Algorithmic Problems for Linear Recurrence Sequences

UNIVERSITÄT
DES
SAARLANDES

A dissertation submitted towards the degree
Doctor of Natural Sciences
of the Faculty of Mathematics and Computer Science
of Saarland University

by

Joris Nieuwveld

Saarbrücken, 2025

Day of Colloquium	September 5th, 2025
Dean of the Faculty	Prof. Dr. Roland Speicher
Chair of the Committee	Prof. Dr. Martina Maggio
Reporters	Prof. Dr. James Maynard Prof. Dr. Kurt Mehlhorn Prof. Dr. Jeffrey Shallit Prof. Dr. Joël Ouaknine
Academic Assistant	Dr. Toghrul Karimov

Abstract

Linear recurrence sequences (LRS) are among the most fundamental and easily definable classes of number sequences, encompassing many classical sequences such as polynomials, powers of two, and the Fibonacci numbers. They also describe the dynamics of iterated linear maps and arise naturally in numerous contexts within computer science, mathematics, and other quantitive sciences. However, despite their simplicity, many easy-to-state decision problems for LRS have stubbornly remained open for decades despite considerable and sustained attention. Chief among these are the *Skolem problem* and the *Positivity problem*, which ask to determine, for a given LRS, whether it contains a zero term and whether it contains only positive terms, respectively. For both problems, decidability is currently open, i.e., whether they are algorithmically solvable.

In this thesis, we present the following results. For the Skolem problem, we introduce an algorithm for simple LRS whose correctness is unconditional but whose termination relies on two classical, widely-believed number-theoretic conjectures. This algorithm is implementable in practice, and we report on experimental results. For the Positivity problem, we introduce the notion of *reversible* LRS, which enables us to carve out a large decidable class of sequences. We also examine various expansions of classical logics by predicates obtained from LRS. In particular, we study expansions of *monadic second-order logic* of the natural numbers with order and present major advances over the seminal results of Büchi, Elgot, and Rabin from the early 1960s. Finally, we investigate fragments of *Presburger arithmetic*, where, among others, we establish the decidability of the existential fragment of Presburger arithmetic expanded with powers of 2 and 3.

Zusammenfassung

Lineare rekursive Folge (LRF) gehören zu den grundlegendsten und am einfachsten zu definierenden Klassen von Zahlenfolgen und umfassen viele klassische Folgen wie Polynome, Zweierpotenzen und die Fibonacci-Zahlen. Sie beschreiben auch die Dynamik iterierter linearer Abbildungen und tauchen in zahlreichen Kontexten der Informatik, Mathematik und anderer quantitativer Wissenschaften auf. Trotz ihrer Einfachheit bleiben jedoch viele leicht zu formulierendes Entscheidungsprobleme für LRF trotz erheblicher und anhaltender Aufmerksamkeit seit Jahrzehnten hartnäckig offen. Dazu gehören vor allem das Skolem-Problem und das Positivitätsproblem, bei denen es darum geht, für eine gegebene LRF zu bestimmen, ob sie einen Nullterm enthält bzw. ob eine LRF nur positive Terme enthält. Für beide Probleme ist die Entscheidbarkeit derzeit offen, d.h., ob sie algorithmisch lösbar sind.

In dieser Doktorarbeit präsentieren wir die folgenden Ergebnisse. Für das Skolem-Problem stellen wir einen Algorithmus für einfache LRF vor, dessen Korrektheit nicht an Bedingungen geknüpft ist, dessen Terminerung aber von zwei weithin akzeptierten zahlentheoretischen Vermutungen abhängt. Dieser Algorithmus ist in der Praxis implementierbar, und wir berichten über experimentelle Ergebnisse. Für das Positivitätsproblem führen wir den Begriff der *umkehrbaren* LRF ein, der es uns ermöglicht, eine große, entscheidbare Klasse von Sequenzen zu bestimmen. Wir untersuchen auch verschiedene Erweiterungen klassischer Logiken durch Prädikate, die aus LRF gewonnen werden. Insbesondere untersuchen wir Erweiterungen von *monadische Prädikatenlogik zweiter Stufe* der natürlichen Zahlen mit Ordnung und präsentieren wichtige Fortschritte gegenüber den bahnbrechenden Ergebnissen von Büchi, Elgot und Rabin aus den frühen 1960er Jahren. Schließlich untersuchen wir Fragmente der *Presburger Arithmetik*, wo wir unter anderem die Entscheidbarkeit des existentiellen Fragments der Presburger Arithmetik, erweitert mit Potenzen von 2 und 3, beweisen.

Acknowledgements

First and foremost, I am deeply grateful to my advisor Joël Ouknine. His scientific ideas guided me toward intriguing problems, exceptional collaborators, and innovative proofs, while his unwavering optimism and encouragement were essential to achieving the results presented in this thesis. Although the Positivity Problem is among his favourite problems, positivity is never one of Joël's open problems.

I also want to thank my senior collaborators and advisers, Valérie Berthé, Florian Luca, and James Worrell, for sharing their knowledge and ideas. Next, I want to thank everyone whom I worked with in Joël's Foundations of Algorithmic Foundations group in Saarbücken, Armand, David, Edon, Erfan, Faraz, Filip, Gorav, Isa, Jean, Mihir, Quintin, Piotr, Sandra, and Toghrul, my office mates in Saarbücken, Diego, Janine, Jimmy, Karam, Kingsley, Liam, Lovita, Melis, Rohan, Thomas, Vasilis, and Yuchen, and all the other wonderful I met at the MPI: Abby, Andrew, Jonathan, Kimaya, Lennard, Vabuk, Yiğit, and Zayd.

I am also grateful to my examiners, Georg Zetzsche (Qualifying and Area Exams), Moses Ganardi (Area Exams), and Holger Herrmans (Qualifying Exam), as well as to my examiners James Maynard, Kurt Mehlhorn, and Jeffrey Shallit for their detailed feedback on this thesis. Furthermore, I am thankful to Wieb Bosma, Bernd Souvingnier, and Wadim Zudilin for teaching me during my undergraduate studies at Radboud University in the Netherlands and helping me to find this exciting field of research.

Lastly, I am grateful to my parents, Peter and Minke, and my brother Niels, who have supported me throughout my life.

Contents

Introduction

A central objective of theoretical computer science is to explore the limits of computing: what is and is not algorithmically solvable (and at what computational cost)? Although the roots of this endeavour go back to the foundational works of Turing and Gödel from the 1930s, such questions remain of critical relevance to this day. One of the grand challenges of modern computer science, as articulated by Tony Hoare some two decades ago, is, broadly speaking, to devise algorithmic methods to help ensuring that computer programs meet their formal specifications [79]. While such problems are computationally unsolvable in general, numerous theoretical and practical open questions remain, particularly concerning simple program loops. As an example, consider the class of *linear loops*, which are program fragments of the form depicted in Figure 1. This loop terminates if and only if there is a natural number n such that $\mathbf{s}^\top M^n \mathbf{t} = 0$. Determining whether such an n exists is a deceptively easy-looking problem equivalent to the *Skolem problem*, whose decidability has been open for some 90 years![1]

Linear loops (and countless variants) are ubiquitous in real-world programs, and their analysis falls within the field of *discrete-time linear dynamical systems*. In its simplest form, a discrete-time linear dynamical system consists of a matrix $M \in \mathbb{Q}^{d\times d}$ together with a vector $\mathbf{t} \in \mathbb{Q}^d$; one then investigates properties of its orbit $(M^n\mathbf{t})_{n=0}^{\infty}$. One of the central decision problems concerning orbits of linear dynamical systems is determining whether the orbit ever reaches a given region $T \subseteq \mathbb{Q}^d$. Classes of regions

[1] Terence Tao has memorably characterized the Skolem problem as *"the halting problem for linear automata"*, and remarked that *"it is faintly outrageous that [it] is still open"* [150].

```
input: M ∈ ℚ^{d×d}, s, t ∈ ℚ^d
x := t
while s · x ≠ 0 do
    x := Mx
```

Figure 1: A linear loop.

for which this problem has been extensively studied include single points, hyperplanes, and half-spaces. The reachability problem for a single point is decidable, as shown by Kannan and Lipton in 1980 [82]. Deciding whether a hyperplane or half-space is reachable is more challenging. For a hyperplane, reachability corresponds to the Skolem problem: for a given $\mathbf{s} \in \mathbb{Q}^d$, is there an $n \in \mathbb{N}$ such that $\mathbf{s}^\top M^n \mathbf{t} = 0$? For a half-space, we obtain the *Positivity problem*: for a given $\mathbf{s} \in \mathbb{Q}^d$, is it the case that, for all $n \in \mathbb{N}$, we have $\mathbf{s}^\top M^n \mathbf{t} \geq 0$?[2]

The Skolem and Positivity problems are intimately connected to various fundamental topics in program analysis and automated verification, such as the termination and model checking of simple while loops [7, 8, 21, 38, 85, 123, 126] or the algorithmic analysis of stochastic systems [2, 3, 15, 48, 124, 125, 155]. They also appear in a variety of other contexts, such as formal power series [133, 147], control theory [34, 68], and even theoretical biology [98] (see [121] for further references). Both the Skolem and Positivity problems are often used as references to establish hardness of other open decision problems; see, e.g., [13, 50]. The Skolem and the Positivity problems often constitute the main bottlenecks for many decision problems related to discrete-time linear dynamical systems: Both can be encoded in many interesting problems. To sharpen one's understanding of these challenges, one observes that $(\mathbf{s}^\top M^n \mathbf{t})_{n=0}^\infty$ forms a linear recurrence sequence using the Cayley-Hamilton theorem.

The Skolem and Positivity problems can then equivalently be formulated in terms of linear recurrence sequences. A *linear recurrence sequence (LRS)* is a sequence of rational numbers $(u_n)_{n=0}^\infty$ that obeys a linear relation of the form

$$\forall n \in \mathbb{N}\colon\ u_{n+d} = c_1 u_{n+d-1} + c_2 u_{n+d-2} + \cdots + c_d u_n\,, \tag{1}$$

where $c_1, \ldots, c_d \in \mathbb{Q}$ are constants. Examples of linear recurrence sequences are geometric progressions $(a \cdot b^n)_{n=0}^\infty$, the Fibonacci numbers $(F_n)_{n=0}^\infty$ (defined by $F_0 = 0$, $F_1 = 1$ and $F_{n+2} = F_{n+1} + F_n$), and values $(P(n))_{n=0}^\infty$ of a fixed univariate polynomial $P \in \mathbb{Z}[X]$ at consecutive integers. In the language of linear recurrence sequences, we reframe the Skolem and the Positivity problems as follows.

Problem 1 (Skolem problem)**.** For a given LRS $(u_n)_{n=0}^\infty$, determine whether $u_n = 0$ for some $n \in \mathbb{N}$.

Problem 2 (Positivity problem)**.** For a given LRS $(u_n)_{n=0}^\infty$, determine whether $u_n \geq 0$ for all $n \in \mathbb{N}$.

[2] In keeping with established terminology, "positivity" is interpreted throughout this thesis in the non-strict sense of "non-negativity".

As noted earlier, the decidability of these problems has been open for many decades, a state of affairs described as a "mathematical embarrassment" by Richard Lipton [100]. We present partial progress on both problems in Chapters 2 and 3, respectively.

It is worth pointing out that in our formulation of the Skolem and Positivity problems as corresponding halting problems for linear loops (as per Figure 1), we are concerned with a *single* initial configuration. Researchers have also investigated termination for *all* initial configurations (i.e., all choices of initial vector $\mathbf{t}$), with positive decidability results [39, 80, 154]. Yet different lines of inquiry concern themselves with *robustness* questions, i.e., whether arbitrarily small perturbations of the initial configuration $\mathbf{t}$ affect reachability [5, 55]. Continuous-time analogues of the Skolem and Positivity problems have also been considered [20, 51]. Note, however, that the inclusion of additional program features such as conditional branching, rounding, or nondeterminism often leads to undecidability. Notable examples include the mortality problem [75] and the use of floating-point rounding [95]. Another related question consists in determining whether, for a given collection of LRS $(u_n^{(1)})_{n=0}^{\infty}, \ldots, (u_n^{(d)})_{n=0}^{\infty}$, there are $n_1, \ldots, n_d \in \mathbb{N}$ such that $\sum_{i=1}^{d} u_{n_i}^{(i)} = 0$; this was however shown to be undecidable by Derksen and Masser [57]. In [73], Guilmant et al. identified a rare instance in which a problem involving multiple update matrices turns out to be decidable.

In Chapters 4 and 5, we then turn our attention to the decidability of expansions of two classical logics by predicates consisting of value sets $P = \{u_n \colon n \in \mathbb{N}\}$ of certain linear recurrence sequences.

We begin with the *monadic second-order logic* (*MSO*) of the natural numbers equipped with order. This logical formalism allows quantification over integer variables (first-order quantification) and sets of integers (second-order quantification). In his 1962 influential paper, Büchi [43] demonstrated the decidability of this theory by uncovering a profound connection between automata and logic. Building upon Büchi's foundation, Elgot and Rabin [62] established the decidability of the MSO theory of $\langle \mathbb{N}; <, P \rangle$ for certain unary predicates P. This method, now known as *Elgot-Rabin contraction*, has been successfully applied to various predicates, such as the sets of powers of two ($2^{\mathbb{N}}$), set of square numbers (Sq), the set of Fibonacci numbers, etc. Accordingly, the MSO theories of $\langle \mathbb{N}; <, 2^{\mathbb{N}} \rangle$ and $\langle \mathbb{N}; <, \mathsf{Sq} \rangle$ are decidable. Elgot and Rabin's approach was nevertheless limited both by the range of predicates that they could consider and by the fact that predicates had to be handled in isolation; in particular, they remained resolutely silent on the decidability of the MSO theories of $\langle \mathbb{N}; <, 2^{\mathbb{N}}, 3^{\mathbb{N}} \rangle$ and $\langle \mathbb{N}; <, 2^{\mathbb{N}}, \mathsf{Sq} \rangle$, for example. In fact, no progress was

recorded on such questions for some 60 years, despite a vast amount of research and a corresponding voluminous literature on the MSO theory of $\langle \mathbb{N}; < \rangle$ and expansions thereof. We present substantial progress on these and related questions in Chapter 4.

Finally, we investigate expansions of *Presburger arithmetic*, the first-order theory of $\langle \mathbb{Z}; 0, 1, <, + \rangle$. More precisely, as in the preceding chapter, we consider expansions of this formalism by value sets of various specific linear recurrence sequences. Presburger arithmetic was originally shown to be decidable by Presburger in 1929 via a quantifier-elimination procedure [128]. Büchi [42] subsequently extended Presburger's result and established the decidability of Presburger arithmetic expanded with the set of powers of 2 (that is, the first-order theory of $\langle \mathbb{Z}; 0, 1, <, +, 2^{\mathbb{N}} \rangle$), or the set of powers of any given base (in isolation). Further work in the ensuing decades demonstrated, among others, that one can alternatively expand Presburger arithmetic with the set of Fibonacci numbers or the powering function $n \mapsto 2^n$ and retain decidability. Nevertheless, formidable complications arise when expanding the base theory by two sets of powers simultaneously (e.g., by powers of 2 and powers of 3). Indeed, Hieronymi and Schulz recently showed that this expansion of Presburger arithmetic is, in fact, undecidable [78], answering what at the time had been an open question for some 40-odd years. They further conjectured that the *existential* fragment of this expansion would be decidable, a result which we establish in Chapter 5.

Structure of the thesis

We now summarize the content of this thesis. This overview also serves as an outline, as the results correspond to the chapters of this thesis.

Skolem problem Chapter 2 is devoted to the Skolem problem, where we achieve progress by employing the onomastically related Skolem conjecture. An LRS $(u_n)_{n=0}^{\infty}$ can be uniquely extended to a bi-infinite sequence $(u_n)_{n=-\infty}^{\infty}$ taking values over the rational numbers. We refer to this bi-infinite sequence as the *bi-completion* of $(u_n)_{n=0}^{\infty}$. The bi-completion has a more global, and thus richer, structure than the one-sided LRS. The *Skolem conjecture* posits that when the LRS is simple (see Section 1.2 for the exact definition), 0 does not occur in its bi-completion if and only if there is an $M \geq 1$ such that $(u_n \bmod M)_{n=-\infty}^{\infty}$ is well-defined and does not contain 0. Thus, it is conjectured that there is a local-global principle for simple LRS.

In greater detail, we explore the notion of *modularity* for an LRS $(u_n)_{n=0}^{\infty}$, which means that there is an M such that the sequence (and not necessarily the bi-infinite sequence) is ultimately non-zero modulo M. Next, we study the Skolem conjecture

and the bi-completion of LRS and construct an algorithm that solves the Skolem problem for simple LRS when assuming this conjecture and a p-adic version of Schanuel's conjecture. This algorithm is of practical significance, as it can resolve many instances of the Skolem problem that are intractable using previous techniques within reasonable computational time. Importantly, this algorithm only relies on the conjectures to guarantee termination. When our algorithm terminates, it provides a certificate of correctness independent of the conjectures. We also describe our implementation of this algorithm in the SKOLEM-tool. The tool also includes a method for solving certain fragments of the Skolem problem using Baker's theorem on linear forms in logarithms.

The results presented in this chapter are based on the works [31, 99] and the tool [11].

Positivity problem In Chapter 3, we study the Positivity problem.

The chapter begins with a review of the current state of the art regarding the Positivity problem. We present the current techniques to decide fragments of the Positivity problem and identify the key obstacles that hinder further progress. We use these insights to study a restricted class of LRS: the *reversible linear recurrence sequences*. These are precisely the LRS whose bi-completion is integer-valued. This defining property imposes strong structural constraints on the sequence, which we unearth using Galois-theoretic methods. Therefore, we can push the borders of the Positivity problem for reversible LRS far beyond the state-of-the-art for general LRS. We also show that when a reversible LRS is sufficiently complicated, we encounter the same barriers for general LRS and give explicit examples of LRS for which we cannot decide positivity.

In the second part of the chapter, we discuss the Positivity problem for real algebraic LRS (LRS taking value in the ring $\mathbb{R} \cap \overline{\mathbb{Q}}$). It is folklore that the Skolem problem for integer LRS is Turing-equivalent to the Skolem problem for algebraic LRS, but the analogous question for the Positivity problem was open. We establish that the Positivity problem for $(\mathbb{R} \cap \overline{\mathbb{Q}})$-LRS is Turing-equivalent to the Positivity problem for $\mathbb{Z}$-LRS by showing that $(\mathbb{R} \cap \overline{\mathbb{Q}})$-LRS can be approximated closely by $\mathbb{Q}$-LRS.

The results presented in this chapter are based on the works [84, 89, 99].

Monadic second-order logic In Chapter 4, we discuss the monadic second-order theory (MSO) of the natural numbers with order, expanded by unary predicates $P_i = \{u_n^{(i)} : n \in \mathbb{N}\}$, where each $(u_n^{(i)})_{n=0}^{\infty}$ is an LRS. Thus we consider $\mathsf{MSO}_{\mathbb{N};<}(P_1, \ldots, P_d)$,

i.e., the MSO theory of $\langle \mathbb{N}; <, P_1, \ldots, P_d \rangle$. We open the chapter by establishing a version of the Elgot-Rabin contraction method and apply it to three different variants.

The Elgot-Rabin contraction method translates predicates $P_1, \ldots, P_d$ and an MSO formula φ into an infinite word β_φ and an ω-automaton $\mathcal{A}_\varphi$ such that $\mathcal{A}_\varphi$ accepts β_φ if and only if φ is satisfied. The central objective is to ensure that β_φ belongs to a class of infinite words for which the acceptance problem for ω-automata is decidable. For the predicates that Elgot and Rabin studied, this word β_φ is ultimately periodic, guaranteeing decidability.

Our first variant concerns multiple simple LRS $(u_n^{(i)})_{n=0}^{\infty}$ with a single dominant root (see Section 1.2 for definitions) such as the set $a^{\mathbb{N}}$ of powers of a natural number a. As a snapshot of our results, we prove that $\mathsf{MSO}_{\mathbb{N};<}(2^{\mathbb{N}}, 3^{\mathbb{N}})$ and $\mathsf{MSO}_{\mathbb{N};<}(2^{\mathbb{N}}, 3^{\mathbb{N}}, 6^{\mathbb{N}})$ are decidable and that $\mathsf{MSO}_{\mathbb{N};<}(2^{\mathbb{N}}, 3^{\mathbb{N}}, 5^{\mathbb{N}})$ is decidable when assuming Schanuel's conjecture. In these cases, the word β_φ obtained from the Elgot-Rabin contraction method is *toric*, a notion recently introduced by Berthé et al. [25], as β_φ can be generated by a rotation on a torus. For $\mathsf{MSO}_{\mathbb{N};<}(2^{\mathbb{N}}, 3^{\mathbb{N}})$ and $\mathsf{MSO}_{\mathbb{N};<}(2^{\mathbb{N}}, 3^{\mathbb{N}}, 6^{\mathbb{N}})$, this structure proves that β_φ is *effectively almost-periodic*, establishing the decidability of the acceptance problem for ω-automata. For $\mathsf{MSO}_{\mathbb{N};<}(2^{\mathbb{N}}, 3^{\mathbb{N}}, 5^{\mathbb{N}})$, decidability depends on the linear independence of $\frac{1}{\log(2)}$, $\frac{1}{\log(3)}$, and $\frac{1}{\log(5)}$ over the rationals, which Schanuel's conjecture implies.

Our second variant combines arithmetic progressions and polynomials that we connect to the base expansions of real algebraic numbers. Consider the expansion generated by the sets of powers of two and the squares, $\mathsf{MSO}_{\mathbb{N};<}(2^{\mathbb{N}}, \mathsf{Sq})$. We show that the decidability of this theory is Turing-equivalent to determining whether a given ω-automaton accepts the binary expansion of $\sqrt{2} = 1.01101010000\cdots$. While this equivalence does not settle decidability, it highlights the core difficulty. It is conjectured that $\sqrt{2}$ is a normal number, which implies that, in the binary expansion of $\sqrt{2}$, finite binary strings should occur with a uniform frequency. However, since ω-automata can detect whether a subword occurs finitely or infinitely often but not compute their frequency, we introduce the weaker notion of *disjunctive words*, where every finite subword occurs infinitely often. Because we can prove that the acceptance problem for ω-automata of a disjunctive word is decidable, the theory $\mathsf{MSO}_{\mathbb{N};<}(2^{\mathbb{N}}, \mathsf{Sq})$ is decidable when assuming that the binary expansion of $\sqrt{2}$ is disjunctive.

Our third variant concerns theories $\mathsf{MSO}_{\mathbb{N};<}(P)$, where $P = \{u_n \colon n \in \mathbb{N}\} \cap \mathbb{N}$ is the value set of a simple LRS $(u_n)_{n=0}^{\infty}$ with two dominant roots. For example, consider $u_n = (1+2i)^n + (1-2i)^n$, which is neither monotonic nor has only positive values. To understand the value set P of such an LRS, we have to sort P. Let $(p_m)_{m=0}^{\infty}$

enumerate P such that p_m is the $(m+1)$th smallest number in P. Our version of the Elgot-Rabin contraction method reduces the decidability of $\mathsf{MSO}_{\mathbb{N};<}(P)$ to whether for all $M \geq 1$, the acceptance problem for ω-automata of $(p_m \bmod M)_{m=0}^{\infty}$ is decidable. The main challenge lies in showing that $(p_m)_{m=0}^{\infty}$ is *effective prodisjunctive*, which is another concept we introduce. Informally, this means that for all $M \geq 1$, the sequence $(p_m \bmod M)_{m=0}^{\infty}$ behaves like a disjunctive word, and we know that the acceptance problem for ω-automata is decidable for disjunctive words. These sequences also produce explicit examples of disjunctive words, which are of independent interest.

The results presented in this chapter are based on the works [24, 119].

Presburger arithmetic Finally, in Chapter 5, we study Presburger arithmetic expanded with two sets of powers. Thus, in particular, we investigate the first-order theory of $\langle \mathbb{Z}; 0, 1, <, 2^{\mathbb{N}}, 3^{\mathbb{N}} \rangle$, denoted $\mathcal{PA}(2^{\mathbb{N}}, 3^{\mathbb{N}})$. Recent work by Hieronymi and Schulz [78] showed that this theory, and its $\forall\exists\forall\exists$-fragment in particular, is undecidable. Our contributions are threefold.

The chapter's main result states that the existential fragment of $\mathcal{PA}(2^{\mathbb{N}}, 3^{\mathbb{N}})$ is decidable. This proof has two technical steps. In the first step, we solve exponential Diophantine equations of the form $\sum_{i=1}^{\ell} c_i z_i^{n_i} = d$, where $c_i, d \in \mathbb{Z}$ are given and each base $z_i \in \{2, 3\}$ is fixed. Owing to the structured nature of the solution sets of such equations, we reduce the decidability of the existential fragment of $\mathcal{PA}(2^{\mathbb{N}}, 3^{\mathbb{N}})$ to determining the solvability of systems of homogeneous inequalities of the form $A\mathbf{z} > 0$, where $\mathbf{z} = (z_1^{n_1}, \ldots, z_\ell^{n_\ell})$. By applying classical results from Diophantine approximation, we can effectively compute whether such a system has a solution.

Our second contribution refines the argument of Hieronymi and Schulz. By reducing from counter machines instead of general Turing machines, we obtain a simpler argument that shows that the $\exists\forall\exists$-fragment of $\mathcal{PA}(2^{\mathbb{N}}, 3^{\mathbb{N}})$ is already undecidable. Hence, we sharpen the upper bound for decidability by one block of quantifiers.

For our final contribution, we study the existential fragment of $\mathcal{PA}(n \mapsto 2^n, n \mapsto 3^n)$, where we consider functions $n \mapsto a^n$ rather than the less expressive unary predicates $a^{\mathbb{N}}$. While we cannot resolve the decidability of this fragment, we encode a decision problem concerning the base-2 expansion of $\log_2(3)$. This encoding highlights a central obstacle: any proof of decidability would require a breakthrough in understanding the distribution of digits in such irrational constants.

The results presented in this chapter are based on [86].

Chapter 1

Preliminaries

Let $\mathbb{N} = \{0, 1, \dots\}$ denote the set of natural numbers, $\mathbb{Z}$ the set of integers, $\mathbb{Q}$ the set of rational numbers, $\mathbb{R}$ the set of real numbers, and $\mathbb{C}$ the set of complex numbers.

We use boldface letters to denote (column) vectors and occasionally write d-dimensional column vectors in the form $\mathbf{x} = (x_1, \dots, x_d)$. In particular, $\mathbf{0}$ denotes a vector containing solely zeros whose dimension should be clear from the context. For vectors $\mathbf{x} = (x_1, \dots, x_d)$ and $\mathbf{y} = (y_1, \dots, y_d)$ and a relation $\sim$, we write $\mathbf{x} \sim \mathbf{y}$ as shorthand for $x_i \sim y_i$ for all $1 \leq i \leq d$. For a ring R, an *R-linear form* is a function of the form $h(x_1, \dots, x_\ell) := c_1 x_1 + \cdots + c_\ell x_\ell$, where $c_1, \dots, c_\ell \in R$.

1.1 Number theory

A number $\alpha \in \mathbb{C}$ is *algebraic* if there is a non-zero polynomial $P \in \mathbb{Q}[X]$ such that $P(\alpha) = 0$. For an algebraic number α, there is a unique non-zero polynomial $P \in \mathbb{Z}[X]$ of least degree with coprime coefficients and a positive leading coefficient, the *minimal polynomial* of α. A minimal polynomial is always irreducible over the rationals, and the *degree* of α is the degree of its minimal polynomial. If the minimal polynomial of α is monic, meaning its leading coefficient is 1, then α is an *algebraic integer*. An algebraic number is a *unit* if $1/\alpha$ is also an algebraic integer, or equivalently when α has a monic minimal polynomial with constant coefficient ± 1. The *Galois conjugates* of α are the roots of the minimal polynomial of α. The set of algebraic numbers forms a field, denoted by $\overline{\mathbb{Q}}$, and the set of algebraic integers forms a ring.

A number ζ is a *root of unity* if $\zeta^d = 1$ for some natural number $d \geq 1$. The *order* of a root of unity ζ is the smallest natural number $d \geq 1$ such that $\zeta^d = 1$, and ζ is a primitive dth root of unity, denoted ζ_d if ζ is a root of unity of order d. For example, the imaginary unit i is a primitive root 4th root of unity. We recall an old result of Kronecker [91].

Theorem 1.1.1. *Let $P \in \mathbb{Z}[X]$ be a monic polynomial such that $P(0) \neq 0$ and all roots of P have absolute value at most 1. Then all the roots of P are roots of unity.*

A *canonical representation* of an algebraic number α consists of its minimal polynomial P and sufficiently accurate rational approximations of its real and imaginary parts to distinguish it from the other roots of P. All arithmetic operations can be performed effectively on canonical representations of algebraic numbers [53, Section 4.2].

We denote the natural logarithm by log. Thus, $\log(e) = 1$. To avoid an overload of brackets, we often write $\log|a|$ instead of $\log(|a|)$.

1.1.1 Number fields and Galois theory

A *number field* K is a field extension of $\mathbb{Q}$ such that K is a finite-dimensional vector space over the rationals. This finite dimension of K as a $\mathbb{Q}$-vector space is called the *degree* of K. If $\alpha_1, \ldots, \alpha_d$ are algebraic numbers, we can effectively compute the degree of $\mathbb{Q}(\alpha_1, \ldots, \alpha_d)$.

If K is a field and $P \in K[X]$, then the *splitting field* of P is the smallest field that contains K and all roots of P. When $K \subseteq L$ are fields, then $L : K$ denotes the *field extension* L of K, and $[L : K]$ denotes the *degree* of $L : K$ (the dimension of L of as a K-vector space). The set of algebraic integers in K forms a ring, the *ring of integers of* K, denoted $\mathcal{O}_K$.

A polynomial $P \in K[X]$ is *separable* if its roots are distinct in an algebraic closure of K. If K has characteristic zero, equivalently, the polynomial P has no double roots. A field extension $L : K$ is *Galois* if L is the splitting field of a separable polynomial $P \in K[X]$. A *Galois automorphism* of $L : K$ is a field automorphism (a map that is compatible with both addition and multiplication) that fixes K. The *Galois group* $\mathrm{Gal}_K(L)$ is the group of Galois automorphisms of the extension $L : K$. The fundamental theorem of Galois theory states that for a Galois extension $L : K$, the order of $\mathrm{Gal}_K(L)$ equals $[L : K]$.

A finite group G is said to act *transitively* on a finite set X if for each pair $x, y \in X$ there is a $g \in G$ such that $g(x) = y$ and *faithfully* if $g(x) = x$ for all $x \in X$ implies that g is the unit of G. The *stabilizer* G_x of an element x in X is defined as the set $\{g \in G \colon gx = x\}$. The Orbit-Stabilizer theorem (see, for example, [132, Theorem 3.19]) implies that if G acts transitively on X, the cardinality of G_x is the same for each $x \in X$. Further, when G acts transitively and faithfully on X, the number $\#\{g \in G \colon gx = y\}$ does not depend on the choice of $x, y \in G$. If K is a field

and $P \in K[X]$ is irreducible and separable with splitting field L, then $\mathrm{Gal}_K(L)$ acts transitively and faithfully on the set of roots of P.

Example 1.1.2. Let $K = \mathbb{Q}$ and L be the splitting field of $X^3 - 2$. Then L is the number field $\mathbb{Q}(\sqrt[3]{2}, \zeta_3\sqrt[3]{2}, \zeta_3^2\sqrt[3]{2})$, which is also a Galois extension of $\mathbb{Q}$. Then $[L : \mathbb{Q}] = 6$ and the Galois group $\mathrm{Gal}_{\mathbb{Q}}(L)$ is generated by two elements: complex conjugation and the unique automorphism defined by $\sqrt[3]{2} \mapsto \zeta_3\sqrt[3]{2}$ and $\zeta_3 \mapsto \zeta_3$. Thus, $\mathrm{Gal}_{\mathbb{Q}}(L)$ is isomorphic to S_3, the symmetric group acting on three elements. □

The p-adic numbers

Let $x \in \mathbb{Z}$ and $p \in \mathbb{N}$ be a prime number (we sometimes refer to prime numbers as *rational primes*). Then the *p-adic valuation* of x, denoted $\nu_p(x)$, is the largest integer n such that p^n divides x, whereas p^{n+1} does not. By convention, $\nu_p(0) = +\infty$. Thus, $\nu_3(12) = 1$ and $\nu_2(12) = 2$. For integers x and y and a prime p, we have that $\nu_p(xy) = \nu_p(x) + \nu_p(y)$ and $\nu_p(x+y) \geq \min\{\nu_p(x), \nu_p(y)\}$. One can extend this notion of the p-adic valuation to number fields.

Let K be a Galois extension of $\mathbb{Q}$ with a ring of integers $\mathcal{O}_K$. Then each non-zero ideal of $\mathcal{O}_K$ admits a unique factorization as a product of non-zero prime ideals. More specifically, fix a rational prime p. By the Dedekind-Kummer theorem, the ideal (p) factors into prime ideals in $\mathcal{O}_K$ as $\mathfrak{p}_1^e \cdots \mathfrak{p}_f^e$, where $\mathfrak{p}_1, \ldots, \mathfrak{p}_f$ (the prime ideals *above* p) are distinct. The number e is called the *ramification index* of $\mathfrak{p}_i$ over p. For $x \in \mathcal{O}_K$, the $\mathfrak{p}$-valuation $\nu_\mathfrak{p} : \mathcal{O}_K \to \mathbb{N} \cup \{\infty\}$ is defined as the number of times $\mathfrak{p}$ appears in the factorisation of (x), the principal ideal generated by x. This definition can be extended from $\mathcal{O}_K$ to K by writing $\alpha \in K$ as x/n for some $x \in \mathcal{O}_K$ and $n \in \mathbb{N}$. Then, setting $\nu_\mathfrak{p}(\alpha) = \nu_\mathfrak{p}(x) - \nu_\mathfrak{p}(n)$ gives a well-defined definition of the $\mathfrak{p}$-adic valuation. Moreover, we have again that $\nu_\mathfrak{p}(x_1 x_2) = \nu_\mathfrak{p}(x_1) + \nu_\mathfrak{p}(x_2)$ and $\nu_\mathfrak{p}(x_1 + x_2) \geq \min(\nu_\mathfrak{p}(x_1), \nu_\mathfrak{p}(x_2))$ for all $x_1, x_2 \in K$.

One defines an absolute value $|\cdot|_\mathfrak{p}$ (the $\mathfrak{p}$-*norm*) on K by $|x|_\mathfrak{p} = p^{-\nu_\mathfrak{p}(x)/e}$. Let $K_\mathfrak{p}$ be the completion of K with respect to this absolute value and $\mathcal{O}_{K_\mathfrak{p}}$ denote the ring of integers of $K_\mathfrak{p}$. Further, let $\exp : \mathfrak{p}^k\mathcal{O}_{K_\mathfrak{p}} \to 1 + \mathfrak{p}^k\mathcal{O}_{K_\mathfrak{p}}$ and $\log : 1 + \mathfrak{p}^{k'}\mathcal{O}_{K_\mathfrak{p}} \to \mathfrak{p}^k\mathcal{O}_{K_\mathfrak{p}}$ denote the $\mathfrak{p}$-adic exponential function and $\mathfrak{p}$-adic logarithm, respectively, for some numbers k, k' for the functions to be well-defined. These two functions are defined by their usual power series:

$$\exp(x) = \sum_{n=0}^{\infty} \frac{x^n}{n!} \quad \text{and} \quad \log(x+1) = \sum_{n=1}^{\infty} (-1)^{n+1} \frac{x^n}{n}\,.$$

The $\mathfrak{p}$-adic exponential function and logarithm are inverses of each other, and for all x, y where the functions are defined, we have $\exp(x+y) = \exp(x)\exp(y)$ and $\log(xy) = \log(x) + \log(y)$.

When $K = \mathbb{Q}$ and $\mathfrak{p} = (p)$, let $\mathbb{Q}_p$ and $\mathbb{Z}_p$ denote $K_\mathfrak{p}$ and $\mathcal{O}_{K_\mathfrak{p}}$, respectively. Moreover, we often write p-adic instead of $\mathfrak{p}$-adic when discussing the valuation, exponential function, or logarithm.

Multiplicative relationships

By a *multiplicative relation* of $\alpha_1, \ldots, \alpha_d \in \mathbb{C}^*$, we mean $\mathbf{k} = (k_1, \ldots, k_d) \in \mathbb{Z}^d$ such that $\alpha_1^{k_1} \cdots \alpha_d^{k_d} = 1$. Let $G := G_M(\alpha_1, \ldots, \alpha_d)$ be the set of all multiplicative relations of $(\alpha_1, \ldots, \alpha_d)$. Then G forms a free abelian group. If $\mathbf{0}$ is the sole element in G, then $\alpha_1, \ldots, \alpha_d$ are called *multiplicatively independent*. The *rank* of G is the cardinality of the largest multiplicatively independent subset of X. If $\operatorname{rank}(G) = m$, then G has a *basis* $B = \{\mathbf{v}_1, \ldots, \mathbf{v}_m\} \subseteq \mathbb{Z}^d$ that is linearly independent over $\mathbb{Q}$ with the property that every $\mathbf{z} \in G$ can be written as an integer linear combination of $\mathbf{v}_1, \ldots, \mathbf{v}_m$. For non-zero algebraic $\alpha_1, \ldots, \alpha_d$, we can compute a basis of G using a deep result of Masser [112].

Theorem 1.1.3 ([112]). *Given $\alpha_1, \ldots, \alpha_d \in \overline{\mathbb{Q}}^*$, one can compute a basis for the group $G_M(\alpha_1, \ldots, \alpha_d)$.*

Similarly, we write $G_A(\alpha_1, \ldots, \alpha_d)$ for the group of additive relations of $\alpha_1, \ldots, \alpha_d$ over the integers. That is,

$$G_A(\alpha_1, \ldots, \alpha_d) = \left\{(k_1, \ldots, k_d) \in \mathbb{Z}^d \colon k_1\alpha_1 + \cdots + k_d\alpha_d = 0\right\}.$$

1.1.2 Transcendence and Baker's theorem on linear forms in logarithms

A *transcendental number* is a complex number that is not algebraic. A finite set $S = \{\alpha_1, \ldots, \alpha_d\}$ of complex numbers is *algebraically independent* over a field R if for all non-zero polynomials $P \in R[X_1, \ldots, X_d]$, we have $P(\alpha_1, \ldots, \alpha_d) \neq 0$. The *transcendence degree* of S is the cardinality of the largest subset of S that is algebraically independent over $\mathbb{Q}$. Below, we state Schanuel's conjecture, a classical conjecture in transcendental number theory with far-reaching implications [92].

Conjecture 1.1.4 (Schanuel's conjecture). *If $\alpha_1, \ldots, \alpha_d \in \mathbb{C}$ are linearly independent over $\mathbb{Q}$, the transcendence degree of $\{\alpha_1, \ldots, \alpha_d, \exp(\alpha_1), \ldots, \exp(\alpha_d)\}$ is at least d.*

Many versions of Schanuel's conjecture exist, including the following weaker version [46, Conjecture 3.9].

Conjecture 1.1.5 (Weak Schanuel conjecture.). *If $\alpha_1, \ldots, \alpha_d \in \overline{\mathbb{Q}}^*$ are multiplicatively independent, then $\log(\alpha_1), \ldots, \log(\alpha_d)$ are algebraically independent over $\mathbb{Q}$.*

Similarly to complex numbers, we can define the notions of transcendence, algebraic independence, and transcendence degree for p-adic numbers and the algebraic closure of the p-adic numbers. Then, Calegari and Mazur [46, Conjecture 3.10] formulated the following p-adic version of the weak Schanuel conjecture.

Conjecture 1.1.6 (Weak p-adic Schanuel conjecture.). *If $\alpha_1, \ldots, \alpha_d \in \overline{\mathbb{Q}}^*$ are in some finite extension of $\mathbb{Q}_p$, the p-adic logarithm $\log(\alpha_i)$ is well-defined for $1 \leq i \leq d$, and $\alpha_1, \ldots, \alpha_d$ are linearly independent over $\mathbb{Q}$. Then $\{\log(\alpha_1), \ldots, \log(\alpha_d)\}$ has transcendence degree d over $\mathbb{Q}$.*

For simplicity, we will often refer to Conjecture 1.1.6 as the p-adic Schanuel conjecture.

Baker's theorem on linear forms in logarithms (which we colloquially call Baker's theorem) is a crucial tool for many results in this thesis, and we will often rely on an effective version of Baker's theorem due to Matveev. Baker's work encompassed many breakthroughs in transcendental number theory, for which he received a Fields medal in 1970 [113]. We start by giving a non-effective version of Baker's theorem.

Theorem 1.1.7 (Theorem 1.6 in [158]). *If $\alpha_1, \ldots, \alpha_d \in \overline{\mathbb{Q}}^*$ are multiplicatively independent, the numbers $1, \log(\alpha_1), \ldots, \log(\alpha_d)$ are linearly independent over $\overline{\mathbb{Q}}$.*

An effective version of Baker's theorem also gives a lower bound on how fast such linear combinations approach zero (in terms of their coefficients).

Lemma 1.1.8. *Given $\alpha_1, \ldots, \alpha_d \in \overline{\mathbb{Q}}$ and $a_0, \ldots, a_d \in \mathbb{R} \cap \overline{\mathbb{Q}}$, we can effectively determine the sign of $a_0 + \sum_{i=1}^{d} a_i \log(\alpha_i)$.*

Proof. After computing the multiplicative relationships among the algebraic numbers α_i using Theorem 1.1.3, we can rewrite this expression as $b_0 + \sum_{i=1}^{e} b_i \log(\alpha_i)$. Here, we relabelled the α_i such that $\alpha_1, \ldots, \alpha_e$ is a maximum multiplicatively independent subset of $\alpha_1, \ldots, \alpha_d$, and $b_0, \ldots, b_e \in \mathbb{R} \cap \mathbb{Q}$ are explicitly computed. By Theorem 1.1.7, this expression is zero if and only if all b_i are 0. If this expression is non-zero, we can compute it to arbitrary precision and determine whether it is positive. □

A p-adic version of Theorem 1.1.7 is due to Brumer [41].

Theorem 1.1.9 (Brumer)**.** *Let p be a rational prime, K_p a finite extension of $\mathbb{Q}_p$, and $\alpha_1, \ldots, \alpha_d$ algebraic units in K_p whose p-adic logarithms are defined and linearly independent over $\mathbb{Q}$, then $1, \log(\alpha_1), \ldots, \log(\alpha_d)$ are linearly independent over $\overline{\mathbb{Q}}$.*

For $\mathfrak{p}$-adic number fields, we have the following generalization. See for example [159].

Theorem 1.1.10. *Let K be a number field whose ring of integers has a prime ideal $\mathfrak{p}$, $K_\mathfrak{p}$ the completion of K with respect to the absolute value induced by $\mathfrak{p}$, and $\alpha_1, \ldots, \alpha_d \in \mathcal{O}_K$ not in $\mathfrak{p}$ whose $\mathfrak{p}$-adic logarithms are defined and linearly independent over $\mathbb{Q}$, then $1, \log(\alpha_1), \ldots, \log(\alpha_d)$ are linearly independent over $\overline{\mathbb{Q}}$.*

Before we can give Matveev's effective version of Baker's theorem, we first introduce the notion of the height of an algebraic number. Let $\alpha \in \overline{\mathbb{Q}}^*$ have degree d and minimal polynomial $a_0 \prod_{i=1}^{d}(X - \alpha_i)$. The *absolute logarithmic Weil height* of α is defined as

$$h(\alpha) = \frac{1}{d}\left(\log|a_0| + \sum_{i=1}^{d} \log\big(\max(|\alpha_i|, 1)\big)\right). \tag{1.1}$$

Furthermore, set $h(0) = 0$. For all algebraic numbers $\alpha_1, \ldots, \alpha_k$ and $n \in \mathbb{Z}$, we have the following properties:

$$\begin{aligned} h(\alpha_1 + \cdots + \alpha_k) &\leq \log(k) + h(\alpha_1) + \cdots + h(\alpha_k)\,, \\ h(\alpha_1\alpha_2) &\leq h(\alpha_1) + h(\alpha_2)\,, \quad \text{and} \\ h(\alpha_1^n) &= |n| h(\alpha_1)\,. \end{aligned}$$

In particular, for $n \in \mathbb{Z} \setminus \{0\}$, we have that $h(n) = \log|n|$. For a rational number p/q with $\gcd(p, q)$, we have that $h(p/q) = \max(\log|p|, \log|q|)$. For simplicity, we call $h(\alpha)$ the *height* of α.

Let $\alpha_1, \ldots, \alpha_k \in \overline{\mathbb{Q}}^*$ lie in a number field K of degree D, $b_1, \ldots, b_n \in \mathbb{Z}$, and

$$\Lambda = \alpha_1^{b_1}\alpha_2^{b_2}\cdots\alpha_k^{b_k} - 1\,.$$

Further, set κ to 1 if $K \subset \mathbb{R}$ and to 2 otherwise, and let

$$B = \max\big(|b_1|, \ldots, |b_k|\big)\,.$$

Lastly, for $1 \leq i \leq k$, let $h'(\alpha_i) \geq \max(Dh(\alpha_i), \log|\alpha_i|, 0.16)$. Then Matveev showed the following.

Theorem 1.1.11 ([115]). *Using the notation above, if $\Lambda \neq 0$, then*

$$\log|\Lambda| > -C(k,\kappa,D)h'(\alpha_1)\cdots h'(\alpha_k)\big(1+\log(B)\big)\,,$$

where

$$C(k,\kappa,D) = D^2\big(1+\log(D)\big)\min\left(\frac{1}{\kappa}\left(\frac{1}{2}ek\right)^{\kappa}30^{k+3}k^{3.5},\,2^{6k+20}\right)\,.$$

Thus, the constant C only depends on k, κ, and D while $h'(\alpha_i)$ only depends on D and α_i. Thus, C and all $h'(\alpha_i)$ only depend on $\alpha_1,\ldots,\alpha_k$. When $B \geq 2$, we have $1+\log(B) < 3\log(B)$, which gives us the following corollary.

Corollary 1.1.12. *There is a computable constant C only depending on $\alpha_1,\ldots,\alpha_k$ such that $|\Lambda| > B^{-C}$ whenever $\Lambda \neq 0$.*

Kronecker's theorem on Diophantine approximation

We need two consequences of Kronecker's theorem in Diophantine approximation [71].

Let $\mathbb{T}$ denote the additive group $\mathbb{R}/\mathbb{Z}$, which we can identify with the interval $[0,1)$. For $x \in \mathbb{R}$, let $\{x\}$ denote the fractional part of x, i.e., $\{x\} = x - \lfloor x \rfloor$. Then we study the behaviour of rotating a point on the higher-dimensional torus $\mathbb{T}^d$.

Theorem 1.1.13. *Let $\boldsymbol{\delta} = (\delta_1,\ldots,\delta_d) \in \mathbb{R}^d$ such that $1,\delta_1,\ldots,\delta_d$ are multiplicatively independent, define $g : \mathbb{T}^d \to \mathbb{T}^d$ by $g((x_1,\ldots,x_d)) = (\{x_1+\delta_1\},\ldots,\{x_d+\delta_d\})$, and let $g^{(n)}$ be the nth iterate of g. Then the orbit $\big(g^{(n)}((0,\ldots,0))\big)_{n=0}^{\infty}$ is dense in $\mathbb{T}^d$. Thus, for every non-empty open subset O of $\mathbb{T}^d$ there exist infinitely many $n \in \mathbb{N}$ such that $g^{(n)}((0,\ldots,0)) \in O$.*

Our second application only concerns a 1-dimensional torus.

Theorem 1.1.14. *Let $\alpha,\beta \in \mathbb{N}_{>1}$ be multiplicatively independent and $I \subseteq \mathbb{R}_{>0}$ be a non-empty open interval. Then there exist infinitely many $n_1,n_2 \in \mathbb{N}$ such that $\alpha^{n_1}/\beta^{n_2} \in I$.*

Proof. By multiplicative independence, the number $\log_\beta(\alpha)$ is irrational, and by Theorem 1.1.13, the sequence $(\{n\log_\beta(\alpha)\})_{n=0}^{\infty}$ is dense in $[0,1)$. That is,

$$\{n_1\log_\beta(\alpha) - n_2 \colon n_1,n_2 \in \mathbb{N}\} \cap (0,1)$$

is dense in $(0,1)$. Equivalently, $\{\alpha^{n_1}/\beta^{n_2} \colon n_1,n_2 \in \mathbb{N}\} \cap (1,\beta)$ is dense in $(1,\beta)$. Thus, for all $k \in \mathbb{Z}$, we have that $\{\alpha^{n_1}/\beta^{n_2-k} \colon n_1,n_2 \in \mathbb{N}\}$ is dense in (β^k,β^{k+1}) and so $\{\alpha^{n_1}/\beta^{n_2} \colon n_1,n_2 \in \mathbb{N}\}$ is dense in $(0,\infty)$. □

1.2 Linear recurrence sequences

Linear recurrence sequences are the primary object of study in this thesis. In this section, we give their fundamental properties and prove several preliminary lemmas. Much of this section is based on Section 1.1 in [64].

A *linear recurrence sequence* over a ring R (an *R-LRS* for short) is an R-valued sequence $(u_n)_{n=0}^{\infty}$ for which there are $c_1, \ldots, c_d$ such that

$$\forall n \in \mathbb{N}\colon u_{n+d} = c_1 u_{n+d-1} + c_2 u_{n+d-2} + \cdots + c_d u_n \,, \tag{1}$$

In other words, an LRS $(u_n)_{n=0}^{\infty}$ has to satisfy a *linear recurrence*, and a tuple $(u_0, \ldots, u_{d-1}, c_1, \ldots, c_d) \in R^{2d}$ uniquely defines an LRS $(u_n)_{n=0}^{\infty}$. Here, $u_0, \ldots, u_{d-1}$ are called the *initial terms* of the sequence. In (1), we refer to d as the *order* of the linear recurrence. Some other works, use the terminology depth instead of order.

Applying linear algebra, one can show that every LRS satisfies a unique linear recurrence of minimal order. This recurrence is the *linear recurrence* of the LRS, and we call its order the *order* of the LRS. The *zero-LRS* is the unique R-LRS of order 0 and is constantly 0.

We will show that we can assume that c_d is non-zero (where c_d is as in (1)). Indeed, if $c_d = 0$, then $(u_n)_{n=1}^{\infty}$ is an LRS of order $d-1$ as it satisfies the linear recurrence

$$\forall n \in \mathbb{N}_{\geq 1}\colon u_{n+d-1} = c_1 u_{n+d-2} - c_2 u_{n+d-3} - \cdots - c_{d-1} u_n$$

of order $d-1$. Hence, any linear recurrence sequence can be reduced to a concatenation of a finite sequence and a linear recurrence sequence whose linear recurrence satisfies $c_d \neq 0$.

In this thesis, we will primarily consider rings R that are integral domains containing the integers, like number fields, the complex numbers, and the p-adic numbers. However, we need one result for rings $\mathbb{Z}/M\mathbb{Z}$ for integers $M \geq 1$. A sequence $(u_n)_{n=0}^{\infty}$ is *ultimately periodic* with *period* P and *preperiod* N if $u_n = u_{n+P}$ for all $n \geq N$.

Lemma 1.2.1. *Let $M \geq 1$ be a natural number, $R = \mathbb{Z}/M\mathbb{Z}$, and $(u_n)_{n=0}^{\infty}$ an R-LRS satisfying (1) for some given $(u_0, \ldots, u_{d-1}, c_1, \ldots, c_d)$. Then $(u_n)_{n=0}^{\infty}$ is ultimately periodic, and both the preperiod and period can be effectively computed.*

Proof. It is sufficient to note that for each $n \in \mathbb{N}$, one can compute u_n, and each $(u_n, \ldots, u_{n+d-1})$ is in the finite set $(\mathbb{Z}/M\mathbb{Z})^d$. Hence, we can enumerate $(u_n)_{n=0}^{\infty}$ until we find $N \geq 0$ and $P \geq 1$ such that $(u_N, \ldots, u_{N+d-1}) = (u_{N+P}, \ldots, u_{N+P+d-1})$. Then, inductively, it follows that $u_{N+n} = u_{N+n+P}$ for all $n \in \mathbb{N}$, and so $(u_n)_{n=0}^{\infty}$ is periodic with period P and preperiod N. □

In the remainder of this section, we will restrict ourselves to rings R that are integral domains of characteristic 0 that contain 1. Let K be the fraction field of R.

The *characteristic polynomial* $P \in R[X]$ of the R-LRS $(u_n)_{n=0}^{\infty}$ is given by

$$P(X) := X^d - c_1 X^{d-1} - \cdots - c_{d-1} X - c_d \,,$$

where $c_1, \ldots, c_d \in R$ are the coefficients of the linear recurrence (1) of $(u_n)_{n=0}^{\infty}$. The *characteristic roots* of $(u_n)_{n=0}^{\infty}$ are the roots of the characteristic polynomial P in a suitable, finite-dimensional extension L of K. The *multiplicity* of a characteristic root λ is the multiplicity of λ as a root of the characteristic polynomial P. Thus, if $(u_n)_{n=0}^{\infty}$ has characteristic roots $\lambda_1, \ldots, \lambda_k$ with respective multiplicities $m_1, \ldots, m_k$, then its characteristic polynomial is equal to

$$P(X) = \prod_{i=1}^{k} (X - \lambda_i)^{m_i} \,.$$

Then, $d = \sum_{i=1}^{k} m_i$. As $c_d \neq 0$, we have that 0 is never a characteristic root.

If λ_i is a characteristic root and $\sigma \in \mathrm{Gal}_K(L)$ a Galois automorphism $\sigma \in \mathrm{Gal}_K(L)$, then $\sigma(\lambda_i)$ is also a characteristic root of $(u_n)_{n=0}^{\infty}$ of the same multiplicity as λ_i. For example, if $(u_n)_{n=0}^{\infty}$ is an $\mathbb{R}$-LRS with a characteristic root λ, then $\overline{\lambda}$ is also a characteristic root. A linear recurrence sequence is called *simple* if all of its characteristic roots have multiplicity one and *degenerate* if it has characteristic roots $\lambda_i \neq \lambda_j$ such that λ_i/λ_j is a root of unity. An LRS that is not degenerate is called *non-degenerate*.

Assume that $(u_n)_{n=0}^{\infty}$ is a $\mathbb{C}$-LRS. Then a characteristic root λ_i is called *dominant* if $|\lambda_i| \geq |\lambda_j|$ for all $1 \leq j \leq k$. For brevity, we often refer to a dominant characteristic root as a dominant root. If λ is a dominant root, we call $|\lambda|$ the *spectral radius* of the LRS. Similarly, when R is some finite extension of $\mathbb{Q}_{\mathfrak{p}}$ for some prime ideal $\mathfrak{p}$, we call a characteristic root of an R-LRS $(u_n)_{n=0}^{\infty}$ $\mathfrak{p}$-*dominant* if its $\mathfrak{p}$-norm is at least as large as the $\mathfrak{p}$-norm of any other characteristic root.

By convention, when we talk about the number of dominant roots *we do not count multiplicity*; e.g., a recurrence sequence that satisfies the relation $u_{n+2} = 2u_{n+1} - u_n$ with characteristic polynomial $(X-1)^2$ has only one dominant root.

We also apply these notions of simple, (non-)degenerate, and ($\mathfrak{p}$-)dominant roots to roots of polynomials.

1.2.1 The structure of a linear recurrence sequence

We need to understand the structure of linear recurrence sequences in greater detail. Again, R is an integer domain of characteristic zero that contains 1.

The *exponential-polynomial form* is a useful direct formula for LRS.

Proposition 1.2.2 ([64, Section 1.1.6]). *Let $(u_n)_{n=0}^{\infty}$ be an R-LRS that has characteristic roots $\lambda_1, \ldots, \lambda_k$ of respective multiplicities $m_1, \ldots, m_k$. Further, let $L = R(\lambda_1, \ldots, \lambda_k)$. Then there are univariate polynomials $Q_i \in L[X]$ of degree $m_i - 1$ such that, for all $n \in \mathbb{N}$,*

$$u_n = Q_1(n)\lambda_1^n + \cdots + Q_k(n)\lambda_k^n \,. \tag{1.2}$$

We call (1.2) the *exponential-polynomial form* of $(u_n)_{n=0}^{\infty}$, and the polynomials Q_i are called the *polynomial coefficients* of $(u_n)_{n=0}^{\infty}$. Sometimes, we also use the terminology *exponential-polynomial formula*, and in the literature, the name *Binet formula* is also used.

Note that an LRS is simple if and only if all polynomial coefficients are constant.

Using the notation of Proposition 1.2.2, we can carefully count and conclude that all R-LRS are exactly those such that for all Galois automorphisms $\sigma \in \mathrm{Gal}_K(L)$ and characteristic roots λ_i and $\lambda_j = \sigma(\lambda_i)$ we have that $\sigma(Q_i) = Q_j$. That is, the expression (1.2) is closed under Galois automorphisms. For example, if $R \subset \mathbb{R}$ and λ_i is a characteristic root of an R-LRS with polynomial coefficient Q_i, then $\overline{\lambda_i}$ is also a characteristic root with polynomial coefficient $\overline{Q_i}$. From the exponential-polynomial formula, we can deduce that for R-LRS $(u_n)_{n=0}^{\infty}$ and $(v_n)_{n=0}^{\infty}$, we have that $(u_n+v_n)_{n=0}^{\infty}$, $(u_n v_n)_{n=0}^{\infty}$, and $(\sum_{k=0}^{n} u_k)_{n=0}^{\infty}$ are all R-LRS.

Let $a \geq 1$ and b be natural numbers. Then for an R-LRS $(u_n)_{n=0}^{\infty}$, the sequence $(u_{an+b})_{n=0}^{\infty}$ is again an R-LRS which we call a *subsequence* of $(u_n)_{n=0}^{\infty}$. If $(u_n)_{n=0}^{\infty}$ is simple (respectively, non-degenerate), the subsequence $(u_{an+b})_{n=0}^{\infty}$ is again simple (respectively, non-degenerate). Mignotte showed one can algorithmically decompose a $\overline{\mathbb{Q}}$-LRS into non-degenerate LRS.

Theorem 1.2.3 ([64, Theorem 1.2]). *Let $(u_n)_{n=0}^{\infty}$ be a $\overline{\mathbb{Q}}$-LRS. Then one can compute $b \in \mathbb{N}_{\geq 1}$ such that $(u_{a+bn})_{n=0}^{\infty}$ is non-degenerate for all $0 \leq a \leq b-1$.*

We sometimes *normalise* a $\mathbb{C}$-LRS and study the $\mathbb{C}$-LRS $(u_n/|\lambda|^n)_{n=0}^{\infty}$ where λ is dominant. We call the (dominant) characteristic roots of $(u_n/|\lambda|^n)_{n=0}^{\infty}$ *normalised (dominant) characteristic roots*. When $R \subset \mathbb{C}$, the *dominant part* of an R-LRS $(u_n)_{n=0}^{\infty}$ is defined as

$$\sum_{\substack{i=1 \\ \lambda_i \text{ is dominant}}}^{k} Q_i(n)\lambda_i^n$$

and the *non-dominant part* as

$$\sum_{\substack{i=1 \\ \lambda_i \text{ is not dominant}}}^{k} Q_i(n)\lambda_i^n \,.$$

Both the dominant and non-dominant parts are again LRS but not necessarily R-LRS.

The following example illustrates this phenomenon and the preceding definitions

Example 1.2.4. Let $R = \mathbb{Z}$. Then the *Fibonacci sequence* $(F_n)_{n=0}^{\infty}$ is defined by $F_0 = 0$, $F_1 = 1$, and the linear recurrence

$$F_{n+2} = F_{n+1} + F_n \,.$$

Then, $(F_n)_{n=0}^{\infty} = (0, 1, 1, 2, 3, 5, 8, 13, 21, 34, \dots)$. The characteristic polynomial of $(F_n)_{n=0}^{\infty}$ is given by $X^2 - X - 1 = (X - \varphi)(X - \widetilde{\varphi})$. Here, $\varphi := \frac{1+\sqrt{5}}{2} = 1.618\cdots$ (the *golden ratio*) and $\widetilde{\varphi} = \frac{1-\sqrt{5}}{2} = -0.618\cdots$ are the two characteristic roots of the Fibonacci sequence. Thus, $(F_n)_{n=0}^{\infty}$ is a simple, non-degenerate LRS whose only dominant root is φ. As both φ and $\widetilde{\varphi}$ are units, both are $\mathfrak{p}$-dominant for every prime ideal $\mathfrak{p}$ of the ring of integers of $\mathbb{Q}(\varphi, \widetilde{\varphi})$. The exponential-polynomial form of $(F_n)_{n=0}$ is also known as the *Binet formula* and given by

$$F_n = \frac{1}{\sqrt{5}}\varphi^n - \frac{1}{\sqrt{5}}\widetilde{\varphi}^n \,.$$

The dominant part of $(F_n)_{n=0}^{\infty}$ is the order-1 $\overline{\mathbb{Q}}$-LRS $\left(\frac{1}{\sqrt{5}}\varphi^n\right)_{n=0}^{\infty}$ and the non-dominant part of $(F_n)_{n=0}^{\infty}$ is the order-1 $\overline{\mathbb{Q}}$-LRS $\left(-\frac{1}{\sqrt{5}}\widetilde{\varphi}^n\right)_{n=0}^{\infty}$. Both of these $\overline{\mathbb{Q}}$-LRS are not $\mathbb{Q}$-LRS. □

This thesis primarily focuses on $\mathbb{Z}$-LRS, but many results extend naturally to $\mathbb{Q}$-LRS due to the following lemma.

Lemma 1.2.5. *Let $(u_n)_{n=0}^{\infty}$ be a $\mathbb{Q}$-LRS of order d. Then one can compute positive integers a and b such that $(ab^n u_n)_{n=0}^{\infty}$ is a $\mathbb{Z}$-LRS.*

Proof. Let $P \in \mathbb{Q}[X]$ be the characteristic polynomial of $(u_n)_{n=0}^{\infty}$, which has degree d by assumption. Then let b be a positive integer such that $b^d P(X/b) \in \mathbb{Z}[X]$. (One can surely compute b. For example, let b be the least common multiple of the denominators of the coefficients of P. Then indeed $b^d P(X/b) \in \mathbb{Z}[X]$.) Next, let a be the least common multiple of the denominators of $b^0 u_0$, $b^1 u_1$, …, $b^{d-1} u_{d-1}$. Then $(ab^n u_n)_{n=0}^{\infty}$ has an integer-valued linear recurrence (as its characteristic polynomial is in $\mathbb{Z}[X]$). As also its d initial values are in $\mathbb{Z}$, the sequence $(ab^n u_n)_{n=0}^{\infty}$ is a $\mathbb{Z}$-LRS. □

Linear recurrence bi-sequences

Every LRS $(u_n)_{n=0}^{\infty}$ satisfying (1) admits a unique extension to a K-valued bi-infinite sequence $(u_n)_{n=-\infty}^{\infty}$ that satisfies the same linear recurrence, where K is the fraction field of R. Intuitively, this is accomplished by reversing the linear recurrence:

$$u_n = \frac{1}{c_d}(u_{n+d} - c_1 u_{n-d-1} - \cdots - c_{d-1}u_{n+1})\,.$$

The resulting bi-infinite sequence $(u_n)_{n=-\infty}^{\infty}$ is the *bi-completion* of $(u_n)_{n=0}^{\infty}$ and is called a *linearly recurrent bi-sequence* (or *LRBS* for short). More precisely, writing $S := \{1, c_d, c_d^2, \dots\}$ and denoting by $S^{-1}R$ the localisation of R by S (the smallest ring containing all elements of the form r/s with $r \in R$ and $s \in S$), then the bi-completion of $(u_n)_{n=0}^{\infty}$ takes values in $S^{-1}R$. In particular, the bi-completion of a $\mathbb{Z}$-LRS takes values in $\mathbb{Q}$.

For all $a, b \in \mathbb{Z}$, the sequence $(u_{an+b})_{n=0}^{\infty}$ is again a K-LRS of order at most d, and in particular, an $S^{-1}R$-LRS. As a can be negative, the LRBS $(u_n)_{n=-\infty}^{\infty}$ can also be interpreted as the union of two LRS: $(u_n)_{n=0}^{\infty}$ and $(u_{-n-1})_{n=0}^{\infty}$.

Reversible LRS We will sporadically discuss the class of *reversible* LRS, consisting of all $\mathbb{Z}$-LRS $(u_n)_{n=0}^{\infty}$ whose characteristic roots are all algebraic units. Equivalently, a $\mathbb{Z}$-LRS is reversible if the constant coefficient of its characteristic polynomial is ± 1, and another equivalent definition is that a $\mathbb{Z}$-LRS is reversible if and only if its bi-completion is contained in the integers.

Linear dynamical systems

Linear recurrence sequences have a strong relationship with matrix powering. If $M \in R^{d\times d}$, $\mathbf{s}, \mathbf{t} \in R^d$, and $P(X) = X^d - c_1 X^{d-1} - \cdots - c_d$ is the characteristic polynomial of M, then the Cayley-Hamilton theorem gives that $P(M) = 0$. Setting $u_n = \mathbf{s}^\top M^n \mathbf{t}$, we get

$$\begin{aligned} &u_{n+d} - c_1 u_{n+d-1} - \cdots - c_d u_n \\ &= \mathbf{s}^\top M^n \big(M^d - c_1 M^{d-1} - \cdots - c_d I_d\big)\mathbf{t} = 0\,. \end{aligned}$$

Therefore, $(u_n)_{n=0}^{\infty}$ satisfies the linear recurrence (1). Conversely, for every R-LRS, there are a matrix $M \in R^{d\times d}$ and vectors $\mathbf{s}, \mathbf{t} \in R^m$ such that $u_n = \mathbf{s}^\top M^n \mathbf{t}$. Define

the *companion matrix* as

$$M = \begin{pmatrix} c_1 & c_2 & \cdots & c_{d-1} & c_d \\ 1 & 0 & \cdots & 0 & 0 \\ 0 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 1 & 0 \end{pmatrix} \tag{1.3}$$

and set $\mathbf{s} = (0, \ldots, 0, 1)$ and $\mathbf{t} = (u_{d-1}, u_{d-2}, \ldots, u_0)$. Then, for any $n \in \mathbb{N}$, we have $M^n\mathbf{t} = (u_{n+d-1}, u_{n+d-2}, \ldots, u_n)$, and so $u_n = \mathbf{s}^\top M^n \mathbf{t}$. Due to this matrix form, other works use the terminology eigenvalue instead of characteristic root.

More generally, for a given $M \in R^{d\times d}$ and $\mathbf{t} \in R^d$ as above, the sequence $(M^n\mathbf{t})_{n=0}^\infty$ is called a *linear dynamical system* for which there is an entire research program to model-check it [83].

1.2.2 The growth of linear recurrence sequences

For many decision problems involving R-LRS with $R \subset \mathbb{C}$, understanding their growth is essential. An upper bound follows directly from the exponential-polynomial form.

Lemma 1.2.6. *Let $(u_n)_{n=0}^\infty$ be a $\mathbb{C}$-LRS with spectral radius ρ and $C \in \mathbb{R}_{>|\lambda|}$. Then one can compute $r \in \mathbb{Q}_{>0}$ and $\rho \in (|\lambda|, C) \cap \mathbb{Q}$ such that $|u_n| < r\rho^n$ for all $n \in \mathbb{N}$.*

Proof. Assume that $(u_n)_{n=0}^\infty$ satisfies $u_n = \sum_{i=1}^k Q_i(n)\lambda_i^n$. If $m = \max_{i=1}^k(\deg(Q_i))$ and Q is larger than the absolute value of any coefficient of any polynomial Q_i, then for all $1 \le i \le k$ and $n \in \mathbb{N}$, we have $|Q_i(n)| < (m+1)Q \cdot n^m$ and thus $|u_n| < k(m+1)Q \cdot n^m|\lambda|^n$. Thus, $|u_n| < k(m+1)Qn^m|\lambda|^n$. Then, for computably large $N \in \mathbb{N}$, we have that $\rho^n \ge n^m|\lambda|^n$ for all $n \ge N$. Then take $r \ge k(m+1)Q$ such that $|u_n| < r\rho^n$ for all $n \le N$. The lemma follows. □

Finding a lower bound for the growth of a LRS is more involved. We cite one version found in [84, Theorem 2], which is a consequence of the works of Evertse [65] and van der Poorten [127].

Theorem 1.2.7. *Let $(u_n)_{n=0}^\infty$ be a non-degenerate $\overline{\mathbb{Q}}$-LRS satisfying the polynomial exponential formula $u_n = \sum_{i=1}^k Q_i(n)\lambda_i^n$ with spectral radius ρ. For every $0 < r < \rho$, there exists $N \in \mathbb{N}$ such that $|u_n| > r^n$ for all $n > N$.*

Since the proof of Theorem 1.2.7 relies on S-units, the result is non-effective, i.e., no implicit method is provided to compute this number N. In restricted cases, effective results can be proven, like the following results of Mignotte, Shorey, and Tijdeman [116] and Vereshchagin [156].

Theorem 1.2.8. *Let $(u_n)_{n=0}^{\infty}$ be a non-degenerate $\overline{\mathbb{Q}}$-LRS with two dominant roots of magnitude $|\lambda|$. Then there are computable positive constants C_1 and C_2 such that*

$$|u_n| \geq |\lambda|^n \cdot n^{-C_1 \log(n)}$$

whenever $n \geq C_2$.

The second of their results does not concern the growth of the LRS itself but the distance between its terms.

Theorem 1.2.9. *Let $(u_n)_{n=0}^{\infty}$ be a non-degenerate $\overline{\mathbb{Q}}$-LRS with two dominant roots of magnitude $|\lambda|$. Then there are computable positive constants C_3 and C_4 such that*

$$|u_{n_1} - u_{n_2}| \geq |\lambda|^{n_1} \cdot {n_1}^{-C_3 \log(n_1) \log(n_2+2)}$$

whenever $n_1 > n_2$ and $n_1 \geq C_4$.

1.3 Logic and automata

1.3.1 Words and automata

Words

By an *alphabet* Σ, we mean a finite non-empty set of letters. For an alphabet Σ and $w_1, \ldots, w_d \in \Sigma$, let $w = w_1 \cdots w_d$ denote the *finite word* w and Σ^* the set of all finite words over the alphabet Σ. Sometimes, we also write w as a tuple $(w_1, \ldots, w_d)$. In this case, we call d the *length* of the word w, which we denote by $|w|$. With ε, we denote the *empty word*, the unique word in Σ^* of length 0. The set Σ^+ denotes all words of strictly positive length, i.e., $\Sigma^+ = \Sigma^* \setminus \{\varepsilon\}$.

Similarly, we can treat infinite words. Let $w = w_0 w_1 \cdots$ denote an *infinite word* over the alphabet Σ. The set of all infinite words over the alphabet Σ is denoted by Σ^ω, and sometimes we write $w = (w_0, w_1, \ldots)$ and call w a sequence.

We refer to elements in the alphabet Σ as *letters* and if w is a word and $v = v_1 \cdots v_d \in \Sigma^+$, then v is a *factor* of $w \in \Sigma^* \cup \Sigma^\omega$ if there is an $n \in \mathbb{N}$ such that $v = w_n \cdots w_{n+d-1}$. In this case, the letter v *occurs* at position n in w. As an example of this terminology, we say that v occurs infinitely often in w if there are infinitely many $n \in \mathbb{N}$ such that v occurs at position n in w

An infinite word $w \in \Sigma^\omega$ is called *recursive* or *effective* if one can compute w_j for all $j \geq 0$. Formally, there is a Turing machine which, at input j, outputs w_j.

A *factorisation* of $w \in \Sigma^\omega$ is a sequence $(u_n)_{n=0}^{\infty}$ such that $w = u_0 u_1 \cdots$ and $u_n \in \Sigma^+$ for all $n \in \mathbb{N}$. A factorisation is uniquely determined by a strictly increasing sequence $(k_n)_{n=0}^{\infty}$ over $\mathbb{N}$ such that $u_n = w_{k_n} \cdots w_{k_{n+1}-1}$ for all $n \in \mathbb{N}$.

Automata and transducers

A *deterministic Muller automaton* $\mathcal{A} = (\Sigma, Q, q_{\text{init}}, \delta, \mathcal{F})$ consists of an alphabet Σ, a finite set of states Q, an initial state $q_{\text{init}} \in Q$, a transition function $\delta : Q \times \Sigma \to Q$, and an accepting family of sets $\mathcal{F} \subseteq \mathcal{P}(Q)$. The *run* of $w \in \Sigma^* \cup \Sigma^\omega$ on $\mathcal{A}$ is the sequence of states visited while reading w starting in q_{init} and repeatedly updating the state using the transition function while reading the word w. We say that $\mathcal{A}$ *accepts* $w \in \Sigma^\omega$ if the set of states visited infinitely often upon reading w belongs to $\mathcal{F}$. The *acceptance problem* of a recursive word $w \in \Sigma^\omega$ is the question of determining whether $\mathcal{A}$ accepts w for any given deterministic Muller automaton $\mathcal{A}$ with alphabet Σ. We denote the acceptance problem by Acc_w.

A Muller automaton is an example of an *ω-automaton*. Instead of using a Muller automaton, one can use Büchi or parity automata, which are equally expressive.

A *deterministic finite transducer* $\mathcal{B} = (\Sigma_{\text{in}}, \Sigma_{\text{out}}, Q, q_{\text{init}}, \delta)$ is given by an input alphabet Σ_{in}, output alphabet Σ_{out}, set of states Q, initial state $q_{\text{init}} \in Q$, and transition function $\delta : Q \times \Sigma_{\text{in}} \to Q \times \Sigma_{\text{out}}^*$. The transducer $\mathcal{B}$ starts in state q_{init} and reads a word $w \in \Sigma_{\text{in}}^* \cup \Sigma_{\text{in}}^\omega$ and upon reading the letter a whilst in state q, computes $(q', w') = \delta(q, a)$, moves to state q', and concatenates w' to the output string. Write $\mathcal{B}(w) \in \Sigma_{\text{out}}^* \cup \Sigma_{\text{out}}^\omega$ to denote the output word computed upon reading $w \in \Sigma_{\text{in}}^* \cup \Sigma_{\text{in}}^\omega$.

Disjunctive words and normal numbers

In this thesis, we frequently encounter disjunctive words. A word $\alpha \in \Sigma^\omega$ is *disjunctive* if every $v \in \Sigma^+$ occurs infinitely often as a factor in α, i.e., if $v = v_1 \cdots v_d$, there is an n such that $v = \alpha_n \cdots \alpha_{n+d-1}$. Equivalently, there are infinitely many such n: If $|\Sigma| = 1$, then α is constant. If $|\Sigma| \geq 2$ and $N \geq 1$, one can construct $w_1, \ldots, w_N \in \Sigma^+$ such that v is a prefix of each w_i while these w_i are not prefixes of each other. Hence, if each factor appears at least once, then v appears at least N times and thus v occurs infinitely often as a factor of α as N was arbitrary.

Some other works use the terminology *rich* instead of disjunctive, and in previous works [24, 119], we used the terminology *weakly normal*.

If $\widetilde{\Sigma} \subset \Sigma$, then $\alpha \in \Sigma^\omega$ is *disjunctive with respect to* $\widetilde{\Sigma}$ if each $u \in \widetilde{\Sigma}^+$ appears at least once (or, equivalently, infinitely often) in α.

The motivation of disjunctive words lies in the base-b expansions of numbers, first studied by Borel in 1909 [37]. We can interpret a number $x \in \mathbb{R}$ as a word $w \in \{0, \ldots, b-1\}^\omega$ by considering its base-b expansion modulo 1. Here, a number

is *normal in base-b* if each factor appears with a frequency one would expect from a 'random' number. That is, for every $v = v_1 \cdots v_d \in \{0, \ldots, b-1\}^+$,

$$\lim_{N\to\infty} \frac{\#\{0 \le n < N \colon v = w_n \cdots w_{n+d-1}\}}{N} = \frac{1}{b^d}\,. \tag{1.4}$$

Then x is *normal* if x is normal in base-b for all $b \ge 2$.

Although rational numbers are not normal (as their base-b expansion is ultimately periodic for every $b \ge 2$), normal numbers are abundant. Borel showed that normal numbers have Lebesgue measure 1 in $(0, 1)$, but little is known about most other natural constants. For example, $\sqrt{2}$, e, $\log(2)$, and π are all conjectured to be normal [14]. For a detailed discussion of normal numbers, see surveys [77, 129] and the book [45].

In particular, [77] states the following conjecture.

Conjecture 1.3.1. *Irrational real algebraic numbers are normal in any base $b \ge 2$.*

So far, the strongest result towards this conjecture is due to Adamczewski and Bugeaud [1]. Let $p(n)$ denote the *factor complexity* of α, i.e., the number of distinct factors in α of length n.

Theorem 1.3.2. *If $b \ge 2$ and α is the base-b expansion of an irrational algebraic number, then*

$$\liminf_{n\to\infty} \frac{p(n)}{n} = +\infty\,.$$

1.3.2 Logical theories

Let U be a universe. Then a *predicate* P is a subset of U^μ for some $\mu \in \mathbb{N}_{\ge 1}$, and predicate P is *unary* if $\mu = 1$. Then, we can interpret P as a function $P : U^\mu \to \{\text{true, false}\}$ and as a set such that $P(\mathbf{x})$ holds if and only if $\mathbf{x} \in P$.

A *structure* $\mathbb{M}$ consists of a universe U, constants $c_1, \ldots, c_k \in U$, predicates $P_1, \ldots, P_\ell$ where each $P_i \subseteq U^{\mu(i)}$ for some $\mu(i) \ge 1$, and functions $f_1, \ldots, f_m$ where each f_i has the type $f_i \colon U^{\delta(i)} \to U$ for some $\delta(i) \ge 1$. We denote this structure by $\langle U; c_1, \ldots, c_k, P_1, \ldots, P_\ell, f_1, \ldots, f_m\rangle$. From the context, it should be clear which items are constants, predicates, and functions.

By the *language* of the structure $\mathbb{M}$ we mean the set of all well-formed first-order formulas constructed from symbols denoting the constants $c_1, \ldots, c_k$, predicates $P_1, \ldots, P_\ell$, and functions $f_1, \ldots, f_m$, as well as the symbols $\forall, \exists, \wedge, \vee, \neg$, and $=$. A *term* is a well-formed expression constructed from constant, function, and variable symbols. Terms represent elements of the universe. A *(first-order) theory* is simply

a set of sentences, i.e., formulas without free variables. The theory of the structure $\mathbb{M}$ is the set of all sentences in the language of $\mathbb{M}$ that are true in $\mathbb{M}$. A formula is *existential* if it is of the form $\exists x_1 \cdots \exists x_m \colon \varphi(x_1, \ldots, x_m)$ for a quantifier-free sentence φ. The *existential fragment* of a theory $\mathcal{T}$, which itself is a theory, is the set of all existential formulas belonging to $\mathcal{T}$. Finally, a theory $\mathcal{T}$ is *decidable* if there exists an algorithm that takes a sentence φ and determines whether $\varphi \in \mathcal{T}$. A set $X \subseteq U^d$ is *definable* in a structure $\mathbb{M}$ if there exists a formula φ in the language of $\mathbb{M}$ with d free variables such that for all $x_1, \ldots, x_d \in U$, we have that $\varphi(x_1, \ldots, x_d)$ holds if and only if $(x_1, \ldots, x_d) \in X$.

The real ordered field with the exponential function

Let $\mathbb{R}_{\exp} := \langle \mathbb{R}; <, +, -, \cdot, \exp(\cdot), 0, 1 \rangle$ denote the structure of the real ordered field with exponentiation. The first-order theory of $\mathbb{R}_{\exp}$ is the set of all well-formed first-order sentences in a suitable language $\mathcal{L}_{\exp}$ that are true in $\mathbb{R}_{\exp}$. In [109], Macintyre and Wilkie showed that the first-order theory of the structure $\mathbb{R}_{\exp}$ is decidable assuming Schanuel's conjecture (Conjecture 1.1.4) for termination.

Theorem 1.3.3. *Assuming Schanuel's conjecture for termination, given a sentence $\varphi \in \mathcal{L}_{\exp}$, one can determine whether φ holds in $\mathbb{R}_{\exp}$.*

Presburger arithmetic

We denote *Presburger arithmetic*, the theory of the structure $\langle \mathbb{Z}; 0, 1, <, + \rangle$, with $\mathcal{PA}$. If $P_1, \ldots, P_\ell$ and $f_1, \ldots, f_m$ denote predicates and functions over the universe $\mathbb{Z}$, respectively, then $\mathcal{PA}(P_1, \ldots, P_\ell, f_1, \ldots, f_m)$ denotes the theory of the structure $\langle \mathbb{Z}; 0, 1, <, +, P_1, \ldots, P_\ell, f_1, \ldots, f_m \rangle$ (Presburger arithmetic *expanded* with these predicates and function).

A set $X \subseteq \mathbb{Z}^d$ is *semilinear* if it is definable in the structure $\langle \mathbb{Z}; 0, 1, <, + \rangle$. By the result of Presburger [128] that Presburger arithmetic admits *quantifier elimination* if we allow a divisibility predicate [74], such X can be defined by a formula of the form

$$\bigvee_{i \in I} \Bigg(\bigwedge_{j = J_i} t_j(x_1, \ldots, x_d) \equiv 0 \pmod{D_j} \wedge \bigwedge_{k \in K_j} h_k(x_1, \ldots, x_d) \sim_k c_k \Bigg), \tag{1.5}$$

where $D_j \geq 1$ and each t_j, h_j is a $\mathbb{Q}$-linear form, $c_k \in \mathbb{Z}$, and $\sim_k \in \{>, =\}$.

Monadic Second-Order Logic

Monadic second-order logic (MSO), contrary to first-order logic, allows quantification over both elements and subsets of the universe. We will only be interpreting MSO formulas over expansions of the structure $\langle \mathbb{N}; < \rangle$. For a general perspective on MSO, see [36].

Let $\mathbb{S} := \langle \mathbb{N}; <, P_1, \ldots, P_m \rangle$ be a structure where each predicate $P_i \subseteq \mathbb{N}$ is unary. We associate a language $\mathcal{L}_{\mathbb{S}}$ of terms and formulas with $\mathbb{S}$ as follows. The terms of $\mathcal{L}_{\mathbb{S}}$ are the countably many constant symbols $\{0, 1, 2, \ldots\}$, lowercase variable symbols that stand for elements of $\mathbb{N}$, and uppercase variable symbols that denote subsets of $\mathbb{N}$. The formulas of $\mathcal{L}_{\mathbb{S}}$ are the well-formed statements constructed from the built-in equality ($=$) and membership ($\in$) symbols, logical connectives, quantification over elements of $\mathbb{N}$ (written Qx for a quantifier Q), and quantification over subsets (written QX for a quantifier Q). The MSO theory of the structure $\mathbb{S}$ is the set of all sentences belonging to $\mathcal{L}_{\mathbb{S}}$ that are true in $\mathbb{S}$. The MSO theory of $\mathbb{S}$ is *decidable* if there exists an algorithm that, given a sentence $\varphi \in \mathcal{L}_{\mathbb{S}}$, determines whether φ belongs to the MSO theory of $\mathbb{S}$.

As an example, consider $\mathbb{S} = \langle \mathbb{N}; <, P \rangle$ where P is the set of all primes. Let $s(\cdot)$ be the successor function defined by $s(x) = y$ if and only if

$$x < y \quad \wedge \quad \forall z \colon x < z \Rightarrow y \leq z \,.$$

That is, $s(x) = x + 1$. Further let

$$\varphi(X) := 1 \in X \wedge 0, 2 \notin X \wedge \forall x \colon x \in X \iff s(s(s(x))) \in X,$$
$$\psi := \exists X \colon \varphi(X) \wedge \forall y . \exists z > y \colon z \in X \wedge P(z) \,.$$

The formula φ defines the subset $\{n \colon n \equiv 1 \pmod 3\}$ of $\mathbb{N}$, and ψ is the sentence "there are infinitely many primes congruent to 1 modulo 3", which is the case.

Fourier-Motzkin elimination

Fourier-Motzkin elimination is an algorithm to eliminate variables from a system of linear inequalities. Let $\Phi(x_1, \ldots, x_m)$ be a Boolean combination of atomic formulas of the form $h(x_1, \ldots, x_m) \sim 0$, where h is a $\mathbb{Q}$-affine form and $\sim$ is a (strict or non-strict) inequality symbol. Let $1 \leq \ell \leq m$, and consider the formula $\exists x_1, \ldots, x_\ell \in \mathbb{R} \colon \Phi(x_1, \ldots, x_m)$. Using the Fourier-Motzkin elimination [54], we can compute a formula $\Psi(x_{\ell+1}, \ldots, x_m) = \bigvee_{j \in J} \bigwedge_{k \in K} h_{j,k}(x_{\ell+1}, \ldots, x_m) \sim_{j,k} 0$ such that

(a) each $\sim_{j,k}$ is an inequality and $h_{j,k}$ is a $\mathbb{Q}$-affine form, and

(b) for all $z_{\ell+1}, \ldots, z_m \in \mathbb{R}$, the sentence

$$\exists x_1, \ldots, x_\ell \in \mathbb{R} \colon \Phi(x_1, \ldots, x_\ell, z_{\ell+1}, \ldots, z_m)$$

holds if and only if $\Psi(x_{\ell+1}, \ldots, x_m)$ holds.

Chapter 2

The Skolem problem

2.1 Introduction and main results

In this chapter, we discuss the Skolem problem, which asks one to determine whether a $\mathbb{Z}$-LRS has a zero; i.e., does there exist an $n \in \mathbb{N}$ such that $u_n = 0$? We can generalize the Skolem problem to rings other than $\mathbb{Z}$.

Problem 2.1.1 (Skolem problem for R-LRS)**.** For an integral domain R, the R-Skolem problem asks if it is decidable whether, for a given R-LRS $(u_n)_{n=0}^{\infty}$, there is an $n \in \mathbb{N}$ such that $u_n = 0$.

The most prominent open cases of the Skolem problem concern the rings $R = \mathbb{Z}$ and $R = \mathbb{Q}$. Due to Lemma 1.2.5, the Skolem problems for these two rings are Turing-equivalent, and the reductions in both directions preserve properties like order, degeneracy, and simplicity. Consequently, throughout this chapter, we refer to $\mathbb{Z}$-LRS simply as LRS, and the $\mathbb{Z}$-Skolem problem as the Skolem problem. The $\overline{\mathbb{Q}}$-Skolem problem is also Turing equivalent to the $\mathbb{Z}$-Skolem problem, but this reduction does not preserve order (see Section 3.5).

Before presenting the current state of the Skolem problem, our technical contributions, we briefly review the history of the Skolem problem.

Early motivation of the Skolem problem The origins of the Skolem problem can be traced back to Hilbert in the 1920s. In his famous program, Hilbert identified decidability as a central objective, envisioning the development of an algorithm that can determine whether any mathematical statement is true or false [160]. Among his famous list of problems, Hilbert's tenth problem sought an algorithm to decide whether a given Diophantine equation has an integer solution. Formally, for a polynomial $P \in \mathbb{Z}[X_1, \ldots, X_d]$, determine whether there are $n_1, \ldots, n_d \in \mathbb{Z}$ such that

$P(n_1, \ldots, n_d) = 0$. However, Hilbert's vision could not be fulfilled. Gödel's incompleteness theorems demonstrated that no algorithm can decide the truth of all mathematical statements. Matiyasevich, utilizing the previous work of Davis, Putnam, and Robinson, underscored this limitation by proving in 1970 that no algorithm exists that determines the solutions of all Diophantine equations [114].

Thoralf Skolem was a Norwegian mathematician primarily interested in logic but also contributed to the fields of Diophantine equations, lattice theory, and group theory. He is the namesake of the Skolem problem, and his involvement with the Skolem problem began during his study of Thue equations [143]. Thue equations are Diophantine equations of the form

$$a_0X^3 + a_1X^2Y + a_2XY^2 + a_3Y^3 = d\,,$$

where a_0, a_1, a_2, a_3, and d are given integers. For many such instances, one can construct a linear recurrence sequence that contains zero if and only if the corresponding Thue equation has a solution. This is exactly an instance of the Skolem problem! To solve these cases of the Skolem problem, Skolem relied on p-adic techniques, and so this method of solving Thue equations is nowadays known as *Skolem's (p-adic) method*. Skolem's method is still sporadically used (e.g., in [149]) but has mostly fallen into obscurity since Baker developed his methods. In general, methods employing Baker's theorem are more practical and more often applicable. As Brock, Elkies, and Jordan [40] note "Skolem's method does not always apply in diophantine problems, and it does not always work even when it applies (...)".

The Skolem-Mahler-Lech theorem Skolem authored several papers on linear recurrence sequences, the most notable being [144], which contains one of the most foundational results on linear recurrence sequences: the Skolem-Mahler-Lech theorem. This theorem states, under mild assumptions, that for an LRS $(u_n)_{n=0}^{\infty}$, there are only finitely many indices $n \in \mathbb{N}$ such that $u_n = 0$.

Theorem 2.1.2 (Skolem-Mahler-Lech)**.** *Let $(u_n)_{n=0}^{\infty}$ be a $\mathbb{C}$-LRS and define $Z = \{n \in \mathbb{N} \colon u_n = 0\}$. Then Z is the union of finitely many arithmetical progressions and a finite set. In particular, if $(u_n)_{n=0}^{\infty}$ is non-degenerate and not the zero-LRS, then Z is finite.*

Skolem [144] initially proved this result for $\mathbb{Q}$-LRS, and Mahler [110] subsequently generalised it to $\overline{\mathbb{Q}}$-LRS. Two decades later, Lech [94] proved the general case for R-LRS, where R is an integral domain of characteristic 0.

Using Theorem 1.2.3, one can decompose an LRS $(u_n)_{n=0}^{\infty}$ into non-degenerate subsequences $(u_{Sn+i})_{n=0}^{\infty}$. If one can decide the Skolem problem for all these S subsequences, it can be decided for the original sequence. Therefore, Theorem 1.2.3 reduces the Skolem problem to non-degenerate LRS. Then, according to the Skolem-Mahler-Lech theorem, the set $Z := \{n \in \mathbb{N} \colon u_n = 0\}$ is finite. However, the Skolem-Mahler-Lech theorem does not provide a method to compute Z, i.e., the theorem is non-effective. The Skolem problem may be viewed as the search for an effective Skolem-Mahler-Lech theorem.

In rings of characteristic $p > 0$, the Skolem-Mahler-Lech theorem as above does not hold. For instance, Lech [94] observed that for the $\mathbb{F}_p(t)$-LRS defined by $u_n = (t+1)^n - t^n + 1$, then $u_n = 0$ if and only if n is a power of p. Thus, the zero set $Z = \{p^m \colon m \in \mathbb{N}\}$ is infinite while $(u_n)_{n=0}^{\infty}$ is non-degenerate. The Skolem-Mahler-Lech theorem seems to fail. Nevertheless, Z still exhibits a high degree of structure: it is *p-automatic*. In 2007, Derksen [56] developed an analogue of the Skolem-Mahler-Lech theorem in positive characteristic by proving that the zero-set Z is always p-automatic. Moreover, in the same paper, he showed that the Skolem problem is decidable for fields R that are finitely generated over $\mathbb{F}_p$.

The MSTV class In the 1960s, Alan Baker made groundbreaking contributions to the study of (exponential) Diophantine equations and number theory in general with his study of linear forms in logarithms. The methods stemming from Baker's ideas proved far more effective, general, and reliable than Skolem's methods and solved many Diophantine equations and other open problems in number theory. Among Baker's most celebrated results is what is now known as *Baker's theorem on linear forms in logarithms*. As a consequence, Mignotte, Shorey, and Tijdeman [116] and Vereshchagin [156] independently applied Baker's theorem to linear recurrence sequences in 1984 and 1985.

Theorem 2.1.3 (Mignotte, Shorey, and Tijdeman and Vereshchagin)**.** *The Skolem problem is decidable for all $\mathbb{Q}$-LRS $(u_n)_{n=0}^{\infty}$ with at most three dominant roots or at most two $\mathfrak{p}$-dominant roots for some prime ideal $\mathfrak{p}$.*

We refer to the class of LRS covered by this theorem as follows.

Definition 2.1.4. The *MSTV class* comprises all $\mathbb{Q}$-LRS with at most three dominant roots, or at most two $\mathfrak{p}$-dominant roots for some prime ideal $\mathfrak{p}$.

Any LRS with at most four characteristic roots belongs to the MSTV class, which includes all LRS of order up to 4, and one can construct explicit order-5 LRS outside of the MSTV class. More subtly, the Skolem problem is also decidable for all $\overline{\mathbb{Q}}$-LRS with at most three dominant roots or at most two $\mathfrak{p}$-dominant roots, and that again all $\overline{\mathbb{Q}}$-LRS with at most four characteristic roots have one of these two properties as shown recently by Bacik [10] (see also Bilu [30]).

Other recent results The previously mentioned results of Mignotte, Shorey, and Tijdeman and Vereshchagin appear to have exhausted the power of Baker's theorem. As a result, recent research shifted to alternative approaches, focusing on certain fragments of the Skolem problem or working under the assumption of certain conjectures.

In one line of research [104, 105, 106, 108], Luca et al. introduce the concept of a *universal Skolem set*. A set $S \subset \mathbb{N}$ is a universal Skolem set if, for all LRS $(u_n)_{n=0}^{\infty}$, one can determine whether $u_n = 0$ for some $n \in S$. From this perspective, the classical Skolem problem amounts to establishing that $\mathbb{N}$ itself forms a universal Skolem set. So far, Luca et al. have produced a universal Skolem set S of density at least 0.29. That is,

$$\liminf_{N\to\infty} \#\{0 \leq s < N \colon s \in S\}/N \geq 0.29\,.$$

The Bateman-Horn conjecture implies that this set S has density 1. Additionally, a variant of the Cramér conjecture implies that the Skolem problem is decidable. Both the Bateman-Horn conjecture and the variant of the Cramér conjecture are deep number-theoretic conjectures related to the distribution of prime numbers.

Furthermore, Blondel and Portier [35] showed that the Skolem problem is NP-hard, and Akshay et al. [4] gave an alternative proof of this fact. In both cases, the authors depend on degenerate LRS to obtain NP-hardness. However, to settle the decidability of the Skolem problem, it is sufficient to consider non-degenerate LRS and no complexity hardness result is known for non-degenerate LRS.

Main results

To demonstrate our approach to solving certain fragments of the Skolem problem, we introduce the *Bi-Skolem problem*. Recall that any LRS $(u_n)_{n=0}^{\infty}$ has a bi-completion $(u_n)_{n=-\infty}^{\infty}$ which forms a linear recurrence bi-sequence (LRBS) that satisfies the linear recurrence (1) for all $n \in \mathbb{Z}$ by running the recurrence backwards:

$$u_n = \frac{1}{c_d}\left(u_{n+d} - c_1 u_{n+d-1} - \cdots - c_{d-1}u_{n+1}\right)\,. \tag{2.1}$$

The Bi-Skolem problem is the natural analogue of the Skolem problem for LRBS.

Problem 2.1.5 (Bi-Skolem problem)**.** For a $\mathbb{Q}$-LRBS $(u_n)_{n=-\infty}^{\infty}$, determine whether $u_n = 0$ for some $n \in \mathbb{Z}$.

As an LRBS is the union of the two sequences $(u_n)_{n=0}^{\infty}$ and $(u_{-n-1})_{n=0}^{\infty}$, the Bi-Skolem problem reduces to the Skolem problem. It is unknown whether there is a reduction in the other direction, but in Section 2.3, we give a reduction of the Bi-Skolem problem to the Skolem problem assuming a p-adic version of Schanuel's conjecture (Conjecture 1.1.6).

Theorem 2.1.6. *Assuming the p-adic Schanuel conjecture, the Skolem problem and the Bi-Skolem problem are Turing-interreducible. Moreover, the p-adic Schanuel conjecture is solely required to prove termination.*

Hence, our focus shifts from LRS and the Skolem problem to LRBS and the Bi-Skolem problem. The advantage of this perspective is that LRBS are a more global object as they are closed under shifts of the sequence. If $(u_n)_{n=0}^{\infty}$ is a $\mathbb{Z}$-LRS satisfying (1), then (2.1) implies that the LRBS takes values in $\mathbb{Z}[1/c_d]$, and the denominators c_d^k of terms in the LRBS are arbitraryly large powers of c_d. Thus, for $M \geq 1$ coprime with c_d, the bi-sequence $(u_n \bmod M)_{n=-\infty}^{\infty}$ is well-defined.

The more global structure of an LRBS seems to imply a local-global principle for LRBS. Skolem first conjectured this in 1937 [145], so we refer to this conjecture as the *Skolem conjecture*. The Skolem conjecture is also known as the *exponential local-global principle*, but this name also sometimes refers to a more general conjecture.

Conjecture 2.1.7 (Skolem conjecture)**.** *Let $(u_n)_{n=0}^{\infty}$ be a simple $\mathbb{Z}$-LRS whose characteristic polynomial has constant coefficient c_d. Then its bi-completion $(u_n)_{n=-\infty}^{\infty}$ is non-zero if and only if for some $M \geq 1$ coprime with c_d, the bi-sequence $(u_n \bmod M)_{n=-\infty}^{\infty}$ is non-zero.*

In other words, if a simple LRBS is never zero, this is witnessed by some modulus M for which $(u_n \bmod M)_{n=-\infty}^{\infty}$ is well-defined.

There exists a substantial body of literature on the Skolem conjecture, including proofs for a variety of special cases. In particular, the Skolem conjecture has been shown to hold for simple LRBS of order 2 [16] and for certain families of LRBS of order 3 [135, 136]. In a different but related vein, Bertók and Hajdu have shown that, in some sense, the Skolem conjecture is valid in 'almost all' instances [26, 27].

The Skolem conjecture directly gives a method to solve the Bi-Skolem problem for simple LRS by running two semi-algorithms in parallel:

1. Enumerate $n \in \mathbb{Z}$ until $u_n = 0$.

2. Enumerate $M \geq 2$ until finding an M coprime with c_d such that $(u_n \bmod M)_{n=-\infty}^{\infty}$ does not contain 0.

These bi-sequences are enumerable, so the first semi-algorithm exists. For the second semi-algorithm, we can use Lemma 1.2.1. The Skolem conjecture implies that exactly one of these two semi-algorithms terminates. If the first terminates, then 0 is in $(u_n)_{n=-\infty}^{\infty}$ and if the second terminates 0 is not in $(u_n)_{n=-\infty}^{\infty}$. Thus, the Skolem conjecture implies that the Bi-Skolem problem is decidable for simple LRS. The observation above and Theorem 2.1.6 imply the main result of the chapter: Under mild assumptions, the Skolem problem is decidable for simple LRS.

Theorem 2.1.8. *For all simple LRS, the Skolem problem is decidable when assuming the p-adic Schanuel conjecture and the Skolem conjecture. Moreover, both conjectures are only needed to ensure termination.*

Because both conjectures are only used to prove termination, we obtain an independent certificate of correctness that does not depend on these conjectures when our procedure terminates.

Organization of the chapter

We start with studying the behaviour of an LRS modulo integers M in Section 2.2. Specifically, we call an LRS $(u_n)_{n=0}^{\infty}$ *modular* if there is an $M \geq 1$ such that $(u_n \bmod M)_{n=0}^{\infty}$ is ultimately non-zero, i.e., $u_n \equiv 0 \pmod{M}$ for only finitely many $n \in \mathbb{N}$. We use our results from modular LRS to give an alternate proof of Theorem 2.1.8 for LRS of order 5, where we do not rely on the p-adic Schanuel conjecture.

Theorem 2.1.9. *Let $(u_n)_{n=0}^{\infty}$ be an LRS of order at most 5. Then, assuming the Skolem conjecture for termination, the Skolem problem is decidable for $(u_n)_{n=0}^{\infty}$.*

Next, in Section 2.3, we prove Theorem 2.1.8, and in Section 2.4, we strengthen this result for low-order recurrences, where we show the following result.

Theorem 2.1.10. *Let $(u_n)_{n=0}^{\infty}$ be a simple LRS of order at most 7. Then, assuming the Skolem conjecture for termination, the Skolem problem is decidable for $(u_n)_{n=0}^{\infty}$.*

Although Theorem 2.1.10 absorbs Theorem 2.1.9, we provide separate proofs and algorithms for both as an algorithm for Theorem 2.1.9 is significantly simpler.

We also present an implementation of our procedure for Theorem 2.1.8, the SKOLEM-tool. We also partially implemented the (often much faster) methods of Mignotte, Shorey, and Tijdeman and Vereshchagin using Baker's theorem on linear forms in logarithms, so we also describe how to use these tools in practice.

2.2 Modularity

An LRS is *modular* if there exists an $M \geq 1$ such that $u_n \equiv 0 \pmod{M}$ for only finitely many $n \geq 0$. For a fixed LRS $(u_n)_{n=0}^{\infty}$, when one knows that such a number M exists, the Skolem problem for $(u_n)_{n=0}^{\infty}$ is decidable: Enumerate $M \geq 1$ and compute the preperiod and the repeating block of the ultimately periodic LRS $(u_n \pmod{M})_{n=0}^{\infty}$. When the repeating block does not contain 0 (which by assumption eventually occurs), computing the finite set $\{n \in \mathbb{N} \colon u_n \equiv 0 \bmod M\}$ is straightforward using this ultimately periodic representation. Hence, to solve the Skolem problem, one only needs to compute u_n for n in this set.

Due to such straightforward proofs, we study the notion of modularity for LRS in this section. We describe a criterion for a simple LRS to be modular and use this criterion to show the decidability of the Skolem problem for LRS of order 5, subject to the Skolem conjecture. The following is a simple but useful observation about modular LRS.

Observation 2.2.1. If an LRS $(u_n)_{n=0}^{\infty}$ is not modular and $q > 0$, then for some residue class $r \in \{0, 1, \ldots, q-1\}$ one of the subsequences $(u_{qn+r})_{n=0}^{\infty}$ is not modular.

Proof. Assume that all subsequences $(u_{qn+r})_{n=0}^{\infty}$ are not modular. Then for some $M_0, \ldots, M_{q-1}$ we have that for all $0 \leq r < q$, there are only finitely many $n \in \mathbb{N}$ such that $u_{qn+r} \equiv 0 \pmod{M_r}$. Now let $M = \operatorname{lcm}(M_0, \ldots, M_{q-1})$ and $Z = \{n \in \mathbb{N} \colon u_n \equiv 0 \pmod{M}\}$. Then Z has a finite intersection with every congruence class modulo M. Thus, Z is finite. Hence, $(u_n)_{n=0}^{\infty}$ is modular, contradicting our assumption that $(u_n)_{n=0}^{\infty}$ is modular. □

Write a simple LRS $(u_n)_{n=0}^{\infty}$ in its exponential-polynomial form (1.2). Then for a prime ideal $\mathfrak{p}$ above a rational prime, define the $\mathfrak{p}$-*dominant part* $(u_n^{\mathfrak{p}})_{n=0}^{\infty}$ as

$$u_n^{\mathfrak{p}} := \sum_{\substack{i=1 \\ \lambda_i \text{ is } \mathfrak{p}\text{-dominant}}}^{k} Q_i(n)\lambda_i^n \,. \tag{2.2}$$

When $(u_n)_{n=0}^{\infty}$ is simple, all polynomials $Q_i(n)$ are constant. While $(u_n^{\mathfrak{p}})_{n=0}^{\infty}$ is a $\overline{\mathbb{Q}}$-LRS, it is not necessarily a $\mathbb{Z}$-LRS.

Proposition 2.2.2. *Suppose that $(u_n)_{n=0}^{\infty}$ is a simple and non-degenerate LRS and that $\mathfrak{p}$ a prime ideal that does not contain all the characteristic roots of the LRS. If there is no $n \in \mathbb{Z}$ such that $u_n = u_n^{\mathfrak{p}} = 0$, the Skolem conjecture implies that $(u_n)_{n=0}^{\infty}$ is modular.*

Proof. Let $\mathfrak{p}$ be a prime ideal that does not contain all the characteristic roots. Suppose that $(u_n)_{n=0}^{\infty}$ is not modular. We argue that $u_n = u_n^{\mathfrak{p}} = 0$ for some $n \in \mathbb{N}$.

Let K be the splitting field of $(u_n)_{n=0}^{\infty}$ and $\mathcal{O}_K$ its ring of integers. Then the residue field $\mathcal{O}_K/\mathfrak{p}$ is finite of order p^f, where p is a rational prime and f the *residual degree* of $\mathfrak{p}$. From the assumption that $(u_n)_{n=0}^{\infty}$ is not modular, by Observation 2.2.1, there exists an $r \in \{0, \ldots, p^f - 2\}$ such that $(w_n)_{n=0}^{\infty} := (u_{(p^f-1)n+r})_{n=0}^{\infty}$ is not modular. Since $(w_n)_{n=0}^{\infty}$ is not modular, iteratively applying Observation 2.2.1 we obtain a sequence of numbers $r_0, r_1, \ldots$ such that $r_k \in \{0, \ldots, p^k - 1\}$, $r_k \equiv r_{k+1} \pmod{p^k}$, and $(w_{p^k n + r_k})_{n=0}^{\infty}$ is not modular. By non-modularity of these subsequences, for all $k \in \mathbb{N}$, there exists an $n_k \in r_k + p^k \mathbb{Z}_{\geq 1}$ such that $w_{n_k} \equiv 0 \pmod{p^k}$.

As each of these subsequences $(w_{p^k n + r_k})_{n=0}^{\infty}$ is simple and non-modular, each contains a zero term by the Skolem conjecture. Since the original sequence $(u_n)_{n=0}^{\infty}$ has only finitely many zeros by the Skolem-Mahler-Lech theorem (Theorem 2.1.2), there exists $x \in \mathbb{Z}$ such that $w_x = 0$ and $n_k \equiv x \pmod{p^k}$ for infinitely many $k \geq 1$. The characteristic roots of $(w_n)_{n=0}^{\infty}$ are $\lambda_1^{p^f-1}, \ldots, \lambda_k^{p^f-1}$ and so the $\mathfrak{p}$-dominant roots of $(w_n)_{n=0}^{\infty}$ are exactly the $(p^f - 1)$th powers of the $\mathfrak{p}$-dominant roots of $(u_n)_{n=0}^{\infty}$. Since $\mathfrak{p}$ divides the non-$\mathfrak{p}$-dominant roots and $n_k \geq p^k \geq k$, we have that $\nu_{\mathfrak{p}}(w_{n_k} - w_{n_k}^{\mathfrak{p}}) \geq (p^f - 1)k$. Thus, for all $k \geq 1$,

$$0 = w_{n_k} \equiv w_{n_k}^{\mathfrak{p}} \pmod{\mathfrak{p}^k}.$$

Since $\mathfrak{p}$ does not divide the $\mathfrak{p}$-dominant roots of $(u_n)_{n=0}^{\infty}$ and since $\mathcal{O}_K/\mathfrak{p}$ is a field of order p^f, all $\mathfrak{p}$-dominant roots $\lambda_i^{p^f-1}$ of $(w_n)_{n=0}^{\infty}$ lie in $1 + \mathfrak{p}$. As $n_k \equiv x \pmod{p^k}$ for infinitely many k, we have that $\lambda_i^{(p^f-1)n_k} \equiv \lambda_i^{(p^f-1)x} \pmod{\mathfrak{p}^k}$ holds for infinitely many k. Thus, $w_x^{\mathfrak{p}} \equiv 0 \pmod{\mathfrak{p}^k}$ for infinitely many k and so $w_x^{\mathfrak{p}} = 0$.

We conclude that $(w_n)_{n=0}^{\infty}$ and $(w_n^{\mathfrak{p}})_{n=0}^{\infty}$ share a zero and hence so do $(u_n)_{n=0}^{\infty}$ and $(u_n^{\mathfrak{p}})_{n=0}^{\infty}$, contradicting the hypothesis. We conclude that $(u_n)_{n=0}^{\infty}$ is modular. □

We generalize this proposition to a statement about all prime ideals.

Theorem 2.2.3. *Let $(u_n)_{n=0}^{\infty}$ be a simple LRS and assume the Skolem conjecture. Then $(u_n)_{n=0}^{\infty}$ is modular if and only if for all $n \in \mathbb{Z}$ for which $u_n = 0$ there exists a prime ideal $\mathfrak{p}$ with $u_n^{\mathfrak{p}} \neq 0$ that does not contain all characteristic roots.*

Proof. Because we can decompose $(u_n)_{n=0}^{\infty}$ into non-degenerate subsequences using Lemma 1.2.3, we can assume that $(u_n)_{n=0}^{\infty}$ is non-degenerate. It also suffices to consider the case that $(u_n)_{n=-\infty}^{\infty}$ is not identically zero and hence, by the Skolem-Mahler-Lech theorem, has finitely many zeros.

Assume the 'only if' condition holds. Then there is an $\ell \geq 1$ such that the LRS $(u_n^{(i)})_{n=-\infty}^{\infty} := (u_{n\ell+i})_{n=-\infty}^{\infty}$ contains at most one zero for all $0 \leq i \leq \ell - 1$. Let $0 \leq i \leq \ell - 1$. If $(u_n^{(i)})_{n=-\infty}^{\infty}$ contains no zero, then by the Skolem conjecture, there exists an $M_i \geq 1$ such that $(u_n^{(i)} \bmod M_i)_{n=-\infty}^{\infty}$ is non-zero. Otherwise $(u_n^{(i)})_{n=-\infty}^{\infty}$ contains a zero—say $u_{n'}^{(i)} = 0$. By construction, $u_n^{(i)} \neq 0$ for all $n \in \mathbb{Z} \setminus \{n'\}$. By assumption, $u_{n'}^{\mathfrak{p}} \neq 0$ for some prime ideal $\mathfrak{p}$ for which Proposition 2.2.2 applies. Thus $(u_n^{(i)})_{n=0}^{\infty}$ is modular and there is an $M_i \geq 1$ such that $(u_n^{(i)} \bmod M_i)_{n=-\infty}^{\infty}$ has finitely many zeros. Letting $M = \operatorname{lcm}(M_0, \ldots, M_{\ell-1})$, we have that $(u_n \bmod M)_{n=0}^{\infty}$ is zero finitely often and so $(u_n)_{n=-\infty}^{\infty}$ is indeed modular.

Now, assume that the 'only if' condition does not hold. Then, there exists $n \in \mathbb{Z}$ such that $u_n = u_n^{\mathfrak{p}} = 0$ for all prime ideals $\mathfrak{p}$ that do not contain all characteristic roots. Shifting the sequence, we can assume that $n = 0$. We want to show that for $M \geq 1$, we have $u_n \equiv 0 \pmod{M}$ for infinitely many $n \in \mathbb{N}$. Assume, for a rational prime p, that p divides M exactly $\ell \geq 1$ times. We claim that there is an $S_{p^\ell} \geq 1$ such that $u_{nS_{p^\ell}} \equiv 0 \pmod{p^\ell}$ for all $n \in \mathbb{N}$.

As the extension $K : \mathbb{Q}$ is Galois, the $\mathfrak{p}$-valuation of the $\mathfrak{p}$-dominant roots is equal for all prime ideals $\mathfrak{p}$ above p. In particular, they are all zero or bigger than zero. Assume $\mathfrak{p}$ has ramification index e. Then $(\nu_{\mathfrak{p}}(u_n^{\mathfrak{p}}))_{n=0}^{\infty}$ is purely periodic modulo $\mathfrak{p}^{e\ell}$, say of period R, and so $u_{Rn}^{\mathfrak{p}} \equiv 0 \pmod{\mathfrak{p}^{e\ell}}$ for all $n \geq 0$ as $u_0^{\mathfrak{p}} = 0$. As $\nu_{\mathfrak{p}}(u_n - u_n^{\mathfrak{p}})$ grows linearly, for large enough n, say larger than P, we have that $\nu_{\mathfrak{p}}(u_n^{\mathfrak{p}}) \geq e\ell$. Hence, taking S_{p^ℓ} to be a multiple of R larger than P we have that $\nu_p(u_{Sn}) = \frac{1}{e}\nu_{\mathfrak{p}}(u_n^{\mathfrak{p}} + (u_n - u_n^{\mathfrak{p}})) \geq \ell$. This proves our claim. Let S be the least common multiple of all S_{p^ℓ} such that p^ℓ divides M. Then $u_{Sn} \equiv 0 \pmod{m}$ for all $n \geq 0$ and so $(u_n)_{n=0}^{\infty}$ is not modular. □

Because the Skolem problem is straightforward for modular LRS, we naturally want to decide whether one can determine whether a given $\mathbb{Q}$-LRS is modular.

Theorem 2.2.4. *For simple LRS $(u_n)_{n=0}^{\infty}$, there is a procedure to determine whether $(u_n)_{n=0}^{\infty}$ is modular assuming the Skolem conjecture.*

Proof. Decompose the LRS into non-degenerate subsequences (Lemma 1.2.3). If a subsequence is constantly zero, the LRS $(u_n)_{n=0}^{\infty}$ is not modular. Otherwise, perform the following on all subsequences.

Run two semi-algorithms in parallel. For the first semi-algorithm, search for a witness $M \geq 1$ that the given subsequence is modular. If one finds such an M, return that this subsequence of $(u_n)_{n=0}^{\infty}$ is modular. For the second semi-algorithm, iterate through the values of $(u_n)_{n=0}^{\infty}$ until one finds $n \in \mathbb{Z}$ such that $u_n = 0$. Having found such an n, one searches for a prime ideal $\mathfrak{p}$ that divides at least one but not all characteristic roots and such that $u_n^{\mathfrak{p}} = 0$. (The required check can be done in finite time since any prime ideal that lies above some characteristic root lies above a rational prime divisor of c_d, of which there are only finitely many.) If such a prime ideal is found, apply Theorem 2.2.3 to conclude that $(u_n)_{n=0}^{\infty}$ is not modular and halt. Otherwise, search for the next value of n such that $u_n = 0$.

By Theorem 2.2.3, one of these two semi-algorithms will terminate for each subsequence, and if the first semi-algorithm terminates for all subsequences, $(u_n)_{n=0}^{\infty}$ is modular. □

We move on to the Skolem problem for LRS of order 5. First, we introduce a technical lemma, which will simplify the problem slightly by reducing to the case of LRS for which all $\mathfrak{p}$-dominant roots have $\mathfrak{p}$-valuation 0.

Lemma 2.2.5. *The Skolem problem reduces to the special case in which the $\mathfrak{p}$-dominant roots of an LRS have $\mathfrak{p}$-valuation 0 for all prime ideals $\mathfrak{p}$ in the ring of integers of its splitting field. Moreover, this reduction preserves simplicity and does not increase the order of the LRS.*

Proof. The result trivially holds for the zero-LRS, so we can assume that the LRS $(u_n)_{n=0}^{\infty}$ satisfying (1) is not zero. As $c_d \neq 0$, only finitely many prime ideals divide $\pm c_d$. Hence, there are only finitely many prime ideals $\mathfrak{p}$ for which some characteristic root has a positive p-adic valuation. For such a prime ideal $\mathfrak{p}$ of ramification index $e_{\mathfrak{p}}$, there is a natural number $d_{\mathfrak{p}}$ such that $\nu_{\mathfrak{p}}(\lambda)d_{\mathfrak{p}}$ is divisible by $e_{\mathfrak{p}}$ for all $\mathfrak{p}$-dominant roots λ. Then let d be the least common multiple of all these $d_{\mathfrak{p}}$ and $S = \prod_{p|c_d} p^{\nu_{\mathfrak{p}}(\lambda)d/e_{\mathfrak{p}}}$. Then, by construction, for $0 \leq i < d-1$, we have that $(u_{dn+i}/S^n)_{n=0}^{\infty}$ is a $\mathbb{Q}$-LRS such that for all prime ideals $\mathfrak{p}$, the $\mathfrak{p}$-dominant roots have $\mathfrak{p}$-valuation 0. Hence, the sequence $(Tu_{dn+i}/S^n)_{n=0}^{\infty}$ is a $\mathbb{Z}$-LRS for some $T \geq 1$. Now solve the Skolem problem for the $\mathbb{Z}$-LRS $(Tu_{dn+i}/S^n)_{n=0}^{\infty}$ for $0 \leq i \leq d-1$.

If $\lambda_1, \ldots, \lambda_k$ are the characteristic roots of $(u_n)_{n=0}^{\infty}$ with respective multiplicities $m_1, \ldots, m_k$, then $\lambda_1^d/S, \ldots, \lambda_k^d/S$ also have respective multiplicities in $(Tu_{dn+i}/S^n)_{n=0}^{\infty}$ for $0 \leq i \leq d-1$. As these are possible characteristic roots of $(Tu_{dn+i}/S^n)_{n=0}^{\infty}$, the order does not increase, but when $(u_n)_{n=0}^{\infty}$ is degenerate, some λ_j^d/S might be equal. □

We can now prove the main result of this section, which resolves Theorem 2.1.9.

Lemma 2.2.6. *Let $(u_n)_{n=0}^{\infty}$ be a non-degenerate LRS of order at most 5 LRS outside of the MSTV class. Then $(u_n)_{n=0}^{\infty}$ has the exponential-polynomial form*

$$u_n = \alpha_1\lambda_1^n + \overline{\alpha_1}\,\overline{\lambda_1}^n + \alpha_2\lambda_2^n + \overline{\alpha_2}\,\overline{\lambda_2}^n + b\rho^n\,, \tag{2.3}$$

where $\alpha_1, \lambda_1, \alpha_2, \lambda_2 \in \overline{\mathbb{Q}}^*$ and $b, \rho \in \mathbb{R} \cap \overline{\mathbb{Q}}^*$.

Proof. Being outside the MSTV class, the LRS $(u_n)_{n=0}^{\infty}$ has at least 5 characteristic roots (and thus precisely 5 that are all simple), of which at least four are dominant. If there are two dominant real roots, their quotient would be ± 1, contradicting non-degeneracy. Thus, there are at least three non-real dominant roots and, as the set of characteristic roots is closed under complex conjugation, there are two pairs of complex conjugate dominant roots, say λ_1, $\overline{\lambda_1}$, and λ_2, $\overline{\lambda_2}$. The remaining root ρ is real (as its complex conjugate is among these five roots), and again, by the Galois closure of the polynomial exponential form, the polynomial coefficient of ρ, b, is also real. As all characteristic roots and polynomial coefficients are non-zero, the result follows. □

Proof of Lemma 2.1.9. Let $(u_n)_{n=0}^{\infty}$ be an LRS of order at most 5. We can assume that $(u_n)_{n=0}^{\infty}$ is non-degenerate by Lemma 1.2.3 and lies outside the MSTV class as the Skolem problem is decidable for such LRS. By Lemma 2.2.6, $(u_n)_{n=0}^{\infty}$ has the exponential-polynomial form (2.3). By Lemma 2.2.5, we can assume that for all prime ideals $\mathfrak{p}$ of $\mathcal{O}_K$, the $\mathfrak{p}$-dominant roots have $\mathfrak{p}$-valuation 0.

By Theorem 3.1.6, $(u_n)_{n=0}^{\infty}$ is not reversible and hence λ_1 is not a unit. Thus, for some prime ideal $\mathfrak{p}$, we have $\overline{\lambda_1} \in \mathfrak{p}$. As $\lambda_1\overline{\lambda_1} = \lambda_2\overline{\lambda_2}$, we have $\lambda_2 \in \mathfrak{p}$ or $\overline{\lambda_2} \in \mathfrak{p}$, and without loss of generality, $\overline{\lambda_2} \in \mathfrak{p}$. As $(u_n)_{n=0}^{\infty}$ is not in the MSTV class, at least three characteristic roots are $\mathfrak{p}$-dominant and have thus $\mathfrak{p}$-valuation 0. Thus, ρ, λ_1, and λ_2 have $\mathfrak{p}$-valuation 0. Therefore,

$$u_n^{\mathfrak{p}} = \alpha_1\lambda_1^n + \alpha_2\lambda_2^n + b\rho^n\,.$$

If $u_n = u_n^{\mathfrak{p}} = 0$ for some $n \in \mathbb{Z}$, then $0 = u_n - u_n^{\mathfrak{p}} - \overline{u_n^{\mathfrak{p}}} = -b\rho^n$, implying $b = 0$ or $\rho = 0$, which are both not allowed. As such, by Proposition 2.2.2, the LRS $(u_n)_{n=0}^{\infty}$ is modular.

Thus, we can search for an $M \geq 1$ and $N \geq 0$ such that $u_n \not\equiv 0 \pmod{M}$ for all $n \geq N$ and check whether $u_n = 0$ for some $0 \leq n \leq N$ to solve the Skolem problem. □

The result cannot be extended beyond order 5. First, there are non-simple order-6 LRS outside the MSTV class (and so the Skolem conjecture does not apply). Secondly, there are non-modular simple order-6 LRS outside of the MSTV class.

Example 2.2.7. Write

$$\lambda_1 = 1 + 2i,\ \lambda_2 = \frac{3}{2} + \frac{1}{2}\sqrt{-11},\ \lambda_3 = \frac{1}{2} + \frac{1}{2}\sqrt{-19}\,,$$

and let

$$u_n = \lambda_1^n + \overline{\lambda_1}^n + \lambda_2^n + \overline{\lambda_2}^n - 2\lambda_3^n - 2\overline{\lambda_3}^n\,.$$

Equivalently, $(u_n)_{n=0}^{\infty}$ is defined by the linear recurrence

$$u_{n+6} = 6u_{n+5} - 26u_{n+4} + 66u_{n+3} - 130u_{n+2} + 150u_{n+1} - 125u_n$$

and the initial values $0, 3, 11, -12, -125, -177$ for $n = 0, \ldots, 5$. Showing that $(u_n)_{n=0}^{\infty}$ is non-degenerate is straightforward. We prove that $(u_n)_{n=0}^{\infty}$ is not modular and does not belong to the MSTV class.

We start with the latter. Let $\mathcal{O}_K$ be the ring of integers of the splitting field of $(u_n)_{n=0}^{\infty}$. All six characteristic roots are dominant with absolute values $\sqrt{5}$. If $\mathfrak{p}$ is a prime ideal in $\mathcal{O}_K$, all roots are $\mathfrak{p}$-dominant when $\mathfrak{p}$ is not above 5. When $\mathfrak{p}$ is above 5 and $i = 1, 2, 3$, exactly one of λ_i and $\overline{\lambda_i}$ is in $\mathfrak{p}$. Thus, there are three $\mathfrak{p}$-dominant roots, and $(u_n)_{n=0}^{\infty}$ is not in the MSTV class.

By the previous remarks, we have that $(u_n)_{n=0}^{\infty} = (u_n^{\mathfrak{p}})_{n=0}^{\infty}$ when a prime ideal $\mathfrak{p}$ is not above 5. When $\mathfrak{p}$ is above (5), exactly one of λ_i and $\overline{\lambda_i}$ is in $\mathfrak{p}$ for $i = 1, 2, 3$. Thus,

$$u_0 = 1 + 1 + 1 + 1 - 2 - 2 = 0 = 1 + 1 - 2 = u_0^{\mathfrak{p}}\,.$$

We conclude that $(u_n)_{n=0}^{\infty}$ is not modular using Theorem 2.2.3. □

2.3 From bi-infinite sequences to sequences

In this section, we discuss the relationship between the Bi-Skolem problem and the Skolem problem and prove Theorems 2.1.6 and 2.1.8. To solve the fragments of the Skolem problem, we will provide a method to compute the set

$$Z = \{n \in \mathbb{Z} \colon u_n = 0\}$$

of zeros of a given non-degenerate LRBS $(u_n)_{n=-\infty}^{\infty}$, which is finite by the Skolem-Mahler-Lech theorem. To solve the Skolem problem, it is sufficient to compute whether Z has a positive entry.

We first show how to compute Z given access to two abstract subroutines. Then, we implement these subroutines assuming the p-adic Schanuel conjecture and Skolem conjecture, respectively.

Proposition 2.3.1. *Let $Z \subseteq \mathbb{Z}$ be a finite recursive set. Then Z can be computed by the following two subroutines:*

- *Subroutine 1: subsequence search. For $a \in \mathbb{Z}$ and $b \geq 1$, return* False *if the set $Z \cap \{bn+a \colon n \in \mathbb{Z}\}$ is empty and else return an element from $Z \cap \{bn+a \colon n \in \mathbb{Z}\}$;*
- *Subroutine 2: leapfrogging. For $z \in Z$, return $S \geq 1$ such that z is the only element in $Z \cap \{z + Sn \colon n \in \mathbb{Z}\}$.*

We postpone the proof of the proposition. First, we explain our application of this result and give an example. Recall that an *arithmetic progression* is a set $\{cn+d \colon n \in \mathbb{Z}\}$ for some $c, d \in \mathbb{Z}$ with $c \geq 1$.

In our application of Proposition 2.3.1, the set Z has the form $\{n \in \mathbb{Z} \colon u_n = 0\}$. Subroutine 1 is thus equivalent to determining whether $u_{bn+a} = 0$ for some $n \in \mathbb{Z}$. If such an n exists, it is returned, and False is returned otherwise. Therefore, we call this routine the *subsequence search*. Determining whether there exists n such that $u_{bn+a} = 0$ is straightforward as $\mathbb{Q}$-LRBS are enumerable. Meanwhile, the Skolem conjecture allows us to certify that $u_{bn+a} \neq 0$ for all $n \in \mathbb{Z}$: we find a number M such that u_{bn+a} is never congruent to 0 modulo M. The Skolem conjecture ensures such an M exists.

For Subroutine 2, we want to isolate one zero of the LRBS. If we discover $z \in \mathbb{Z}$ such that $u_z = 0$ then we find an arithmetic progression $\{Sn + z \colon n \in \mathbb{Z}\}$ such that $u_{Sn+z} = 0$ implies that $n = 0$. By this choice of S, we 'leapfrog' over all other zeros in the sequence, hence we name this subroutine *leapfrogging*. Implementation of this subroutine relies on the p-adic Schanuel conjecture.

Example 2.3.2. Let $(u_n)_{n=-\infty}^{\infty}$ be the LRBS defined by $(u_0, u_1, u_2) = (0, 3, 1)$ and the recurrence $u_{n+3} = 2u_{n+2} - u_n$ for all $n \in \mathbb{Z}$. Then, $u_n = -F_{n-3} - 2$, where F_k is the kth Fibonacci number (which are defined in Example 1.2.4).

To compute $Z = \{n \in \mathbb{Z} \colon u_n = 0\}$, we first apply Subroutine 1 (subsequence search) to discover that $u_0 = 0$. Using Subroutine 2 (leapfrogging), we obtain $S = 8$ such that $u_{Sn} = 0$ if and only if $n = 0$. Now we apply our method recursively on the remaining seven subsequences:

$$(u_{8n+1})_{n=-\infty}^{\infty}, (u_{8n+2})_{n=-\infty}^{\infty}, \ldots, (u_{8n+7})_{n=-\infty}^{\infty}.$$

Using the subsequence search, we find that $(u_{8n+1})_{n=-\infty}^{\infty}$ is always non-zero modulo 4. Similarly, $(u_{8n+2})_{n=-\infty}^{\infty}$, $(u_{8n+3})_{n=-\infty}^{\infty}$, $(u_{8n+4})_{n=-\infty}^{\infty}$, $(u_{8n+5})_{n=-\infty}^{\infty}$, $(u_{8n+7})_{n=-\infty}^{\infty}$ are all non-zero modulo 3. Hence, these six subsequences cannot contain 0. For $(u_{8n+6})_{n=-\infty}^{\infty}$, such a modulus does not exist as the subroutine will find that $u_6 = 0$.

Using the leapfrogging method on the LRBS $(u_{8n+6})_{n=-\infty}^{\infty}$, we find $S = 3$. Thus, $u_{8 \cdot S \cdot n+6} = u_{24n+6} = 0$ if and only if $n = 0$. The only subsequences which could still contain unknown zeros are $(u_{8 \cdot 3 \cdot n+8+6})_{n=-\infty}^{\infty}$ and $(u_{8 \cdot 3 \cdot n+2 \cdot 8+6})_{n=-\infty}^{\infty}$, for which the subsequence search shows that their terms are never zero modulo 2: They are always odd and thus never 0.

In this way, we can conclude that $Z = \{0, 6\}$. □

Now we will prove Proposition 2.3.1.

Proof of Proposition 2.3.1. We use a recursive method.

Assume we have to compute $Z \cap \{bn + a \colon n \in \mathbb{Z}\}$ for some given integers a and $b \geq 1$. Then, using Subroutine 1 (subsequence search), we either conclude that $Z \cap \{bn + a \colon n \in \mathbb{Z}\}$ is empty or find a specific z such that $bz + a \in Z$. In the first case, we are done. In the second case, we use Subroutine 2 (leapfrogging) to compute an $S \geq 1$ such that $Sbn + bz + a \in Z$ if and only if $n = 0$. Hence, we found all indices n where $Sbn + (bz + a) \in Z$ and only need to focus on the arithmetic progressions

$$\{Sbn + (b + bz + a) \colon n \in \mathbb{Z}\}, \ldots, \{Sbn + ((S-1)b + bz + a) \colon n \in \mathbb{Z}\}.$$

We apply this recursive approach to each of these arithmetic progressions.

Starting from $\mathbb{Z}$, we partition $\mathbb{Z}$ into finitely many arithmetic progressions when encountering an element from Z and refine this partition every time we encounter an element in Z. Because we refine our partition only once for every element in the finite set Z and Z is finite, this recursive method ultimately terminates, giving us finitely many arithmetic progressions that form a partition of $\mathbb{Z}$ such that for each, we either know they contain no element from Z or exactly one element of Z (which we know). Therefore, we can compute Z. □

2.3.1 The subsequence search

As mentioned earlier, the Skolem conjecture will be our main tool for Subroutine 1.

Proposition 2.3.3. *For simple $\mathbb{Z}$-LRBS, a procedure for Subroutine 1 exists that only depends on the Skolem conjecture for termination.*

Proof. Let $(u_n)_{n=-\infty}^{\infty}$ be an LRBS, which is the bi-completion of a simple $\mathbb{Q}$-LRS $(u_n)_{n=0}^{\infty}$. Then, using Lemma 1.2.5, we can assume that $(u_n)_{n=0}^{\infty}$ is a $\mathbb{Z}$-LRS. We have to show that for all $a \in \mathbb{Z}$ and $b \geq 1$, we can show whether the subsequence $(u_{bn+a})_{n=-\infty}^{\infty}$ contains a zero term. As each LRBS $(u_{bn+a})_{n=-\infty}^{\infty}$ is again simple, it is sufficient to show the statement for $(u_n)_{n=-\infty}^{\infty}$. Let c_d be the constant coefficient of the characteristic polynomial of $(u_n)_{n=-\infty}^{\infty}$. To do this, run two semi-algorithms in parallel:

- Semi-algorithm 1: Enumerate $M = 1, 2, 3, 4, \ldots$ until $\gcd(M, c_d) = 1$ and $(u_n \bmod M)_{n=-\infty}^{\infty}$ does not contain 0. Then return False.
- Semi-algorithm 2: Enumerate the sequence $u_0, u_1, u_{-1}, u_2, u_{-2}, u_3, \ldots$ until $u_z = 0$ holds. Then return z.

Semi-algorithm 1 exists because one can compute the integer c_d and for each M, the LRBS $(u_n \bmod M)_{n=-\infty}^{\infty}$ is periodic, and this period can be computed due to Lemma 1.2.1. So if one computes the LRBS modulo M and finds an $n \geq 1$ such that $u_0 \equiv u_n \pmod{M}$,...,$u_{d-1} \equiv u_{n+d-1} \pmod{M}$, then the period of the LRBS modulo M divides n. Then, if $u_0, \ldots, u_{n-1}$ are all non-zero modulo M, the LRBS $(u_n)_{n=-\infty}^{\infty}$ does not contain 0.

Semi-algorithm 2 exists because an LRBS is recursive.

If Semi-algorithm 1 does not terminate, the Skolem conjecture (Conjecture 2.1.7) implies that $(u_n)_{n=-\infty}^{\infty}$ contains a zero, say $u_z = 0$ for some $z \in \mathbb{Z}$. In that case, Semi-algorithm 2 will terminate. If Semi-algorithm 1 does terminate, the LRBS $(u_n)_{n=-\infty}^{\infty}$ does not contain any number congruent to 0 modulo M, and so certainly has no zero term.

Thus, assuming the Skolem conjecture, exactly one of the semi-algorithms terminates, giving Subroutine 1. If the Skolem conjecture is false for a particular instance, both semi-algorithms will fail to terminate. □

In our proof, we assumed the Skolem conjecture in full generality to prove that the procedure for Subroutine 1 terminates. This can be made more precise. Let $\mathcal{C}$ be a subset of LRBS. Then call $\mathcal{C}$ *closed* if for all $(u_n)_{n=-\infty}^{\infty} \in \mathcal{C}$, $b \in \mathbb{Z}_{\geq 1}$, and $0 \leq a \leq b-1$, we have that $(u_{bn+a})_{n=-\infty}^{\infty}$ is in $\mathcal{C}$. These closed subsets form a topology. In particular, the union and intersection of closed subsets of $\mathbb{Q}$-LRBS are closed. Many natural classes of LRBS are closed: non-degenerate LRBS, simple LRBS, LRBS of order at most d, LRBS with a characteristic root that is a power of 2, etc.

If the Skolem conjecture holds for a class $\mathcal{C}$ of LRBS and one also has a leapfrogging procedure for all $\mathbb{Q}$-LRBS in $\mathcal{C}$, then the Skolem problem is decidable for all LRS whose bi-completion is in $\mathcal{C}$.

The Skolem conjecture does not hold for some non-simple LRS such as $u_n = (2n+1)2^n$. In this case, $u_n \neq 0$ for all $n \in \mathbb{Z}$ but $(u_n)_{n=0}^{\infty}$ is also not modular: If $M = 2^t s$ for some odd number s, then there are infinitely many $n \in \mathbb{N}$ such that $s \mid (2n+1)$ while $2^t \mid 2^n$ for large enough n. Thus, there are infinitely many $n \in \mathbb{N}$ such that $u_n \equiv 0 \pmod{M}$.

2.3.2 Leapfrogging

In this section, we discuss various methods to implement Subroutine 2 (leapfrogging) and obtain Theorem 2.1.6 as a corollary.

Recall that for a given $z \in Z$, we want to compute $S \geq 1$ such that z is the only element in $Z \cap \{Sn + z \colon n \in \mathbb{Z}\}$. In the context of LRBS, this means that for a given LRS $(u_n)_{n=0}^{\infty}$ such that $u_z = 0$ for some $z \in \mathbb{Z}$, we have to compute $S \geq 1$ such that $u_{nS+z} = 0$ if and only if $n = 0$. By shifting the sequence to $(u_{n-z})_{n=-\infty}^{\infty}$, we can assume that $z = 0$.

Using Lemma 3.5.1, we can assume that $(u_n)_{n=-\infty}^{\infty}$ is non-degenerate and non-zero. Let $(u_n)_{n=0}^{\infty}$ satisfy the order-d linear recurrence (1) ($u_{n+d} = c_1 u_{n+d-1} + \cdots + c_d u_n$) and has characteristic polynomial P and characteristic roots $\lambda_1, \ldots, \lambda_d$. As $(u_n)_{n=0}^{\infty}$ is integer-valued, the numbers $u_0, \ldots, u_{d-1}, c_1, \ldots, c_d$ are all integers.

Let us sketch our method. For all our methods, we rely on p-adic numbers and the p-adic valuation. In particular, we find a prime number p, integers a and b, and $S \geq 1$ such that

$$\nu_p(u_{Sn}) = a + b\nu_p(n)\,.$$

In that case, we have the following implications:

$$\begin{aligned} u_{Sn} = 0 &\iff \nu_p(u_{Sn}) = +\infty \\ &\iff a + b\nu_p(n) = +\infty \\ &\iff \nu_p(n) = +\infty \\ &\iff n = 0. \end{aligned}$$

Hence, Sn is in Z if and only if $n = 0$. This argument is exactly the leapfrogging subroutine.

To compute S, a, and b, we construct a p-adic power series $f(X) = \sum_{j=0}^{\infty} a_j X^j$. For the sake of exposition, assume that p does not divide c_d such that we can use the

p-adic valuation in some field extension K_p of $\mathbb{Q}_p$. We will compute a natural number $L \geq 1$ such that $f(Ln) = u_{Ln}$ for all $n \in \mathbb{Z}$ with computable p-adic coefficient a_j such that $\nu_p(a_j) \geq r$ for some $r \in \mathbb{Z}$.

Let b be the smallest number such that $a_b \neq 0$. Then, for a computable number t, we have that

$$\nu_p(a_b \cdot (Lnp^t)^b) = \nu_p(a_b) + bt + b\nu_p(Ln) < r + jt + j\nu_p(Ln) \geq \nu_p(a_j \cdot (Lnp^t)^j)$$

for all $j > b$ and $n \in \mathbb{Z} \setminus \{0\}$. Therefore, if we take $S = Lp^t$, we have that

$$\nu_p(u_{Sn}) = \nu_p(a_b(np^t)^b) = \nu_p(a_b) + bt + \nu_p(n^b) = a + b\nu_p(n)\,.$$

We will show how to compute the step size S for each prime p. We exhibit multiple methods that differ in practicality and generality.

A simple criterion

Assume that $(u_n)_{n=0}^{\infty}$ is simple, non-zero, and integer-valued. Let $\mathcal{D}$ be the discriminant of P and $\mathcal{C}$ the *circular determinant*, the determinant of the *circular matrix*:

$$\mathcal{C} = \det \begin{pmatrix} u_0 & u_1 & \cdots & u_{d-1} \\ u_1 & u_2 & \cdots & u_d \\ \vdots & \vdots & \ddots & \vdots \\ u_{d-1} & u_d & \cdots & u_{2d-2} \end{pmatrix}.$$

Because $(u_n)_{n=0}^{\infty}$ is simple and has minimal order d, the numbers $\mathcal{C}$ and $\mathcal{D}$ are non-zero integers.

Let $p \geq 3$ be a rational prime and L the period of $(u_n \bmod p)_{n=-\infty}^{\infty}$.

Theorem 2.3.4. *Assume that p does not divide c_d, $\mathcal{C}$, and $\mathcal{D}$ and that p^2 does not divide u_L. Then $u_{nL} \neq 0$ for all $n \in \mathbb{Z} \setminus \{0\}$.*

Proof. Let K_p be the splitting field of the characteristic polynomial of $(u_n)_{n=-\infty}^{\infty}$ over $\mathbb{Q}_p$. As $(u_n)_{n=-\infty}^{\infty}$ is simple, we can write the polynomial exponential form (1.2) as

$$u_n = \alpha_1\lambda_1^n + \cdots + \alpha_d\lambda_d^n$$

for some non-zero numbers $\alpha_i \in K_p$. As $(u_n)_{n=0}^{\infty}$ is a $\mathbb{Z}$-LRS, the numbers $\lambda_1, \ldots, \lambda_d$ are in $\mathcal{O}_{K_p}$. Next,

$$\begin{pmatrix} 1 & 1 & \cdots & 1 \\ \lambda_1 & \lambda_2 & \cdots & \lambda_d \\ \vdots & \vdots & \ddots & \vdots \\ \lambda_1^{d-1} & \lambda_2^{d-1} & \cdots & \lambda_d^{d-1} \end{pmatrix} \begin{pmatrix} \alpha_0 \\ \alpha_1 \\ \vdots \\ \alpha_{d-1} \end{pmatrix} = \begin{pmatrix} u_0 \\ u_1 \\ \vdots \\ u_{d-1} \end{pmatrix},$$

and the square matrix is a Vandermonde matrix with determinant $\pm\sqrt{\mathcal{D}}$. As $p \nmid \mathcal{D}$, we know that all the α_i are in $\mathcal{O}_{K_p}$. If any α_i were in $p\mathcal{O}_{K_p}$, then $(u_n \bmod p)_{n=-\infty}^{\infty}$ would have order less than d, and so the circular matrix would not be invertible modulo p. Hence, p divides $\mathcal{C}$, which we explicitly assumed not to be the case. Thus all α_i are in $\mathcal{O}_{K_p}^*$.

We claim that $\lambda_i^L \equiv 1 \pmod{p}$ for all $1 \le i \le d$. To see this, set $x_i = \alpha_i(\lambda_i^L - 1)$. Then,

$$\begin{pmatrix} 1 & 1 & \cdots & 1 \\ \lambda_1 & \lambda_2 & \cdots & \lambda_d \\ \vdots & \vdots & \ddots & \vdots \\ \lambda_1^{d-1} & \lambda_2^{d-1} & \cdots & \lambda_d^{d-1} \end{pmatrix} \begin{pmatrix} x_0 \\ x_1 \\ \vdots \\ x_{d-1} \end{pmatrix} = \begin{pmatrix} u_L - u_0 \\ u_{L+1} - u_1 \\ \vdots \\ u_{L+d-1} - u_{d-1} \end{pmatrix} \equiv \begin{pmatrix} 0 \\ 0 \\ \vdots \\ 0 \end{pmatrix} \pmod{p}$$

as $u_n \equiv u_{n+L} \pmod{p}$ for all $n \in \mathbb{Z}$. Again, as p does not divide $\mathcal{D}$, the square matrix is invertible modulo p, and so $x_i \equiv 0 \pmod{p}$. As $\alpha_i \in \mathcal{O}_{K_p}^*$, the claim follows.

Because $p \ge 3$, we have that $\nu_p(\log(\lambda_i^L)) \ge 1 > 1/p - 1$, and so the p-adic analytic function

$$f(z) = \sum_{i=1}^{d} \alpha_i \exp\left(z \log(\lambda_i^L)\right)$$

is well-defined on $\mathcal{O}_{K_p}$. Moreover, $f(n) = u_{Ln}$ for all $n \in \mathbb{Z}$. Setting

$$a_j = \frac{1}{j!} \sum_{i=1}^{d} \alpha_i \big(\log(\lambda_i^L)\big)^j ,$$

gives that $f(z) = \sum_{j=0}^{\infty} a_j z^j$, where $a_0 = u_0 = 0$.

By Legendre's formula [72, Chapter 5], we have that for $j \ge 1$,

$$\nu_p(j!) = \sum_{j=1}^{\infty} \left\lfloor \frac{j}{p^j} \right\rfloor < \sum_{j=1}^{\infty} \frac{j}{p^j} = \frac{j}{p-1} . \tag{2.4}$$

Therefore,

$$\nu_p(a_j) \ge \min_{1 \le i \le d} \big(\nu_p(\alpha_i) + j\,\nu_p(\log(\lambda_i^L))\big) - \nu_p(j!) > j - \frac{j}{p-1}$$

as $\nu_p(\log(\lambda_i^L)) \ge 1$. Thus, as $p \ge 3$, we have $\nu_p(a_1) \ge 1$ and $\nu_p(a_j) \ge 2$ for all $j \ge 2$. Only, $\nu_p(a_1) = 1$ as

$$0 \not\equiv U(N) \equiv f(1) \equiv \sum_{j=1}^{\infty} a_j \equiv a_1 \pmod{p^2} .$$

The result follows as now for any non-zero integer n,

$$\nu_p(u_{Ln}) = \nu_p(f(n)) = \nu_p(a_1 n) \ne +\infty . \qquad \square$$

Using the exponential function

In this subsection, we will give a more general method than the one described in the previous section. In particular, we will prove the following.

Proposition 2.3.5. *Let $(u_n)_{n=0}^{\infty}$ be a non-zero and non-degenerate $\mathbb{Z}$-LRS such that $u_0 = 0$ and let p be a rational prime. Then one can construct a power series $f(X) = \sum_{j=0}^{\infty} a_j X^j$ with computable coefficients a_j in a finite extension K_p of $\mathbb{Q}_p$ when $p \nmid c_d$ or when $p \mid c_d$ and $\mathfrak{p}$ is a prime ideal above p in $K_{\mathfrak{p}}$.*

We start with the easiest case: p does not divide c_d. We generalise our argument later. Let K_p be an extension of Q_p that contains all characteristic roots, and their polynomial coefficients Q_i are in $K_p[X]$.

Write $(u_n)_{n=0}^{\infty}$ in its polynomial exponential form (1.2): $u_n = \sum_{i=1}^{k} Q_i(n)\lambda_i^n$. By our earlier assumptions, each characteristic roots λ_i is in the ring of integers of K_p, $\mathcal{O}_{K_p}$, but not in $p\mathcal{O}_{K_p}$ as $p \nmid c_d$. By multiplying the sequence with an appropriate power of p, we can assume that all polynomial coefficients Q_i are in $\mathcal{O}_{K_p}[X]$. Now, compute $L \geq 1$ such that $\log(\lambda_i^L)$ is defined for all $1 \leq i \leq k$. Then,

$$u_{Ln} = \sum_{i=1}^{k} Q_i(Ln)\lambda_i^{Ln} = \sum_{i=1}^{k} Q_i(Ln)\exp(n\log(\lambda_i^L)) = \sum_{j=0}^{\infty}\left(\sum_{i=1}^{k} Q_i(Ln)\log(\lambda_i^L)^j\right)\frac{n^j}{j!} \tag{2.5}$$

gives the p-adic power series $f(X) = \sum_{j=0}^{\infty} a_j X^j$ we had to construct. Using the fact that $a_j = \frac{1}{j!}f^{(j)}(0)$, we can compute an explicit formula for the coefficients a_j using differentiation that also shows these numbers are computable:

$$a_j = \frac{1}{j!}\sum_{\ell=0}^{j}\binom{j}{\ell}L^{\ell}\sum_{i=1}^{d} Q_i^{(\ell)}(0)\log(\lambda_i^L)^{j-\ell}. \tag{2.6}$$

When the LRS is simple, say $u_n = \sum_{i=1}^{d}\alpha_i\lambda_i^n$, we have that $a_j = \frac{1}{j!}\sum_{i=1}^{d}\alpha_i\log(\lambda_i^L)^j$.

As the power series converges for all $n \in \mathbb{N}$, the p-adic valuation of the coefficients a_j tends to infinity. Using (2.4) and a lower bound on $\log(\lambda_i^L)^j$, we can give a linear lower bound on the p-adic valuation of a_j.

Thus, in this restricted case, we have proved Proposition 2.3.5.

For primes p that divide c_d, one has to be more careful. By computing S_+ and S_- such that $u_{S_+n} \neq 0$ and $u_{-S_-n} \neq 0$ for $n \geq 1$, we can take $S = \operatorname{lcm}(S_-, S_+)$. By symmetry, we consider S_+ as one can compute S_- analogously by considering the sequence $(u_{-n})_{n=0}^{\infty}$. Let $\mathfrak{p}$ be a prime ideal above p in the splitting field of f. By Lemma 2.2.5, we can assume that the $\mathfrak{p}$-dominant roots have $\mathfrak{p}$-valuation 0. Then

$\nu_{\mathfrak{p}}(u_n - u_n^{\mathfrak{p}})$ grows faster than some computable linear function. Then, again for a computable number L, the expression $\exp(n\log(\lambda_i^L))$ is well-defined for all $\mathfrak{p}$-dominant λ_i. Hence, a_j can be computed up to arbitrary precision modulo powers of $\mathfrak{p}$, giving the remainder of Proposition 2.3.5.

Using the companion matrix

Here, we present a different method to compute the coefficients a_j. This other method applies to general primes p, but we again begin by assuming that $p \nmid c_d$ and that p is odd. Recall that the *companion matrix* (1.3) of $(u_n)_{n=0}^{\infty}$ is defined as

$$A := \begin{pmatrix} c_1 & \cdots & c_{d-1} & c_d \\ 1 & \cdots & 0 & 0 \\ \vdots & \ddots & \vdots & \vdots \\ 0 & \cdots & 1 & 0 \end{pmatrix}.$$

If

$$\mathbf{s} = \begin{pmatrix} 0 & \cdots & 0 & 1 \end{pmatrix} \quad \text{and} \quad \mathbf{t} = \begin{pmatrix} u_{d-1} & \cdots & u_0 \end{pmatrix}^{\top}$$

then $u_n = \mathbf{s}A^n\mathbf{t}$ for all $n \in \mathbb{Z}$. Modulo p, the matrix A is reversible as $p \nmid c_d = \pm\det(A)$. Interpreting A as an element of $\mathrm{GL}_d(\mathbb{F}_p)$, the group of invertible matrices modulo p, there is a minimal $L \geq 1$ such that $A^L \equiv I_d \pmod{p}$ because the group of invertible matrices over the finite field $\mathbf{F}_p$ ($\mathrm{GL}_d(\mathbb{F}_p)$) is finite. Then, setting $B = \frac{1}{p}(A^L - I_d)$, we have that

$$\begin{aligned}
u_{Ln} &= \mathbf{s}A^{Ln}\mathbf{t} \\
&= \mathbf{s}(I_d + pB)^n\mathbf{t} \\
&= \sum_{\ell=0}^{n} \binom{n}{\ell} p^\ell \mathbf{s}B^k\mathbf{t} \\
&= \sum_{\ell=0}^{n} \frac{n(n-1)\cdots(n-\ell+1)}{k!} p^\ell \mathbf{s}B^\ell\mathbf{t} \\
&= \sum_{\ell=0}^{\infty} \frac{n(n-1)\cdots(n-\ell+1)}{\ell!} p^\ell \mathbf{s}B^\ell\mathbf{t} \\
&= \sum_{\ell=0}^{\infty}\sum_{j=0}^{\infty} c_{\ell,j} n^j \frac{p^\ell}{\ell!} \quad \text{for certain } c_{\ell,j} \in \mathbb{Z} \text{ with } c_{\ell,j} = 0 \text{ for } j > \ell \\
&= \sum_{j=0}^{\infty} \left(\sum_{\ell=j}^{\infty} c_{\ell,j} \frac{p^\ell}{\ell!} \right) n^j .
\end{aligned}$$

The last step is allowed by [72, Proposition 4.1.4] as (2.4) implies that $\nu_p(c_{\ell,j}p^\ell/\ell!) \geq \frac{(p-2)\ell}{p-1}$ and so the terms $c_{\ell,j}p^\ell/\ell!$ converge to 0 p-adically as j goes to infinity and uniformly to 0 as ℓ goes to infinity. Similarly, the double sum in the last line converges.

Therefore, the coefficients $a_j := \sum_{\ell=j}^{\infty} c_{\ell,j}\frac{p^\ell}{\ell!}$ of the power series $\widetilde{f}(X) = \sum_{n=0}^{\infty} a_j X^n$ are well-defined in $\mathbb{Z}_p$. Moreover, by the previous considerations, we have $\nu_p(a_j) \geq \nu_p(\frac{p^j}{j!}) \geq \frac{(p-2)j}{p-1}$ and so $\lim_{j\to\infty} \nu_p(a_j) = \infty$.

Next, if $p = 2$, then we search for a minimal $L \geq 1$ such that $A^L \equiv I_d \pmod{2^2}$, which exists as $\mathrm{GL}_d(\mathbb{Z}/4\mathbb{Z})$ is finite. The same computation applies, but now we get that for some integers $c_{\ell,j}$,

$$u_{Ln} = \sum_{\ell=0}^{\infty}\sum_{j=0}^{\infty} c_{\ell,j} n^j \frac{2^{2\ell}}{\ell!} = \sum_{j=0}^{\infty}\left(\sum_{\ell=j}^{\infty} c_{\ell,j}\frac{2^{2\ell}}{\ell!}\right) n^j\,.$$

Again, we claim that we can invert the summation, that $a_j := \sum_{\ell=j}^{\infty} c_{\ell,j}\frac{2^{2\ell}}{\ell!}$ gives a converging sum, and that $\lim_{j\to\infty} a_j = \infty$. This will follow in the same manner as in the odd case, but now $\nu_2(c_{\ell,j}\frac{2^{2\ell}}{\ell!}) \geq 2\ell - \frac{\ell}{2-1} = \ell$, which surely goes to infinity when ℓ does. Thus, $\nu_j(a_j) \geq \nu_j(2^{2j}/j!) \geq j$ and so $\lim_{j\to\infty} a_j = \infty$ as well. As such, our claim follows.

2.3.3 Computing the step

In the previous two sections, we discussed two methods to produce a power series: The exponential function gives a power series $f_{\exp}(X)$ and the companion matrix method gives $f_{\mathrm{comp}}(X)$. As $f_{\exp}(n) = u_n = f_{\mathrm{comp}}(X)$ for infinitely many numbers n, Strassman's theorem [72, Chapter 5] implies that $f_{\exp}(X) = f_{\mathrm{comp}}(X)$. Thus, the companion matrix method lets us compute a_j to arbitrary precision without using extensions of the p-adics, the p-adic logarithms, or p-adic exponential functions explicitly.

Thus, to compute S, we have a power series

$$f(X) = \sum_{j=0}^{\infty} a_j X^n\,, \tag{2.7}$$

where a_j is defined by (2.6) and can be computed up to arbitrary precision. For $j \geq 0$, one can thus compute a polynomial $F \in \overline{\mathbb{Q}}[X_1, \ldots, X_k]$ such that $a_j = F(\log(\lambda_1^L), \cdots, \log(\lambda_k^L))$. First, we reduce the number of variables with Theorem 1.1.3, which we apply to $\lambda_1, \ldots, \lambda_k$. With possible reordering, we can then compute $t \in \mathbb{N}$,

$s_{t+1}, \ldots, s_k \in \mathbb{Z}_{\neq 0}$, and $s_{i,j} \in \mathbb{Z}$ such that $\lambda_1, \ldots, \lambda_t$ are multiplicatively independent and

$$\lambda_1^{s_{i,1}} \cdots \lambda_t^{s_{i,t}} = \lambda_i^{s_i} .$$

We use these relationships to construct a polynomial $\widetilde{F} \in \overline{\mathbb{Q}}[X_1, \ldots, X_t]$ such that $F(\log(\lambda_1), \ldots, \log(\lambda_k)) = \widetilde{F}(\log(\lambda_1), \ldots, \log(\lambda_t))$.

When p does not divide c_d, the characteristic roots $\lambda_1, \ldots, \lambda_t$ lie in an finite extension of $\mathbb{Q}_p$, and so the (weak) p-adic Schanuel conjecture says that $\log(\lambda_1), \ldots, \log(\lambda_t)$ are algebraically independent over $\mathbb{Q}$. So either the polynomial F is the zero polynomial (which we can check), or a_j is non-zero.

When p divides the constant coefficient of the characteristic polynomial, we take a $\mathfrak{p}$-adic logarithm in the field $\mathbb{Q}_\mathfrak{p}$. To the best of our knowledge, no version of Schanuel's conjecture exists for such fields in the literature, although one can easily formulate one.

But we do not know whether the p-adic Schanuel conjecture is true! Yet, when the p-adic Schanuel conjecture implies that a_j is non-zero, we can still compute a_j to arbitrary precision. If a_j is truly non-zero, we will encounter a non-zero term in its p-adic expansion. Otherwise, this method will not terminate. As such, we have proven the following.

Theorem 2.3.6. *Using the notation above, there is a procedure that will return whether $a_j = 0$, which terminates subject to the p-adic Schanuel conjecture when p does not divide c_d.*

Therefore, we indeed have an algorithm to compute the step S, which gives us an algorithm for the leapfrogging procedure. We conclude Theorems 2.1.6 and Theorem 2.1.8.

Simple recurrences

Using Baker's theorem on linear forms in logarithms, we can remove the dependence on the p-adic Schanuel conjecture when determining zeroness of low-order coefficients of the power series (2.7). In particular, we have the following result.

Proposition 2.3.7. *For simple LRS, using the above notation, we can determine whether $a_0 = 0$ and whether $a_1 = 0$.*

Proof. As the LRS is simple, we can write it in its polynomial-exponential form $u_n = \sum_{i=1}^{k} \alpha_i \lambda_i^n$. When p does not divide c_d, we have $a_j = \sum_{i=1}^{k} \alpha_i \log(\lambda_i^L)^j$. When p divides c_d, we have that

$$a_j = \sum_{\substack{i=1 \\ \lambda_i \text{ is } \mathfrak{p}\text{-dominant}}}^{k} \alpha_i \log_{\mathfrak{p}}(\lambda_i^L)^j \,.$$

Hence, a_0 is always an algebraic number for which we can test whether it is zero. When $j = 1$, the proof of this proposition follows from Masser's result (Theorem 1.1.3), where we rewrite the linear combination of logarithms into one where all logarithms are linearly independent over $\mathbb{Q}$ and then apply one of the two versions of Baker's theorem on linear forms in logarithms for p-adic logarithms: Theorem 1.1.9 or 1.1.10. Hence, either the linear combination is trivially zero (as all coefficients are zero) or is non-zero.

The result follows. □

2.4 Simple low-order recurrences

For simple, low-order LRS, we can eliminate the dependence on the p-adic Schanuel conjecture and thus improve on Theorem 2.1.8. Namely, we want to prove Theorem 2.1.10: The Skolem problem is decidable for simple LRS up to order 7 under the assumption of the Skolem conjecture.

Thus, assume that $(u_n)_{n=0}^{\infty}$ is a simple, non-degenerate LRS satisfying $u_n = \sum_{i=1}^{k} \alpha_i \lambda_i^k$ with all α_i non-zero. Let c_d denote the constant coefficient of the characteristic polynomial of $(u_n)_{n=0}^{\infty}$.

By the discussion in the previous section, Theorem 2.1.10 follows from showing certain coefficients a_j of the power series can be zero-tested explicitly sufficient to show that we do not have to rely on the p-adic Schanuel conjecture to decide the Skolem problem.

As in the proof of Proposition 2.3.5, we only need to compute $S_- \geq 1$ and $S_+ \geq 1$ such that $u_{S_-,-n}$ and u_{S_+n} are both non-zero for all integers $n \geq 1$. By symmetry, it is sufficient to compute S_+.

By Theorem 2.1.3, we can assume that $(u_n)_{n=0}^{\infty}$ is not in the MSTV class. Hence, $(u_n)_{n=0}^{\infty}$ has at least four dominant roots and at least three $\mathfrak{p}$-dominant roots for each prime ideal $\mathfrak{p}$ of the ring of integers $\mathcal{O}_K$ of the splitting field of the LRS. By Lemma 2.2.5 we may assume that the $\mathfrak{p}$-dominant roots have $\mathfrak{p}$-valuation 0 for all prime ideals $\mathfrak{p}$ of $\mathcal{O}_K$. Hence, by Theorem 3.1.6, the characteristic roots are not all

units, and so there is a prime ideal $\mathfrak{p}$ above a rational prime p containing at least one dominant root but not all dominant roots.

Let q be a rational prime not dividing c_d (which exists as $c_d \neq 0$) and $\mathfrak{p}$ a prime ideal of $\mathcal{O}_K$. For clarity, we will write $\log_q$ (respectively, $\log_{\mathfrak{p}}$) instead of log for the q-adic logarithm (respectively, $\log_{\mathfrak{p}}$). For a prime ideal $\mathfrak{p}$, we write $a_j^{\mathfrak{p}}$ for the jth coefficient of the power series $f(X) = \sum_{j=0} a_j^{\mathfrak{p}} X^j$.

By splitting into subsequences, we can assume that $\log_q(\lambda_i)$ is defined for all $1 \leq i \leq k$ and that $\log_{\mathfrak{p}}(\lambda_i)$ is defined for all $\mathfrak{p}$-dominant roots λ_i. Using the formula (2.6), we have that

$$a_j^q = \sum_{i=1}^{k} \alpha_i \log_q(\lambda_i)^j \quad \text{and} \quad a_j^{\mathfrak{p}} = \sum_{\substack{i=1 \\ \lambda_i \text{ is } \mathfrak{p}\text{-dominant}}}^{k} \alpha_i \log_{\mathfrak{p}}(\lambda_i)^j \,. \tag{2.8}$$

Using multiple applications of Baker's theorem, we can conclude the following.

Lemma 2.4.1. *Assume $\mathfrak{p}$ is a prime ideal and q a prime not dividing c_d, then*

$$a_1^{\mathfrak{p}} = 0 \Longleftrightarrow \widetilde{a}_1^q := \sum_{\substack{i=1 \\ \lambda_i \text{ is } \mathfrak{p}\text{-dominant}}}^{k} \alpha_i \log_q(\lambda_i) = 0 \,.$$

Proof. If $a_1^{\mathfrak{p}} = 0$, Theorem 1.1.10 implies that the multiplicative relations of $\lambda_1, \ldots, \lambda_k$ are sufficient to show that $a_1^{\mathfrak{p}} = 0$. Hence, $\widetilde{a}_1^q = 0$ using the same relationships. The other direction follows similarly, using Theorem 1.1.9 instead. □

The observation above severely limits the possible cases.

Lemma 2.4.2. *Let $\mathfrak{p}$ and q as above and assume $(u_n)_{n=0}^{\infty}$ has d' characteristic roots that are $\mathfrak{p}$-dominant. Then,*

1. *if $d' = 3$, at least one of $a_0^{\mathfrak{p}}$, $a_1^{\mathfrak{p}}$, and $a_2^{\mathfrak{p}}$ is non-zero;*
2. *if $d' = k - 2$, at least one of $a_0^{\mathfrak{p}}$, a_0^q, $a_1^{\mathfrak{p}}$, and a_1^q is non-zero.*

Proof.

1. Assume $\lambda_1, \lambda_2, \lambda_3$ are $\mathfrak{p}$-dominant. Then (2.8) implies that

$$\begin{pmatrix} 1 & 1 & 1 \\ \log(\lambda_1) & \log(\lambda_2) & \log(\lambda_3) \\ \log(\lambda_1)^2 & \log(\lambda_2)^2 & \log(\lambda_3)^2 \end{pmatrix} \begin{pmatrix} \alpha_1 \\ \alpha_2 \\ \alpha_3 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix} .$$

   As $\alpha_1, \alpha_2, \alpha_3 \neq 0$, the square Vandermonde matrix has determinant 0. Thus, $\lambda_i = \lambda_j$ for some distinct $i, j \in \{1, 2, 3\}$, which contradicts that $(u_n)_{n=0}^{\infty}$ is non-degenerate.

2. If $d' = k-2$, assume λ_{k-1} and λ_k are not $\mathfrak{p}$-dominant and that $a_0^{\mathfrak{p}}$, a_0^q, $a_1^{\mathfrak{p}}$, and a_1^q are all zero. Then, $0 = a_0^q - a_0^{\mathfrak{p}} = \alpha_{k-1} + \alpha_k$, and $a_1^{\mathfrak{p}} = a_1^q = 0$ by Lemma 2.4.1. Thus, $\alpha_{k-1}\log_q(\lambda_{k-1}) + \alpha_k \log_q(\lambda_k) = 0$. Thus, $\log_q(\lambda_{k-1}) = \log_q(\lambda_k)$, again contradicting non-degeneracy. □

We are in the position to prove Theorem 2.1.10.

Proof of Theorem 2.1.10. Let q be a prime number that does not divide c_d and assume that $a_0^q = a_1^q = 0$. By Lemma 2.2.5, for any prime ideal $\mathfrak{p}$, the $\mathfrak{p}$-dominant roots are not in $\mathfrak{p}$.

As $(u_n)_{n=0}^{\infty}$ is outside of the MSTV class and non-degenerate, the LRS $(u_n)_{n=0}^{\infty}$ has at least four dominant roots. By non-degeneracy, at most one dominant root is real, and so there are at least two conjugate pairs of non-real dominant roots: $\lambda_1, \overline{\lambda_1}, \lambda_2$, and $\overline{\lambda_2}$ are dominant roots. We claim that these dominant roots are not units. Indeed, restricting $(u_n)_{n=0}^{\infty}$ to the LRS

$$v_n := \sum_{\substack{i=1, \\ \text{there is a Galois automorphism } \sigma \\ \text{such that } \sigma(\lambda_i) \in \{\lambda_1, \overline{\lambda_1}, \lambda_2, \overline{\lambda_2}\}}}^{k} \alpha_i \lambda_i^n$$

gives a reversible LRS with at least four dominant roots and order at most 7. This contradicts Theorem 3.3.2, and so these dominant roots are indeed not units.

Assume that $\mathfrak{p}$ is a prime ideal such that $\overline{\lambda_1} \in \mathfrak{p}$ and thus $\overline{\lambda_1}$ is not $\mathfrak{p}$-dominant. As $\lambda_1\overline{\lambda_1} = \lambda_2\overline{\lambda_2}$ is in $\mathfrak{p}$, without loss of generality, we have that $\overline{\lambda_2}$ is also in $\mathfrak{p}$ and thus not $\mathfrak{p}$-dominant. Assume that $a_0^{\mathfrak{p}} = a_1^{\mathfrak{p}} = a_2^{\mathfrak{p}} = a_0^{\overline{\mathfrak{p}}} = a_1^{\overline{\mathfrak{p}}} = 0$.

By Lemma 2.4.2, there are at least two characteristic roots that are not $\mathfrak{p}$ dominant, and the number of $\mathfrak{p}$-dominant roots is not three. Thus, there are at least three characteristic roots that are not $\mathfrak{p}$-dominant and at least four $\mathfrak{p}$-dominant roots. As there are at most $k = 7$ roots, we conclude that these lower bounds are sharp. Let the other three characteristic roots be λ_3, λ_4, and λ_5. Then exactly one of λ_1, λ_2, λ_3, λ_4, and λ_5 is not $\mathfrak{p}$-dominant, which by symmetry, is λ_1 or λ_3.

If λ_1 is not $\mathfrak{p}$-dominant, applying Lemma 2.4.1 for the prime ideals $\mathfrak{p}$ and $\overline{\mathfrak{p}}$ gives

$$\begin{aligned}
0 = a_0^{\mathfrak{p}} &= \alpha_1 + \alpha_2 + \overline{\alpha_2} \\
= a_0^{\overline{\mathfrak{p}}} &= \overline{\alpha_1} + \alpha_2 + \overline{\alpha_2} \quad \text{and} \\
0 = a_1^{\mathfrak{p}} &= \alpha_1 \log_q(\lambda_1) + \alpha_2 \log_q(\lambda_2) + \overline{\alpha_2}\log_q(\overline{\lambda_2}) \\
= a_1^{\overline{\mathfrak{p}}} &= \overline{\alpha_1}\log_q(\overline{\lambda_1}) + \alpha_2 \log_q(\lambda_2) + \overline{\alpha_2}\log_q(\overline{\lambda_2}).
\end{aligned}$$

The first two lines give that $\alpha_1 = \overline{\alpha_1}$ and the second that $\log_q(\lambda_1) = \log_q(\overline{\lambda_1})$ as $\alpha_1 \neq 0$. This contradicts the non-degeneracy condition. Thus, without loss of generality, the characteristic root λ_3 is not $\mathfrak{p}$-dominant.

If λ_3 is real, using Lemma 2.4.1 on the difference of a_1^q and $a_1^{\overline{\mathfrak{p}}}$ gives that

$$\alpha_1 \log_q(\lambda_1) + \alpha_2 \log_q(\lambda_2) + \alpha_3 \log_q(\lambda_3) = 0\,.$$

Hence, by Theorem 1.1.9, the characteristic roots λ_1, λ_2, and λ_3 are multiplicatively dependent. Say the multiplicative relationship is $\lambda_1^{s_1}\lambda_2^{s_2}\lambda_3^{s_3} = 1$. As $\nu_{\mathfrak{p}}(\lambda_3) > 0 = \nu_{\mathfrak{p}}(\lambda_1) = \nu_{\mathfrak{p}}(\lambda_2)$, it follows that $s_3 = 0$ and thus λ_1 and λ_2 are multiplicatively dependent. We have that $|\lambda_1| > 1$ because λ_1 is dominant, and the characteristic roots multiply to a non-zero integer larger than 1 as the LRS is not reversible. Thus, as $|\lambda_1| = |\lambda_2| > 1$, we have $s_1 = -s_2$, and so λ_1/λ_2 is a root of unity, contradicting non-degeneracy.

Thus, λ_3 is not real. Then, without loss of generality, $\lambda_4 = \overline{\lambda_3}$. Then, as $a_0^q = a_0^{\mathfrak{p}} = a_0^{\overline{\mathfrak{p}}} = 0$, we get

$$0 = \alpha_5 + \sum_{i=1}^{3}(\alpha_i + \overline{\alpha_i}) = \alpha_5 + \sum_{i=1}^{3}\overline{\alpha_i} = \alpha_5 + \sum_{i=1}^{3}\alpha_i\,,$$

and so $\alpha_5 = 0$, a contradiction.

Hence at least one of $a_0^{\mathfrak{p}}$, $a_1^{\mathfrak{p}}$, $a_2^{\mathfrak{p}}$, $a_0^{\overline{\mathfrak{p}}}$, $a_1^{\overline{\mathfrak{p}}}$, a_0^q, and a_1^q is non-zero. For all, except for $a_2^{\mathfrak{p}}$, zeroness can be decided, and hence if all of $a_0^{\mathfrak{p}}$, $a_1^{\mathfrak{p}}$, $a_0^{\overline{\mathfrak{p}}}$, $a_1^{\overline{\mathfrak{p}}}$, and a_0^q are zero, $a_2^{\mathfrak{p}}$ is non-zero, suggesting we can simply compute its p-adic expansion until a non-zero digit appears.

As such, we can apply Proposition 2.3.5 to find an S such that $u_{Sn} \neq 0$ for all $n \geq 1$. □

2.5 The Baker-Davenport method

All $\mathbb{Z}$-LRS with at most three dominant roots or at most two $\mathfrak{p}$-dominant roots for some prime ideal $\mathfrak{p}$ are in the MSTV class, and thus, in principle, the Skolem problem is decidable for such LRS. The SKOLEM-tool contains a procedure that solves (unconditionally) the Skolem problem for a large subset of the MSTV class (Section 2.6), namely all LRS with at most three dominant roots, all of which are simple. To the best of our knowledge, this is the first implementation of the method of [116, 156].

However, our actual method will differ from the works of Mignotte, Shorey, and Tijdeman [116] and Vershchagin [156] because nowadays, we can rely on the convenient tools of Matveev [115] and Yu [159]. These tools simplify their proofs greatly.

To start, we make a brief observation.

Proposition 2.5.1. *Let $(u_n)_{n=0}^{\infty}$ be a $\overline{\mathbb{Q}}$-LRS with dominant roots $\lambda_1, \ldots, \lambda_k$ having respective multiplicities $m_1, \ldots, m_k$. If $m_1 > m_2, \ldots, m_k$, then the Skolem problem is decidable for $(u_n)_{n=0}^{\infty}$.*

Proof. Writing $(u_n)_{n=0}^{\infty}$ in its exponential-polynomial form (1.2) gives

$$u_n = Q_1(n)\lambda_1^n + \cdots + Q_k(n)\lambda_k^n + r(n)\,,$$

where $|\lambda_1| = \cdots = |\lambda_k|$ and $r(n)$ is an LRS whose dominant roots have modulus strictly smaller than $|\lambda_1|$. We see that λ_1 is real as the exponential-polynomial form is invariant under complex conjugation. Dividing by n^{m_1-1} gives

$$u_n = \frac{Q_1(n)}{n^{m_1-1}}\lambda_1^n + \cdots + \frac{Q_k(n)}{n^{m_1-1}}\lambda_k^n + \frac{r(n)}{n^{m_1-1}}\,,$$

where one easily can compute constants $N \in \mathbb{Z}_{\geq 0}$ and $C > 0$ such that for all $n \geq N$, we have

$$\left|\frac{Q_1(n)}{n^{m_1-1}}\right| > C \quad \text{and} \quad \left|\frac{r(n)}{n^{m_1-1}|\lambda_1|^n}\right| + \sum_{i=2}^{k}\left|\frac{Q_i(n)}{n^{m_1-1}}\right| < C\,.$$

Then, by construction, $u_n > 0$ for all $n \geq N$. □

Proposition 2.5.1 shows in particular that the Skolem problem is decidable when $k = 1$, i.e., there is a single dominant root.

The remainder of our implementation is based on [102], which relies on Matveev's version of Baker's theorem on linear forms in logarithms (Theorem 1.1.11) and a reduction method due to Dujella and Pethő [61].

2.5.1 Establishing a bound using Baker's theorem

Assume $(u_n)_{n=0}^{\infty}$ is a non-degenerate $\mathbb{Z}$-LRS with two or three dominant roots and that all these dominant roots are simple. Then u_n can be expressed as

$$u_n = \alpha\lambda^n + \overline{\alpha}\,\overline{\lambda}^n + b\rho^n + r(n)\,,$$

where ρ is real algebraic, $|\lambda| = |\rho|$, and $r(n)$ is an LRS with dominant roots strictly smaller than $|\rho|$. By scaling, let $\lambda_1 = \lambda/\rho$, $b_1 = b/|\alpha|$, and $\alpha_1 = \alpha/|\alpha|$. Then $|\lambda_1| = |b_1| = |\alpha_1| = 1$. Then, for $r_1(n) = r(n)/|\alpha\lambda^n|$, we have that

$$\alpha\lambda^n + \overline{\alpha}\,\overline{\lambda}^n + b|\lambda|^n + r(n) = |\alpha|\rho^n\left(\alpha_1{\lambda_1}^n + \overline{\alpha_1}\,\overline{\lambda_1}^n + b_1 + r_1(n)\right)\,.$$

Next, using Lemma 1.2.6, we compute an integer $N' \geq 0$ and rational numbers $c_1 > 0$ and $0 < d < 1$ such that for all $n \geq N'$, we have $|r_1(n)| < c_1 d^n$. We will split into two cases: $|b_1| > 2$ and $|b_1| \leq 2$.

We first deal with the case $|b_1| > 2$. If $u_n = 0$, then $\left|b_1 - (\alpha_1 {\lambda_1}^n + \overline{\alpha_1}\,\overline{\lambda_1}^n)\right| \leq c_1 d^n$. However, using the triangle inequality, we have that

$$c_1 d^n \geq \left|b_1 - (\alpha_1 {\lambda_1}^n + \overline{\alpha_1}\,\overline{\lambda_1}^n)\right| \geq |b_1| - |\alpha_1 {\lambda_1}^n + \overline{\alpha_1}\,\overline{\lambda_1}^n| \geq |b_1| - 2\,.$$

As $d < 1$ and $|b_1| - 2 > 0$, the number n can be bounded effectively.

Now let $|b_1| \leq 2$ and write $\lambda_1 = e^{i\varphi}$. We want to find $N \geq 0$ such that for all $n \geq N$,

$$|\alpha_1 e^{in\varphi} + \overline{\alpha_1} e^{-in\varphi} + b_1| > c_1 d^n\,. \tag{2.9}$$

As $|e^{in\varphi}| = 1$ for all $n \geq 0$, (2.9) is equivalent to

$$\left|\alpha_1 e^{2in\varphi} + b_1 e^{in\varphi} + \overline{\alpha_1}\right| > c_1 d^n\,.$$

Treating the term inside the absolute value as a quadratic equation in $e^{in\varphi}$ and observing that $\alpha_1\overline{\alpha_1} = 1$ gives that

$$\left|e^{in\varphi} - \frac{-b_1 + \sqrt{{b_1}^2 - 4}}{2\alpha_1}\right| \cdot \left|e^{in\varphi} - \frac{-b_1 - \sqrt{{b_1}^2 - 4}}{2\alpha_1}\right| > c_1 d^n\,.$$

As b_1 is real and $|b_1| \leq 2$, we have ${b_1}^2 \leq 4$ by our earlier assumption; the two values for $-b_1 \pm \sqrt{{b_1}^2 - 4}$ are complex conjugates and have equal absolute value. Hence, the absolute value of $\gamma := (-b_1 \pm \sqrt{{b_1}^2 - 4})/2\alpha_1$ equals the absolute value of $(-b_1 \mp \sqrt{{b_1}^2 - 4})/2\alpha_1$. As $\overline{\alpha_1}$ is the product of these two complex conjugate numbers and $|\alpha_1| = 1$, they both lie on the unit circle. Thus, $\left|e^{in\varphi} - (-b_1 \pm \sqrt{{b_1}^2 - 4})/(2\alpha_1)\right| \leq 2$. Therefore, we have to show for all $n \geq N$ that

$$\Lambda := \left|\gamma^{-1} e^{in\varphi} - 1\right| > \frac{1}{2} c_1 d^n \tag{2.10}$$

for γ equal to $\frac{-b_1 \pm \sqrt{{b_1}^2 - 4}}{2\alpha_1}$.

We apply Matveev's theorem (Theorem 1.1.11) to (2.10). Then, when $\Lambda \neq 0$, we obtain

$$\log|\Lambda| > -C h'(\lambda_i) h'(\gamma^{-1})\big(1 + \log(B)\big)\,,$$

where $C \geq 1$ is a computable constant and $B = \max(|-1|, |n|)$. When demanding that $n \geq 1$, we have that $B = n$. Recall that for $\beta \in \overline{\mathbb{Q}}$, we have $h'(\beta) \geq \max(Dh(\beta), \log|\beta|, 0.16)$. Computing D (the degree of the extension $\mathbb{Q}(\lambda_1, \gamma)$ is relatively expensive, and so we estimate it from above. The scaled numbers λ_1 and α_1 lie

in quadratic extensions of the Galois closure of $\mathbb{Q}(\lambda)$ as $\lambda/|\lambda| = \sqrt{\lambda^2/(\lambda\overline{\lambda})} = \sqrt{\lambda/\overline{\lambda}}$ and similarly for α. Thus, if λ has a minimum polynomial of degree d_1, the number field $Q(\lambda_1, \alpha_1)$ has degree at most $2 \cdot 2 \cdot d_1!$. Similarly, b lies in the Galois closure of $\mathbb{Q}(\rho)$, and thus $b_1 = b/|\alpha|$ lies in an extension of degree at most $2d_1!d_2!$, where d_2 denotes the degree of the minimum polynomial of ρ. As γ lies in a quadratic extension of $\mathbb{Q}(\alpha_1, b_1)$, the number field $\mathbb{Q}(\lambda_1, \gamma)$ has degree $D \leq 2^3 d_1!d_2!$.

If $\Lambda = 0$, then $\lambda_1^n = \gamma$. Thus, $h(\gamma) = h(\lambda_1^n) = nh(\lambda_1)$. If γ is a root of unity, we have $n = 0$ as λ_1 is not a root of unity by non-degeneracy. Thus, we only need to find a lower bound for $h(\lambda_1)$ and an upper bound for $h(\gamma)$ when γ is not a root of unity. We start with γ. As γ has to be in at least one prime ideal, we have that $h(\gamma) \geq \log(2)/D$ using (1.1).

For an algebraic number x, we have that

$$h(x/|x|) = h\left(x/\sqrt{x\overline{x}}\right) \leq h(x) + \frac{1}{2}\big(h(x) + h(\overline{x})\big) = 2h(x)\,.$$

As we have the minimum polynomial of λ, estimating $h(\lambda_1)$ is easy: $h(\lambda_1) \leq 2h(\lambda)$. Thus, when $\Lambda = 0$, we have that $n \leq 2h(\lambda)D/\log(2)$.

Computing $h'(\gamma)$ exactly is expensive, so we will also estimate it from above. First, we estimate $h(\gamma)$ as follows

$$\begin{aligned}
h(\gamma) &= h\left(\frac{-b \pm \sqrt{{b_1}^2 - 4\alpha_1\overline{\alpha_1}}}{2\alpha_1}\right) \\
&\leq h\left(-b_1 \pm \sqrt{{b_1}^2 - 4\alpha_1\overline{\alpha_1}}\right) + h(2\alpha_1) \\
&\leq 2\log(2) + h(b_1) + \frac{1}{2}h({b_1}^2 - 4\alpha_1\overline{\alpha_1}) + h(\alpha_1) \\
&\leq 3\log(2) + h(b_1) + \frac{1}{2}h({b_1}^2) + \frac{1}{2}h(4\alpha_1\overline{\alpha_1}) + h(\alpha_1) \\
&\leq 4\log(2) + 2h(b_1) + 2h(\alpha_1) \leq 4(\log(2) + h(b) + h(\alpha))\,.
\end{aligned}$$

From here, we can compute an estimate for $h'(\gamma)$ directly.

Finally, we have to estimate $\log|\lambda_1|$ and $\log|\gamma|$. As both lie on the unit circle, they are at most $\log(\pi)$.

We have now computed a constant $c_3 > 0$ such that $\log|\Lambda| < c_3 \cdot (1 + \log(n))$ when n is large enough. Taking logarithms on both sides of (2.10) gives

$$\log|\Lambda| > \log(c_2) + n\log(d)\,. \tag{2.11}$$

As our n are large, we can assume that $B = n$. Then we want to find n for which

$$c_3 \cdot \big(1 + \log(n)\big) > \log(c_2) + n\log(d)\,.$$

As this difference is monotone, a computer can quickly determine for which N this inequality fails for all $n \geq N$.

The typical case is $b = 0$, i.e., there are two dominant roots. In this case, the formula for γ is much simpler, giving sharper estimates for $h(\gamma)$.

2.5.2 The Baker-Davenport reduction

The method described above successfully computes an N such that $u_n \neq 0$ for all $n \geq N$. Unfortunately, this N is fairly large, for example 10^{20}. Practically speaking, determining whether $u_n = 0$ for some $n < N$ is impossible by simply iterating the LRS. As such, we need to reduce N to something more manageable. One possibility is to use the famous LLL-algorithm [97], but in our use case with only two logarithms, a continued fractions method is sufficient. A specialized method for LRS using the p-adic zeros of a sequence is described by Bacik et al. in the very recent work [12]. The described method is due to Dujella and Pethő [61] and uses a version of a lemma by Baker and Davenport.

The *continued fraction* of a number $\alpha = \alpha_0 \in \mathbb{R} \setminus \mathbb{Q}$ is the infinite sequence $[a_n]_{n=0}^{\infty}$ defined by $a_n = \lfloor \alpha_n \rfloor$ and $\alpha_{n+1} = \left\lfloor \frac{1}{\alpha_n - a_n} \right\rfloor$ for all $n \geq 0$. As such, $a_0 \in \mathbb{Z}$, and $a_n \in \mathbb{Z}_{\geq 1}$ for all $n \geq 1$. For example, the continued fraction of the golden ratio, $\frac{1+\sqrt{5}}{2}$, is the constant sequence $[1, 1, 1, \dots]$. For each $n \geq 0$, put

$$\frac{p_n}{q_n} = a_0 + \cfrac{1}{a_1 + \cfrac{1}{a_2 + \cfrac{1}{\ddots + \cfrac{1}{a_{n-1} + \cfrac{1}{a_n}}}}}$$

where $p_n \in \mathbb{Z}$ and $q_n \in \mathbb{Z}_{\geq 1}$ are coprime. Then p_n/q_n is called the *nth convergent* of α. A fraction $p'/q' \in \mathbb{Q}$ is a convergent of α if and only if $p/q = p_n/q_n$ for some $n \geq 0$. Moreover, the sequences $(p_n)_{n=0}^{\infty}$ and $(q_n)_{n=0}^{\infty}$ satisfy for $n \geq 2$,

$$\begin{aligned} p_0 &= a_0 & p_1 &= a_0 a_1 + 1 & p_n &= a_n p_{n-1} + p_{n-2} \\ q_0 &= 1 & q_1 &= a_1 & q_n &= a_n q_{n-1} + q_{n-2}. \end{aligned}$$

Recall that we tried to solve the equation $|\gamma \lambda_1^n - 1| > c_2 d^n$ and that $\lambda_1 = e^{i\varphi}$. As we also know that $|\gamma| = 1$, we write $\gamma = e^{i\theta}$. Then we need to solve the equation

$$|e^{i(\theta + n\varphi)} - 1| < c_2 d^n \,. \tag{2.12}$$

If $||x||$ denotes the distance from x to the nearest integer, then using geometry and the Taylor expansion of the sine function, we see that if we assume that $|\theta+n\varphi| < \pi$, then the bound

$$|e^{i(\theta+n\varphi)} - 1| \geq |\sin(\theta + n\varphi)| \geq \left|\frac{\theta + n\varphi}{2\pi}\right| \geq \left|\left|\frac{\theta + n\varphi}{2\pi}\right|\right|$$

holds. Combining this bound with (2.12) gives

$$\left|\left|(\theta + n\varphi)/2\pi\right|\right| < c_2 d^n \,. \tag{2.13}$$

Now we can apply the theorem of Dujella and Pethő [61].

Theorem 2.5.2. *Let N be a positive integer, $\psi \in \mathbb{R}$, and p/q a convergent of the continued fraction of $x \in \mathbb{R} \setminus \mathbb{Q}$ such that $q > 6N$. Set $\varepsilon = ||\psi q|| - N||xq||$. If $\varepsilon > 0$, then there is no solution to*

$$0 < |nx - m + \psi| < CD^{-n}$$

in positive integers n and m such that

$$\frac{\log(Cq/\varepsilon)}{\log(D)} \leq n \leq N \,.$$

In our context we have that $x = \pm\frac{\varphi}{2\pi}$, $\psi = \pm\frac{\theta}{2\pi}$, $C = c_2$, $D = d$, and N is the upper bound found in the previous subsection. Here, the two $\pm$ have equal sign, and we need both (together with the inclusion of m in this theorem) to transform (2.13) into the pure absolute value, i.e., the inequality $|\pm(\theta + n\varphi)/2\pi + m| < c_2 d^n$ has no solution.

Often, this method successfully reduces N to $\log(Cq/\varepsilon)/\log(D)$, which is much smaller. One can simply check whether $u_n = 0$ for some $n \leq \log(Cq/\varepsilon)/\log(D)$ by iterating the sequence.

However, the reduction does not always succeed. In certain specific cases, there exists no convergent p/q such that $||\psi q|| > N||xq||$ holds. For example, if $\psi = 0$, then $||\psi q|| = 0$ for every convergent. This scenario is not impossible as it occurs when α is real. Moreover, the reduced N may still be very large. For example, this occurs when D is very close to 1. In that case, iterating the sequence is still too laborious. Hence, we need another reduction.

Another reduction using local methods As mentioned previously, the Baker-Davenport reduction cannot always be applied. When the LRS is simple, we can use another approach. As explained in Section 1.2, the sequence $(u_n \bmod p)_{n=0}^{\infty}$ is periodic for almost all primes p. For $m \geq 2$ such that $(u_n \bmod p)_{n=0}^{\infty}$ is periodic, let L_m be the period of $(u_n \bmod p)_{n=0}^{\infty}$ and

$$Z_m = \left\{0 \leq n < L_m \colon u_n \equiv 0 \pmod{m}\right\}.$$

If $m' \geq 2$ as well, the set $Z_{\mathrm{lcm}(m,m')}$ can be computed by applying the Chinese Remainder theorem repeatedly. Then, $\#Z_{\mathrm{lcm}(m,m')} \leq \#Z_m \#Z_{m'}$. We can use this trick when Z_m and $Z_{m'}$ are not too big and m and m' have few common factors.

In our approach, we use primes p for which the characteristic polynomial of $(u_n)_{n=0}^{\infty}$ splits over $\mathbb{F}_p$. Then, $L_p \mid p-1$ because we assumed that our LRS is simple. Hence, if we assume that the numbers $u_n \bmod p$ are randomly distributed for each $0 \leq n < L_p$, we have that $u_n \equiv 0 \pmod{p}$ with a probability of $1/p$. Thus, we expect that $\#Z_p \approx L_p/p < 1$. If L_p is smaller than $p-1$, this expected value is even lower. Computing Z_p and L_p for primes $p = p_1, \ldots, p_k$, lets us to compute Z_m and L_m for $m = p_1 \cdots p_k$ using the Chinese remainder theorem. When $m \geq N$ exceeds N, only a few values remain in Z_m. For each element in Z_m below N, one can compute u_n and check whether it is zero. As we are only interested in $Z_p \cap \{0, \ldots, N-1\}$, adding a few more primes p_i will eliminate many large values in Z_m, leaving only true in zeros of $(u_n)_{n=0}^{\infty}$ in Z_m, most of the time. By explicitly computing these values, we can check whether they are truly zero.

2.5.3 The non-simple case

For non-simple LRS with two or three dominant roots, the method to compute an upper bound N such that $u_n \neq 0$ for all $n \geq N$ still applies. Here,

$$\alpha(n)\lambda^n + \overline{\alpha}(n)\overline{\lambda}^n + b(n)\rho^n + r(n) = 0$$

where $\alpha(n)$, $\overline{\alpha}(n)$, and $b(n)$ are now polynomials in n. Following the same recipe, when $b(n) \leq 2|\alpha(n)|$ for large enough n, we first find a bound N such that $\alpha(n) \neq 0$ for all $n \geq N$ and write $\alpha_1(n) = \alpha(n)/|\alpha(n)| = \alpha(n)/\sqrt{\alpha(n)\overline{\alpha}(n)}$ etc. Thus, $\alpha_1(n)$, $\overline{\alpha_1}(n)$, and $b_1(n)$ are now algebraic functions. Hence,

$$\gamma(n) = \frac{-b_1(n) \pm \sqrt{b_1(n)^2 - 4}}{2\alpha_1(n)}$$

is also an algebraic function. Using the rules for the height, we can estimate the height of $\gamma(n)$ as a linear function in both $\alpha(n)$ and $b(n)$. As $\alpha(n)$ and $b(n)$ are known polynomials in n, their height can be bounded effectively by a linear function in $\log(n)$. Hence, $h(\gamma(n)) < C_1 + C_2 \log(n)$ for some computable constants C_1 and C_2. Thus, showing that Λ is non-zero, that is, $e^{in\varphi} \neq \gamma(n)$ can be done as

$$nh(e^{i\varphi}) = h(e^{in\varphi}) = h(\gamma(n)) \leq C_1 + C_2 \log(n)$$

does not hold for computably large n. Meanwhile, in the application of Baker's theorem, we have that $\log|\Lambda| > -C\log^2(n)$ due to this extra factor $\log(n)$. As $\log^2(n)$ still grows slowly, the number n can be bounded by some constant N.

Unfortunately, this bound N is again often large, and the Baker-Davenport reduction Section 2.5.2) cannot be used. Moreover, the local methods of Section 2.5.2 can also fail. Even for primes p such that the characteristic polynomial splits over $\mathbb{F}_p$, the period L_p of the LRS modulo p is still only bounded by $(p-1)p$, and so $\#Z_p$ has size approximately $\approx (p-1)p/p = p-1$. Hence, the size of this set is too large for practical purposes, and one has to resort to other algorithms like the LLL-algorithm or compute the p-adic zeros up to a decent precision using the recent result [12].

2.6 The SKOLEM-tool

We have implemented many of the methods described in Sections 2.3 and 2.5 in the SKOLEM-tool, which finds all the zeros of a given LRBS. Specifically, the leapfrogging algorithm from Section 2.3 is implemented for simple LRS, and the Baker-Davenport method from Section 2.5 is implemented for all described classes, including all LRS with at most three dominant roots, which are all simple. No guarantees are provided.

The tool supports non-degenerate LRS, and, by default, rejects degenerate inputs. Degenerate sequences could be manually decomposed into several non-degenerate subsequences. The check for non-degeneracy can be skipped, yielding a modest speed-up, but may lead to non-termination if the given LRS contains infinitely many zeros.

For both the Baker-Davenport method and the Leapfrogging method, there are several modes in which the algorithm can run. In the web version of the tool, they can be enabled with a slider. One such example is that the user can choose whether to automatically reduce the (sub)sequences by their greatest common divisor using the *Use GCD reduction* setting. Doing so will slightly improve the bounds found by the Baker-Davenport method and will significantly speed up the Leapfrogging algorithm.

Leapfrogging

The Leapfrogging algorithm of Section 2.3 is implemented only for primes p that do not divide the constant coefficient of the characteristic polynomial of the LRS. For the sake of efficiency, Theorem 1.1.3 is not implemented and checking whether a coefficient a_j of the power series (2.7) is zero is done by computing its p-adic expansion to high precision. If the p-adic valuation of a_j is sufficiently large, then a_j is assumed to be zero. Hence, the output of the tool may be unsound in very specific circumstances.

For the Leapfrogging algorithm, the option *factor subcases* reduces the number of subsequences analyzed. For example, when $u_0 = 0$, and the step S is computed to be 24, normally the 23 subsequences $(u_{24n+1})_{n=-\infty}^{\infty}, \ldots, (u_{24n+23})_{n=-\infty}^{\infty}$ are analyzed separately. *Factor subcases* finds a more optimal partition of the remaining subsequences. In this case, $(u_{3n+1})_{n=-\infty}^{\infty}$, $(u_{3n+2})_{n=-\infty}^{\infty}$, $(u_{6n+3})_{n=-\infty}^{\infty}$, $(u_{12n+6})_{n=-\infty}^{\infty}$, and $(u_{24n+12})_{n=-\infty}^{\infty}$. Thus, only five instead of 23 sequences need to be analyzed. In practice, enabling this setting can slow the algorithm as higher numbers M may be needed, although the length of the certificate is reduced.

Secondly, we have two approaches to searching for a suitable M:

minimal M. This option will simply start from $M = 1$ and search upwards to seek M such that the sequence is non-zero mod M. Whilst this method guarantees the minimal M is found, the period such that it is non-zero mod m can be quite large.

minimal period. Our second approach finds the shortest period. Iterating k, we test whether there is an M such that $(u_n \bmod M)_{n=-\infty}^{\infty}$ has period k and is non-zero. In practice, despite obtaining far larger numbers M found than in the minimal M method, the algorithm runs faster with approach.

Baker-Davenport

In the implementation of the Baker-Davenport method from Section 2.5, complex numbers are treated numerically rather than symbolically, as this speeds up the algorithm immensely. The resulting approximations do not affect the correctness of the algorithm but might give slightly higher bounds for N. When the dominant roots and the largest non-dominant root have absolute values which are almost equal, the algorithm returns an error. Currently, the minimum allowed distance is 2^{-950}, but this number can be changed manually in the code. Still, when this distance is small,

Order	Number of LRS	zero LRS	Degenerate LRS	Non-simple LRS
2	9250	6	358	50
3	8995	0	74	2
4	9195	0	35	2
5	9188	0	15	3
6	9172	0	10	6
7	9213	0	12	0
8	9157	0	10	0
9	9143	0	4	3
10	9047	0	8	1

Table 2.1: The distribution of the LRS in our randomly generated dataset.

the Baker-Davenport method can return an enormous bound N, making the method impractical.

For the Baker-Davenport method, there is the choice of whether to only compute the bound or whether to also check for zeros up to this bound. The method is in principle only for LRS (not LRBS), but we provide an option to search in both directions, giving bounds N^- and N^+ such that $u_n \neq 0$ for $n \leq N^-$ and $n \geq N^+$. Therefore, all the zeros of the corresponding LRBS will be found (alternatively, either of the two directions can be run independently).

2.6.1 Testing results

We tested the SKOLEM-tool on a suite of random LRS generated in [31]. This dataset comprises 82367 LRS of orders uniformly distributed between 2 and 10. The coefficients and initial values are uniformly chosen between -20 and 20, with the only restriction being that the constant coefficient of the characteristic polynomial is non-zero. That is, c_d, as in (1), is non-zero.

A subset of these LRS have all initial values set to zero, yielding identically zero sequences. Additionally, some LRS are non-degenerate, yet our implementation does not currently incorporate a decomposition algorithm to extract the non-degenerate components. We list the number of occurrences of these phenomena in Table 2.1 by order. Some of these recurrences may not be minimal, which we have not tested. But we believe this should only involve a negligible number of LRS).

For each LRS, we attempted to solve the Skolem problem using five different variations of our algorithm, each with a 60-second timeout. Testing was conducted using SageMath in Docker on a Dell PowerEdge M620 blade equipped with 2×3.3

Order	Leapfrog A	Leapfrog B	Leapfrog C	Leapfrog D	Baker-Davenport
2	100%	100%	100%	100%	100%
3	100%	100%	100%	100%	100%
4	99.92%	99.90%	99.97%	99.93%	100%
5	84.6%	98.0%	83.6%	97.4%	100%
6	36.5%	62.1%	34.0%	58.9%	99.99%
7	8.0%	16.4%	7.0%	14.4%	100%
8	2.1%	4.4%	1.9%	3.8%	99.97%
9	0.42%	0.83%	0.38%	0.68%	99.95%
10	0.14%	0.32%	0.13%	0.21%	99.88%

Table 2.2: The success rate for each variant, by order.

Order	Leapfrog A	Leapfrog B	Leapfrog C	Leapfrog D	Baker-Davenport
2	0.18	0.18	0.18	0.18	0.28
3	0.21	0.21	0.21	0.21	0.34
4	0.38	0.29	0.40	0.30	0.39
5	5.8	3.1	6.0	3.4	0.45
6	7.8	9.0	8.9	9.5	0.51
7	10.0	11.3	10.1	11.4	0.58
8	7.7	10.8	8.4	9.6	0.72
9	4.5	8.3	4.1	7.1	0.88
10	14.8	6.0	11.8	8.1	1.03

Table 2.3: The mean time in seconds for successful runs for each variant, by order.

GHz Intel Xeon E5-2667 v2 (2×8 cores, 32 with hyper-threading) and 256GB RAM. Testing was restricted to 16 parallel threads (50% of the computer's resources).

The five variations of our algorithm comprise of four variants of the leapfrogging algorithm (labelled A–D) and the Baker-Davenport method. All variants use the *GCD-reduction method*, and all Leapfrogging variants use the *exponential-polynomial method* in Section 2.3.2 to compute the step. Variants A and C use the *minimal M method*, and variants B and D use the *minimal period method*. Variants C and D use the *factor subcases method* while variants A and B do not. The Baker-Davenport method computes the zeros in both the positive and the negative direction and iterates the zeros up to the bounds found by the Baker-Davenport reduction. Tables 2.2 and 2.3 present the success rates and mean runtimes of these experiments, broken down by order. For the leapfrogging variants, the experiments were restricted to non-simple LRS.

In our experiments, the Baker-Davenport method is by far the fastest and most

successful. However, it does not apply to LRS with more than three dominant roots, which were not present in our randomly generated dataset. Moreover, for rare instances, the leapfrogging methods are indeed faster than the Baker-Davenport method. For example, when $(u_0, \ldots, u_4) = (-1, -17, -3, -6, 8)$ and

$$u_{n+5} = -11u_{n+4} - 4u_{n+3} + 5u_{n+2} + 12u_{n+1} + 13u_n\,,$$

Our Leapfrogging B experiment runs in 0.23 seconds while the Baker-Davenport method needs 14.9 seconds because the bound N^- in the negative direction after the Baker-Davenport reduction is -196851. Among the leapfrogging variants, the *minimal period method* is substantially faster and more successful than the *minimal M method* while the *factor subcases method* slows down the algorithm to a lesser extent. Note that for higher orders, the success rate of the leapfrogging variants is too low to allow for a direct comparison of run times.

Surprisingly, in this large random set of LRS, zeros do not occur often. For example, for the Baker-Davenport method (which almost always succeeded), the largest n such that at least one LRS $(u_n)_{n=0}^{\infty}$ satisfies $u_n = 0$ is $n = 13$. Thus, despite the enormous number of LRS, for all these LRS $(u_n)_{n=0}^{\infty}$, we have that $u_n \neq 0$ for all $n > 13$.

2.7 Concluding remarks

The proof of Theorem 2.1.8 only depends on the Skolem and p-adic Schanuel conjecture for termination and provides a certificate when terminating. The following theorem gives a shortened version of this certificate.

Theorem 2.7.1. *Let $(u_n)_{n=0}^{\infty}$ be a simple LRS. Then, assuming the Skolem conjecture and the p-adic Schanuel conjecture for termination, one can compute a number $Q \geq 1$ such that for all $0 \leq i < Q$ either*

1. *$(u_{Qn+i})_{n=-\infty}^{\infty}$ is constantly zero;*
2. *one can compute integers $b_i > 0$, a_i, and $z_i \equiv i \pmod{Q}$ and a prime number p_i such that $\nu_{p_i}(u_{Qn-z_i}) = a_i + b_i\nu_{p_i}(n)$ for all $n \in \mathbb{Z}$; or*
3. *one can compute $M_i \geq 1$ coprime with the constant coefficient of the characteristic polynomial of $(u_n)_{n=0}^{\infty}$ such that $u_{Qn+i} \not\equiv 0 \pmod{M_i}$ for all $n \in \mathbb{Z}$.*

Thus, we partition the LRBS $(u_n)_{n=0}^{\infty}$ into Q LRBS that are either constantly 0 (case 1), contain 0 exactly once as proven by the p_i-adic valuation of the sequence (case 2), or the subsequence is non-zero as demonstrated by considering the bi-infinite sequence modulo M_i (case 3). Case 1 can only occur when $(u_n)_{n=0}^{\infty}$ is degenerate or is the zero-LRS. Theorem 1.2.3 allows us to decompose the LRBS into non-degenerate LRBS. When $(u_n)_{n=0}^{\infty}$ is non-degenerate and not constantly zero, only cases 2 and 3 occur.

The statement of Theorem 2.7.1 is complex due to the many different parameters a_i, b_i, z_i, p_i, and M_i. Only finitely many primes p_i are excluded, and we can choose them freely otherwise, so we can choose all p_i to be equal to, say, a single prime p. As the numbers M_i also are describe the local behaviour of the subsequence, we would like to relate M_i to p. That is, we want that all M_i to be powers of p. Then (at the cost of needing to choose a bigger number Q), we can compute a_i such that a_i, $\nu_{p_i}(u_n) = a_i$ for all $n \equiv i \pmod{Q}$. After [111], we refer to such a formula as a *Marques-Lengyel formula*.

Definition 2.7.2 (Marques-Lengyel formula)**.** If $(u_n)_{n=-\infty}^{\infty}$ is an LRBS and p a prime number, $(u_n)_{n=-\infty}^{\infty}$ allows a Marques-Lengyel formula for the prime p if there are $Q \geq 1$ and integers $b_i \neq 0$, a_i, and $z_i \equiv i \pmod{Q}$ for $0 \leq i < Q$ such that either

a. for all $n \in \mathbb{Z}$, $\nu_p(u_{Qn-z_i}) = a_i + b_i\nu_p(n)$; or

b. for all $n \in \mathbb{Z}$, $\nu_p(u_{Qn+i}) = a_i$.

Encouraged by Lengyel's result that the Fibonacci numbers allow a Marques-Lengyel formula for every prime number [96], Marques and Lengyel were optimistic about the existence of Marques-Lengyel formulas in [111]. They studied the Tribonacci numbers, which are defined by $T_0 = 0$, $T_1 = T_2 = 1$, and the recurrence

$$T_{n+3} = T_{n+2} + T_{n+1} + T_n\,.$$

They show that the Tribonacci numbers allow a Marques-Lengyel formula for the prime 2:

$$\nu_2(T_n) = \begin{cases} 0 & \text{if } n \equiv 1, 2 \pmod{4} \\ 1 & \text{if } n \equiv 3, 11 \pmod{16} \\ 2 & \text{if } n \equiv 4, 8 \pmod{16} \\ 3 & \text{if } n \equiv 7 \pmod{16} \\ \nu_2(n) - 1 & \text{if } n \equiv 0 \pmod{16} \\ \nu_2(n+4) - 1 & \text{if } n \equiv 12 \pmod{16} \\ \nu_2(n+1) + 1 & \text{if } n \equiv 31 \pmod{32} \\ \nu_2(n+17) + 1 & \text{if } n \equiv 15 \pmod{32} \end{cases}$$

and conjectured that such a formula exists for every prime p. In [32], we showed that this conjecture holds for $p = 3$, 83, and 397 and fails for all other primes $p < 600$ except for $p = 11$, 103, and 163 for which one needs more precise arguments to reach a conclusion.

For some primes, a more general phenomenon occurs: z_i can be rational. For example, let $p = 269$ and $Z = \{0, -1, -4, -17, 1/3, -5/3\}$. Then,

$$\nu_{269}(T_n) = \begin{cases} 0 & \text{if } n \not\equiv z_i \pmod{268} \text{ for all } z_i \in Z \\ \nu_{269}(n - z_i) & \text{if } n \equiv z_i \pmod{268} \text{ for some } z_i \in Z\,. \end{cases}$$

Besides the integer zeros $0, -1, -4, -17$, there are also *rational zeros* $1/3$ and $-5/3$. For at least 1/12th of all primes, these rational zeros occur, contradicting the conjecture of Marques and Lengyel.

Thus, the valuation argument can also apply when the zero is outside the subsequence. Bacik et al. formalised this idea as follows [12]. For an LRBS $(u_n)_{n=-\infty}^{\infty}$ and a prime p such that $(u_n \bmod p)_{n=-\infty}^{\infty}$ is well-defined and periodic, call $z \in \mathbb{Z}_p$ a *p-adic zero* if for some sequence $(z_n)_{n=0}^{\infty} \in \mathbb{Z}^{\omega}$, the p-adic limits $\lim_{n\to\infty} z_n = z$ and $\lim_{n\to\infty} u_{z_n} = 0$ hold. We revisit the Tribonacci sequence with $p = 257$ and $Z = \{0, -1, -4, -17, 1/3, -5/3\}$:

$$\nu_{257}(T_n) = \begin{cases} 0 & \text{if } n \not\equiv z_i \pmod{256} \text{ for all } z_i \in Z \cup \{56\} \\ \nu_{257}(n - z_i) & \text{if } n \equiv z_i \pmod{256} \text{ for some } z_i \in Z \\ \nu_{257}(\frac{n-56}{256} - z_{257}) & \text{if } n \equiv 56 \pmod{256}\,, \end{cases}$$

where $z_{257} = 20 + 95\cdot 257 + 199\cdot 257^2 + 234\cdot 257^3 + 165\cdot 257^4 + \cdots$ is a 257-adic zero. Our other formulas for the p-adic valuations assert z_{257} does not correspond to an integer zero of the Tribonacci sequence. Bacik et al. showed that one can compute p-adic zeros to arbitrary precision, and when one assumes the p-adic Schanuel conjecture, one can also show that two zeros are distinct and thus enumerate the p-adic zeros. This method provides limited help for the Skolem problem. It allows one to search for zeros quite quickly, but one still has to determine whether a p-adic zero corresponds to a zero of the LRS. As these p-adic zeros can correspond to transcendental numbers that are very hard to describe, this is not an easy task.

Therefore, we investigate p-adic zeros that can be explained algebraically. Such an algebraic explanation shows that an entire subsequence does not contain zero. The earlier-mentioned rational zeros, and more generally, *twisted rational zeros*.

Definition 2.7.3 ([33]). A non-degenerate LRBS $(u_n)_{n=-\infty}^{\infty}$ with an exponential-polynomial form $u_n = \sum_{i=1}^{k} Q_i(n)\lambda_i^n$ has a *twisted rational zero* at $q \in \mathbb{Q}$ if for some roots of unity $\zeta_1, \ldots, \zeta_k$ and some definition of $\lambda_1^q, \ldots, \lambda_k^q$,

$$\zeta_1 Q_1(q)\lambda_1^q + \zeta_2 Q_2(q)\lambda_2^q + \cdots + \zeta_k Q_k(q)\lambda_k^q = 0\,.$$

In other words, twisted rational zeros imply there is a non-degenerate LRBS $(v_n)_{n=-\infty}^{\infty}$ such that $v_0 = 0$ (the twisted rational zero) and the LRBS $(u_n)_{n=-\infty}^{\infty}$ and $(v_n)_{n=-\infty}^{\infty}$ share a subsequence. That is, there are integers a, b, c, d with $a, c \geq 1$ such that $v_{an+b} = u_{cn+d}$ for all $n \in \mathbb{N}$.

The Skolem-Mahler-Lech theorem extends to twisted rational zeros.

Theorem 2.7.4 (Theorem 1.9 in [33]). *A non-degenerate, non-zero $\mathbb{Q}$-LRS has finitely many twisted rational zeros.*

To end the chapter, we give an explicit example of an LRS for which the sole known method of solving the Skolem problem depends on twisted rational zeros.

Example 2.7.5. Define the non-degenerate LRBS $(u_n)_{n=-\infty}^{\infty}$ by

$$u_n = (-4+7i)^n + (-4-7i)^n + 2(8+i)^n + 2(8-i)^n - n\,.$$

The first four characteristic roots have modulus $\sqrt{65}$, and for every prime deal $\mathfrak{p}$, there are either three or five $\mathfrak{p}$-dominant roots. Thus, $(u_n)_{n=0}^{\infty}$ is outside the MSTV class, and Luca showed that $(u_n)_{n=0}^{\infty}$ is not modular [103]. We will show that $u_n \neq 0$ for all $n \in \mathbb{Z}$ using the technology developed in this chapter and a twisted rational zero at $q = 0$:

$$1 \cdot (-4+7i)^0 + (-1) \cdot (-4-7i)^0 + 1 \cdot 2(8+i)^0 + (-1) \cdot 2(8-i)^0 - 0 = 0\,.$$

Using modular arithmetic, $u_n \equiv 0 \pmod 4$ if and only if $n \equiv 2 \pmod 4$. Then define the $\overline{\mathbb{Q}}$-LRBS $(v_n)_{n=-\infty}^{\infty}$ by

$$\begin{aligned} v_n &= (-4+7i)^n - i^n(-4-7i)^n + 2(8+i)^n - 2i^n(8-i)^n - n \\ &= (-4+7i)^n - (7-4i)^n + 2(8+i)^n - 2(1+8i)^n - n \end{aligned}$$

such that $u_{4n+2} = v_{4n+2}$ for all $n \in \mathbb{Z}$. We can exploit the fact that $v_0 = 0$ to deduce that for $n \equiv 2 \pmod 4$,

$$\nu_3(u_n) = \nu_3(v_n) = \begin{cases} 1 & \text{if } n \equiv 3 \pmod 6 \\ \nu_3(n) & \text{if } n \equiv 0 \pmod 6\,. \end{cases}$$

We conclude that for all $n \in \mathbb{Z}$:

$$\begin{aligned} u_n &\not\equiv 0 \pmod 4 \quad \text{if } n \not\equiv 2, 6, 10 \pmod{12}\,; \\ u_n &\not\equiv 0 \pmod 9 \quad \text{if } n \equiv 2, 10 \pmod{12}\,; \\ \nu_3(u_n) &= \nu_3(n) \quad \text{if } n \equiv 6 \pmod{12}\,, \end{aligned}$$

which is sufficient to conclude that $u_n \neq 0$ for all $n \in \mathbb{Z}$. □

The method above does not seem to apply to all non-simple order-6 LRS outside the MSTV class. An explicit example for which we cannot solve the Skolem problem is the non-simple, non-degenerate, and non-modular sequence $(u_n)_{n=0}^{\infty}$ defined by

$$u_n = 2(-4+7i)^n + 2(-4-7i)^n + 4(8+i)^n + 4(8-i)^n + n\,.$$

Chapter 3

The Positivity problem

3.1 Introduction and main results

In this chapter, we study the *Positivity problem* for linear recurrence sequences. The Positivity problem shares many similarities with the Skolem problem and, as discussed in the introduction, shows up in many different contexts. The Positivity problem asks to determine whether a $\mathbb{Q}$-LRS only contains positive numbers:

Problem 2 (Positivity problem)**.** For a given LRS $(u_n)_{n=0}^{\infty}$, determine whether $u_n \geq 0$ for all $n \in \mathbb{N}$.

Like the Skolem problem, we can generalise this problem to other rings than the integers, but for 'positivity' to make sense, the ring has to be a subring of $\mathbb{R}$.

Problem 3.1.1 (Positivity problem for R-LRS)**.** For a ring $R \subseteq \mathbb{R}$, the R-Positivity problem asks to determine for a given R-LRS $(u_n)_{n=0}^{\infty}$ whether $\forall n \in \mathbb{N}$: $u_n \geq 0$.

Similarly to the Skolem problem, the Positivity problems for $R = \mathbb{Z}$ and $R = \mathbb{Q}$ are Turing-interreducible: If $(u_n)_{n=0}^{\infty}$ is a $\mathbb{Q}$-LRS, one can compute integers $a, b \geq 1$ such that $(ab^n u_n)_{n=0}^{\infty}$ is an integer-valued LRS. As a and b are positive, the numbers u_n and $ab^n u_n$ have the same sign. The case $R = \mathbb{Z}$ is also the most prominent case of the Positivity problem, so, in this chapter, assume that all LRS are $\mathbb{Z}$-LRS unless stated otherwise.

The Positivity problem is known to be Skolem-hard. If $(u_n)_{n=0}^{\infty}$ is a $\mathbb{Z}$-LRS, then $u_n = 0$ if and only if $u_n^2 - 1$ is negative. Hence, as $(u_n^2 - 1)_{n=0}^{\infty}$ is again a $\mathbb{Z}$-LRS.

However, the distribution of negative terms in a sequence is far more complicated than the distribution of zeros. For the Skolem problem, the Skolem-Mahler-Lech theorem describes the set of zeros of an LRS as a union of finitely many arithmetic progressions and a finite set. Moreover, by Lemma 1.2.3, we can explicitly compute

all these arithmetic progressions. There is no analogue of the Skolem-Mahler-Lech theorem for the Positivity problem. In fact, it is still open to decide whether an LRS has finitely many negative terms. This problem is called the *Ultimate Positivity problem*, which asks to determine whether an LRS is ultimately positive or, in other words, contains only finitely many negative terms.

Problem 3.1.2 (Ultimate Positivity problem for R-LRS)**.** The Ultimate Positivity problem asks to determine for a given R-LRS $(u_n)_{n=0}^{\infty}$ whether $(u_n)_{n=0}^{\infty}$ is ultimately positive, i.e., $\exists N \geq 0\ \forall n \geq N \colon u_n \geq 0$.

The case $R = \mathbb{Z}$ is again the most prominent case of the Ultimate Positivity problem, and so when the ring R is not specified, we assume that $R = \mathbb{Z}$.

The (Ultimate) Positivity problem has a much shorter history than the Skolem problem. In the 1970s, around a decade before Mignotte, Shorey, and Tijdeman, and Vereshchagin published their powerful results on the Skolem problem, the Positivity problem emerged in computational biology, and the study of Lindenmayer systems in particular [98, 134].

Since, various authors have made incremental progress on the Positivity problem. In 1981, Burke and Webb solved the Ultimate Positivity problem for LRS up to order 2, and in 1990, Nagasaka and Shiue did the same for LRS up to order 3 with repeated roots. In 2006, Halava et al. [76] showed that the Positivity problem was decidable for LRS of order at most 2, and Laohakosol and Tangsupphathawat [93] solved the order-3 case three years later while also solving the general order-3 case for the Ultimate Positivity problem. All of these solutions relied on relatively elementary tools.

In 2014 and 2015, Ouaknine and Worrell pushed the boundaries of decidability of the Positivity and Ultimate Positivity problems to their current limit in a series of papers.

Theorem 3.1.3 ([121])**.** *The Positivity problem and the Ultimate Positivity problem are decidable for all LRS of order at most 5. Moreover, both problems are Diophantine-hard for order-*6 *LRS.*

Theorem 3.1.4 ([120])**.** *The Positivity problem is decidable for all simple LRS of order at most 9.*

Theorem 3.1.5 ([122])**.** *The Ultimate Positivity problem is decidable for all simple LRS.*

For the first two results, Ouaknine and Worrell rely on Baker's theorem on linear forms and for the latter two results, they use Tarski's result that the first-order theory of $\langle \mathbb{R}; 0, 1, +, \cdot \rangle$ is decidable. In Section 3.2, we will give intuition about these three theorems and also explain what "Diophantine hardness" as in Theorem 3.1.3 means.

Main results and organization of this chapter

In Section 3.2, we break down the current state-of-the-art of the Positivity and Ultimate Positivity problems. We sketch the proofs of Theorems 3.1.3, 3.1.4, and 3.1.5 and explain why these techniques are insufficient to make further progress. We also give explicit LRS for which we cannot decide (ultimate) positivity and show that the Skolem problem for order 5 LRS reduces to the Positivity problem for simple LRS of order 10, sharpening a bound of Ouaknine and Worrell.

In Section 3.3, we study a special class of linear recurrence sequences, namely *reversible linear recurrence sequences*. Recall that an LRS is reversible if its bi-completion is integer-valued. We will show that the Skolem, Positivity, and Ultimate Positivity problems are decidable for LRS of a much higher order when they are in this class by exploiting their structural properties.

Theorem 3.1.6. *For reversible LRS, the Skolem problem is decidable up to order* 11.

Theorem 3.1.7. *For reversible LRS, the Positivity and Ultimate Positivity problems are both decidable up to order* 11.

Theorem 3.1.8. *For simple reversible LRS, the Positivity problem is decidable up to order* 17.

To prove these theorems, we will analyse the potential possible ways the characteristic roots of a reversible LRS are distributed, i.e., the roots of a monic polynomial $P \in \mathbb{Z}[X]$ whose constant coefficient is ± 1. We rely on Galois-theoretic techniques to show that such polynomials cannot have 'many' dominant roots.

In Section 3.4, we show that we cannot push these Galois-theoretic techniques any further: we will inevitably encounter the same type of problem as one encounters for the general Skolem and (Ultimate) Positivity problems. We construct reversible LRS for which the Positivity problem is open. In particular, we construct a reversible LRS of order 12 and a simple reversible LRS of order 18 for which we do not know how to solve the Positivity problem and a reversible LRS of order 7 for which we do not know how to decide the Skolem problem. Thus, Theorems 3.1.6, 3.1.7, and 3.1.8 are sharp.

In Section 3.5, we discuss the Positivity problem for real algebraic linear recurrence sequences. For the Skolem problem, it is folklore [107, Lemma 9] that the $\mathbb{Z}$- and $\overline{\mathbb{Q}}$-Skolem problems are Turing-interreducible. However, such a result was not known for the Positivity problem. We prove the following result:

Theorem 3.1.9. *The $\mathbb{R} \cap \overline{\mathbb{Q}}$-Positivity problem and $\mathbb{Q}$-Positivity problem are Turing-interreducible. Moreover, the Positivity problem for simple $\mathbb{R} \cap \overline{\mathbb{Q}}$-LRS is Turing equivalent to the Positivity problem for simple $\mathbb{Q}$-LRS.*

3.2 Overview of the Positivity problem

In this section, we discuss obstacles to extending our results for deciding the (Ultimate) Positivity problem at higher orders.

As stated in the introduction, the (Ultimate) Positivity problem is decidable for all LRS up to order 5. However, extending this result to LRS of order 6 is challenging because it would resolve certain number-theoretical problems in Diophantine approximation. Hence, Ouaknine and Worrell referred to the (Ultimate) Positivity problem as *Diophantine-hard* in Theorem 3.1.3. This hardness depends on the LRS being non-simple. In contrast, for non-simple LRS, this Diophantine-hardness obstacle does not exist. Ouaknine and Worrell established decidability for simple LRS for the Ultimate Positivity problem and decidability for simple LRS of order at most 9 for the Positivity problem. For simple order-10 LRS, one encounters another type of hardness. Before exploring these two forms of hardness, we present a lemma that addresses most cases.

Lemma 3.2.1. *Let $(u_n)_{n=0}^{\infty}$ be a non-degenerate LRS. Then the (ultimate) positivity for $(u_n)_{n=0}^{\infty}$ is decidable if the following condition is not satisfied: Among the dominant roots of $(u_n)_{n=0}^{\infty}$ of maximum multiplicity, there is a non-real dominant root and a positive real dominant root.*

Proof. As $(u_n)_{n=0}^{\infty}$ is non-degenerate, the LRS $(u_n)_{n=0}^{\infty}$ has at most one real dominant root. Let λ be the modulus of a dominant root. Then we scale $(u_n)_{n=0}^{\infty}$ with $|\lambda|^n$ to obtain the following normalised exponential-polynomial form:

$$\frac{u_n}{|\lambda|^n} = Q(n)(\pm 1)^n + \sum_{i=1}^{k} Q_i(n)\lambda_i^n + \overline{Q_i}(n)\overline{\lambda_i}^n + r_n\,.$$

Here, $Q(n)$ is the zero polynomial if $(u_n)_{n=0}^{\infty}$ has no real dominant root, all λ_i are normalised dominant roots, and $(r_n)_{n=0}^{\infty}$ is the normalised non-dominant part

of $(u_n)_{n=0}^{\infty}$. As the dominant roots of the normalized sequence have modulus 1, we have $\lim_{n\to\infty} r_n = 0$. Moreover, by Lemma 1.2.6, we can compute rational numbers $r > 1$ and $0 < R < 1$ such that $|r_n| < rR^n$. Let $d := \max_{i=1}^{k} \deg(Q_i)$.

If $\deg(Q) > d$, then

$$\frac{u_n}{|\lambda|^n n^{d+1}} = \frac{Q(n)(\pm 1)^n}{n^{d+1}} + O(1/n)\,,$$

where the constant in the big O-notation can be controlled. For computably large n, we thus have that $Q(n)(\pm 1)^n$ determines the sign of u_n.

If $\deg(Q) < d$ or $\deg(Q) = d$ and the normalised real dominant root is -1, we apply a lemma due to Braverman [39].

Lemma 3.2.2. *Let $\gamma_1, \ldots, \gamma_k \in \{z \in \mathbb{C} \colon |z| = 1, z \neq 1\}$ be distinct complex numbers, $\alpha_1, \ldots, \alpha_k \in \mathbb{C}\backslash\{0\}$, and $w_n = \sum_{i=1}^{k} \alpha_i \gamma_i^n$. Then there is a $c < 0$ such that $\mathrm{Re}(w_n) < c$ for infinitely many n.*

In our case, we can write that

$$\frac{u_n}{|\lambda|^n n^d} = \alpha(\pm 1)^n + \sum_{i=1}^{k} \alpha_i \lambda_i^n + \overline{\alpha_i}\,\overline{\lambda_i}^n + O(1/n)\,,$$

where $\alpha = 0$ when the normalised real dominant root is 1 and at least one α_i is non-zero. Then Lemma 3.2.2 says that for some $c < 0$, there are infinitely many n such that $u_n < c + O(1/n)$, and thus $(u_n)_{n=0}^{\infty}$ is not (ultimately) positive.

As the degrees of the polynomials Q and Q_i correspond to the multiplicities of their characteristic roots, the lemma follows. □

Hardness at order 6 From the lemma above, we can conclude the following: If we cannot decide positivity for an LRS $(u_n)_{n=0}^{\infty}$ with non-simple dominant roots, then the LRS has at least three characteristic roots of multiplicity 2, and hence, order at least 6. Ouaknine and Worrell use such LRS to obtain their Diophantine hardness.

We give a simplified version of their method. Let

$$u_n = \frac{1}{2} n\big(2 - \lambda^n - \overline{\lambda}^n\big) - r\big(i\lambda^n - i\overline{\lambda}^n\big) = n\big(1 - \cos(2\pi\varphi n)\big) - r\sin(2\pi\varphi n)\,, \qquad (3.1)$$

where $\lambda = e^{i2\pi\varphi}$ is an algebraic number in $\mathbb{Q}(i)$ of modulus 1 which is not a root of unity. If u_n were negative, then $\cos(2\pi\varphi n)$ is close to 1. Let $m \in \mathbb{Z}$ such that $n\varphi - m$ is in $(-1/2, 1/2]$. Then, using the Taylor series expansion,

$$1 - \cos(2\pi\varphi n) = \frac{1}{2}(2\pi\varphi n - 2\pi m)^2 + O\big((2\pi\varphi n - 2\pi m)^4\big) \quad \text{and}$$
$$\sin(2\pi\varphi n) = (2\pi\varphi n - 2\pi m) + O\big((2\pi\varphi n - 2\pi m)^4\big)\,.$$

Hence, when $2\pi\varphi n - m$ decreases, we have that $u_n < 0$ roughly when $n(2\pi\varphi n - 2\pi m)^2 < r(n\varphi - 2\pi m)$. The latter is equivalent to

$$\left|\varphi - \frac{m}{n}\right| < \frac{r(2\pi)^{-1}}{n^2}\,.$$

For a real number x, define the *Lagrange constant* (or *homogeneous Diophantine approximation constant*)

$$L(x) = \inf\left\{c \in \mathbb{R}\colon \left|x - \frac{m}{n}\right| < \frac{c}{n^2} \text{ for some } n, m \in \mathbb{Z}\right\}\,,$$

and the *type* (or *homogeneous Diophantine approximation type*)

$$L_\infty(x) = \inf\left\{c \in \mathbb{R}\colon \left|x - \frac{m}{n}\right| < \frac{c}{n^2} \text{ for infinitely many } n, m \in \mathbb{Z}\right\}\,.$$

Now assume that (Ultimate) Positivity is decidable for LRS of the shape (3.1). Then, except the error from the Taylor approximation, we can vary r and apply Positivity oracles (respectively, Ultimate Positivity oracles) to compute the Lagrange constant (respectively, type) of φ up to arbitrary precision. However, computing $L_\infty(\varphi)$ and $L(\varphi)$ is a problem in Diophantine approximation that appears very difficult, and this shows that the (Ultimate) Positivity problem is 'Diophantine-hard'.

Hardness for simple LRS of order 10 Ouaknine and Worrell [120] also achieved the state-of-the-art for the Positivity problem for simple LRS from which they conclude Theorem 3.1.4.

Theorem 3.2.3. *Let $(u_n)_{n=0}^\infty$ be a non-degenerate simple $(\mathbb{R} \cap \overline{\mathbb{Q}})$-LRS with characteristic polynomial $f \in \mathbb{Z}[X]$ and a positive dominant root. If $f \in \mathbb{Z}[X]$ has either at most eight dominant roots or precisely nine roots, then we can determine whether $u_n \geq 0$ for each $n \in \mathbb{N}_0$.*

We will recall the techniques employed in Theorem 3.2.3. For the sake of brevity, we shall give only a brief outline here; we direct the interested reader to the full argument given in [120].

Let $(u_n)_{n=0}^\infty$ be a simple, non-degenerate $(\mathbb{R} \cap \overline{\mathbb{Q}})$-LRS satisfying the assumptions of Theorem 3.2.3 (again, we note that the study of the Positivity problem reduces to the study of non-degenerate LRS). By Lemma 3.2.1, we know that any hard instance has a positive real dominant root.

We then normalise $(u_n)_{n=0}^\infty$ such that the dominant roots have modulus 1. Thus, the non-real dominant roots $\lambda_1, \overline{\lambda_1}, \ldots, \lambda_k, \overline{\lambda_k}$ lie on the unit circle and the positive

real characteristic root of $(u_n)_{n=0}^{\infty}$ is equal to 1, and so $|\lambda_i| = 1$ for all $1 \leq i \leq k$. For each $n \in \mathbb{N}$, the normalised sequence $(u_n)_{n=0}^{\infty}$ satisfies the exponential-polynomial form

$$u_n = \alpha + \sum_{i=1}^{k} \alpha_i \lambda_i^n + \overline{\alpha_i}\,\overline{\lambda_i}^n + \sum_{i=1}^{k'} \beta_i \xi_i^n \,,$$

where $\xi_1, \ldots, \xi_{k'}$ are the normalised non-dominant roots of $(u_n)_{n=0}^{\infty}$ satisfying $0 < |\xi_i| < 1$ and $\alpha_1, \ldots, \alpha_k, \alpha, \beta_1, \ldots, \beta_{k'}$ are non-zero algebraic numbers. Thus, if write $\lambda_j = e^{2\pi i \varphi_j}$ and $\alpha_j = \frac{1}{2} a_1 e^{2\pi i \psi_j}$, we get that $\alpha_j \lambda_j^n + \overline{\alpha_j}\,\overline{\lambda_j}^n = a_j \cos(2\pi(\varphi_j n + \psi_n))$. Therefore,

$$u_n = a_1 \cos\big(2\pi(\varphi_1 n + \psi_1)\big) + \cdots + a_k \cos\big(2\pi(\varphi_k n + \psi_k)\big) + \alpha + O(\xi^n) \tag{3.2}$$

where $\xi < 1$. Before giving the general argument, we illustrate it through an example.

Example 3.2.4. Define the $\mathbb{Q}$-LRS $(u_n)_{n=0}^{\infty}$ by

$$\begin{aligned} u_n = &\left(\frac{3+4i}{5}\right)^n + \left(\frac{3-4i}{5}\right)^n - 2\left(\frac{5+12i}{65}\right)^n - 2\left(\frac{5-12i}{65}\right)^n \\ &+ 3\left(\frac{63-16i}{13}\right)^n + 3\left(\frac{63+16i}{13}\right)^n + 12 - (1/2)^n \,. \end{aligned}$$

Then, $\lambda_1 = \frac{3+4i}{5} = e^{i2\pi\varphi_1}$, $\lambda_2 = \frac{5+12i}{13} = e^{i2\pi\varphi_2}$, $\lambda_3 = \frac{63-16i}{65} = e^{i2\pi\varphi_3}$ are representatives of the complex conjugate pairs of the non-real dominant roots, $\alpha_1 = 1 = \frac{1}{2} 2 e^{2\pi 0}$, $\alpha_2 = 2 = \frac{1}{2} 4 e^{2\pi \frac{1}{2}}$, $\alpha_3 = 3 = \frac{1}{2} 6 e^{2\pi 0}$, 12, and 1 are the polynomial coefficients of the characteristic roots λ_1, λ_2, λ_3, 1, and $\xi = 1/2$, respectively. Then, $\lambda_3 = \overline{\lambda}_1 \lambda_2$, and λ_1 and λ_2 are multiplicatively independent. Thus, in the notation of (3.2), we get that

$$u_n = 2\cos(2\pi\varphi_1 n) + 4\cos(2\pi(\varphi_2 n + 1/2)) + 6\cos(2\pi(-\varphi_1 n + \varphi_2 n)) + 12 - (1/2)^n \,.$$

Using the trigonometric identities, we can rewrite the latter as

$$\begin{aligned} u_n = &\, 2\cos(2\pi\varphi_1 n) - 4\cos(2\pi\varphi_2 n) + 6\cos(2\pi\varphi_1 n)\cos(2\pi\varphi_2 n) \\ &- 6\sin(2\pi\varphi_1 n)\sin(2\pi\varphi_2 n) + 12 - (1/2)^n \,. \end{aligned}$$

By Kronecker's theorem (Theorem 1.1.13), $((\varphi_1 n \bmod 1, \varphi_2 n \bmod 1))_{n=0}^{\infty}$ is dense in the torus $\mathbb{T}^2$ (recall that $\mathbb{T} = [0,1)$). Thus,

$$\begin{aligned} &\liminf_{n\to\infty} 2\cos(2\pi\varphi_1 n) - 4\cos(2\pi\varphi_2 n) + 6\cos(2\pi\varphi_1 n)\cos(2\pi\varphi_2 n) \\ &\qquad - 6\sin(2\pi\varphi_1 n)\sin(2\pi\varphi_2 n) \\ &= \min_{(x_1,x_2)\in\mathbb{T}^2} 2\cos(2\pi x_1) - 4\cos(2\pi x_2) + 6\cos(2\pi x_1)\cos(2\pi x_2) \\ &\qquad - 6\sin(2\pi x_1)\sin(2\pi x_2) = -12 \,. \end{aligned}$$

where the minimum is achieved at $(x_1, x_2) = (1/2, 0)$. Hence, if $(u_n)_{n=0}^{\infty}$ is negative, $\varphi_1 n \bmod 1$ has to be (exponentially) close to $1/2$ and $\varphi_2 n \bmod 1$ has to be (exponentially) close to 0. Whether $\varphi_1 n \bmod 1$ is (exponentially) close to $1/2$ for some n can be verified with Baker's theorem on linear forms in logarithms, and positivity for $(u_n)_{n=0}^{\infty}$ is decidable. □

We return to the general case. In general, we can find all the multiplicative relationships between the normalised dominant roots $\lambda_1, \ldots, \lambda_k$ with Masser's theorem (Theorem 1.1.3). By renumbering the roots, we can assume that $\lambda_1, \ldots, \lambda_\ell$ form a maximal multiplicatively independent subset of $\{\lambda_1, \ldots, \lambda_k\}$ and that for $\ell + 1 \leq j \leq k$, we have that $\lambda_j = \lambda_1^{a_{j,1}} \cdots \lambda_\ell^{a_{j,\ell}}$ for some integers $a_{j,i}$. Then, $\varphi_1, \ldots, \varphi_\ell, 1$ are linearly independent over the rationals and so by Theorem 1.1.13 the points $\big((\varphi_1 n + \psi_1) \bmod 1, \ldots, (\varphi_\ell n + \psi_\ell) \bmod 1\big)$ are dense in the torus $\mathbb{T}^\ell$.

This induces a function $f : \mathbb{T}^\ell \to \mathbb{R}$ defined by

$$f(x_1, \ldots, x_\ell) = \sum_{j=1}^{k} 2a_j \cos\big(2\pi(x_j + \psi_j)\big) + \sum_{j=1}^{k} 2a_j \cos\Big(2\pi\Big(\sum_{i=1}^{\ell} a_{j,i} x_i + \psi_j\Big)\Big) + \alpha\,.$$

Then, $u_n = f(\varphi_1 n, \ldots, \varphi_\ell n) + O(\xi^n)$. Using trigonometric identities, we can rewrite the cosines to obtain a polynomial $P \in (\mathbb{R} \cap \overline{\mathbb{Q}})[X_1, \ldots, X_{2\ell}]$ such that

$$f(x_1, \ldots, x_\ell) = P\big(\cos(2\pi x_1), \sin(2\pi x_1), \ldots, \cos(2\pi x_\ell), \sin(2\pi x_\ell)\big)\,.$$

Then, using a density argument, we get

$$\begin{aligned}\mu := \liminf_{n\to\infty} u_n &= \min_{(x_1,\ldots,x_\ell)\in\mathbb{T}^\ell} P\big(\cos(2\pi x_1), \sin(2\pi x_1), \ldots, \cos(2\pi x_\ell), \sin(2\pi x_\ell)\big) \\ &= \min_{\substack{(y_1,\widetilde{y}_1,\ldots,y_\ell,\widetilde{y}_\ell)\in\mathbb{R}^{2\ell} \\ \forall j\colon y_j^2+\widetilde{y}_j^2=1}} P\big(y_1, \widetilde{y}_1, \ldots, y_\ell, \widetilde{y}_\ell\big)\big)\,.\end{aligned}$$

The last minimum is computable as it can be expressed in the theory of the reals, which Tarski [151] famously showed to be decidable. Moreover, we have that μ is a real algebraic number. If $\mu \neq 0$, the term $O(\xi^n)$ cannot influence the sign of u_n (and we can compute from which point onward this does not occur). Thus, for a hard instance, we can assume that $\mu = 0$. Let

$$X = \Big\{(x_1, \ldots, x_\ell) \in \mathbb{T}^\ell \colon P\big(\cos(2\pi x_1), \sin(2\pi x_1), \ldots, \cos(2\pi x_\ell), \sin(2\pi x_\ell)\big) = 0\Big\}\,.$$

Then, if u_n is negative, the vector $(\varphi_1 n, \ldots, \varphi_\ell n)$ has to be exponentially close (in n) to X such that the $O(\xi^n)$ can potentially force a negative term. That is, we compute a rational number $0 < r < 1$ such that $u_n < 0$ implies that $(\varphi_1 n, \ldots, \varphi_\ell n)$ is at most

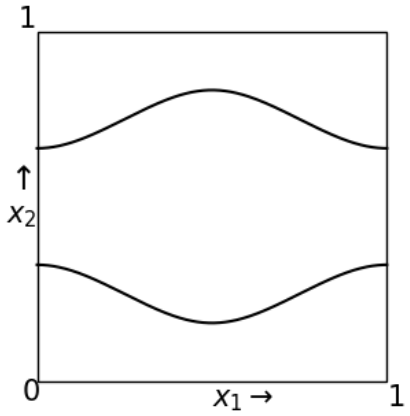

Figure 3.1: The set $X = \left\{(x_1, x_2) \in \mathbb{T}^2 \colon \cos(2\pi x_1) + 2\cos(2\pi x_2) = 0\right\}$.

r^n away from X. This occurs only finitely often.[1] For a single point $\mathbf{x} \in X$, we can compute all $n \in \mathbb{N}$ such that $n(\varphi_1, \ldots, \varphi_\ell)$ and $\mathbf{x}$ have distance at most r^n in $\mathbb{T}^\ell$ for some $n \in \mathbb{N}$, but we do not know how to do this when X is infinite.

The meat of the result of Ouaknine and Worrell is that X is finite when $\ell \leq 1$ or $k - \ell \leq 1$. The smallest pair (k, ℓ) such that their methods do not apply is $(4, 2)$, and the LRS also a real dominant root and some non-dominant root (as $\xi = 0$ forces the LRS to be positive or not ultimately positive). Thus, the hard example $(u_n)_{n=0}^\infty$ has to have at least order 10 (one dominant real root, one non-dominant root, and four complex conjugate pairs of dominant roots). We now give one example for order 10 of the so-called *squaring form*.

Example 3.2.5. We take the simple $\mathbb{Q}$-LRS $(u_n)_{n=0}^\infty$ defined by

$$u_n = \left(\left(\frac{3+4i}{5}\right)^n + \left(\frac{3-4i}{5}\right)^n + 2\left(\frac{5+12i}{13}\right)^n + 2\left(\frac{5-12i}{13}\right)^n\right)^2 - 2^{-n}.$$

Then, $(u_n)_{n=0}^\infty$ has order 10, $\ell = 2$, and

$$X = \left\{(x_1, x_2) \in \mathbb{T}^2 \colon \cos(2\pi x_1) + 2\cos(2\pi x_2) = 0\right\}.$$

We sketch X in Figure 3.1. □

As a by-product of this construction, we can extract an improvement on a result of Ouaknine and Worrell [120]. Instead of showing that the Skolem problem for order-5 LRS reduces to the Positivity problem for simple sequences of order 14, we can reduce to the Positivity problem for simple sequences of order 10.

[1] Ouaknine and Worrell prove Theorem 3.1.5 (the Ultimate Positivity problem is decidable for simple LRS) in this manner.

Theorem 3.2.6. *The Skolem problem at order* 5 *reduces to the Positivity problem for simple LRS of order* 10.

Proof. Let $(u_n)_{n=0}^{\infty}$ be an LRS of order 5. Using Lemma 1.2.3, we can assume that $(u_n)_{n=0}^{\infty}$ is non-degenerate, and as we can solve the Skolem problem for LRS in the MSTV class, we assume that $(u_n)_{n=0}^{\infty}$ is outside the MSTV class. Thus, by Lemma 2.2.6, the LRS $(u_n)_{n=0}^{\infty}$ has the exponential-polynomial form (2.3). Due to Theorem 3.3.2, the dominant roots are not units. Further, we can assume by Lemma 2.2.5 that for any prime ideal $\mathfrak{p}$ the $\mathfrak{p}$-dominant roots have $\mathfrak{p}$-valuation 0.

We claim that $|\rho| < |\lambda_1|$. As λ_1 is dominant, we have $|\rho| \leq |\lambda_1|$, so presume that equality holds. Further, let $\mathfrak{p}$ be a prime ideal containing λ_1 (which exists as λ_1 is not a unit). Then, $\lambda_1\overline{\lambda_1} = \lambda_2\overline{\lambda_2} = \rho^2 \in \mathfrak{p}$, and so ρ and at least one of λ_2 and $\overline{\lambda_2}$ is in $\mathfrak{p}$. As the $\mathfrak{p}$-dominant roots have $\mathfrak{p}$-valuation 0, there are thus at most two $\mathfrak{p}$-dominant roots, and hence $(u_n)_{n=0}^{\infty}$ would be in the MSTV class, contradicting our earlier assumptions. The claim follows.

We claim that ρ is an integer. If not, ρ would have a Galois conjugate among the other characteristic roots, say σ is a Galois automorphism such that $\sigma(\lambda_1) = \rho$. Then applying σ on $\lambda_1\overline{\lambda_1} = \lambda_2\overline{\lambda_2}$ and taking absolute values gives that

$$|\rho||\sigma(\overline{\lambda_1})| = |\sigma(\lambda_2)||\sigma(\overline{\lambda_2})|$$

where within each pair of absolute value bars, there is a different characteristic root. Thus, ρ is also a dominant root, contradicting our previous claim, and so ρ has no non-trivial Galois conjugates. As ρ is an algebraic integer, it is a non-zero integer.

By the Galois closure, $b \in \mathbb{Q}$. When $u_n = 0$,

$$\alpha_1\lambda_1^n + \overline{\alpha_1}\,\overline{\lambda_1}^n + \alpha_2\lambda_2^n + \overline{\alpha_2}\,\overline{\lambda_2}^n = -b\rho^n\,,$$

where both sides of the equation are rational. If the above holds, then clearly

$$v_n := \left(\alpha_1\lambda_1^n + \overline{\alpha_1}\,\overline{\lambda_1}^n + \alpha_2\lambda_2^n + \overline{\alpha_2}\,\overline{\lambda_2}^n\right)^2 - 2b^2\rho^{2n} < 0,$$

where $(v_n)_{n=0}^{\infty}$ is an order-10 LRS which by Theorem 1.2.7 and the construction above, is only 0 finitely often. Thus, apply Positivity oracles to $(v_n)_{n=N}^{\infty}$ (which is simple and has order 10) until one returns this sequence is positive. Then, $u_n = 0$ implies that $n < N$, and we can enumerate these finitely many numbers. □

Beyond orders 6 and 10 So far, we have seen two different obstacles that prevent progress on the Positivity problem:

1. For φ such that $e^{i2\pi\varphi}$ is algebraic, $r, q \in \mathbb{Q}$, can one compute all solutions to $|\varphi - \frac{m}{n}| < rn^{-q}$?

2. When X is an infinite curve as defined in the previous section, can one compute all $n \in \mathbb{N}$ such that $n(\varphi_1, \ldots, \varphi_\ell)$ come expectationally close to a point in X?

When investigating LRS of higher orders, one can encounter instances where these two problems are combined: can one compute all $n \in \mathbb{N}$ such that $n(\varphi_1, \ldots, \varphi_\ell)$ come polynomially close to a point in X?

3.3 The Positivity problem for reversible linear recurrence sequences

In this section, we will add an extra restriction (reversibility) to the linear recurrence sequences and show that this minor restriction already enlarges the realms of decidability for these decision problems for LRS. Recall the definition of a reversible LRS.

Definition 3.3.1. A $\mathbb{Z}$-LRS $(u_n)_{n=0}^{\infty}$ is *reversible* if its bi-completion $(u_n)_{n=-\infty}^{\infty}$ is integer-valued.

By an old result of Fatou [66] (see also [23, Chapter 7]), an LRS is reversible if and only if its constant coefficient is equal to ± 1. That is, reversible LRS are exactly the LRS whose characteristic roots are all units. Similarly, a polynomial $f(X) \in \mathbb{Z}[X]$ *reversible* if it is monic and its constant coefficient is ± 1, or, equivalently, all its roots are units.

The LRS $(\mathbf{t}^\top M^n \mathbf{s})_{n=0}^{\infty}$ is reversible when M is an integer-valued matrix with determinant ± 1 (that is, M is *uni-modular*). Examples of reversible LRS are polynomials (which are LRS that only have the characteristic root 1) and the Fibonacci numbers introduced in Example 1.2.4 whose bi-completion is $(\ldots, -3, 2, -1, 1, 0, 1, 1, 2, 3, \ldots)$.

Our goal for this section is to prove Theorems 3.1.6, 3.1.7, and 3.1.8. Theorem 3.1.6 was first proven by Lipton et al. [99] and was later reproven by Kenison [88]. In the same paper, Kenison also showed that the Positivity problem is decidable for simple reversible LRS of order 10 or less [88]. However, in both of these papers, the authors focused on these relatively low-order LRS, and we will deduce the result from a more general study of reversible sequences.

3.3.1 Reducing reversible LRS to reversible polynomials

We are going to deduce Theorems 3.1.6, 3.1.7, and 3.1.8 from the following result, which is proven in Section 3.3.2.

Theorem 3.3.2. *Let f be a non-degenerate reversible polynomial such that more than half of its roots are dominant. Then either f is linear or f is cubic with two dominant roots.*

For the Skolem, Positivity, and Ultimate Positivity problems, we can use Theorem 1.2.3 to reduce to non-degenerate LRS. This reduction does not increase the order and preserves simplicity and reversibility. The latter follows because the characteristic roots of any of these non-degenerate subsequences are powers of units and thus units.

This enables us to instantly prove Theorem 3.1.6: The Skolem problem is decidable for reversible LRS up to order 7.

Proof of Theorem 3.1.6. Assume that $(u_n)_{n=0}^{\infty}$ is a reversible, non-degenerate LRS of order $d \leq 7$. Due to Theorem 2.1.3, we can assume that $(u_n)_{n=0}^{\infty}$ has at least four dominant roots. Thus, the characteristic polynomial of $(u_n)_{n=0}^{\infty}$ is reversible and non-degenerate, has at least four dominant roots, and its degree is at most 7. This contradicts Theorem 3.3.2. □

Next, we turn our attention to the (Ultimate) Positivity problem. As a consequence of Lemma 3.2.1, we can reduce the problem of deciding positivity to LRS that possess a positive dominant root. Next, we deal with Theorem 3.1.8: The Positivity problem is decidable for reversible LRS up to order 17.

Proof of Theorem 3.1.8. As previously noted, we can assume that the LRS (and thus its characteristic polynomial f) is simple, reversible, non-degenerate), has a dominant root in $\mathbb{R}_{>0}$, and its degree is at most 17. Due to Theorem 3.3.2, we can assume that f is cubic, linear, or that at most half of its characteristic roots are dominant. Thus, f has at most eight dominant roots. Hence, by Theorem 3.2.3, positivity is decidable. □

Lastly, we tackle Theorem 3.1.7: positivity and ultimate positivity are decidable for reversible LRS up to order 11.

Proof of Theorem 3.1.7. Let $(u_n)_{n=0}^{\infty}$ be a reversible LRS of order at most 11 for which we have to prove that (ultimate) positivity is decidable. Then, by the same arguments earlier in this section, we can assume that $(u_n)_{n=0}^{\infty}$ is non-degenerate. Moreover, by Lemma 3.2.1, we can assume that $(u_n)_{n=0}^{\infty}$ has a real positive dominant root ρ and a pair of non-real complex conjugate dominant roots λ_i and $\overline{\lambda_i}$.

Claim 3.3.3. *We have that $(u_n)_{n=0}^{\infty}$ has a non-simple dominant root.*

Proof. If $(u_n)_{n=0}^{\infty}$ would be simple, then positivity and ultimate positivity would be decidable by Theorem 3.1.5 and Theorem 3.1.8, respectively as $(u_n)_{n=0}^{\infty}$ has order at most 11. Thus, $(u_n)_{n=0}^{\infty}$ has a non-simple characteristic root.

Otherwise, due to the Galois closure of the polynomial form, we can split $(u_n)_{n=0}^{\infty}$ into two reversible LRS $(v_n)_{n=0}^{\infty}$ and $(w_n)_{n=0}^{\infty}$ where $(v_n)_{n=0}^{\infty}$ is simple and contains all the non-dominant roots, $(w_n)_{n=0}^{\infty}$ is non-simple and its characteristic roots have modulus strictly smaller than ρ, and $u_n = v_n + w_n$ for all $n \in \mathbb{N}$.

By Lemma 1.2.6, there are positive rational numbers r, R such that $R < \lambda$ and $|w_n| < rR^n$ for all $n \in \mathbb{N}$. Let $(u_n^-)_{n=0}^{\infty}$ and $(u_n^+)_{n=0}^{\infty}$ be the simple $\mathbb{Q}$-LRS defined by $u_n^- = u_n - rR^n$ and $u_n^+ = u_n + rR^n$ such that $u_n^- < u_n < u_n^+$ for all $n \in \mathbb{N}$. As $(w_n)_{n=0}^{\infty}$ has a non-simple characteristic root, we know that $(w_n)_{n=0}^{\infty}$ has order at least two. Thus, $(v_n)_{n=0}^{\infty}$ has order at most $9 = 11 - 2$. By Theorem 3.3.2, $(v_n)_{n=0}^{\infty}$ has at most four dominant roots.

Theorem 1.2.7 implies that the signs of u_n^- and u_n^+ differ for only finitely $n \in \mathbb{N}$, and thus the same holds for $(u_n)_{n=0}^{\infty}$ and $(u_n^-)_{n=0}^{\infty}$. Hence, $(u_n)_{n=0}^{\infty}$ is ultimately positive if and only if $(u_n^-)_{n=0}^{\infty}$ is ultimately positive, where the latter is decidable due to Theorem 3.1.5. Thus, ultimate positivity is decidable for $(u_n)_{n=0}^{\infty}$.

If $(u_n)_{n=0}^{\infty}$ is not ultimately positive, the LRS $(u_n)_{n=0}^{\infty}$ is not positive. Otherwise, $(u_n^-)_{n=0}^{\infty}$ is also ultimately positive, and so there are only finitely many n such that $u_n^- < 0$. As $(u_n^-)_{n=0}^{\infty}$ has the same dominant roots as $(v_n)_{n=0}^{\infty}$, the sequence $(u_n^-)_{n=0}^{\infty}$ has at most four dominant roots and is simple, and so, for every $k \in \mathbb{N}$, we can determine whether $(u_n^-)_{n=k}^{\infty}$ is positive due to Theorem 3.2.3. As $(u_n^-)_{n=0}^{\infty}$ is ultimately positive, the LRS $(u_n^-)_{n=k}^{\infty}$ is positive for some $k \in \mathbb{N}$ (which we can decide for all $k \in \mathbb{N}$). Then $(u_n)_{n=0}^{\infty}$ is positive if and only if $u_n \geq 0$ for all $0 \leq n < k$. □

Thus, we can assume that one of the dominant roots of $(u_n)_{n=0}^{\infty}$ is non-simple. By Lemma 3.2.1, both the real dominant root and a pair of complex conjugate roots are non-simple. We can assume this complex conjugate pair is λ and $\overline{\lambda}$.

Let f be the monic integer-valued polynomial of the smallest degree with ρ and λ as roots. Then, f is non-degenerate and reversible. By Theorem 3.3.2, it follows that

at most half of the roots of f are dominant if f is neither linear nor cubic with two dominant roots. As such, f has degree at least 6 and, additionally, as each of these roots is non-simple (being a Galois conjugate of either ρ or λ), the sequence $(u_n)_{n=0}^{\infty}$ has order at least 12.

We thus deduce the desired result: Positivity and ultimate positivity are decidable for reversible LRS up to order 11. □

3.3.2 Root analysis of reversible polynomials

The main result of this section is Theorem 3.3.2. Essentially, the theorem says that excepting a number of special cases, no more than half of the roots of such a polynomial can be dominant.

We will prove Theorem 3.3.2 by studying identities between the roots of irreducible polynomials. We employ a powerful result due to Dubickas and Smyth [59], Theorem 3.3.4 below, concerning necessary conditions for an algebraic unit and all its Galois conjugates to lie on two concentric circles centred at the origin. (Theorem 3.3.4 is a specialisation of the general result [59, Theorem 2.1].)

Theorem 3.3.4. *Let $f \in \mathbb{Z}[X]$ be an irreducible, reversible polynomial of degree d whose roots lie on two circles centred at the origin with respective r and R. Without loss of generality, at most half of the roots of f lie on the circle of radius r. Then, either*

1. *d is even, and half of the roots lie on the circle of radius r; or*
2. *d is a multiple of three, and a third of the roots lie on the circle of radius r. Moreover, for every root β on the circle of radius r, there exists $n > 0$ such that $\beta^n \in \mathbb{R}$.*

We also need the following lemma, versions of which were proven by Smyth [146] and Ferguson [67].

Lemma 3.3.5. *Suppose that λ is an algebraic number with Galois conjugates β and γ satisfying $\lambda^2 = \beta\gamma$. Then λ/β is a root of unity.*

The last lemma has the following immediate helpful consequence.

Lemma 3.3.6. *Let $f \in \mathbb{Z}[X]$ be an irreducible, non-degenerate polynomial with a real root ρ. Then f has exactly one root of modulus $|\rho|$.*

Proof. If λ also has modulus ρ, we have $\rho^2 = \lambda\overline{\lambda}$ and hence ρ/λ is a root of unity by Lemma 3.3.5. As f is non-degenerate, $\rho = \lambda$; that is, f has exactly one dominant root. □

Lemma 3.3.7. *Suppose that $f \in \mathbb{Z}[X]$ is irreducible, non-degenerate, and reversible, with $2m$ non-real dominant roots and no real dominant roots. Then f is constant, we have $(\deg(f), m) = (3, 1)$ or $\deg(f) > 3m$.*

Proof. Since f has at least $2m$ roots, it is clear that $\deg(f) \geq 2m$. If $m = 0$, f is a constant polynomial. Thus, we can assume that $m \geq 1$.

We first show $\deg(f) > 2m$. Assume, to get a contradiction, that $\deg(f) = 2m$. Then the roots of f all lie on the circumference of some circle centred at the origin. As f is reversible, the polynomial f is monic and its constant coefficient is ± 1 (which, by Vieta's formulas, equals the product of the roots of f). Thus, all roots of f have modulus 1 and lie on the unit circle. So, by Theorem 1.1.1, they are therefore roots of unity. As $m \geq 1$, the polynomial f has at least two roots, whose quotient is thus a root of unity. We have reached a contradiction: f was assumed to be non-degenerate. Thus $\deg(f) > 2m$.

We study the case $m = 1$. Then, $\deg(f) \geq 3$. If $\deg(f) = 3$, we are the second case of the lemma, and if $\deg(f) > 3 = 3m$, we have the third case. Thus, we assume that $m \geq 2$.

We now show that $m \geq 2$ implies that $\deg(f) \geq 3m$. Let $\lambda_1, \overline{\lambda_1}, \ldots, \lambda_m, \overline{\lambda_m}$ be the $2m$ dominant roots of f. Thus $\lambda_1\overline{\lambda_1} = \lambda_i\overline{\lambda_i}$ for each $i \in \{1, \ldots, m\}$. Since $2m < \deg(f) < 3m$, the polynomial f has somewhere between 1 and $m-1$ non-dominant roots. Let γ be one such non-dominant root. Further, since f is irreducible, there is a Galois automorphism σ such that $\sigma(\lambda_1) = \gamma$. Then, the equation

$$\gamma\sigma(\overline{\lambda_1}) = \sigma(\lambda_2)\sigma(\overline{\lambda_2}) = \cdots = \sigma(\lambda_m)\sigma(\overline{\lambda_m}) \tag{3.3}$$

contains $2m$ different roots of f. As there are at most $m-1$ non-degenerate roots, the roots $\sigma(\lambda_i)$ and $\sigma(\overline{\lambda_i})$ are both not non-dominant for at least one $2 \leq i \leq m$. That is, $\sigma(\lambda_i)$ and $\sigma(\overline{\lambda_i})$ are dominant. But, $|\gamma\sigma(\overline{\lambda_1})| = \sigma(\lambda_i)\sigma(\overline{\lambda_i})|$ cannot hold, as $|\gamma| < |\sigma(\lambda_i)|$ and $|\sigma(\overline{\lambda_1})| \leq |\sigma(\overline{\lambda_i})|$. Thus, we have a contradiction, and so f has at least m non-dominant roots. Thus, $\deg(f) \geq m + 2m = 3m$.

Finally, we eliminate the case that $\deg(f) = 3m$ when $m \geq 2$. Using the same argument as before (with the same choice of γ and σ), we deduce that for all $1 \leq i \leq m$, exactly one of $\sigma(\lambda_i)$ and $\sigma(\overline{\lambda_i})$ is dominant, and the other is non-dominant. By taking absolute values in (3.3), we see that all non-dominant roots have an equal

modulus. Thus, the roots lie on two circles centred at the origin: the m non-dominant roots lie on one, and the $2m$ dominant roots lie on the other. Thus, by Theorem 3.3.4, for each non-dominant root γ_i there is an n_i such that $\gamma_i^{n_i} \in \mathbb{R}$. Then, $|(\gamma_1/\gamma_2)^{n_1 n_2}| = 1$ as $|\gamma_1| = |\gamma_2|$ while $\gamma_1^{n_1 n_2} = (\gamma_1^{n_1})^{n_2}$ and $\gamma_2^{n_1 n_2} = (\gamma_2^{n_2})^{n_1}$ are both real numbers. Thus, $(\gamma_1/\gamma_2)^{n_1 n_2} = \pm 1$ and so γ_1/γ_2 is a root of unity, contradicting our assumption that f is non-degenerate.

Hence, $\deg(f) > 3m$ if $m \geq 2$, from which the desired result follows. □

To improve the bound from $\deg(f) > 3m$ to $\deg(f) \geq 4m$, we shall introduce new and novel techniques for counting symmetries in the roots of f. Let $\lambda_1, \ldots, \lambda_\ell$ be the roots of f. The interesting case occurs when all the dominant roots of f are non-real. We denote the dominant roots of f by $\lambda_1, \overline{\lambda_1}, \ldots, \lambda_m, \overline{\lambda_m}$. Let $\mu_1 := \lambda_1 \overline{\lambda_1}$ and g be the minimal polynomial of μ_1 (hereafter we shall refer to g as the *dominating polynomial of* f). Let $\mu_2, \ldots, \mu_n$ be the Galois conjugates of μ_1 (and thus the other roots of g) and $\sigma_1, \ldots, \sigma_n$ the Galois automorphisms associated with g such that $\sigma_j(\mu_1) = \mu_j$.

Set $K = \mathbb{Q}(\mu_1, \ldots, \mu_n)$ and $L = \mathbb{Q}(\lambda_1, \ldots, \lambda_\ell)$. Clearly, $K \subset L$, and so each σ_j can be lifted to an automorphism $\widetilde{\sigma}_j$ in $\mathrm{Gal}_{\mathbb{Q}}(L)$ such that $\widetilde{\sigma}_j|_K = \sigma_j$. Applying these $\widetilde{\sigma}_j$ on $\lambda_1\overline{\lambda_1} = \cdots = \lambda_m\overline{\lambda_m} = \mu_1$ gives rise to the following n equations:

$$\begin{aligned} \alpha_{1,1,1}\alpha_{1,1,2} = \cdots = \alpha_{m,1,1}\alpha_{m,1,2} &= \mu_1 \\ \vdots \qquad \vdots \qquad\qquad \vdots & \\ \alpha_{1,n,1}\alpha_{1,n,2} = \cdots = \alpha_{m,n,1}\alpha_{m,n,2} &= \mu_n \end{aligned} \tag{3.4}$$

where $\alpha_{i,j,1} = \widetilde{\sigma}_j(\lambda_i)$ and $\alpha_{i,j,2} = \widetilde{\sigma}_j(\overline{\lambda_i})$. Since each $\alpha_{i,j,k}$ is a Galois conjugate of a dominant root of f, each $\alpha_{i,j,k}$ is also a root f. Given a root λ of f, we define the *equation number*

$$E = \#\big\{(i,j,k) \colon \alpha_{i,j,k} = \lambda \text{ for } 1 \leq i \leq m, 1 \leq j \leq n, k = 1,2\big\}.$$

In Lemma 3.3.9, we will show that E is independent of the choice of root λ. It is useful to see the two roots of f in one position in one equation in (3.4) as a *pair*. In other words, $\alpha_{i,j,1}$ and $\alpha_{i,j,2}$ are *paired* for all $1 \leq i \leq m$ and $1 \leq j \leq n$. Further, for $j = 1, \ldots, n$, let $\mathcal{A}_j := \{\alpha_{1,j,1}, \alpha_{1,j,2}, \ldots, \alpha_{m,j,1}, \alpha_{m,j,2}\}$. Note that $\#\mathcal{A}_j = 2m$, as $\widetilde{\sigma}_j$ is a bijection between the set of dominant roots of f and $\mathcal{A}_j$.

We claim that $\mathcal{A}_j$ is independent of the choice of lift $\widetilde{\sigma}_j$ of σ_j. If λ and λ' are roots of f such that $\lambda\lambda' = \mu_j$, then $\widetilde{\sigma}_j^{-1}(\lambda)\widetilde{\sigma}_j^{-1}(\lambda') = \mu_1 = \lambda_1\overline{\lambda_1}$. Thus, $\widetilde{\sigma}_j^{-1}(\lambda)$ and $\widetilde{\sigma}_j^{-1}(\lambda')$ are dominant roots of f. Further, since $\widetilde{\sigma}_j$ is a bijection, $\lambda = \widetilde{\sigma}_j\big(\widetilde{\sigma}_j^{-1}(\lambda)\big) \in \mathcal{A}_j$ (and similarly $\lambda' \in \mathcal{A}_j$). We make two deductions. First, if μ_j is the product of two distinct

roots of f, then those roots are two elements of $\mathcal{A}_j$. Second, we infer our claim that $\mathcal{A}_j$ is independent of the choice of $\widetilde{\sigma}_j$.

In the case of one dominant root, the same construction applies: g is defined as the minimal polynomial of λ_1^2, where λ_1 is the sole dominant root of f. By non-degeneracy, the squares of all roots of f are distinct, and so $\deg(f) = \deg(g)$, $\mu_j = \lambda_j^2$ for $j = 1, \ldots, \deg(f)$, and $E = 2$ for all roots of f (it appears once as a square). Only, $\mathcal{A}_j = \{\lambda_j\}$ consists of exactly one root of f.

Lemma 3.3.8. *Suppose that $f \in \mathbb{Z}[x]$ is reversible, non-degenerate, and irreducible with $2m$ non-real dominant roots and has degree less than $4m$. Write g for the dominating polynomial of f. Then g is also reversible, non-degenerate, and irreducible.*

Proof. By construction, g is irreducible, and as all the roots are products of roots of f (which are units), g is reversible. Assume, to get a contradiction with g being non-degenerate, that a quotient of roots of g, say $\mu_j/\mu_{j'}$, is a root of unity. Both sets of roots $\mathcal{A}_j$ and $\mathcal{A}_{j'}$ have cardinality $2m$. Since $\deg(f) < 4m = \#\mathcal{A}_j + \#\mathcal{A}_{j'}$, we deduce that $\mathcal{A}_j \cap \mathcal{A}_{j'}$ is non-empty. Let $\lambda \in \mathcal{A}_j \cap \mathcal{A}_{j'}$ and κ, κ' be roots of f such that $\lambda\kappa = \mu_j$ and $\lambda\kappa' = \mu_{j'}$. Since $\mu_j \neq \mu_{j'}$, we have $\kappa \neq \kappa'$. It follows that f is degenerate because $\kappa/\kappa' = \mu_j/\mu_{j'}$ is a root of unity. From this contradiction, we deduce that g is non-degenerate. □

Lemma 3.3.9. *Suppose that $f \in \mathbb{Z}[X]$ is reversible, non-degenerate, and irreducible with $2m$ non-real dominant roots dominant polynomial g. Then all the roots of f have the same equation number E and*

$$2m \deg(g) = E \deg(f)\,. \tag{3.5}$$

Proof. We use the notation of $\lambda_i, \mu_j, \sigma_j, \widetilde{\sigma}_j$, $\alpha_{i,j,k}$, K, L, etc. as above.

Set $H = \mathrm{Gal}_{\mathbb{Q}}(K)$ and $G = \mathrm{Gal}_{\mathbb{Q}}(L)$. By the Orbit-Stabilizer theorem, the number of $\sigma \in H$ such that $\sigma(\mu_1) = \mu_j$ is independent of the choice of $1 \leq j \leq n$. Now each $\sigma \in H$ has the same number of lifts to G, and so the number of elements of G that map μ_1 to each μ_j is independent of $1 \leq j \leq n$. Thus, the number of elements of G such that the image of $\mathcal{A}_1$ is $\mathcal{A}_j$ is also independent of the choice of $1 \leq j \leq n$.

We claim there is no pair of distinct j_1 and j_2 for which $\mathcal{A}_{j_1} = \mathcal{A}_{j_2}$. Indeed, we assume, to get a contradiction, that $\mathcal{A}_{j_1} = \mathcal{A}_{j_2}$ for $j_1 \neq j_2$. Then $\mu_{j_1} \neq \mu_{j_2}$ and

$$\mu_{j_1}^m = \prod_{i=1}^{m} \alpha_{i,j_1,1}\alpha_{i,j_1,2} = \prod_{i=1}^{m} \alpha_{i,j_2,1}\alpha_{i,j_2,2} = \mu_{j_2}^m\,.$$

Thus μ_{j_1}/μ_{j_2} is a root of unity. However, by Lemma 3.3.8, the polynomial g is non-degenerate, causing the contradiction.

We also make the following observation. By the Orbit-Stabilizer theorem, for every choice of two roots λ and λ' of f, the number of $\sigma \in G$ such that $\sigma(\lambda) = \lambda'$ is equal. Thus for each root λ of f, the number of $\sigma \in G$ such that one of $\widetilde{\sigma}(\lambda_1), \widetilde{\sigma}(\overline{\lambda_1}), \ldots, \widetilde{\sigma}(\lambda_m), \widetilde{\sigma}(\overline{\lambda_m})$ equals λ is independent of the choice of λ. Thus, the equation number E is independent of the choice of the root λ.

The equation $2m \deg(g) = E \deg(f)$ follows from counting the number of $\alpha_{i,j,k}$. On the one hand, there are $\deg(g)$ equations with $2m$ entries as all $\mathcal{A}_j$ are distinct. On the other hand, the roots of $\deg(f)$ each appear E times. □

The following result increases the bound on the degree of f to $\deg(f) \geq 4m$.

Theorem 3.3.10. *Let $f \in \mathbb{Z}[X]$ be an irreducible, non-degenerate, and reversible polynomial with $2m$ dominant non-real roots and no real dominant roots. Then $(\deg(f), m) = (3, 1)$ or $\deg(f) \geq 4m$.*

Proof. Assume, to get a contradiction, that f is a counterexample of lowest degree.

From Lemma 3.3.7, we deduce that $\deg(f) > 3m$ if we are not in the exceptional case $(\deg(f), m) = (3, 1)$. As f is a counterexample to Theorem 3.3.10, we have $\deg(f) < 4m$ as well. Therefore, $3m < \deg(f) < 4m$, and so $m \geq 2$ as $\deg(f) \in \mathbb{N}$. We shall employ the preceding notation for the dominating polynomial g, the sets of roots $\mathcal{A}_j$ of f, and the equation number E.

In each equation in (3.4), there are $2m$ distinct roots of f. Since $\deg(f) < 4m$ and f has $2m$ dominant roots, the polynomial f has less than $2m$ non-dominant roots, and so each such equation contains at least one dominant root of f. In an equation in (3.4), the smallest root of f appearing in the equation has to be paired with the largest root of f in the equation, which is dominant by the argument above. Let γ be a root of f with minimal absolute value, then in any equation in (3.4) where γ appears, the root γ is paired with a dominant root. Thus, $|\gamma\lambda_1|$ is the minimal absolute value attained by any root of g. We now show that at least half of the roots of g lie on the circle $\{z \in \mathbb{C} \colon |z| = |\gamma\lambda_1|\}$. Using (3.5) and that $\deg(f) < 4m$, we obtain that $E > \deg(g)/2$. Thus, γ is in more than half of the equations in (3.4). Each such equation corresponds to a root of g of minimal absolute value, which has modulus $|\gamma\lambda_1|$. Hence, $\pm X^n g(X^{-1})$ is an irreducible, reversible, non-degenerate such that more than half of its characteristic roots have modulus $|\gamma\lambda_1|$ and are thus dominant.

Hence, g is either a counterexample to the theorem or $\deg(g) = 3$. In the latter case, as γ is not dominant, we have $E < \deg(g) = 3$. Meanwhile, $E > \deg(g)/2$,

and so $E = 2$. Then, by (3.5), we have $\deg(f) = 3m$, which we excluded. As f is a counterexample of lowest degree, we have that $\deg(g) \geq \deg(f)$. Thus, $2m \leq E$ by (3.5).

If λ_1 is paired with a dominant root λ, then $\mu_j = \lambda\lambda_1$ and $|\mu_j| = \mu_1$. As g is non-degenerate by Lemma 3.3.8, Lemma 3.3.6 implies that g has no other roots of modulus μ_1 as μ_1 is real. Thus, $\overline{\lambda_1}$ is the only dominant root with which λ_1 is paired, and so λ_1 is paired with at most $\deg(f) - 2m$ non-dominant roots of f and $\overline{\lambda_1}$. This gives the upper bound $E \leq \deg(f) - 2m + 1 \leq 2m$. Combined with $2m \leq E$, we have that $E = 2m$ and $E = \deg(f) - 2m + 1$ and so $\deg(f) = 4m - 1$. Therefore, $\deg(g) = \deg(f) = 4m - 1$ by (3.4).

Each of the $2m$ dominant roots pair exactly with their respective complex conjugate and all $2m - 1$ non-dominant roots. Let γ and γ' be non-dominant. Both are paired with $2m$ dominant roots and no other roots as $E = 2m$. Further, as there are $4m - 1$ different equations, the roots γ and γ' both appear in at least one equation in (3.4) where they are therefore paired with a dominant root. Thus, $|\gamma| = |\gamma'|$ and so f has exactly $2m - 1$ non-dominant roots with the same modulus. As $2m - 1$ is odd, at least one non-dominant root is real. Hence, Lemma 3.3.6 implies that as f is non-degenerate, we have $2m - 1 = 1$ and so $m = 1$. This contradicts that $m \geq 2$. □

In Theorem 3.3.10, we made the superfluous assumption that f is irreducible. We circumvent the irreducibility assumption with a careful case analysis.

Proof of Theorem 3.3.2. Let f be a counterexample of minimal degree, and factor f into irreducible reversible polynomials $f_1, \ldots, f_k$. For $1 \leq i \leq k$, let m_i' be the number of dominant roots of f_i. Call an irreducible factor *sharp* if $2m_i' = \deg(f_i)$ and *special* if $2m_i' > \deg(f_i)$. From Lemma 3.3.6 and Theorem 3.3.10, it follows that if an irreducible factor is special, then $(\deg(f_i), m_i') = (1, 1)$ or $(3, 2)$. If $k = 1$, the polynomial f is irreducible, and the result follows automatically. Thus, we can assume that $k \geq 2$. Since f is a counterexample of minimal degree, a straightforward proof by contradiction permits us to assume $k = 2$ and that the dominant roots of each f_i are dominant for f. Thus, our argument reduces to the following cases: we need only show that the product of either two special polynomials or a special and sharp polynomial breaks the hypothesis. By renumbering, we can assume f_1 is special and f_2 is either sharp or special. We observe that under our assumptions, the dominant roots of f_1 and f_2 are necessarily equal in absolute value and, as we do not count multiplicity, we have $f_1 \neq f_2$.

We begin our case analysis. First, consider the case where $(\deg(f_1), m_1') = (1, 1)$. Then $f_1(X) = X \pm 1$ as f_1 is reversible. Thus, the dominant roots of f_2 also have modulus 1. As the product of the roots of f_2 is ± 1 by the Vieta formulas, they all lie on the unit circle. Hence, using Theorem 1.1.1, the roots of f_2 are roots of unity. Thus, the quotient of a root of f_1 and f_2 is a root of unity, causing a contradiction.

Second, we suppose that $(\deg(f_1), m_1') = (3, 2)$. Following the argument in the preceding case, either $(\deg(f_2), m_2') = (3, 2)$ or $\deg(f_2) = 2m_2'$. In the former, the non-dominant roots γ_1 and γ_2 of f_1 and f_2 (respectively) are both real and equal in modulus. This is straightforward to see since each f_j is of the form

$$f_j = (x - \gamma_j)(x - R\mathrm{e}^{\mathrm{i}\theta_j})(x - R\mathrm{e}^{-\mathrm{i}\theta_j})$$

with constant coefficient ± 1 for some $R \in \mathbb{R}_{>0}$. Thus, $\gamma_j := \pm R^{-2}$. We cannot have two such irreducible factors since then the ratio $\gamma_1/\gamma_2 = \pm 1$, which contradicts the non-degeneracy assumption on f.

We continue with the latter subcase $(\deg(f_1), m_1') = (3, 2)$ and $\deg(f_2) = 2m_2'$. Since the dominant roots of f_1 and f_2 are dominant roots of f, the dominating polynomials of f_1 and f_2 are one and the same, say g. Let E_1 and E_2 be the respective equation numbers of f_1 and f_2. From (3.5), we obtain that $2\deg(g) = 3E_1$ and that E_1 is even. Since $1 \leq E_1 \leq \deg(f_1) = 3$ (each pairing is distinct), we have that $E_1 = 2$ and, it follows immediately, $\deg(g) = 3$. By (3.5), $2m_2' \deg(g) = E_2 \deg(f_2)$ and as $\deg(g) = 3$ and $\deg(f_2) = 2m_2'$, we deduce that $E_2 = 3/2$ is not an integer.

We have exhausted the possibilities for constructing a minimal counterexample f and find that no such counterexample exists. We have thus proved Theorem 3.3.2. □

3.4 Hard instances of the Positivity problem

In this section, we extend the difficulties found for the Positivity problem for general LRS in Section 3.2 to reversible LRS to show that our methods above are optimal and cannot be improved. Specifically, we construct a simple reversible LRS of order 18 and sketch the construction of a reversible LRS of order 12 that, to the best of our knowledge, for which the same difficulties apply as in Section 3.2.

The calculations involved in preparing these hard instances were performed in SageMath [58].

The Positivity problem for simple, reversible linear recurrence sequences

We start with the most complicated hard instances: constructing a simple, reversible order 18 for which we cannot decide positivity.

Example 3.4.1. Take the irreducible polynomial

$$f(X) = X^8 - 3X^7 + 4X^6 - 4X^5 + 11X^4 - 21X^3 + 19X^2 - 7X + 1\,,$$

which has eight non-real roots $\lambda_1, \overline{\lambda_1}, \lambda_2, \overline{\lambda_2}, \lambda_3, \overline{\lambda_3}, \lambda_4, \overline{\lambda_4}$ such that λ_1 and λ_2 are dominant, the roots λ_3 and λ_4 are both non-dominant, and $|\lambda_4| < 1 < |\lambda_3| \approx 1.143$. Let $\varphi := (1+\sqrt{5})/2$ denote the golden ratio. Then, with a certain labelling of complex conjugates,

$$\lambda_1\overline{\lambda_1} = \lambda_2\overline{\lambda_2} = \varphi^2 \quad \text{and} \quad \lambda_3\lambda_4 = \overline{\lambda_3}\,\overline{\lambda_4} = \varphi^{-2}\,,$$

which, due to the number of relations, severely limits the possible Galois automorphisms. In particular, the Galois group has the form of a *wreath product* $D_4 \wr C_2$. Thus a dihedral group D_4 acts on $\lambda_1, \overline{\lambda_1}, \lambda_2$, and $\overline{\lambda_2}$ and is generated by the elements (written in cycle notation) $(\lambda_1\ \lambda_2\ \overline{\lambda_1}\ \overline{\lambda_2})$ and $(\lambda_1\ \overline{\lambda_1})$. A second dihedral group D_4 acts on $\lambda_3, \overline{\lambda_3}, \lambda_4, \overline{\lambda_4}$ and is generated by $(\lambda_3\ \overline{\lambda_3}\ \lambda_4\ \overline{\lambda_4})$ and $(\lambda_3\ \lambda_4)$. Lastly, there is a cyclic C_2 group acting on these two sets of four roots generated by the permutation $(\lambda_1\ \lambda_3)(\overline{\lambda_1}\ \lambda_4)(\lambda_2\ \overline{\lambda_3})(\overline{\lambda_2}\ \overline{\lambda_4})$.

The terms in the hard sequence $(u_n)_{n=0}^{\infty}$ are given as follows:

$$\begin{aligned} u_n = \frac{1}{\sqrt{5}}\Big((1+\lambda_1)\lambda_1^n + (1+\overline{\lambda_1})\overline{\lambda_1}^n + (1+\lambda_2)\lambda_2^n + (1+\overline{\lambda_2})\overline{\lambda_2}^n\Big)^2 \\ - \frac{1}{\sqrt{5}}\Big((1+\lambda_3)\lambda_3^n + (1+\overline{\lambda_3})\overline{\lambda_3}^n + (1+\lambda_4)\lambda_4^n + (1+\overline{\lambda_4})\overline{\lambda_4}^n\Big)^2 . \end{aligned}$$

By the action of the Galois group, it can be seen that each term u_n is rational and further that $(u_n)_{n=0}^{\infty}$ is simple, reversible, and has exactly order 18. The initial values $u_0, \ldots, u_{17}$ of $(u_n)_{n=0}^{\infty}$ are

$$\begin{aligned} &-11, -8, 0, 240, 704, -20, 192, 5508, 46305, 2625, 13425, 73117, 2469800, 536000, \\ &554151, 77287, 108792361, 66461616\,. \end{aligned}$$

The simple LRS $(u_n)_{n=0}^{\infty}$ satisfies the relation

$$\begin{aligned} u_{n+18} = u_{n+17} - 10u_{n+16} + 6u_{n+15} + 43u_{n+14} - 93u_{n+13} + 672u_{n+12} - 596u_{n+11} \\ + 120u_{n+10} + 3972u_{n+9} - 15345u_{n+8} + 29654u_{n+7} - 36108u_{n+6} + 23847u_{n+5} \\ - 9572u_{n+4} + 2361u_{n+3} - 325u_{n+2} + 26u_{n+1} - u_n\,. \end{aligned}$$

Observe that u_0, u_1, and u_5 are negative, but these are the only negative terms up to $n = 10^5$. Thus, the question is to prove that $u_n \geq 0$ for all $n \geq 6$.

It remains to show that the torus T associated with $(u_n)_{n=0}^{\infty}$ has the prescribed 'squaring form' (as in Example 3.2.5) and that $(u_n)_{n=0}^{\infty}$ is non-degenerate. To start, the numbers u_n and $\frac{u_n}{\varphi^{2n}}$ have the same sign. Moreover, we observe that $|1+\lambda_1| \neq |1+\lambda_2|$ and that both λ_1/φ and λ_2/φ lie on the unit circle. For $a = 1+\lambda_1, b = \lambda_2$ and some $0 < r < 1$, we have that

$$\begin{aligned}\frac{u_n}{\varphi^{2n}} &= \frac{1}{\varphi^{2n}}\left((1+\lambda_1)\lambda_1^n + (1+\overline{\lambda_1})\overline{\lambda_1}^n + (1+\lambda_2)\lambda_2^n + (1+\overline{\lambda_2})\overline{\lambda_2}^n\right)^2 + O(r^n) \\ &= \left(a\left(\frac{\lambda_1}{\varphi}\right)^n + \overline{a}\left(\frac{\lambda_1}{\varphi}\right)^{-n} + b\left(\frac{\lambda_2}{\varphi}\right)^n + \overline{b}\left(\frac{\lambda_2}{\varphi}\right)^{-n}\right)^2 + O(r^n)\end{aligned}$$

is close to the 'squaring form' discussed in Example 3.2.5. To verify that we indeed have such a hardness, we have to verify that λ_1/φ and λ_2/φ are multiplicatively independent.

Proposition 3.4.2. *We have that $\lambda_1/|\lambda_1|$ and $\lambda_2/|\lambda_2|$ are multiplicatively independent.*

Proof. Note that $|\lambda_1| = |\lambda_2| = \varphi$ as $\lambda_1\overline{\lambda_1} = \lambda_2\overline{\lambda_2} = \varphi^2$. By the earlier described Galois action, we see that there are Galois automorphisms σ and τ such that $\sigma(\lambda_1) = \tau(\lambda_1) = \lambda_3$, $\sigma(\lambda_2) = \lambda_4$, and $\tau(\lambda_2) = \overline{\lambda_3}$. Further, by this choice, $\sigma(\varphi) = \tau(\varphi) = -\varphi^{-1}$.

Assume, to get a contradiction, that $\lambda_1/|\lambda_1|$ and $\lambda_2/|\lambda_2|$ are multiplicatively dependent; that is to say, there are $a, b \in \mathbb{Z}$, not both 0, such that $(\lambda_1/|\lambda_1|)^a (\lambda_2/|\lambda_2|)^b = 1$. By applying σ to this identity, we obtain

$$1 = \left(\frac{\lambda_3}{-\varphi^{-1}}\right)^a \left(\frac{\lambda_4}{-\varphi^{-1}}\right)^b = \zeta\left(\frac{|\lambda_3\lambda_4|}{\varphi^{-2}}\right)^a \left(\frac{\lambda_4}{-\varphi^{-1}}\right)^{b-a} = \zeta\left(\frac{\lambda_4}{-\varphi^{-1}}\right)^{b-a}$$

for some ζ on the unit circle. Since $|\lambda_4/-\varphi^{-1}| \neq 1$, we conclude that $a = b$. Then when we apply τ to the identity $(\lambda_1/|\lambda_1|)^a (\lambda_2/|\lambda_2|)^b = 1$ we obtain

$$1 = \left(\frac{\lambda_3}{-\varphi^{-1}}\right)^a \left(\frac{\overline{\lambda_3}}{-\varphi^{-1}}\right)^b = \zeta'\left(\frac{|\lambda_3|}{|\overline{\lambda_3}|}\right)^a \left(\frac{\overline{\lambda_3}}{-\varphi^{-1}}\right)^{b+a} = \zeta'\left(\frac{\overline{\lambda_3}}{-\varphi^{-1}}\right)^{b+a}$$

for some ζ' on the unit circle. Since $|\overline{\lambda_3}/(-\varphi^{-1})| \neq 1$, this implies that $a = -b$. Together with $a = b$, we deduce that $a = b = 0$. Thus $\lambda_1/|\lambda_1|$ and $\lambda_2/|\lambda_2|$ are multiplicatively independent. □

Constructing other hard instances of reversible linear recurrence sequences

In this subsection, we will construct reversible LRS $(u_n)_{n=0}^{\infty}$ of

- order 8 for which we cannot decide the Skolem problem; and
- order 12 for which we cannot decide the (Ultimate) Positivity problem.

We start with the order-8 LRS. Using Theorem 3.3.2, we can conclude that any reversible order-8 LRS outside the MSTV class is simple, and so we can apply the methods from Chapter 2 that depend on the Skolem Conjecture and the p-adic version of Schanuel's conjecture. However, for this example, we care about constructing a class of examples we do not have an unconditional algorithm.

Due to Theorem 2.1.3, a hard example $(u_n)_{n=0}^{\infty}$ has to have at least four dominant roots. As $(u_n)_{n=0}^{\infty}$ is reversible, all of its characteristic roots are units, and thus the p-adic condition implies that $(u_n)_{n=0}^{\infty}$ has at least three dominant roots. This condition is thus subsumed by the condition that $(u_n)_{n=0}^{\infty}$ has at least four dominant roots. Lastly, we have to show that $(u_n)_{n=0}^{\infty}$ is non-degenerate.

Fix non-zero integers a, b with $a \neq \pm b$ and let $\rho = \sqrt{2} + 1$ (more generally, the construction below applies to any real quadratic unit ρ greater than 1). Let k be an even positive integer parameter. Writing

$$g_1(X) := (X^2 - aX + \rho^k)(X^2 - bX + \rho^k)\,,$$

the roots of g_1 are

$$\frac{a \pm \sqrt{a^2 - 4\rho^k}}{2} \quad \text{and} \quad \frac{b \pm \sqrt{b^2 - 4\rho^k}}{2}\,.$$

If k is large enough, the four roots of g_1 are all non-real and have modulus $\rho^{k/2}$.

Now write $g(X) := g_1(X)g_2(X)$, where

$$g_2(X) := (X^2 - aX + \rho^{-k})(X^2 - bX + \rho^{-k})\,.$$

Noting that the Galois conjugate of $\rho = 1 + \sqrt{2}$ is $-\rho^{-1} = 1 - \sqrt{2}$, we obtain that g is an integer polynomial of degree 8 and constant term 1 (and thus g is reversible). Using that k is even, the Galois conjugate of ρ^k is ρ^{-1}, and so the roots of g_2 are

$$\frac{a \pm \sqrt{a^2 - 4\rho^{-k}}}{2} \quad \text{and} \quad \frac{b \pm \sqrt{b^2 - 4\rho^{-k}}}{2}\,.$$

For suitably large k, we have $a^2, b^2 > 4\rho^{-k}$, giving that the roots of g_2 have modulus less than $\max(|a|, |b|) < \rho^{k/2}$ and so g has exactly four dominant roots: the roots of g_2.

It remains to observe that g is non-degenerate. Indeed, the non-dominant roots (in modulus) are real and all pairwise distinct. Thus, if λ/λ' is a root of unity for two distinct roots λ and λ' of g, then both roots are dominant and have degree 4. Therefore there is a Galois automorphism σ that maps λ to a non-dominant root, but $\sigma(\lambda')$ maps to another root of g and thus $|\sigma(\lambda)| \neq |\sigma(\lambda')|$. This gives a contradiction as $\sigma(\lambda/\lambda') = \sigma(\lambda)/\sigma(\lambda')$ is a root of unity of modulus 1. Thus, g is non-degenerate.

We now make the numbers a, b, k more explicit.

Example 3.4.3. Let $\rho = \sqrt{2}+1$ as earlier, and write

$$\lambda_1 = \frac{1+\sqrt{1-4\rho^2}}{2} \quad \text{and} \quad \lambda_2 = \frac{2+\sqrt{4-4\rho^2}}{2}\,.$$

The characteristic roots of maximum modulus will be λ_1, $\overline{\lambda_1}$, λ_2, and $\overline{\lambda_2}$. The other four (real) roots are

$$r_1 = \frac{1+\sqrt{1-4\rho^{-2}}}{2}\,, \quad \widetilde{r}_1 = \frac{1-\sqrt{1-4\rho^{-2}}}{2}\,,$$
$$r_2 = \frac{2+\sqrt{4-4\rho^{-2}}}{2} \quad \text{and} \quad \widetilde{r}_2 = \frac{2-\sqrt{4-4\rho^{-2}}}{2}\,.$$

Let

$$u_n = \sqrt{2}\left(\lambda_1^n + \overline{\lambda_1}^n + 2\lambda_2^n + 2\overline{\lambda_2}^n - r_1^n - \widetilde{r}_1^n - 2r_2^n - 2\widetilde{r}_2^n\right)\,.$$

Equivalently, write:

$$u_{n+8} = 6u_{n+7} - 25u_{n+6} + 66u_{n+5} - 120u_{n+4} + 150u_{n+3} - 89u_{n+2} + 18u_{n+1} - u_n\,,$$

with initial values (for $n = 0, \ldots, 7$) of $(0, 0, -48, -120, 0, 520, 624, -2016)$.

Then $(u_n)_{n=0}^{\infty}$ has zeros at indices 0, 1, and 4. It does not belong to the MSTV class, is non-degenerate and not modular.[2]

With the SKOLEM-tool (see Section 2.6), we can still solve the Skolem problem for $(u_n)_{n=0}^{\infty}$ and confirm that $u_n = 0$ if and only if $n \in \{0, 1, 4\}$. □

For a reversible LRS of order 12 for which we cannot decide positivity and ultimate positivity, we use the same construction. Our example will be related to the discussion in Section 3.1, which is based on the work of Ouaknine and Worrell [121]. We recall

[2]Any reversible LRS that has a zero (or whose bi-completion has a zero) will necessarily fail to be modular since for any given integer $m \geq 2$, the sequence of residues modulo m is always periodic.

the following point from Theorem 3.1.7: a non-simple LRS that is a hard example of (ultimate) positivity possesses (at least) three dominant roots of multiplicity (at least) 2 which one is real and positive. As our LRS has order 12 and is reversible, the methods from Theorem 3.1.7 imply there are six roots each of multiplicity 2.

With a, b, k, ρ as above, we take the polynomial

$$g(X) = (X^2 - aX + \rho^{2k})(X^2 - aX + \rho^{-2k})$$

Let λ and $\overline{\lambda}$ be the roots of $X^2 - aX + \rho^{2k}$ and λ_3 and λ_4 be the roots of $X^2 - aX + \rho^{-2k}$. Let the last two characteristic roots be $\widetilde{\rho}^{\pm 1} := \rho^{\pm k}$.

Then λ, $\overline{\lambda}$, and ρ^k are dominant while the three other roots are not dominant and real. Moreover, g is non-degenerate. Lastly, let $q \in \mathbb{Q}_{>0}$ with denominator d. Then define the reversible LRS $(u_n^{(q)})_{n=0}^{\infty}$ as

$$u_n^{(q)} = (n+\widetilde{\rho})\widetilde{\rho}^n + (n+\widetilde{\rho}^{-1})\widetilde{\rho}^{-n} + q\big((n+\lambda)\lambda^n + (n+\overline{\lambda})\overline{\lambda}^n + (n+\lambda_3)\lambda_3^n + (n+\lambda_4)\lambda_4^n\big)\,.$$

Then, $(du_n^{(q)})_{n=0}^{\infty}$ is a reversible LRS. However, given the current state of the art, it is unknown how to solve the Positivity problem for such LRS as one can again encode the approximation of certain real numbers.

3.5 The Positivity problem for real algebraic linear recurrence sequences

So far, we have mostly covered decision problems involving $\mathbb{Z}$-LRS and $\mathbb{Q}$-LRS. For the Skolem problem, the folklore result [107, Lemma 9] implies that the $\mathbb{Q}$- and $\overline{\mathbb{Q}}$-Skolem problems are Turing-interreducible. Proving this fact is straightforward:

Using the same technique as in Lemma 1.2.3, we can assume we are dealing with a non-degenerate $\overline{\mathbb{Q}}$-LRS $(u_n)_{n=0}^{\infty}$ with polynomial-exponential form (1.2) and splitting field K. Then,

$$v_n = \prod_{\sigma \in \mathrm{Gal}_{\mathbb{Q}}(K)} \sigma\Big(\sum_{i=1}^{k} Q_i(n)\lambda_i^n\Big) = \prod_{\sigma \in \mathrm{Gal}_{\mathbb{Q}}(K)} \sum_{i=1}^{k} \sigma(Q_i(n))\sigma(\lambda_i)^n$$

is a $\mathbb{Q}$-LRS (being closed under Galois automorphisms) while it is the product of non-degenerate $\overline{\mathbb{Q}}$-LRS that by the Skolem-Mahler-Lech theorem are only zero finitely often. Hence, $(v_n)_{n=0}^{\infty}$ has finitely many zeros that we can compute using a Skolem oracle. Then, for these finitely many zeros z, one can test whether $u_z = 0$.

Unfortunately, this reduction does not apply to the Positivity problem as we do not know how many factors in this product are positive. Fortunately, we can find another approach.

Theorem 3.1.9. *The $\mathbb{R}\cap\overline{\mathbb{Q}}$-Positivity problem and $\mathbb{Q}$-Positivity problem are Turing-interreducible. Moreover, the Positivity problem for simple $\mathbb{R}\cap\overline{\mathbb{Q}}$-LRS is Turing equivalent to the Positivity problem for simple $\mathbb{Q}$-LRS.*

We outline our proof of Theorem 3.1.9. Fix a real-algebraic LRS $(u_n)_{n=0}^{\infty}$, which we can assume to be non-degenerate by Lemma 1.2.3.

We first show that a $(\mathbb{R}\cap\overline{\mathbb{Q}})$-LRS is a $(\mathbb{R}\cap\overline{\mathbb{Q}})$-linear combination of $\mathbb{Q}$-LRS. That is, $u_n = \sum_{i=1}^{\ell} n^{\ell_i}\beta_i u_n^{(i)}$ for some $\beta_i \in \mathbb{R}\cap\overline{\mathbb{Q}}$, $\ell_i \in \mathbb{N}$, and $\mathbb{Q}$-LRS $(u_n^{(i)})_{n=0}^{\infty}$. To sketch the proof, when writing $(u_n)_{n=0}^{\infty}$ in its exponential form, it is sufficient to prove this for sequences $(\lambda^n)_{n=0}^{\infty}$. Using the minimal polynomial of λ of degree d, we find $\mathbb{Q}$-LRS $(\widetilde{u}_n^{(0)})_{n=0}^{\infty},\dots,(\widetilde{u}_n^{(d-1)})_{n=0}^{\infty}$ such that $\lambda^n = \sum_{i=1}^{d-1}\lambda^i\widetilde{u}_n^{(i)}$ in Lemma 3.5.1. By complex conjugation, we have $\overline{\lambda}^n = \sum_{i=1}^{d-1}\overline{\lambda}^i\widetilde{u}_n^{(i)}$, and so when considering the entire LRS $(u_n)_{n=0}^{\infty}$, we indeed obtain real algebraic coefficients β_i.

Thus, the constants β_i are the only non-rational part in these formulas $u_n = \sum_{i=1}^{\ell} n^{\ell_i}\beta_i u_n^{(i)}$. In Lemma 3.5.2, we show we can approximate these constants β_i as follows. Given $0<r<1$, we construct simple $\mathbb{Q}$-LRS $(c_n^{(i)})_{n=0}^{\infty}$ and $(d_n^{(i)})_{n=0}^{\infty}$, $b>0$, and $N\in\mathbb{N}$ such that $|\beta - c_n^{(i)}/d_n^{(i)}| < br^n$ for $n\geq N$.

Then set $(D_n)_{n=0}^{\infty}$ to be the product of all LRS $(d_n^{(i)})_{n=0}^{\infty}$, which is a rapidly growing $\mathbb{Q}$-LRS and construct three new $\mathbb{Q}$-LRS $(v_n)_{n=0}^{\infty}$, $(v_n^+)_{n=0}^{\infty}$, and $(v_n^-)_{n=0}^{\infty}$, where $v_n = D_n u_n$, and

$$v_n^{\pm} = \pm\frac{1}{2}\widetilde{r}^n + D_n\sum_{i=1}^{\ell}\frac{c_n^{(i)}}{d_n^{(i)}}n^{\ell_i}u_n^{(i)}$$

are $\mathbb{Q}$-LRS. For some computable $N\in\mathbb{N}$ and $n\geq N$, we have that $D_n>0$ and $v_n^-\leq v_n\leq v_n^+$, and due to their similar growth, Theorem 1.2.7 implies there are only finitely many n such that $v_n^-<0<v_n^+$. If a Positivity Oracle allows us to decide the Positivity problem for the $\mathbb{Q}$-LRS $(v_n^-)_{n=N'}^{\infty}$ and $(v_n^+)_{n=N'}^{\infty}$ for all $N'\in\mathbb{N}$, we can decide the Positivity problem for $(v_n)_{n=N'}^{\infty}$ and thus $(u_n)_{n=0}^{\infty}$.

Our proof of Theorem 3.1.9 starts with the two technical lemmas mentioned above. The first shows that the powers of an algebraic number form a $\overline{\mathbb{Q}}$-linear combination of rational LRS.

Lemma 3.5.1. *Let $\lambda\in\overline{\mathbb{Q}}^*$ have the monic minimal polynomial $P(x) = X^d - c_1X^{d-1} - \cdots - c_d \in \mathbb{Q}[X]$. One can construct simple rational LRS $u_n^{(0)},\dots,u_n^{(d-1)}$ such that for all $n\geq 0$,*

$$\lambda^n = \sum_{i=0}^{d-1}\lambda^i u_n^{(i)}\,.$$

Proof. Let $a_0, \dots, a_{d-1} \in \mathbb{Q}$ and consider $\alpha = a_0 + a_1\lambda + \cdots + a_{d-1}\lambda^{d-1}$. Then, using the minimal polynomial of λ, we have $\alpha\lambda$ is again in the direct sum of $\mathbb{Q}, \lambda\mathbb{Q}, \dots, \lambda^{d-1}\mathbb{Q}$. Using the companion matrix (1.3), we can compute $\alpha\lambda$ as

$$\alpha\lambda = \begin{pmatrix} a_{d-1} & a_{d-1} & \cdots & a_0 \end{pmatrix} \begin{pmatrix} c_1 & c_2 & \cdots & c_{d-1} & c_d \\ 1 & 0 & \cdots & 0 & 0 \\ 0 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 1 & 0 \end{pmatrix} \begin{pmatrix} \lambda^{d-1} \\ \lambda^{d-2} \\ \vdots \\ 1 \end{pmatrix}.$$

Let M denote the square companion matrix above and e_i the ith standard unit vector. Then,

$$\lambda^n = \begin{pmatrix} 0 & \cdots & 0 & 1 \end{pmatrix} M^n \begin{pmatrix} \lambda^{d-1} \\ \lambda^{d-2} \\ \vdots \\ 1 \end{pmatrix}.$$

Hence, $\lambda^n = \sum_{i=0}^{d-1} \lambda^i e_d^\top M^n e_{d-i}$, where each $(e_d^\top M^n e_{d-i-1})_{n=0}^\infty$ is a computable $\mathbb{Q}$-LRS. As P is the minimal polynomial of λ and the characteristic polynomial of M, these $\mathbb{Q}$-LRS are also all simple. □

The second technical lemma shows that every algebraic number can be approximated exponentially fast by the quotient of two $\mathbb{Q}$-LRS.

Lemma 3.5.2. *For $\beta \in \mathbb{R} \cap \overline{\mathbb{Q}}$ and $0 < r < 1$, one can construct simple $\mathbb{Q}$-LRS $(c_n)_{n=0}^\infty$ and $(d_n)_{n=0}^\infty$, $b \in \mathbb{Q}_{>0}$, and $N \in \mathbb{N}$ such that*

$$\forall n \geq N \colon \left|\beta - \frac{c_n}{d_n}\right| < b \cdot r^n \wedge d_n > 0\,.$$

Proof. If $\beta \in \mathbb{Q}$, choosing $c_n = \beta$, $d_n = 1$, $b = 1$, and $N = 0$ suffices. Hence assume β is irrational and let $\beta = \beta_1, \dots, \beta_d$ denote the Galois conjugates of β.

First, we claim that we can find $p/q \in \mathbb{Q}$ and $s \in \{-1, 1\}$ with the property that when setting $f(x) = \frac{s}{x - p/q}$, $f(\beta) \in \mathbb{R}_{>0}$, and $f(\beta) > |f(\beta_i)|$ for $i = 2, \dots, d$.

Let $\delta = \min(|\beta - \beta_2|, \dots, |\beta - \beta_d|)$. As $\mathbb{Q}$ is dense in $\mathbb{R}$, there exists an effectively computable $\frac{p}{q} \in \mathbb{Q}$ such that $|\beta - \frac{p}{q}| < \delta/2$. Let $f_1(x) = x - \frac{p}{q}$. Then, for $i = 2, \dots, d$, the reverse triangle inequality implies that

$$|f_1(\beta_i)| = \left|\beta_i - \frac{p}{q}\right| \geq |\beta_i - \beta| - \left|\beta - \frac{p}{q}\right| \geq \delta - \delta/2 = \delta/2 > |f_1(\beta)|\,.$$

Thus, $f_1(\beta)$ is closer to the origin than $f_1(\beta_i)$, but $f_1(\beta) \neq 0$ as β is irrational.

Let $s \in \{-1, 1\}$ be the sign of $f_1(\beta)$ and $f(x) = \frac{s}{f_1(x)}$. Then $f(\beta) > 0$ and $|f(\beta)| > |f(\beta_i)|$ for all $2 \leq i \leq d$. Therefore, f has the properties stated above.

Define the simple LRS $(c_n)_{n=0}^{\infty}$ and $(d_n)_{n=0}^{\infty}$ as follows:

$$c_n = \sum_{i=1}^{d} \beta_i f(\beta_i)^n \quad \text{and} \quad d_n = \sum_{i=1}^{d} f(\beta_i)^n \,.$$

As any Galois automorphism would permute the terms of the sum in c_n, we deduce that c_n is rational for all $n \geq 0$. Similarly, we have that d_n is rational for all $n \in \mathbb{N}$. Thus, $(c_n)_{n=0}^{\infty}$ and $(d_n)_{n=0}^{\infty}$ are simple $\mathbb{Q}$-LRS. As β is the single dominant root of $(d_n)_{n=0}^{\infty}$ with polynomial coefficient 1 in front of it, the LRS $(d_n)_{n=0}^{\infty}$ is ultimately strictly positive, and one can compute $N \in \mathbb{N}$ such that $d_n > 0$ for all $n \geq N$.

Let $b \geq \frac{1}{2}\sum_{i=2}^{d} |\beta - \beta_i|$ be rational and $\widetilde{r} = \max_{i=2}^{d} |f(\beta_i)|$. Then $\widetilde{r} < f(\beta)$ and update N such that $\left(\frac{f(\beta)}{\widetilde{r}}\right)^N > 2(d-1)$ also holds. Then, for all $n \geq N$,

$$\begin{aligned}
\left|\beta - \frac{c_n}{d_n}\right| &= \left|\frac{\sum_{i=1}^{d}(\beta - \beta_i) f(\beta_i)^n}{\sum_{i=1}^{d} f(\beta_i)^n}\right| \\
&= \left|\frac{\sum_{i=2}^{d}(\beta - \beta_i) f(\beta_i)^n}{f(\beta)^n + \sum_{i=2}^{d} f(\beta_i)^n}\right| \\
&\leq \left|\frac{\frac{1}{2} b \cdot \widetilde{r}^n}{f(\beta)^n - (d-1)\widetilde{r}^n}\right| \\
&= \frac{1}{2} b \left|\left(\frac{f(\beta)}{\widetilde{r}}\right)^n - (d-1)\right|^{-1} \\
&\leq b \left(\frac{\widetilde{r}}{f(\beta)}\right)^n .
\end{aligned}$$

It remains to choose $\ell \in \mathbb{N}$ such that $(\widetilde{r}/f(\beta))^\ell < r$. Then, choosing the LRS $(c_{\ell n})_{n=0}^{\infty}$ and $(d_{\ell n})_{n=0}^{\infty}$ gives that

$$\left|\beta - \frac{c_{\ell n}}{d_{\ell n}}\right| < b \cdot r^n \,. \qquad \square$$

Proof of Theorem 3.1.9. Let $(u_n)_{n=0}^{\infty}$ be a $(\mathbb{R} \cap \overline{\mathbb{Q}})$-LRS with a exponential-polynomial form $u_n = \sum_{i=1}^{k} Q_i(n)\lambda_i^n$. By decomposing into non-degenerate subsequences using Theorem 1.2.3, we can assume that $(u_n)_{n=0}^{\infty}$ is non-degenerate. We assume that $k \geq 1$ as the Positivity problem is trivial for the zero-LRS. As the characteristic polynomial of $(u_n)_{n=0}^{\infty}$ is in $(\mathbb{R} \cap \overline{\mathbb{Q}})[X]$, the non-real characteristic roots come in complex conjugate pairs. That is, for each $1 \leq i \leq k$ such that $\lambda_i \notin \mathbb{R}$, there exists $1 \leq j \leq k$ such that $\lambda_j = \overline{\lambda_i}$. Therefore, by the Galois closure of the polynomial-exponential form, we have $Q_j(n) = \overline{Q_i(n)}$, and we can rewrite u_n as

$$u_n = \sum_{i \in \mathcal{R}} n^{\sigma(i)} \gamma_i \lambda_i^n + \sum_{i \in \mathcal{C}} n^{\sigma(i)} (\gamma_i \lambda_i^n + \overline{\gamma_i}\,\overline{\lambda_i}^n) \tag{3.6}$$

where $\mathcal{R}$ and $\mathcal{C}$ are disjoint finite sets, $\sigma(i) \in \mathbb{N}$ for all $i \in \mathcal{C} \cup \mathcal{R}$, $\lambda_i, \gamma_i \in \mathbb{R} \cap \overline{\mathbb{Q}}$ for $i \in \mathcal{R}$, and $\lambda_i \in \overline{\mathbb{Q}} \backslash \mathbb{R}$ and $\gamma_i \in \overline{\mathbb{Q}}$ for $i \in \mathcal{C}$. Next, for $i \in \mathcal{C} \cup \mathcal{R}$, we apply Lemma 3.5.1 to λ_i and obtain that $\lambda_i^n = \sum_{j=0}^{d-1} \lambda_i^j v_n^{(i,j)}$ for some simple $\mathbb{Q}$-LRS $(v_n^{(i,0)})_{n=0}^{\infty}, \ldots, (v_n^{(i,d_i-1)})_{n=0}^{\infty}$ where d_i is the degree of λ_i. Next, observe that if $i \in \mathcal{C}$,

$$\overline{\lambda_i}^n = \overline{\lambda_i^n} = \sum_{j=0}^{d-1} \overline{\lambda_i}^j v_n^{(i,j)} .$$

For $j \geq 0$ and $i \in \mathcal{R}$, we have $\gamma_i \lambda_i^j \in \mathbb{R}$ and for $i \in \mathcal{C}$, we have $\gamma_i \lambda_i^j + \overline{\gamma_i}\,\overline{\lambda_i}^j \in \mathbb{R}$. Therefore,

$$u_n = \sum_{i \in \mathcal{R}} \sum_{j=0}^{d_i-1} n^{\sigma(i)} \gamma_i \lambda_i^j v_n^{(i,j)} + \sum_{i \in \mathcal{C}} \sum_{j=0}^{d_i-1} n^{\sigma(i)} (\gamma_i \lambda_i^j + \overline{\gamma_i}\,\overline{\lambda_i}^j) v_n^{(i,j)} = \sum_{i=1}^{\ell} n^{\ell_i} \beta_i u_n^{(i)}$$

for some $\ell, \ell_i \in \mathbb{N}$, $\beta_i \in \mathbb{R} \cap \overline{\mathbb{Q}}$, and simple $\mathbb{Q}$-LRS $(u_n^{(i)})_{n=0}^{\infty}$.

Let R_i denote the spectral radius of $(u_n^{(i)})_{n=0}^{\infty}$ and let $r \in \mathbb{Q}$ satisfy $0 < r < \min(1, 1/R_i)$ for all $1 \leq i \leq \ell$. Then, for each $1 \leq i \leq \ell$, invoke Lemma 3.5.2 with β_i and r to obtain simple $\mathbb{Q}$-LRS $(c_n^{(i)})_{n=0}^{\infty}$ and $(d_n^{(i)})_{n=0}^{\infty}$, $b_i \in \mathbb{Q}_{>0}$, and the threshold N_i. Let $\widetilde{r} \in \mathbb{Q}$ satisfy $rR < \widetilde{r} < \min(1, R)$ and compute $N' \in N$ such that for all $n \geq N'$,

$$\frac{1}{2}\widetilde{r}^n \geq r^n \sum_{i=1}^{\ell} b_i n^{\ell_i} |u_n^{(i)}| .$$

This is possible as $(r^n \sum_{i=1}^{\ell} b_i u_n^{(i)} n^{\ell_i})_{n=0}^{\infty}$ is a $\overline{\mathbb{Q}}$-LRS with spectral radius $rR_i < \widetilde{r}$. Define $N = \max(N_1, \ldots, N_\ell, N')$, the simple $\mathbb{Q}$-LRS $(D_n)_{n=0}^{\infty}$ as $D_n = d_n^{(1)} \cdots d_n^{(\ell)}$, and the $\mathbb{Q}$-LRS $(v_n)_{n=0}^{\infty}$, $(v_n^+)_{n=0}^{\infty}$, and $(v_n^-)_{n=0}^{\infty}$ as

$$\begin{aligned}
v_n &= D_n \sum_{i=1}^{\ell} \beta_i n^{\ell_i} u_n^{(i)} = D_n u_n, \\
v_n^+ &= D_n \Big(\frac{1}{2}\widetilde{r}^n + \sum_{i=1}^{\ell} \frac{c_n^{(i)}}{d_n^{(i)}} n^{\ell_i} u_n^{(i)} \Big), \quad \text{and} \\
v_n^- &= D_n \Big(-\frac{1}{2}\widetilde{r}^n + \sum_{i=1}^{\ell} \frac{c_n^{(i)}}{d_n^{(i)}} n^{\ell_i} u_n^{(i)} \Big) .
\end{aligned}$$

Then, by construction, for all $n \geq N$,

$$\begin{aligned}
|v_n^+ - v_n| &\leq |D_n| \Big(\frac{1}{2}\widetilde{r}^n + \sum_{i=1}^{\ell} \Big| \beta_i - \frac{c_n}{d_n} \Big| n^{\ell_i} |u_n^{(i)}| \Big) \\
&< |D_n| \Big(\frac{1}{2}\widetilde{r}^n + r^n \sum_{i=1}^{\ell} b_i n^{\ell_i} |u_n^{(i)}| \Big) \\
&\leq |D_n| \Big(\frac{1}{2}\widetilde{r}^n + \frac{1}{2}\widetilde{r}^n \Big) = |D_n| \widetilde{r}^n .
\end{aligned}$$

Thus $(v_n^+ - v_n^-)_{n=0}^\infty$ has a spectral radius strictly smaller than $(v_n)_{n=0}^\infty$. The same holds for $|v_n^- - v_n|$ and $|v_n^+ - v_n^-|$. Moreover, $v_n^- \leq v_n \leq v_n^+$ holds for all $n \geq N$. As $d_n^{(i)}$ is strictly positive for $n \geq N$, so is D_n. Hence, we have that for all $n \geq N$, we have $v_n \geq 0$ if and only if $u_n \geq 0$.

Therefore, as $\widetilde{r} < R$, Theorem 1.2.7 implies there exists an $N'' \geq N$ such that $|v_n|, |v_n^+|, |v_n^-| > |D_n|\widetilde{r}^n$ for all $n \geq N''$. Thus, the signs of v_n, v_n^+, and v_n^- are identical for all $n \geq N''$. In particular, $(v_n)_{n=0}^\infty$ is ultimately positive if and only if both $(v_n^+)_{n=0}^\infty$ and $(v_n^-)_{n=0}^\infty$ are ultimately positive.

As such, we can decide positivity for $(v_n)_{n=N''}^\infty$ using the Positivity query on $(v_n^+)_{n=N''}^\infty$ and $(v_n^+)_{n=N''}^\infty$. Let $K \geq N''$. Then, $(v_n^-)_{n=K}^\infty$ being positive, implies that $(v_n)_{n=K}^\infty$ is positive while $(v_n^-)_{n=K}^\infty$ not being positive implies that $(v_n)_{n=K}^\infty$ is not positive. As $(v_n^-)_{n=K}^\infty$ is positive if and only if $(v_n^+)_{n=K}^\infty$ is positive, we can iterate natural numbers K until $(v_n^-)_{n=K}^\infty$ are both positive or both not positive. Then $(u_n)_{n=0}^\infty$ is positive if and only if $u_n \geq 0$ for all $0 \leq n \leq K$, which we can check.

This proves the first claim of Theorem 3.1.9. For the second claim, note that if $(u_n)_{n=0}^\infty$ is simple, the LRS $(v_n^+)_{n=0}^\infty$ and $(v_n^-)_{n=0}^\infty$ are also simple. Hence, our reduction only involves Positivity queries for simple rational LRS. □

Chapter 4

Monadic second-order logic

4.1 Introduction and main results

In this chapter, we study expansions of monadic second-order (MSO) logic with predicates derived from linear recurrence sequences. In particular, we study expansions of the structure $\langle \mathbb{N}; < \rangle$.

Büchi, in his seminal 1962 paper [43], established the decidability of the MSO theory of $\langle \mathbb{N}; < \rangle$ (call this theory $\mathsf{MSO}_{\mathbb{N};<}$) and in so doing brought to light the profound connections between mathematical logic and automata theory. Over the ensuing decades, considerable work has been devoted to the question of which expansions of $\langle \mathbb{N}; < \rangle$ retain MSO decidability. In most cases, the resulting theory is undecidable. When adding a single function or non-unary predicate, there is little chance of retaining decidability. For example, the MSO theory of $\langle \mathbb{N}; <, n \mapsto 2n \rangle$ (where $n \mapsto 2n$ denotes the doubling function) is undecidable while the first-order theory of this structure is decidable: It can be encoded in Presburger arithmetic.

Therefore, we study the decidability of the MSO theory of $\langle \mathbb{N}; < P_1, \ldots, P_d \rangle$ for unary predicates $P_1, \ldots, P_d$ (call this theory $\mathsf{MSO}_{\mathbb{N};<}(P_1, \ldots, P_d)$). Here, we interpret a unary predicate P as a fixed set of non-negative integers $P \subseteq \mathbb{N}$ (see Section 1.3.2).

Taking, for example, P to be the set of prime numbers, Büchi and Landweber [44] observed in 1969 that a proof of decidability of $\mathsf{MSO}_{\mathbb{N};<}(P)$ would "seem very difficult", as it would *inter alia* enable one to settle the twin prime conjecture. (Decidability was subsequently established assuming Schinzel's hypothesis H [18].)

The set of prime numbers is, of course, highly intricate. In 1966, Elgot and Rabin [62] considered a large class of simpler predicates of 'arithmetic' origin, such as, for any fixed k, the set $k^{\mathbb{N}} = \{k^n : n \in \mathbb{N}\}$ of powers of k, and the set $\mathsf{N}_k = \{n^k : n \in \mathbb{N}\}$ of kth powers. For any such predicate P, they systematically established decidability of $\mathsf{MSO}_{\mathbb{N};<}(P)$ by using what is now known as the *Elgot-Rabin contraction method*.

In short, they reduce the decidability of the MSO theory to deciding whether an ω-automaton accepts a certain word α, where in their cases, the word α is ultimately periodic.

Although Elgot and Rabin establish separately the decidability of the MSO theories with a single added predicate, for example, of $\mathsf{MSO}_{\mathbb{N};<}(2^{\mathbb{N}})$ and $\mathsf{MSO}_{\mathbb{N};<}(\mathsf{N}_2)$, they remain resolutely silent on the obvious joint expansions $\mathsf{MSO}_{\mathbb{N};<}(2^{\mathbb{N}}, 3^{\mathbb{N}})$ and $\mathsf{MSO}_{\mathbb{N};<}(2^{\mathbb{N}}, \mathsf{N}_2)$. Over the last several decades worth of research on monadic second-order expansions of $\langle \mathbb{N}; < \rangle$, it is fair to say that the bulk of the attention has focused on the addition of a single, well-behaved predicate P. The decidability of such single-predicate expansions of $\langle \mathbb{N}; < \rangle$ can usually be handled with automata-theoretic techniques alone, by reasoning about individual patterns in isolation. This is not the case when multiple predicates are at play simultaneously.

Moreover, usually, a single predicate is given in a convenient form to allow one to reduce to the acceptance of a word α by an automaton $\mathcal{A}$, like an increasing sequence. But, one can also consider a non-increasing sequence (like $u_n = (2+i)^n + (2-i)^n$) or the union of two 'nice' predicates (like $2^{\mathbb{N}} \cup 3^{\mathbb{N}}$). When a sequence is not increasing, one has to 'sort' the predicate first: One takes the sequence's positive values and sorts them to obtain a new sequence $(p_m)_{m=0}^{\infty}$, which destroys the original structure of the predicate. The classical automata-theoretic methods cannot handle the disparity between the sequence $(u_n)_{n=0}^{\infty}$ and its sorted counterpart $(p_m)_{m=0}^{\infty}$. Such collections of predicates can exhibit highly complex interaction patterns, which existing approaches are ill-equipped to handle. However, this is sometimes still possible (see for example also [52]).

Due to these two difficulties, dealing with ultimately periodic words is insufficient, forcing us to employ wider classes of predicates. The first such class are *(effective) almost-periodic words*, as introduced in the 1980s by Semenov [140], which roughly speaking, are words $\alpha \in \Sigma^*$ such that if a factor $w \in \Sigma^+$ appears infinitely often in α, the finite word w appears in every factor of α of a certain length. In particular, almost-periodic words are examples of toric words, which are generated by a dynamical system on a torus [25]. We also use *disjunctive words* (see, for example, [45, Section 4.4]), which are words $\alpha \in \Sigma^{\omega}$ of maximum factor complexity: every $w \in \Sigma^+$ appears infinitely often as a factor in α. In addition, we apply number-theoretic tools to ensure effectiveness at various junctures of our algorithms.

Since, the general theory has been substantially developed and abstracted by, among others, Carton, Rabinovich, and Thomas [47, 130, 131], giving rise to concepts like profinitely ultimately periodic words. In our study, we steer away from

such abstract concepts and use less general, but practical, classes of words to establish results for explicit predicates.

This chapter contains three main results. First, Theorem 4.4.1 considers predicates arising from the value sets of linear recurrence sequences, generalizing the predicates considered above. A simplified version of that result is as follows:

Theorem 4.1.1. *Let $\rho_1, \ldots, \rho_d > 1$ be natural numbers.*

1. $\mathsf{MSO}_{\mathbb{N};<}(\rho_1^{\mathbb{N}}, \ldots, \rho_d^{\mathbb{N}})$ *is decidable, assuming the weak Schanuel conjecture.*
2. *If $1/\log(\rho_1), \ldots, 1/\log(\rho_d)$ are linearly independent over $\mathbb{Q}$, then the decidability is unconditional.*
3. *If each triple of distinct ρ_i, ρ_j, ρ_k is multiplicatively dependent, then the decidability is unconditional.*

Item (3) captures, for example, that $\mathsf{MSO}_{\mathbb{N};<}(2^{\mathbb{N}}, 3^{\mathbb{N}})$ and $\mathsf{MSO}_{\mathbb{N};<}(2^{\mathbb{N}}, 3^{\mathbb{N}}, 6^{\mathbb{N}})$ are decidable. However, Item (3) does not apply $\mathsf{MSO}_{\mathbb{N};<}(2^{\mathbb{N}}, 3^{\mathbb{N}}, 5^{\mathbb{N}})$, so our decidability for this theory is conditional on the weak Schanuel conjecture. More precisely, we need to decide whether $\frac{1}{\log(2)}$, $\frac{1}{\log(3)}$, and $\frac{1}{\log(5)}$ are linear independent over $\mathbb{Q}$. Although these results are novel, we should mention that the decidability of the *first-order* theory of $\langle \mathbb{N}; <, 2^{\mathbb{N}}, 3^{\mathbb{N}} \rangle$ has been known for over forty years, an important result of Semenov [139]. It was shown in [87] that, in a restricted case, termination using our method requires only Schanuel's conjecture, using the decidability results for $\mathbb{R}_{\exp}$ established by Macintyre and Wilkie [109].

When $\eta \in \mathbb{R}$ and $b \geq 2$ is a natural number, then one can represent η as an infinite word α over the finite alphabet $\Sigma = \{0, \ldots, b-1\}$ using its base-b expansion. By the *MSO theory of the base-b expansion of η* we mean the MSO theory of $\langle \mathbb{N}; <, P_0, \ldots, P_{b-1} \rangle$, where each P_i is the unary predicate $\{n \in \mathbb{N} \colon \alpha_n = i\}$. Our second main result is Theorem 4.5.2, restated here:

Theorem 4.1.2. *Let $p, q, d \geq 1$ and $b \geq 2$ be integers and set $\eta = \sqrt[d]{p/q}$, $P_1 = \{qn^d \colon n \in \mathbb{N}\}$, and $P_2 = \{pb^{nd} \colon n \in \mathbb{N}\}$. The decidability of $\mathsf{MSO}_{\mathbb{N};<}(P_1, P_2)$ is Turing-equivalent to that of the MSO theory of the base-b expansion of η.*

The underlying dynamical system here is symbolic in nature: it consists of the base-b expansion of the irrational number η, which is a dth root of a rational number. When η is irrational, such expansions are widely conjectured to be *normal*, and *a*

fortiori disjunctive: every finite pattern of digits occurs infinitely often. As the MSO theory of a recursive and disjunctive word is decidable (Theorem 4.2.8), we obtain a conditional decidability result. When η is rational, the base-b expansion of η is periodic, and so we obtain unconditional decidability. Therefore, $\mathsf{MSO}_{\mathbb{N};<}(4^{\mathbb{N}}, \mathsf{N}_2)$ is decidable and more specifically, we have the following result.

Corollary 4.1.3. *For any positive integers b and d, the theory $\mathsf{MSO}_{\mathbb{N};<}((b^d)^{\mathbb{N}}, \mathsf{N}_d)$ is decidable.*

When a sequence $(u_n)_{n=0}^{\infty}$ is non-increasing, deciding $\mathsf{MSO}_{\mathbb{N};<}(P)$ with a single predicate $P = \{u_n \colon n \in \mathbb{N}\} \cap \mathbb{N}$ already transcends the methods of Elgot and Rabin. In our last main contribution in this chapter, we study LRS with two dominant roots instead of one. In that case, we obtain our third main result.

Theorem 4.1.4. *Let $(u_n)_{n=0}^{\infty}$ be a non-degenerate, simple LRS with two dominant roots and set $P = \{u_n \colon n \in \mathbb{N}\} \cap \mathbb{N}$. Then $\mathsf{MSO}_{\mathbb{N};<}(P)$ is decidable.*

Although the hypothesis of Theorem 4.1.4 seems quite restrictive, the result covers a wide variety of LRS. In fact, it covers almost all LRS as nearly all LRS have a single dominant root (for which $\mathsf{MSO}_{\mathbb{N};<}(P)$ is decidable using Elgot and Rabin's results) or satisfy the hypothesis of Theorem 4.1.4. In particular, the LRS $(u_n)_{n=0}^{\infty}$ for which we can decide $\mathsf{MSO}_{\mathbb{N};<}(P)$ have density 1 in the space of all LRS (see the work of Dubickas and Sha in [60] and their citations).

Next, we summarise our methods. First, in Section 4.2, we study various classes of words (procyclic, almost-periodic, toric, and disjunctive) and the acceptance problem for an ω-automaton for words in these classes.

For the remaining sections in this chapter, we need more space to explain their content.

The Elgot-Rabin contraction method

In Section 4.3, we study the Elgot-Rabin contraction technique. To outline this method, we return to Büchi, who first defined the *characteristic word*. Let $\Sigma = \{0,1\}^d$ and $P_1, \ldots, P_d \subset \mathbb{N}$ be unary predicates.

Definition 4.1.5. The *characteristic word* of $P_1, \ldots, P_d$, written

$$\alpha := \mathrm{Char}(P_1, \ldots, P_d) \in \Sigma^{\omega},$$

is defined by $\alpha_n = (b_{n,1}, \ldots, b_{n,d})$ where $b_{n,i} = 1$ if $n \in P_i$ and $b_{n,i} = 0$ otherwise.

For example, if $P_1 = \mathsf{N}_2$ and $P_2 = 2^{\mathbb{N}}$,

$$\mathrm{Char}(P_1, P_2) = \big((1,0),(1,1),(0,1),(0,0),(1,1),(0,0),(0,0),(0,0),(0,1),(1,0),\dots\big).$$

When there is only a single predicate (that is, $d = 1$), we also write $\alpha_n = b_{n,1}$ instead of $\alpha_n = (b_{n,1})$.

Recall that Acc_α denotes the *acceptance problem* for an infinite word α over an alphabet Σ and asks to determine, given a deterministic Muller automaton $\mathcal{A}$ with an alphabet Σ, whether $\mathcal{A}$ accepts α.[1] The connection between the decidability of $\mathsf{MSO}_{\mathbb{N};<}$ and the acceptance problem is captured by the following theorem of Büchi.

Theorem 4.1.6 ([153, Theorems 5.4 and 5.9])**.** *Decidability of* $\mathsf{MSO}_{\mathbb{N};<}(P_1,\dots,P_d)$ *is Turing-equivalent to* Acc_α*, where* $\alpha = \mathrm{Char}(P_1,\dots,P_d)$*.*

Elgot and Rabin studied the case where $d = 1$ by adding a single infinite predicate P that is 'sparse' and behaves decently well (which we formally define later). Restricting ourselves to infinite predicates is a minor limitation as finite predicates can be defined in $\mathsf{MSO}_{\mathbb{N};<}$. Examples of such predicates are the set of powers of k (denoted by $k^{\mathbb{N}}$), the kth powers (N_k), the factorials, the Fibonacci numbers and many more. Then the crux of Elgot and Rabin's method is that for every deterministic automaton $\mathcal{A}$, they construct an ultimate periodic word that is accepted by $\mathcal{A}$ if and only if $\mathcal{A}$ accepts $\mathrm{Char}(P)$. Deciding the former is straightforward.

We often rely on the *positive value sequence* and the *order word*.

Definition 4.1.7.

- The *positive value sequence* $(p_n)_{n=0}^{\infty} \subset \mathbb{N}$ of a tuple of infinite predicates $(P_1,\dots,P_d)$ is defined by setting p_{n-1} to be the nth smallest number such that $\alpha_n \neq (0,\dots,0)$. That is, we enumerate $\bigcup_{i=1}^{d} P_i$.
- The *order word* $\beta \in (\Sigma \setminus \{(0,\dots,0)\})^\omega$ is defined as $\beta_n = \alpha_{p_n}$. Thus, removing all occurrences of $(0,\dots,0)$ in the characteristic word gives the order word.

If $\mathbf{0} = (0,\dots,0)$ and $(p_m)_{m=0}^{\infty}$ is the positive value sequence of $(P_1,\dots,P_d)$, we have rewritten as α as

$$\begin{aligned}\alpha &= \mathbf{0}^{p_0}\alpha_{p_0}\mathbf{0}^{p_1-p_0-1}\alpha_{p_1}\cdots\mathbf{0}^{p_n-p_{n-1}}\alpha_{p_n}\cdots\\ &= \mathbf{0}^{p_0}\beta_0\mathbf{0}^{p_1-p_0-1}\beta_1\cdots\mathbf{0}^{p_n-p_{n-1}}\beta_n\cdots.\end{aligned}$$

Next, we quantify our notion of sparsity.

[1] Instead of deterministic Muller automata, one could also use other, equally expressible classes of ω-automata like Büchi automata or deterministic parity automata.

Definition 4.1.8 (Effectively Sparse Predicates)**.** Let $d \geq 2$, $P_1, \ldots, P_d \subset \mathbb{N}$ be predicates, and $(p_m^{(i)})_{m=0}^{\infty}$ be the positive value sequence of P_i.

- P_1 is said to be *effectively sparse* if for any $K \in \mathbb{N}$, the inequality $p_{m+1}^{(1)} - p_m^{(1)} \leq K$ has finitely many solutions in m which can be effectively enumerated.
- $(P_1, \ldots, P_d)$ are called *pairwise effectively sparse* if $\bigcup_{i=1}^{d} P_i$ is effectively sparse.

Thus, a predicate is effectively sparse if its entries grow arbitrarily far apart, and we control how far they are at least apart from a certain point onward. A family of predicates is pairwise effectively sparse if their positive value sequence $(p_m)_{m=0}^{\infty}$ satisfies $\lim_{m\to\infty} p_{m+1} - p_m = \infty$ and for each $K \geq 0$, we can effectively compute all the solutions m to $p_m - p_{m-1} \leq K$.

Using this notion of sparsity, we generalize the Elgot-Rabin contraction method in the following two ways, which are proven in Section 4.3. First, we can show the following for a single predicate P.

Theorem 4.1.9. *Let P be an infinite, recursive, and effectively sparse predicate with positive value sequence $(p_m)_{m=0}^{\infty}$. Then the decidability $\mathsf{MSO}_{\mathbb{N};<}(P)$ reduces to the decidability of $\mathrm{Acc}_{(p_m \bmod M)_{m=0}^{\infty}}$ for all $M \geq 1$.*

For the predicates studied by Elgot and Rabin, Theorem 4.1.9 instantly reproduces their result as for every $M \geq 1$, then $(p_n \bmod M)_{n=0}^{\infty}$ will be ultimately periodic. For example, if $P = 2^{\mathbb{N}}$, we have $p_m = 2^m$. Then $(2^m \bmod M)_{m=0}^{\infty}$ is ultimate periodic, and we can compute the lengths of the preperiod and period explicitly. We call predicates with this property *effectively procyclic*. An infinite predicate with a positive value sequence $(p_m)_{m=0}^{\infty}$ is called effectively procyclic if for all $M \geq 1$, then $(p_m \bmod M)_{m=0}^{\infty}$ one can compute $N \geq 0$ and $r \geq 1$ such that $p_{m+r} \equiv p_m \pmod{M}$ for all $m \geq N$. All LRS with a single dominant root give rise to effectively procyclic predicates, which include the powers of 2, squares and Fibonacci numbers), and the set of factorial numbers is effectively procyclic as well.

Secondly, we obtain the following when expanding MSO with multiple predicates.

Theorem 4.1.10. *Assume that $P_1, \ldots, P_d$ are infinite, recursive, pairwise effectively sparse, effectively procyclic predicates with order word β. Then $\mathsf{MSO}_{\mathbb{N};<}(P_1, \ldots, P_d)$ reduces to Acc_{β}.*

Theorem 4.1.10 also recovers the results of Elgot and Rabin. Their predicates were effectively procyclic, and as $d = 1$, the word β is constant.

Multiple LRS with a simple dominant roots

In Section 4.4, we prove Theorem 4.1.1. Next, we present a high-level overview of our approach towards $\mathsf{MSO}_{\mathbb{N};<}(\rho_1^{\mathbb{N}}, \ldots, \rho_d^{\mathbb{N}})$, where each $\rho_i \geq 2$ is a natural number. As alluded to, we apply a variation of the Elgot-Rabin contraction method, and more specifically, Theorem 4.1.10.

These predicates $\rho_i^{\mathbb{N}}$ correspond to strictly increasing LRS and are thus effectively procyclic by Lemma 1.2.1. Moreover, such predicates are infinite, recursive, and effectively sparse.

However, a set $\rho_1^{\mathbb{N}}, \ldots, \rho_d^{\mathbb{N}}$ of predicates is not necessarily pairwise effectively sparse. If ρ_1 and ρ_2 are multiplicatively dependent, the sets $\rho_1^{\mathbb{N}}$ and $\rho_2^{\mathbb{N}}$ have an infinite overlap. In that case, (consider for example the case $\rho_1 = 4$ and $\rho_2 = 8$), we can find a number b (in our case, take $b = 2$), such that the sets $\rho_1^{\mathbb{N}}$ and $\rho_2^{\mathbb{N}}$ can be defined in $\mathsf{MSO}_{\mathbb{N};<}(b^{\mathbb{N}})$. Thus, instead of $\mathsf{MSO}_{\mathbb{N};<}(\rho_1^{\mathbb{N}}, \ldots, \rho_d^{\mathbb{N}})$, we study $\mathsf{MSO}_{\mathbb{N};<}(b^{\mathbb{N}}, \rho_3^{\mathbb{N}}, \ldots, \rho_d^{\mathbb{N}})$. Therefore, we can assume that $\rho_1, \ldots, \rho_d$ are pairwise multiplicatively independent. Then Baker's theorem implies that $\rho_1^{\mathbb{N}}, \ldots, \rho_d^{\mathbb{N}}$ are pairwise effectively sparse.

Thus, the hypothesis of Theorem 4.1.10 is satisfied, and we solely have to consider its order word γ. We devote Section 4.4.2 to proving that Acc_γ is decidable by showing that a suffix of γ can also be generated as a toric word. Toric words enjoy a crucial combinatorial property: they are almost-periodic. It is known that if a word β is *effectively* almost-periodic, then Acc_β is decidable (Theorem 4.2.4). Through number-theoretic arguments, we establish when the order word γ is indeed effectively almost-periodic.

Geometric series and polynomials

In Section 4.5, we again apply the Elgot-Rabin contraction method to expansions derived from multiple LRS. We analyse $\mathrm{Char}(P_1, P_2)$ for predicates $P_1 = \{qn^d \colon n \in \mathbb{N}\}$ and $P_2 = \{pb^{nd} \colon n \in \mathbb{N}\}$ to prove Theorem 4.1.2. The difference is that the underlying dynamical systems are driven by *numeration systems* [101, Chapter 7].

For the sake of exposition, assume we are dealing with $P_1 = \{n^2 \colon n \in \mathbb{N}\}$ and $P_2 = \{2 \cdot 4^n \colon n \in \mathbb{N}\}$. Then, P_1 and P_2 satisfy the conditions of Theorem 4.1.10 (where pairwise sparsity is the only non-trivial condition), and so we reduce to the order word, which enumerates the squares and numbers of the form $2 \cdot 4^n$. There are exactly $1 + \lfloor \sqrt{2 \cdot 4^n} \rfloor = 1 + \lfloor \sqrt{2} \cdot 2^n \rfloor$ squares below $2 \cdot 4^n$. This number is increasing, so the terms $2 \cdot 4^n$ are again effectively sparse within the set $P_1 \cup P_2$. And moreover, $\lfloor \sqrt{2} \cdot 2^n \rfloor \bmod 2$ is exactly the nth digit in the base-2 expansion of

$\sqrt{2}$. Hence, Theorem 4.1.2 follows by two applications of the Elgot-Rabin contraction method.

Linear recurrence sequences with two dominant roots

We now sketch our approach for our third and final main result of this chapter: the theory $\mathsf{MSO}_{\mathbb{N};<}(\{u_n\colon n \in \mathbb{N}\} \cap \mathbb{N})$ is decidable when $(u_n)_{n=0}^{\infty}$ is a simple, non-degenerate LRS with two dominant roots (Theorem 4.1.4).

Using Theorem 4.1.9, one of our versions of the Elgot-Rabin contraction method, the key to Theorem 4.1.4 is to understand the positive value sequence $(p_m)_{m=0}^{\infty}$ of $\{u_n\colon n \in \mathbb{N}\} \cap \mathbb{N}$. In particular, it is sufficient to show that $\mathrm{Acc}_{(p_m \bmod M)_{m=0}^{\infty}}$ is decidable for all $M \geq 1$.

To outline our approach, we defer to an example. Let $(u_n)_{n=0}^{\infty}$ defined by

$$u_{n+3} = 6u_{n+2} - 13u_{n+1} + 10u_n \tag{4.1}$$

and $u_0 = 2, u_1 = 4$, and $u_2 = 7$. Its subsequent values are

$$10, 9, -6, -53, -150, -271, -206, 787, 4690, 15849, 41994, 92827, 169530, 230369,$$
$$106594, -659933, -3041630, -8604711, -18686406, -30673493, -27164790, \ldots .$$

One can readily verify that $(u_n)_{n=0}^{\infty}$ satisfies the exponential-polynomial formula

$$u_n = \frac{1}{2}(2+i)^n + \frac{1}{2}(2-i)^n + 2^n\,. \tag{4.2}$$

From the exponential-polynomial form, we can infer that $(u_n)_{n=0}^{\infty}$ is indeed simple and non-degenerate with two dominant roots and thus satisfies the hypothesis of Theorem 4.1.4. It is straightforward that u_n is both infinitely often positive and infinitely often negative. Moreover, although $\lim_{n\to\infty} |u_n| = +\infty$, the sequence $(|u_n|)_{n=0}^{\infty}$ is *not* monotonically increasing. On the contrary, for all $N \in \mathbb{N}$ there is an $n \in \mathbb{N}$ such that $|u_{n+N}| < |u_n|$.

For our LRS $(u_n)_{n=0}^{\infty}$, let $(p_m)_{m=0}^{\infty}$ be the positive value sequence of $\{u_n\colon n \in \mathbb{N}\} \cap \mathbb{N}$. Then, in our example,

$$(p_m)_{m=0}^{\infty} = (u_0, u_1, u_2, u_4, u_3, u_{10}, u_{11}, u_{12}, u_{13}, u_{14}, u_{17}, u_{15}, u_{16}, u_{24}, u_{25}, u_{26}, u_{27}, u_{28},$$
$$u_{30}, u_{29}, u_{38}, u_{39}, u_{44}, u_{40}, u_{41}, u_{42}, u_{43}, u_{51}, u_{52}, u_{53}, u_{54}, u_{55}, u_{57}, u_{56}, \ldots)\,.$$

One immediately notices that two complications are at play here. The first one is that we are 'throwing away' all the negative values of our LRS (this in turn is necessary since we are working over the domain $\mathbb{N}$ rather than $\mathbb{Z}$). This restriction is, however, entirely benign in view of the following result:

Corollary 4.1.11. *Let $\overline{P} = \{u_n : n \in \mathbb{N}\}$ for $(u_n)_{n=0}^{\infty}$ a non-degenerate, simple, integer-valued linear recurrence sequence with two dominant roots. Then the MSO theory of $\langle \mathbb{Z}; 0, <, \overline{P} \rangle$ is decidable.*

This result is straightforwardly obtained from Theorem 4.1.4 via an application of Shelah's celebrated composition method in model theory [142]. In the case at hand, one can directly invoke, for example, [152, Corollary 6], since the structure $\langle \mathbb{Z}; 0, <, \overline{P} \rangle$ is isomorphic to the ordered sum $\langle \mathbb{Z} - \mathbb{N}; <, P^- \rangle + \langle \mathbb{N}; <, P \rangle$, where the second summand is as per Theorem 4.1.4, and the first structure is obtained by setting $P^- = \{u_n : n \in \mathbb{N}\} \cap (\mathbb{Z} - \mathbb{N})$. Intuitively speaking, MSO sentences over $\langle \mathbb{Z}; <, \overline{P} \rangle$ can be faithfully decomposed into component subformulas dealing exclusively with either positive or negative values of our LRS, with the truth value of the original sentence obtained by appropriately piecing together truth values of each of the sub-sentences within their respective structures.

The second complication is that, as noted earlier, the ordering of the positive values of our LRS does not respect the index ordering of the LRS. This ostensibly precludes the direct application of classical techniques such as Elgot and Rabin's contraction method (or, more generally, Carton and Thomas's effective profinite ultimate periodicity criterion [47]) or Semenov's toolbox of effective almost periodicity.

To prove Theorem 4.1.4, we therefore rely instead on a new concept, that of *(effective) prodisjunctivity*, which we introduce formally in Definition 4.2.9. Prodisjunctivity is itself predicated on the notion of *disjunctivity*. Informally an increasing sequence $(p_m)_{m=0}^{\infty}$ is effective prodisjunctive if for each $M \geq 1$, the suffix of $(p_m \bmod M)_{m=0}^{\infty}$ is disjunctive, and we can control where this suffix starts and which entries in $\{0, \ldots, M-1\}$ it contains. That is, we can compute $S_M \subset \{0, \ldots, M-1\}$ and $N_M \in \mathbb{N}$ such that $(p_m \bmod M)_{m=N_M}^{\infty}$ is a disjunctive sequence with respect to S_M (See Section 1.3.1 for the exact definitions).

We establish the following instrumental result:

Theorem 4.1.12. *Let $(p_m)_{m=0}^{\infty}$ be the positive value sequence of $\{u_n : n \in \mathbb{N}\} \cap \mathbb{N}$, where $(u_n)_{n=0}^{\infty}$ is a non-degenerate, simple, integer-valued LRS with two dominant roots. Then $(p_m)_{m=0}^{\infty}$ is effectively prodisjunctive.*

Theorem 4.1.12 already implies Theorem 4.1.4 because we can determine whether an ω-automaton accepts a recursive and disjunctive word (Theorem 4.2.8).

We return to our example and set $M = 5$. Then, applying Lemma 1.2.1,

$$u_n \bmod 5 = \begin{cases} 0 & \text{if } n \equiv 3 \pmod 4 \\ 2 & \text{if } n = 0, \text{ or if } n \equiv 2 \pmod 4 \\ 4 & \text{otherwise}. \end{cases} \tag{4.3}$$

Thus, it follows that $S_5 = \{0, 2, 4\}$, $N_5 = 0$, and

$$(p_m \bmod 5)_{m=0}^{\infty} = (2, 4, 2, 4, 0, 2, 0, 4, 4, 2, 4, 0, 4, 4, 4, 2, 0, 4, 2, 4, 2, 0, 4, 4, 4, 2, 0, 0, 4, 4, \\ 2, 0, 4, 4, 4, 2, 0, 0, 4, 4, 2, 2, 0, 4, 4, 2, 0, 4, 4, 4, 2, 0, 2, 4, 4, 2, 0, \dots)$$

is in $\{0, 2, 4\}^{\omega}$.

When computing the first million terms of $(p_m \bmod 5)_{m=0}^{\infty}$, one does not encounter the factor $(0, 0, 0)$. Nevertheless, according to Theorem 4.1.12, it should appear infinitely often! Indeed, we prove this in Section 4.6.2 and construct an index of roughly $2.18 \cdot 10^{59}$ where this occurs. However, in Section 4.6.2, we follow the suboptimal construction of our proof of Theorem 4.1.4, overlooking the far smaller number $m = 8479226$ for the index of the first occurrence of $(0, 0, 0)$ in $(p_m \bmod 5)_{m=0}^{\infty}$.

The above discussion indicates that, whilst $(p_m \bmod 5)_{m=0}^{\infty}$ is disjunctive with respect to the alphabet S_5, it is seemingly not normal, i.e., a given factor $w \in \{0, 2, 4\}^*$ does not necessarily appear with frequency $3^{-|w|}$. This is, however, unsurprising in view of (4.3): the residue class 4 should statistically appear approximately twice as often as either of the other two residue classes. Indeed it is possible to prove that this is asymptotically the case.

4.2 Special classes of infinite words

To establish the broad range of results presented in the previous section, we have to discuss a couple of classes of words and sequences.

Procyclic sequences

First, we generalize the notion of periodic words to sequences.

Definition 4.2.1. An effective, strictly increasing integer-valued sequence $(u_n)_{n=0}^{\infty}$ is called *effectively procyclic* if for all $M \geq 1$, the sequence $(u_n \bmod M)_{n=0}^{\infty}$ is ultimately periodic, and given M, one can effectively compute the period and preperiod of $(u_n \bmod M)_{n=0}^{\infty}$. A predicate is *effectively procyclic* if its positive value sequence is effectively procyclic.

The sets of factorial numbers, squares, and powers of two are all examples of effectively procyclic predicates.

Lemma 4.2.2. *Let $(u_n)_{n=0}^{\infty}$ satisfy $\lim_{n\to\infty} u_n = \infty$ and be a polynomial or an LRS with a single dominant root. Further, set $P = \{u_n : n \in \mathbb{N}\} \cap \mathbb{N}$. Then, P is recursive, infinite, effectively sparse, and effectively procyclic.*

Proof. Write $(u_n)_{n=0}^{\infty}$ in its exponential-polynomial form,

$$u_n = Q(n)\rho^n + r_n\,,$$

where $Q(n) \in (\mathbb{R} \cap \overline{\mathbb{Q}})[X]$ is non-zero, $\rho \in \mathbb{R} \cap \overline{\mathbb{Q}}^*$ and $(r_n)_{n=0}^{\infty}$ is a $\overline{\mathbb{Q}}$-LRS of spectral radius smaller than $|\rho|$. Using Lemma 1.2.6, there are $r, R \in \mathbb{Q}$ such that $R < \rho$ and $|r_n| < rR^n$ for all $n \in \mathbb{N}$. Hence, for large enough n, we have $|Q(n)\rho^n| > rR^n$ and so $Q(n)\rho^n$ determines the sign of u_n. As the sign of $Q(n)$ stabilizes for large enough n and $\lim_{n\to\infty} u_n = \infty$, we have $\rho > 1$ and the leading coefficient of $Q(n)$ is strictly positive. Moreover, if $\rho = 1$, we have $\deg(Q) > 1$.

Then $(u_n)_{n=0}^{\infty}$ and $(u_{n+1} - u_n)_{n=0}^{\infty}$ are both ultimately positive and ultimately increasing sequences, where the start of the ultimately positive suffix can be explicitly computed. Thus, P is recursive, infinite, and effectively sparse. To show that P is effectively procyclic, we combine the fact that $(u_n)_{n=0}^{\infty}$ is ultimately increasing with Lemma 1.2.1. □

Almost periodic and toric words

The second class of words we employ are the so-called almost-periodic words as introduced by Semenov in 1984 [140].

Definition 4.2.3. A word $\alpha \in \Sigma^{\omega}$ is *almost-periodic*, if for every $u \in \Sigma^+$, there exists $R(u) \in \mathbb{N}$ such that the word u either

(a) does not occur in $(\alpha_{R(u)}, \alpha_{R(u)+1}, \dots)$, or

(b) occurs in every factor of α of length $R(u)$.

If, moreover, α is recursive and for every $u \in \Sigma^+$ the *return time* $R(u)$ computable, then α is called *effectively almost-periodic*. If for every $u \in \Sigma^+$, only option (b) occurs (i.e., every factor of α occurs infinitely often with bounded intervals), then α is *uniformly recurrent*.

This result leads to the following theorem.

Theorem 4.2.4 (Theorem 3 in [118]). *If $\alpha \in \Sigma^\omega$ is effectively almost-periodic, then* Acc_α *is decidable.*

Although we only rely on the following lemma for uniformly recurrent words, we prove it for the more general class of almost-periodic words.

Lemma 4.2.5. *Let $\alpha \in \Sigma^\omega$ be recursive and almost periodic. Then α is effectively almost-periodic if and only if for all $w \in \Sigma^+$, one can decide whether w occurs in α.*

Proof. For increasing values of $n \geq |w|$, compute $T_n \subseteq \Sigma^n$, the set of all factors of α length n. If for all $u \in T_n$, the (shorter) word w is not a suffix of u, each u appears only finitely often in w and $R \leq n + |w| + 1$ (as else $\alpha_0 \cdots \alpha_{n+|w|+1}$ would be a factor of α of length $n + |w| + 2$ with suffix u). Thus, $R = n$ is a valid return time. If for all $u \in T_n$, the (shorter) word w is a factor of u, then w occurs infinitely often as a factor of w with a return time $R = n$. As one of these two scenarios has to unfold for large enough n (as a return time exists), we compute T_n until one of these scenarios occurs. □

Now, we can introduce the notion of toric words. Recall that $\mathbb{T}$ denotes the abelian group $\mathbb{R}/\mathbb{Z}$, viewed as the interval $[0, 1)$. For $x \in \mathbb{R}$, let $\{x\} := x - \lfloor x \rfloor$ be the fractional part of x. A word $\alpha \in \Sigma^\omega$ is *toric* if there exist a dimension $d > 0$, initial point $\mathbf{s} = (s_1, \ldots, s_d) \in \mathbb{T}^d$, translation $g : \mathbb{T}^d \to \mathbb{T}^d$ given by

$$(x_1, \ldots, x_d) \to \big(\{x_1 + \delta_1\}, \ldots, \{x_d + \delta_d\}\big) \tag{4.4}$$

for $\delta = (\delta_1, \ldots, \delta_d) \in \mathbb{T}^d$, and a collection $\mathcal{S} = \{S_b \colon b \in \Sigma\}$ of open, pairwise disjoint subsets of $\mathbb{T}^d$ such that for all $n \in \mathbb{N}$ and $b \in \Sigma$,

$$\alpha_n = b \iff g^{(n)}(\mathbf{s}) \in S_b\,. \tag{4.5}$$

Here, $g^{(n)}(\mathbf{s})$ denotes the result of iteratively applying g to $\mathbf{s}$ a total of n times. Thus, $g^{(n)}(\mathbf{s}) = (\{s_1 + n\delta_1\}, \ldots, \{s_d + n\delta_d\})$. That is, α is the *coding* (with respect to $\mathcal{S}$) of the trajectory of the discrete-time dynamical system on $\mathbb{T}^d$ defined by $(g, \mathbf{s})$. See [25] for a discussion of various subclasses of toric words.

For $\lambda_1, \ldots, \lambda_d \in \mathbb{T}$, we define the *group of additive relations on the torus* as

$$G_A(\lambda_1, \ldots, \lambda_d) := \big\{(k_1, \ldots, k_d) \in \mathbb{Z}^d \colon \{k_1\lambda_1 + \cdots + k_d\lambda_d\} = 0\big\}\,.$$

As $G_A(\boldsymbol{\delta}) \subset \mathbb{Z}^d$ is a torsion-free, finitely generated abelian group that has thus a basis $\mathbf{v}_1, \ldots, \mathbf{v}_m \in \mathbb{Z}^d$ for some $m \leq d$. Let $\mathbf{s}$, $\boldsymbol{\delta} \in \mathbb{T}^d$, and $g \colon \mathbb{T}^d \to \mathbb{T}^d$ be as in (4.4). Next, we define the following set:

$$\mathbb{T}_{\boldsymbol{\delta}} := \big\{\mathbf{z} \in \mathbb{T}^d \colon \big(\forall 1 \leq i \leq m \colon \mathbf{v}_i \cdot \mathbf{z} \in \mathbb{Z}\big)\big\}\,.$$

Thus, $\mathbb{T}_{\boldsymbol{\delta}}$ is equal to the quotient $\mathbb{T}^d/G_A(\boldsymbol{\delta})$. One can define $\mathbb{T}_{\boldsymbol{\delta}}$ arithmetically as follows. When the basis $\mathbf{v}_1, \ldots, \mathbf{v}_m$ is known, compute $C = \max_i \|\mathbf{v}_i\|_\infty$ such that $|\mathbf{v}_i \cdot \mathbf{z}| \leq Cd$ for all $\mathbf{z} \in \mathbb{T}^d$. We have that $z \in \mathbb{T}_{\boldsymbol{\delta}}$ if and only if $\mathbf{v}_i \cdot \mathbf{z} \in \mathbb{Z}$ for all i, which is equivalent to

$$\bigwedge_{i=1}^{m} \bigvee_{|k| \leq Cd} \mathbf{v}_i \cdot \mathbf{z} = k\,.$$

That is, $\mathbb{T}_{\boldsymbol{\delta}}$, viewed as a subset of $\mathbb{R}^d$, is an intersection of $\mathbb{T}^d$ with a union of affine subspaces of $\mathbb{R}^d$ with integer parameters.

We can now define $\mathbb{T}_{\boldsymbol{\delta},\mathbf{s}} := \{(\{s_1+z_1\}, \ldots, \{s_d+z_d\}) \mid (z_1, \ldots, z_d) \in \mathbb{T}_{\boldsymbol{\delta}}\}$. Applying Theorem 1.1.13 and shifting by $\mathbf{s}$, gives the following result.

Theorem 4.2.6. *The orbit $(g^{(n)}(\mathbf{s}))_{n=0}^{\infty}$ is dense in $\mathbb{T}_{\boldsymbol{\delta},\mathbf{s}}$, and for every non-empty open subset O of $\mathbb{T}_{\boldsymbol{\delta},\mathbf{s}}$ there exist infinitely many $n \in \mathbb{N}$ such that $g^{(n)}(\mathbf{s}) \in O$.*

To prove Theorem 4.1.1, we will need to show that Acc_α is decidable for certain toric words. The following fact will play an important role in this [25].

Theorem 4.2.7. *Every toric word is uniformly recurrent.*

Disjunctive words

The third class of words and sequences we study stems from disjunctive words. Recall from Section 1.3.1 that $\alpha \in \Sigma^\omega$ is *disjunctive* if every $u \in \Sigma^*$ appears infinitely often in α.

Theorem 4.2.8. *If α is recursive and disjunctive, then Acc_α is decidable.*

Proof. Consider a given deterministic Muller automaton $\mathcal{A}$ as a directed graph allowing multiple edges. We partition the graph into its strongly connected components (SCCs) and call an SCC without outgoing edges a bottom SCC. We will show that the set of states visited infinitely often by the run of α on $\mathcal{A}$ is precisely a bottom SCC. Hence, we decide Acc_α by simulating this run until a bottom SCC is inevitably reached. Then α is accepted if and only if this bottom SCC is in the Muller acceptance condition.

We need to show that (a) if an SCC is not a bottom SCC, then the run eventually exits it; and (b) if the run enters a bottom SCC, it visits all its states infinitely often.

For (a), let S be a non-bottom SCC. There thus exist $q_1 \in S, b \in \Sigma$ such that $\delta(q_1, b) \notin S$, i.e., reading the letter b in state q_1 exits S. We will order the states of S as $q_1, \ldots, q_k$, construct words $u_1 = b, u_2, \ldots, u_k \in \Sigma^+$, and inductively prove that

$\delta(q_j, u_i) \notin S$ for all $j \leq i$. Since α is disjunctive, the word u_k will inevitably occur as a factor, and hence, the run will exit S.

We have observed the base case to hold with $u_1 = b$. For the induction step, assume that $\delta(q_j, u_i) \notin S$ for all $j \leq i$. Now, if $\delta(q_{i+1}, u_i) \notin S$, take $u_{i+1} = u_i$. Otherwise, if $\delta(q_{i+1}, u_i) = q \in S$, the strong connectivity of S implies that $\delta(q, v_i) = q_1$ for some $v_i \in \Sigma^*$. Thus, take $u_{i+1} = u_i v_i b$, and observe that $\delta(q_j, u_{i+1}) \notin S$ for all $j \leq i + 1$.

We prove (b) similarly. Fix an order of states $q_1, \ldots, q_k$ of S. By definition, a run entering the bottom SCC S will be confined in S. It thus suffices to prove that for any $q \in S$, we can inductively construct a word $u_k \in \Sigma^+$ such that for all $j \leq k$, the non-empty run of u_k on $\mathcal{A}$ starting from q_j visits q. The induction is similar to the one above. Choose $q_1 \in S$ and $u_1 \in \Sigma^*$ to be such that $\delta(q_1, u_1) = q$. By the induction hypothesis, for every $j \leq i$, the run on u_i starting in q_j visits q. If the run on u_i starting in q_{i+1} visits q, take $u_{i+1} = u_i$. Otherwise, use that S is a bottom SCC to identify $v_i \in \Sigma^+$ such that $\delta(q_{i+1}, u_i v_i) = q$, and take $u_{i+1} = u_i v_i$. We have thus ensured that for all $j \leq i + 1$, the run on u_i starting in q_j visits q. Since α is disjunctive, the word u_k occurs as a factor infinitely often, and hence all $q \in S$ are visited infinitely often. □

Intuitively, the proof uses the abundance of each factor to guarantee that the set of states visited infinitely often is an entire bottom strongly connected component in the graph induced by the automaton.

Similar to the generalization of ultimately periodic words to procyclic predicates, we can lift the notion of disjunctive words to sequences of integers. First, recall that for $\widetilde{\Sigma} \subset \Sigma$, the infinite word $\alpha \in \Sigma^\omega$ is disjunctive with respect to $\widetilde{\Sigma}$ if every $u \in \widetilde{\Sigma}$ appears infinitely often in α.

Definition 4.2.9. Let $(p_m)_{m=0}^{\infty}$ be an increasing integer-valued sequence. For any integer $M \geq 1$, let S_M be the set of residue classes modulo M that appear infinitely often in $(p_m \bmod M)_{m=0}^{\infty}$:

$$S_M = \left\{ s \in \{0, \ldots, M-1\} \colon \left(\exists^\infty m \in \mathbb{N} \colon p_m \equiv s \pmod{M} \right) \right\} \tag{4.6}$$

and N_M a threshold such that for every $m \geq N_M$, we have $p_m \equiv s \pmod{M}$. We say that the sequence $(p_m)_{n=0}^{\infty}$ is *prodisjunctive* if, for all $M \geq 1$, the sequence of residues $(p_m \bmod M)_{m=N_M}^{\infty}$ is disjunctive with respect to S_M, and $(p_m)_{m=0}^{\infty}$ is *effectively prodisjunctive* if, in addition, for each M, the set S_M and threshold N_M are also computable.

The classes of almost-periodic and disjunctive words are orthogonal as the words within the intersection of these classes are ultimately constant as the following proposition shows.

Proposition 4.2.10. *Assume that $|\Sigma| > 1$. Then $\alpha \in \Sigma^\omega$ cannot be almost-periodic and disjunctive.*

Proof. Let α be almost periodic and disjunctive, and let $a, b \in \Sigma$ be distinct. Then, by disjunctivity, a and b appear infinitely often in α. Hence, as α is almost-periodic, there is a number R such that for every $N \geq 0$, the letter a is in $\alpha_N \cdots \alpha_{N+R-1}$. However, as α is disjunctive, the word b^R also appears infinitely often in α. Thus, for some $n \geq 0$, we have $\alpha_n \cdots \alpha_{n+R-1}$ is equal to b^R but also contains the letter a. This causes a contradiction. □

4.3 Reductions to order words

The purpose of this section is to establish Theorems 4.1.9 and 4.1.10, which are our versions of the Elgot-Rabin contraction method. Our tools will be very similar to those we proved in [24] on which part of this chapter is based. For conciseness, we only prove the results necessary for the main theorems of this chapter. In particular, we rely on the following lemma that asserts that the decidability of acceptance of an automaton is closed under transduction. The proof is given in [24, Lemma 4.5].

Lemma 4.3.1 (Transduction)**.** *Let $\alpha \in \Sigma^\omega$, $\mathcal{B}$ be a deterministic finite transducer with input alphabet Σ and output alphabet Γ, and $\beta = \mathcal{B}(\alpha) \in \Gamma^\omega$. Then the problem* Acc_β *reduces to* Acc_α*.*

Recall that we associate to a *characteristic word* $\alpha \in \Sigma^\omega$ to a tuple of predicates $(P_1, \ldots, P_d)$, where $\Sigma = \{0,1\}^d$ and α_n records which of the predicates hold for n. Next to the transduction property stated above, we will rely on the property that the predicates $P_1, \ldots, P_d$ we use are 'sparse'. That is, if we look at the characteristic word α, the vast majority of all letters will be $(0, \ldots, 0)$: Most numbers are in none of the predicates P_i. Hence, we want to 'compress' the characteristic word to make it more convenient. By Büchi's theorem (Theorem 4.1.6), we only need to determine whether a given deterministic Muller automaton $\mathcal{A}$ accepts α, and this automaton has finite memory. Therefore, when $\mathcal{A}$ reads a string of K letters $(0, \ldots, 0)$, the automaton $\mathcal{A}$ can only remember what K is modulo a certain number and whether K exceeds a certain threshold. That is, for every such automaton $\mathcal{A}$, there are numbers M and N

such that for all $K \geq N$, the automaton $\mathcal{A}$ cannot discern between reading $(0,\dots,0)^K$ and $(0,\dots,0)^{K+M}$.

Recall that the positive value sequence $(p_m)_{m=0}^{\infty}$ enumerates $\bigcup_{i=1}^{d} P_i$, i.e., the sequence of indices n for which n is is some P_i, and that the order word β is the word obtained when removing all letters $(0,\dots,0)$ from α. Thus, for all $m \in \mathbb{N}$, we have $\beta_m = \alpha_{p_m}$. We compress α as follows.

Theorem 4.3.2. *Assume that $P_1,\dots,P_d$ are infinite, recursive, and pairwise effectively sparse predicates with positive value sequence $(p_m)_{m=0}^{\infty}$ and order word β. Then $\mathsf{MSO}_{\mathbb{N};<}(P_1,\dots,P_d)$ reduces to $\mathrm{Acc}_{((\beta_m, p_m \bmod M))_{m=0}^{\infty}}$.*

Proof. Using Theorem 4.1.6, $\mathsf{MSO}_{\mathbb{N};<}(P_1,\dots,P_d)$ is decidable if one can determine whether a deterministic Muller automaton $\mathcal{A} = (\{0,1\}^d, Q, q_{\text{init}}, \delta, \mathcal{F})$ accepts the characteristic word α of $(P_1,\dots,P_d)$. Write $\mathbf{0}$ for $(0,\dots,0)$ and restrict $\mathcal{A}$ to a directed graph G with nodes Q and $\mathbf{0}$-transitions as arrows.

By construction, every node in G has outdegree 1, and so each state is in at most one cycle in G. Therefore, we can compute the least common multiple of the cycle lengths in G (call this number M) and the longest path to a cycle (call this number N). Then, for all states q, numbers $n \geq M+N$, and $d \geq 1$, reading $\mathbf{0}^n$ and $\mathbf{0}^{n+dM}$ leads to journeying through the exact same set of states and ending up in the same state.

As the tuple $(P_1,\dots,P_d)$ is pairwise effectively sparse, we can compute a number K such that for all $m \geq K$, we have $p_{m+1} - p_m > M + N$. We construct a deterministic finite transducer $\mathcal{B}$ that hard-codes the initial segment of α as follows: $\alpha_{\text{init}} := \mathbf{0}^{p_0}\beta_0\mathbf{0}^{p_1-p_0-1}\beta_1 \cdots \mathbf{0}^{p_K-p_{K-1}-1}\beta_K$. For $m > K$, after reading $(p_{m-1} \bmod M)$ and $(p_m \bmod M)$, the transducer $\mathcal{B}$ outputs $\mathbf{0}^{k_m}\beta_m$, where $M+N < k_m \leq 2M+N$ is congruent to $p_m - p_{m-1} - 1$ modulo M. Then, by construction, a state q is visited infinitely often upon reading the characteristic word of $(P_1,\dots,P_d)$ if and only if q is visited infinitely often when $\mathcal{A}$ reads $\alpha_{\text{init}}\mathcal{B}((\beta_m, p_m \bmod M))_{m=K+1}^{\infty})$.

Assume $\mathrm{Acc}_{(\beta_m, p_m \bmod M))_{m=0}^{\infty}}$ is decidable. By hard-coding the initial segment, the problem $\mathrm{Acc}_{(\beta_m, p_m \bmod M))_{m=K+1}^{\infty}}$ is also decidable and thus $\mathrm{Acc}_{\mathcal{B}((\beta_m, p_m \bmod M))_{m=K+1}^{\infty})}$ is decidable by Theorem 4.3.1. Then, $\mathrm{Acc}_{\alpha_{\text{init}}\mathcal{B}((\beta_m, p_m \bmod M))_{m=K+1}^{\infty}}$ is decidable (by again hard-coding the initial segment), and so by construction, the automaton $\mathcal{A}$ accepts α if and only if $\mathcal{A}$ accepts $\alpha_{\text{init}}\mathcal{B}((\beta_m, p_m \bmod M))_{m=K+1}^{\infty})$. Hence, Acc_{α} is decidable, as required. □

Due to Theorem 4.3.2, we are now in the position to prove Theorems 4.1.9 and 4.1.10. We start with the former.

Proof of Theorem 4.1.9. The hypothesis of Theorem 4.3.2 is clearly fulfilled. We now observe in the case of a single predicate, we have $\beta \in (\{0,1\} \setminus \{0\})^\omega = \{1\}^\omega$ and thus β is a constant word. Then by relabelling the alphabet from $\{1\} \times \{0, \ldots, M-1\}$ to the $\{0, \ldots, M-1\}$, we obtain the result. □

Proof of Theorem 4.1.10. Again, the hypothesis of Theorem 4.3.2 is clearly fulfilled. Fix $M \geq 1$ and let α and β denote the characteristic- and order word of $(P_1, \ldots, P_d)$, respectively. For $1 \leq i \leq d$, one can compute the preperiod N_i and period M_i of the positive value sequence $(p_m^{(i)} \bmod M)_{m=0}^\infty$ of P_i as each P_i is effectively procyclic.

We construct a transducer $\mathcal{B}$ such that $\mathcal{B}(\beta) = ((\beta_m, p_m \bmod M))_{m=0}^\infty$. Then, Theorem 4.3.1 implies the result. The transducer $\mathcal{B}$ has states $\{0, \ldots, N_1\} \times \{0, \ldots, M_1 - 1\} \times \cdots \times \{0, \ldots, N_d\} \times \{0, \ldots, M_d - 1\}$. The state corresponding to N_i keeps track of the number of entries from P_i the transducer $\mathcal{B}$ has read, with the exact number if there are less than N_i and N_i otherwise. The state corresponding to M_i keeps track how many entries from P_i the transducer $\mathcal{B}$ has read modulo M_i. Now assume we read β_m. If the ith component of β_m is 0, then we do not change the components of the state belonging to N_i and M_i. If the ith component of β_m is 1, then we update the components of the state belonging to N_i and M_i accordingly and output the corresponding value for $(\beta_m, p_m \bmod M)$. This value $p_m \bmod M$ is independent of which component of β_m equals 1 causes it to output, but as $\beta_m \neq (0, \ldots, 0)$ at least one such value outputs it. Hence, when reading β, the automaton indeed outputs $((\beta_m, p_m \bmod M))_{m=0}^\infty$. □

4.4 Multiple linear recurrence sequence with a single dominant root

In this section, we prove Theorem 4.1.1 by proving a more general version.

Theorem 4.4.1. *Consider $\mathbb{Z}$-LRS $(u_n^{(1)})_{n=0}^\infty, \ldots, (u_n^{(d)})_{n=0}^\infty$ with a single dominant root with the following properties for $1 \leq i \leq d$:*

(1) the exponential-polynomial form of $(u_n^{(i)})_{n=0}^\infty$ can be decomposed as $u_n^{(i)} = c_i \rho_i^n + r_n^{(i)}$ such that $(c_i \rho_i^n)_{i=0}^\infty$ and $(r_n^{(i)})_{n=0}^\infty$ are the dominant and non-dominant parts of $(u_n^{(i)})_{n=0}^\infty$, respectively;

(2) $c_i > 0$ and $\rho_i > 1$;

(3) for $1 \leq j \leq d$ not equal to i, there exist only finitely many pairs $(n, m) \in \mathbb{N}^2$ such that $c_i \rho_i^n = c_j \rho_j^m$.

Writing $P_i = \{u_n^{(i)} : n \in \mathbb{N}\}$, *the theory* $\mathsf{MSO}_{\mathbb{N};<}(P_1, \ldots, P_d)$ *is decidable assuming the weak Schanuel conjecture. The decidability is unconditional in either of the following cases:*

(a) If $1/\log(\rho_1), \ldots, 1/\log(\rho_d)$ *are linearly independent over* $\mathbb{Q}$*;*

(b) If $\operatorname{rank}(G_M(\rho_1, \ldots, \rho_d)) \geq d - 2$ *and* $\rho_1, \ldots, \rho_d$ *are pairwise multiplicatively independent.*

Several remarks are in order concerning this theorem. First, we go over conditions (1)–(3).

Condition (1) is equivalent to $(u_n^{(i)})_{n=0}^{\infty}$ having a single dominant root that is not repeated and excludes examples like polynomials and $(n2^n)_{n=0}^{\infty}$.

Condition (2) is not necessary, but we include it as it simplifies the proof. As we are dealing with the minimal representation of an LRS, we have $c_i \neq 0$, and if $c_i < 0$, then $u_n^{(i)}$ is positive for finitely many n that we can compute (this is an instance of the Positivity problem with a single dominant root). Thus, P_i is finite and computable and therefore definable already in the MSO theory of $\langle \mathbb{N}; < \rangle$. As $(u_n^{(i)})_{n=0}^{\infty}$ is a $\mathbb{Z}$-LRS, the ineuqality $|\rho_i| \leq 1$ would imply that $|u_n^{(i)}|$ is bounded, and we can again actually compute the finite set P_i. Finally, if $\rho_i < 1$, then $(u_{2n}^{(i)})_{n=0}^{\infty}$ and $(u_{2n+1}^{(i)})_{n=0}^{\infty}$ are both LRS with a strictly positive dominant root where one LRS satisfies Condition (2) while the other has a negative coefficient c in front of its dominant root. Thus, we can apply the theorem with either the even or odd terms of $(u_n)_{n=0}^{\infty}$ while the other one gives rise to a finite set P_i that we can explicitly compute.

Condition (3) is equivalent to the following: For every $1 \leq i, j \leq d$, the equation $c_i \rho_i^n = c_j \rho_j^m$ has at most one solution $(n, m) \in \mathbb{N}^2$; The proof is elementary. Omitting this condition quickly leads to Positivity-hardness.

Next, we discuss the second half of the statement of Theorem 4.4.1 involving the weak Schanuel conjecture. The main technical difficulty in the proof of Theorem 4.4.1 is to translate the MSO theory into a problem of linear programming on a $d-1$-dimensional torus $\mathbb{T}^{d-1} = [0, 1)^{d-1}$. We carefully partition this torus into pairwise disjoint open subsets $\mathcal{S} := \{S_b : b \in \Sigma\}$ defined by linear inequalities and find a point $\mathbf{s} = (s_1, \ldots, s_{d-1})$ to reduce to understanding the orbit of the map

$$g : \mathbb{T}^{d-1} \to \mathbb{T}^{d-1}, \quad (z_1, \ldots, z_{d-1}) \mapsto \left(\left\{ z_1 + \frac{\log(\rho_2)}{\log(\rho_1)} \right\}, \ldots, \left\{ z_{d-1} + \frac{\log(\rho_d)}{\log(\rho_1)} \right\} \right).$$

That is, if we can decide Acc_γ, where γ is defined by $\gamma_n = b$ if and only if $g^{(n)}(\mathbf{s}) \in S_b$, the theory $\mathsf{MSO}_{\mathbb{N};<}(P_1, \ldots, P_d)$ is decidable.

Of course, γ will be a toric word, and by Theorem 4.2.7, uniformly recurrent and thus almost periodic. We need to prove that γ is *effectively* almost-periodic, which implies that Acc_γ is decidable by Theorem 4.2.4. However, the open subsets S_b and the initial point $\mathbf{s}$ are defined by ugly expressions in logarithms, while we also have to determine the subtorus where the orbit of g is dense. Using the weak Schanuel conjecture, we can circumvent these number-theoretical problems. Condition (a) implies the orbit is dense, and we can avoid relying on the weak Schanuel conjecture. Condition (b) implies Condition (a), as we will see in Lemma 4.4.12.

Theorem 4.4.1 implies Theorem 4.1.1 from this chapter's introduction.

Proof of Theorem 4.1.1. Assume that ρ_1 and ρ_2 are pairwise multiplicatively independent, say $\rho_1^a = \rho_2^b$ for some $a, b \in \mathbb{N}_{\geq 1}$. Then set $\rho = \rho_1^{a/\operatorname{lcm}(a,b)} = \rho_2^{b/\operatorname{lcm}(a,b)}$, which is an integer as $\rho^{\operatorname{lcm}(a,b)}$ is both a power of a and b. The sets $\rho_1^{\mathbb{N}}$ and $\rho_2^{\mathbb{N}}$ can be defined in $\mathsf{MSO}_{\mathbb{N};<}(\rho^{\mathbb{N}})$ and so $\mathsf{MSO}_{\mathbb{N};<}(\rho_1^{\mathbb{N}}, \ldots, \rho_d^{\mathbb{N}})$ reduces to $\mathsf{MSO}_{\mathbb{N};<}(\rho^{\mathbb{N}}, \rho_3^{\mathbb{N}}, \ldots, \rho_d^{\mathbb{N}})$. Thus, without loss of generality, $\rho_1, \ldots, \rho_d$ are pairwise multiplicatively independent. The result is immediate. □

4.4.1 Reduction to the order word

To prove Theorem 4.4.1, we will use the framework described in Section 4.1. Let $(u_n^{(i)})_{n=0}^{\infty}$ and P_i be as in the statement of Theorem 4.4.1. Let α denote the characteristic word of $(P_1, \ldots, P_d)$, and by β the order word of $(P_1, \ldots, P_d)$.

Our first goal is to understand the structure of the linear recurrence sequences satisfying Theorem 4.4.1. We prove the following lemma derived from Matveev's version of Baker's theorem (Theorem 1.1.11).

Theorem 4.4.2. *Let $b_1, b_2, c_1, c_2, \rho_1, \rho_2, R_1, R_2$ be positive real algebraic numbers such that $\rho_1 > R_1, 1$ and $\rho_2 > R_2, 1$. Then, one can compute $N \in \mathbb{N}$ such that for all $n_1, n_2 \geq N$, the inequality*

$$\left|c_1\rho_1^{n_1} - c_2\rho_2^{n_2}\right| \leq b_1 R_1^{n_1} + b_2 R_2^{n_2} \tag{4.7}$$

implies that $c_1\rho_1^{n_1} = c_2\rho_2^{n_2}$.

Proof. First assume that ρ_1 and ρ_2 are multiplicatively dependent, say $\rho_1^{m_1} = \rho_2^{m_2}$. Then let $\rho_3 = \rho_1^{1/m_2}$ and $R_3 = \max(R_1^{1/m_2}, R_2^{1/m_1}, 1)$. Then $R_3 < \rho_3$ and solving

$$\left|c_1\rho_3^{n_1'} - c_2\rho_3^{n_2'}\right| \leq (b_1 + b_2) R_3^{\max(n_1', n_2')} \tag{4.8}$$

gives all solutions to (4.7) by setting $n_1 = m_2 n_1'$ and $n_2 = m_1 n_2'$. When $c_1\rho_1^{n_1} = c_2\rho_2^{n_2}$ has finitely many solutions, the left-hand side of (4.8) can be bounded from below by $C_1\rho_3^{\max(n_1',n_2')}$ for a computable constant $C_1 > 0$. Thus, we can take N large enough such that $C_1\rho_3^n > (b_1+b_2)R_3^n$ for all $n \geq N$.

Now assume that ρ_1 and ρ_2 are multiplicatively independent. It is sufficient to bound n_1, n_2 that satisfy

$$\left|c_1\rho_1^{n_1} - c_2\rho_2^{n_2}\right| \leq 2b_1R_1^{n_1} \quad \text{or} \quad \left|c_1\rho_1^{n_1} - c_2\rho_2^{n_2}\right| \leq 2b_2R_2^{n_2} \tag{4.9}$$

as adding the two cases gives the result. Due to Theorem 1.1.3, the finitely many solutions of $c_1\rho_1^{n_1} = c_2\rho_2^{n_2}$ can be effectively computed, and we can assume that n_1 and n_2 are large enough such that $c_1\rho_1^{n_1} \neq c_2\rho_2^{n_2}$.

Without loss of generality, $\rho_1 < \rho_2$, and if $n_1 > n_2$, we have that for n_1 and n_2 satisfying (4.9), the number n_1 is effectively bounded as $c_1\rho_1^{n_1} < c_2\rho_2^{n_1} + 2b_iR_i^{n_1}$ has finitely many solutions for $i = 1, 2$. Moreover, we can assume that $n_2 \geq 1$. Similarly, let $k \in \mathbb{N}$ be such that $\rho_2^k > \rho_1$. By the same argument, we can assume that $n_2 \leq kn_1$.

After dividing both sides of (4.9) by $c_1\rho_1^{n_1}$, we apply Matveev's result to $\Lambda = c_1^{-1}c_2(\rho_1^{-1})^{n_1}\rho_2^{n_2} - 1$ (which is non-zero by assumption) to find that

$$\log|\Lambda| > -C_1\big(1 + \log(\max(1, n_1, n_2))\big) = -C_1\big(1 + \log(n_2)\big) \tag{4.10}$$

for some computable constant $C_1 > 0$. Thus, combining (4.9) and (4.10), we need to bound $n_1 \leq n_2$ that satisfy

$$\begin{aligned} C_1(1 + \log(n_2)) &> n_1\log(\rho_1/R_1) - \log(2b_1/c_1) \quad \text{or} \\ C_1(1 + \log(n_2)) &> n_1\log(\rho_1) - n_2\log(R_2) - \log(2b_2/c_1)\,. \end{aligned}$$

Thus,

$$\begin{aligned} C_1\big(1 + \log(n_2)\big) &> \frac{1}{k}n_2\log(\rho_1/R_1) - \log(2b_1/c_1) \quad \text{or} \\ C_1\big(1 + \log(n_2)\big) &> \frac{1}{k}n_2\log(\rho_1) - n_2\log(R_2) - \log(2b_2/c_1)\,. \end{aligned}$$

In both cases, we have an inequality of the form $C_1(1 + \log(n_2)) > n_2c_1 + c_3$ for some computable constants c_2 and c_3. This effectively bounds n_2 (and thus n_1 as well). The result follows. □

Corollary 4.4.3. *With the notation from Theorem 4.4.1, if $1 \leq i, j \leq d$ are distinct, the predicates P_i and P_j have a finite overlap which we can effectively compute, are effectively pairwise sparse, and for some computable number N, we have $u_{n_1}^{(i)} < u_{n_2}^{(j)}$ if and only if $c_i\rho_i^{n_1} < c_j\rho_j^{n_2}$.*

Proof. Let $K \in \mathbb{N}$. Using Lemma 1.2.6, we compute $r_1, r_2, R_1 < \rho_1$ and $R_2 < \rho_2$ such that $|r_n^{(i)}| < r_1 R_1^n - K$ and $|r_n^{(j)}| < r_2 R_2^n$ for all $n \in \mathbb{N}$. Then, applying the lemma above, we have that

$$|u_{n_1}^{(i)} - u_{n_2}^{(i)}| - K > |c_i \rho_i^{n_1} - c_j \rho_i^{n_1}| - |r_{n_1}^{(i)}| - |r_{n_2}^{(j)}| - K > |c_i \rho_i^{n_1} - c_j \rho_i^{n_1}| - r_1 R_1^{n_1} - r_2 R_2^{n_2}$$

has only finitely many solutions, and we can compute $N \in \mathbb{N}$, such that there are no solutions with $n_1, n_2 \geq N$. □

Due to the previous corollary, we prefer to rely on the sequences $(c_i \rho_i^n)_{n=0}^\infty$ instead of the sequences $(u_n^{(i)})_{n=0}^\infty$. To do this, we need to generalize the notion of order words.

Definition 4.4.4. Let $(v_n^{(1)})_{n=0}^\infty, \dots, (v_n^{(d)})_{n=0}^\infty$ be a family of real-valued, pairwise disjoint, and strictly increasing sequences. Further let $Z = \bigcup_{i=1}^d \{v_n^{(i)} \colon n \in \mathbb{N}\}$. We define the word

$$\xi := \mathsf{Ord}\big((v_n^{(1)})_{n=0}^\infty, \dots, (v_n^{(d)})_{n=0}^\infty\big) \in \{1, \dots, d\}^\omega$$

by

$$\xi_n = i \iff \exists z \in (v_n^{(i)})_{n=0}^\infty \colon \big|\{y \in Z \colon y < z\}\big| = n\,.$$

If predicates $P_1, \dots, P_d$ are pairwise disjoint, then the two definitions of the order word almost coincide, with the sole difference being the alphabet being $\{1, \dots, d\}$ or $\{0,1\}^d$.

We now prove the following.

Theorem 4.4.5. *With the notation as in Theorem 4.4.1, let α be the characteristic word of $(P_1, \dots, P_d)$. Then there are computable real algebraic numbers $r_2, \dots, r_d$ such that when setting*

$$\xi := \mathsf{Ord}\left(\mathbb{N}, \left(\frac{\log(r_2)}{\log(\rho_1)} + n\frac{\log(\rho_2)}{\log(\rho_1)}\right)_{n=0}^\infty, \dots, \left(\frac{\log(r_d)}{\log(\rho_1)} + n\frac{\log(\rho_d)}{\log(\rho_1)}\right)_{n=0}^\infty\right), \tag{4.11}$$

the problem Acc_α reduces to Acc_ξ. Moreover, we can assume that $\rho_1 \leq \rho_2, \dots, \rho_d$, that $0 < \log(r_i) < \log(\rho_i)$ for $2 \leq i \leq d$, and the sets $\mathbb{N}$, $\left(\frac{\log(r_2)}{\log(\rho_1)} + n\frac{\log(\rho_2)}{\log(\rho_1)}\right)_{n=0}^\infty$, $\dots$, $\left(\frac{\log(r_d)}{\log(\rho_1)} + n\frac{\log(\rho_d)}{\log(\rho_1)}\right)_{n=0}^\infty$ have pairwise empty intersections.

Proof of Theorem 4.4.5. As we have not assumed anything about the relative sizes of $\rho_1, \dots, \rho_d$ yet, we can indeed assume that $\rho_1 \leq \rho_2, \dots, \rho_d$. Let β be the order word of $(P_1, \dots, P_d)$. Using Lemma 4.2.2, we have that each P_i is infinite, recursive, effectively

sparse, and effectively procyclic. Due to Corollary 4.4.3, the predicates $P_1, \dots, P_d$ are also pairwise effectively sparse. Hence, by Theorem 4.1.10, the problem Acc_α reduces to Acc_β.

Due to the same Corollary 4.4.3, one can compute a number $N \in \mathbb{N}$ such that for all $n_1, n_2 \geq N$ and distinct $1 \leq i, j \leq d$, we have $c_i\rho_i^{n_1} \neq c_j\rho_j^{n_2}$ and $u_{n_1}^{(i)} < u_{n_2}^{(j)}$ if and only if $c_i\rho_i^{n_1} < c_j\rho_j^{n_2}$. Thus,

$$\xi' = \mathsf{Ord}\big((u_n^{(1)})_{n=N}^\infty, \dots, (u_n^{(d)})_{n=N}^\infty\big) = \mathsf{Ord}\big((c_1\rho_1^n)_{n=N}^\infty, \dots, (c_d\rho_d^n)_{n=N}^\infty\big)\,.$$

By the comments above, the problem Acc_β reduces to $\mathrm{Acc}_{\xi'}$ as (by the comments above) changing the alphabet of ξ' to $\{0,1\}^d$ by mapping i to the unit vector e_i and adding a finite, computable prefix gives β. Now let $N_1 \geq N$ be the smallest number such that for all $2 \leq i \leq d$, we have $c_i\rho_i^{N-1} < c_1\rho_1^{N_1}$. Next, for $2 \leq i \leq d$, let N_i be the smallest number such that $c_1\rho_1^{N_1} < c_1\rho_i^{N_i} < \rho_i c_1\rho_1^{N_1}$, which exists as $N_i, N_1 \geq N$ and thus $c_1\rho_1^{N_1} \neq c_i\rho_i^{N_1}$ would have been impossible. Then set $r_i = \frac{c_i\rho_i^{N_i}}{c_1\rho_1^{N_1}}$ and

$$\xi'' = \mathsf{Ord}\big((c_1\rho_1^n)_{n=N_1}^\infty, \dots, (c_d\rho_d^n)_{n=N_d}^\infty\big)\,.$$

Then, as ξ'' and ξ' only have different, computable prefixes, we have that $\mathrm{Acc}_{\xi'}$ reduces to $\mathrm{Acc}_{\xi''}$. Meanwhile,

$$\begin{aligned}
\xi'' &= \mathsf{Ord}\big((c_1\rho^{N_1}\rho_1^n)_{n=0}^\infty, \dots, (c_d\rho^{N_d}\rho_d^n)_{n=0}^\infty\big) \\
&= \mathsf{Ord}\big((\rho_1^n)_{n=0}^\infty, (r_2\rho_2^n)_{n=0}^\infty, \dots, (r_d\rho_d^n)_{n=0}^\infty\big) \\
&= \mathsf{Ord}\big((n\log(\rho_1))_{n=0}^\infty, (\log(r_2) + n\log(\rho_2))_{n=0}^\infty, \dots, (\log(r_d) + n\log(\rho_d))_{n=0}^\infty\big) = \xi\,,
\end{aligned}$$

because dividing by $c_i > 0$, taking logarithms, and dividing by $\log(\rho_1)$ are all concave functions. The result follows. □

4.4.2 Effective almost periodicity of the order word

In this section, we use the notation from the earlier section. For $1 \leq i \leq d$, let $r_i, \rho_i \in \mathbb{R} \cap \overline{\mathbb{Q}}$ with $r_i > 0$ and $\rho_i > 1$. Suppose that for all $1 \leq i < j \leq d$ and $n, m \in \mathbb{N}$, we have $r_i\rho_i^n \neq r_j\rho_j^m$. Let $\xi := \mathsf{Ord}((r_1\rho_1^n)_{n=0}^\infty, \dots, (r_d\rho_d^n)_{n=0}^\infty) \in \{1, \dots, d\}^\omega$ as in Definition 4.4.4.

We prove the following.

Theorem 4.4.6.

(a) The word ξ is almost-periodic.

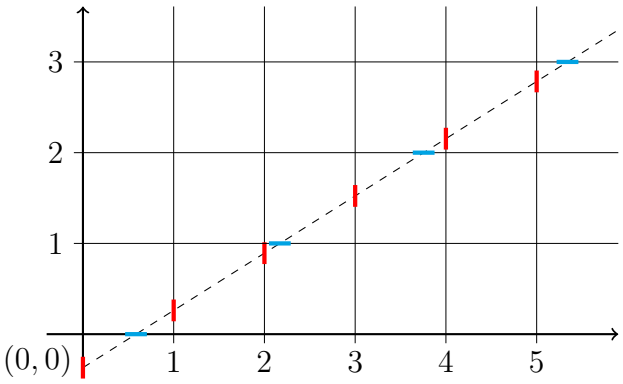

Figure 4.1: The cutting sequence ξ.

(b) Assuming the weak Schanuel conjecture, the word ξ is effectively almost-periodic.

(c) If $1/\log(\rho_1), \ldots, 1/\log(\rho_d)$ are linearly independent over $\mathbb{Q}$, then ξ is unconditionally effectively almost-periodic.

(d) If $\operatorname{rank}(G_M(\rho_1, \ldots, \rho_d)) \geq d - 2$, and $\rho_1, \ldots, \rho_d$ are pairwise multiplicatively independent, then ξ is unconditionally effectively almost-periodic.

This result, together with Theorem 4.4.5, will prove Theorem 4.4.1. By scaling and reordering the d sequences and Theorem 4.4.5, we can equivalently consider (4.11):

$$\xi := \mathsf{Ord}\left(\mathbb{N}, \left(\frac{\log(r_2)}{\log(\rho_1)} + n\frac{\log(\rho_2)}{\log(\rho_1)}\right)_{n=0}^{\infty}, \ldots, \left(\frac{\log(r_d)}{\log(\rho_1)} + n\frac{\log(\rho_d)}{\log(\rho_1)}\right)_{n=0}^{\infty}\right),$$

where $\rho_1 \leq \rho_2, \ldots, \rho_d$ and $0 < \log(r_i) < \log(\rho_i)$ for $2 \leq i \leq d$.

To prove Theorem 4.4.6, we will show that ξ is derived from a toric word; recall from Section 4.2 that toric words are almost-periodic. To prove *effective* almost-periodicity, we rely on Baker's theorem or Schanuel's conjecture.

Example 4.4.7. Suppose $r_1 = 2, \rho_1 = 2, r_2 = 3, \rho_2 = 3$. Then

$$2 \cdot 2^0 < 3 \cdot 3^0 < 2 \cdot 2^1 < 2 \cdot 2^2 < 3 \cdot 3^1 < 2 \cdot 2^3 < 3 \cdot 3^2 < 2 \cdot 2^4 < \cdots$$

and equivalently,

$$0 < \frac{\log(3/2)}{\log(2)} < 1 < 2 < \frac{\log(3/2)}{\log(2)} + \frac{\log(3)}{\log(2)} < 3 < \frac{\log(3/2)}{\log(2)} + 2\frac{\log(3)}{\log(2)} < 4 < \cdots$$

and hence $\xi = 12112121121\cdots$.

Now consider the line $\ell(t) = \{(t, \frac{\log(2)}{\log(3)}t - 1 + \frac{\log(2)}{\log(3)}) : t \in \mathbb{R}_{\geq 0}\}$ and the grid $\mathbb{N} \times \mathbb{N}$ as pictured in Figure 4.1. Then, we can also generate ξ as follows: start at $t = 0$ and

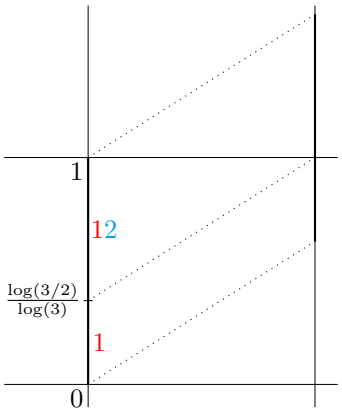

Figure 4.2: Generating ξ with a 1-dimensional system. Start at $y = 1 + \frac{\log(2/3)}{\log(3)}$ and at each step, maps y to $\{y + \frac{\log(2)}{\log(3)}\}$ and output 1 if $y \in (0, \frac{\log(3/2)}{\log(3)})$ and 12 if $y \in (\frac{\log(3/2)}{\log(3)}, 1)$.

follow the line in the positive direction. Then, every time one crosses a vertical line $y = n$ for some $n \in \mathbb{N}$, write a 1 and when crossing a horizontal line $x = n$ for some $n \in \mathbb{N}$, write a 2.

This construction obtains a *cutting sequence* (equivalently, a *billiard word*). (See for example [69, Chapter 4.1.2] and [9, 17] for more on billiard words.)

We aim to transform this 2-dimensional system into a 1-dimensional system. We achieve this by first noting that only the intersections of ℓ with the $\mathbb{N} \times \mathbb{N}$-grid matter and not which line is intersected, except for whether it is horizontal or vertical. If ℓ intersects a vertical line at height y, the next intersection of a vertical line occurs at height $y + \frac{\log(2)}{\log(3)}$. By translation invariance, only the fractional part of y, $\{y\}$, matters. Meanwhile, this fractional part of y also determines whether a horizontal line is intersected before another vertical line is intersected. If $\{y\} \in (0, 1 - \frac{\log(2)}{\log(3)})$, then no horizontal line is crossed, and when $\{y\} \in (1 - \frac{\log(2)}{\log(3)}, 1)$, exactly one horizontal line is crossed. Here, $1 - \frac{\log(2)}{\log(3)}$ can be simplified to $\frac{\log(3/2)}{\log(3)}$.

Therefore, we can generate ξ with the following system: Let $s = 1 + \frac{\log(2/3)}{\log(3)}$ and define $g : \mathbb{T} \to \mathbb{T}$ (recall that $\mathbb{T} = [0, 1)$ by $g(y) = \{y + \frac{\log(2)}{\log(3)}\}$). Further, let $\Sigma = \{1, 12\}$, $S_1 = (0, \frac{\log(3/2)}{\log(3)})$, $S_{12} = (\frac{\log(3/2)}{\log(3)}, 1)$, and $\mathcal{S} = \{S_1, S_{12}\}$. Thus, iterating the map g produces a toric word $\beta \in \Sigma^\omega$, which is an obvious morphism away from ξ. In our case, $\zeta = (12)(1)(12)(12)(1)(12)\cdots$. This construction is pictured in Figure 4.2.

This 1-dimensional method to generate ζ is sufficient to show that ζ is a toric word, and in particular, ζ is almost-periodic. But we can do better: ζ is effectively

almost-periodic. Let $w_1 \cdots w_r \in \Sigma^*$ and use linear programming on the torus $\mathbb{T}$ to either find an open interval $(a, b) \subset (0, 1)$ where for each $x \in (a, b)$ and $1 \leq i \leq r$,

$$g^{(i)}(s) \in S_b \iff w_i = b$$

or no such point exists. In the latter case, we are done. In the former case, we can use that $\frac{\log(2)}{\log(3)}$ and Theorem 4.2.6 to find a number t such that

$$\mathbb{T} = \bigcup_{i=0}^{t-1} g^{(i)}\big((a, b)\big)\,.$$

Of course, one has to be careful with these logarithms, but using Fourier-Motzkin elimination and Lemma 1.1.8, we can avoid such problems. Then, $w_1 \cdots w_r$ has a return time of $R := t + r$ in ζ (that is, $w_1 \cdots w_r$ appears in every factor of ζ of length R) and of $2R$ in ξ. We note that cutting sequences generated by a line on the plane with irrational slope (as in Figure 4.1) are exactly the Sturmian words [101]. □

We continue our proof of Theorem 4.4.6, and as in Example 4.4.7, we will begin by generating ξ with a cutting sequence.

Let $\ell : \mathbb{R}_{\geq 0} \to \mathbb{R}^d$ be a line such that for no $t \in \mathbb{R}_{>0}$, we have that two coordinates of $\ell(t)$ that are integer-valued. Then the *cutting sequence* of ℓ is the sequence $(x_n)_{n=0}^{\infty} \in \{1, \ldots, d\}^{\omega}$ generated by the positions of the integer-valued components of $\ell(t)$. That is, $x_n = i$ if and only if for the $n-1$th smallest $t \geq 0$ such that $\ell(t)$ has an integer component, this is the ith coordinate.

For $2 \leq i \leq d$, let $\delta_i = \frac{\log(\rho_1)}{\log(\rho_i)}$. Let $\ell : \mathbb{R}_{\geq 0} \to \mathbb{R}^d$ be the line

$$\ell(t) = \left(0, -\frac{\log(r_2)}{\log(\rho_1)}, \ldots, -\frac{\log(r_d)}{\log(\rho_1)}\right) + t\,(1, \delta_2, \ldots, \delta_d)$$

and consider the cutting sequence generated by ℓ. We have that $\ell(t)_i = n$ for some integer n if and only if $t = n$, and for $2 \leq i \leq d$, we have $\ell(t)_i = n$ for some integer n if and only if $t = \frac{\log(r_i)}{\log(\rho_1)} + n\frac{\log(\rho_i)}{\log(\rho_1)}$. As $\mathbb{N}$, $\left(\frac{\log(r_2)}{\log(\rho_1)} + n\delta_2\right)_{n=0}^{\infty}, \ldots, \left(\frac{\log(r_d)}{\log(\rho_1)} + n\delta_d\right)_{n=0}^{\infty}$ are pairwise disjoint by Theorem 4.4.5, the cutting sequence of ℓ indeed generates ξ.

Now we want to squish this d-dimensional method of generating ξ into a $d-1$-dimensional method. We achieve this by constructing a toric word defined on the torus $\mathbb{T}^{d-1}$ and finding an appropriate factorization ζ of ξ.

First, note that each $\delta_i = \frac{\log(\rho_1)}{\log(\rho_i)}$ is smaller than 1 and so each number $2 \leq i \leq d$ appears at once between two consecutive occurrences of 1 in ξ. Let

$$\Sigma = \big\{1w_1, \ldots, w_r \colon w_i \in \{2, \ldots, d\}, w_i \neq w_j \text{ for all } i \neq j\big\}\,.$$

Then, by the previous comment, we can factor ξ into a word $\zeta \in \Sigma^\omega$ such that there is an obvious morphism that maps ζ to ξ.

We formalize this construction. Let $(z_m)_{m=0}^\infty \in \mathbb{N}^\omega$ be the sequence such that $z_m = n$ if and only if ξ_n is the mth occurrence of 1 in ξ. For example, using the sequence in Example 4.4.7, we have $\xi = 12112121121\cdots$, and so $(z_m)_{m=0}^\infty = (0, 2, 3, 5, 7, 8, 10, \dots)$. Then, as $\xi_0 = 1$ and 1 appears infinitely often in ξ, we can factor ξ as

$$\zeta_m = \xi_{z_m}\xi_{z_m+1}\cdots\xi_{z_{m+1}-1} = 1\xi_{z_m+1}\cdots\xi_{z_{m+1}-1}\,.$$

Thus, by interpreting the finite words in this factorization as single letters, the word ζ is in Σ^ω and ξ is the image of ζ under the application of the morphism $\mu\colon \Sigma^* \to \{1, \dots, d\}^*$ defined by $\mu(w) = w_0\cdots w_{|w|-1}$ for $w \in \Sigma^*$. Since effectively almost-periodic words are closed under applications of morphisms (provided that the image word is also infinite) and finite modifications [118], the word ζ is (effectively) almost-periodic if and only if ξ is. We will next show that ζ is toric. Recall that $\{x\} := x - \lfloor x \rfloor$ denotes the fractional part of x.

Theorem 4.4.8. *Using the notation above let* $s_i = \left\{\frac{-\log(r_i)}{\log(\rho_i)}\right\} \in \mathbb{T}$ *for* $2 \le i \le d$. *Then* ζ *is the toric word generated by* $\boldsymbol{\delta} = (\delta_2, \dots, \delta_d) \in \mathbb{T}^{d-1}$ *and* $\mathbf{s} = (s_2, \dots, s_d) \in \mathbb{T}^{d-1}$, *as well as a collection of open subsets of* $\mathbb{T}^{d-1}$ *defined by linear inequalities (in variables* $x_2, \dots, x_d$*) of the form* $(1 - x_i)/\delta_i < (1 - x_j)/\delta_j$ *or* $(1 - x_i)/\delta_i \sim 1$, *where* $\sim \in \{>, <\}$ *and* $2 \le i, j \le d$ *are distinct.*

Proof. By the hypothesis, we have that $g : \mathbb{T}^d \to \mathbb{T}^d$ is the map defined by

$$g\big((x_2, \dots, x_d)\big) = \big(\{x_2 + \delta_2\}, \dots, \{x_d + \delta_d\}\big)\,.$$

By construction, we have that for all $n \in \mathbb{N}$ that $g^{(n)}(\mathbf{s}) = (\{s_2 + n\delta_2\}, \dots, \{s_d + n\delta_d\})$ and $\ell(n) = (n, s_2 + n\delta_2, \dots, s_d + n\delta_d)$, which are equal when restricted to the fractional part of the last $d - 1$ coordinates.

Assume that we want to compute $\zeta_n = \xi_{z_n}\cdots\xi_{z_{n+1}-1}$. As stated before, each $2 \le i \le d$ appears at most once in ζ_n. Let $\ell(n) = (n, x_2, \dots, x_d) \in \mathbb{R}^d$. By Theorem 4.4.5, we have that $\{x_i\} \neq 0$ for $2 \le i \le d$. Thus, $\{x_i\} \in (0, 1)$.

Then, tracing the direction of the line ℓ, the first coordinate of $\ell(t)$ is again integer-valued for $n + 1$. Then i is in ζ_n if and only if there is a $t_i \in (0, 1)$ such that $\{x_i + t_i\delta_i\} = 0$. Using that $\{x_i\}, \delta_i, t_i \in (0, 1)$ we have that $\{x_i + t_i\delta_i\} = 0$ if and only if $t_i = \{-x_i\}/\delta_i \in (0, 1)$. Note that as $\{x_i\} \in (0, 1)$, we have $\{-x_i\} = 1 - \{x_i\}$. As $\{-x_i\} > 0$ and $\delta_i > 0$, the inequality $t_i > 0$ is trivially satisfied, but we also need that $\{-x_i\}/\delta_i < 1$. Thus, $i \in \zeta_n$ if and only if $\{-x_i\}/\delta_i < 1$. Moreover, i appears

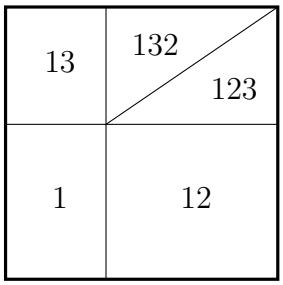

Figure 4.3: The torus for $\rho_1 = 2$, $\rho_2 = 3$, and $\rho_3 = 5$.

before $j \in \{2, \dots, d\} \setminus \{i\}$ in ζ_n if and only if $0 < t_i < t_j$, which occurs exactly when $\{-x_i\}/\delta_i < \{-x_j\}/\delta_j < 1$. The latter is equivalent to $(1-\{x_i\})/\delta_i < (1-\{x_j\})/\delta_j < 1$.

Now let $b = 1b_1 \cdots b_m \in \Sigma$, where $b_i \in \{2, \dots, d\}$ for all $1 \le i \le m$. Then, we set

$$S_b = \left\{ (x_2, \dots, x_d) \in \mathbb{T}^{d-1} \colon \frac{1-x_{b_1}}{\delta_{b_1}} < \cdots < \frac{1-x_{b_m}}{\delta_{b_m}} < 1 \wedge \bigwedge_{\substack{1 \le j \le d \\ j \ne b_1, \dots, b_m}} \frac{1-x_j}{\delta_j} > 1 \right\}. \tag{4.12}$$

These sets S_b are open and pairwise disjoint. Let these inequalities define S_b, which gives 2^{d-1} open and pairwise disjoint sets of $\mathbb{T}^{d-1}$.

By construction, we have that $\zeta_n = b$ if and only if $\ell(n) = (n, x_2, \dots, x_d)$ has the property that $(\{x_2\}, \dots, \{x_d\}) \in S_b$ while $(\{x_2\}, \dots, \{x_d\}) = g^{(n)}(\mathbf{s})$. Thus, ζ is a toric word defined by $\mathbf{s}$, $\boldsymbol{\delta}$, and the sets S_b. □

In the remainder of this section, let $\mathbf{s} = (s_2, \dots, s_d) \in \mathbb{T}^{d-1}$, $\boldsymbol{\delta} := (\delta_2, \dots, \delta_d) \in \mathbb{T}^{d-1}$, and $g \colon \mathbb{T}^{d-1} \to \mathbb{T}^{d-1}$ be defined as above. Figure 4.3 illustrates the target sets $\{S_b \colon b \in \Sigma\}$ constructed in Theorem 4.4.8 for the sequences $(2^n)_{n=1}^{\infty}$, $(3^n)_{n=1}^{\infty}$, and $(5^n)_{n=1}^{\infty}$. Figure 4.3 can also be viewed as follows. Consider a dynamical system on $\mathbb{N}$ that starts at 2 and jumps to the next power of 2 at each step. At each step, a letter from $\{1, 12, 13, 123, 132\}$ is written depending on whether the point jumped over a power of 3 or a power of 5 in the last step (and in what order). The exact letter to be written is determined by keeping track of the fractional part of $\log_3(2^n)$ and $\log_5(2^n)$; this gives rise to the linear inequalities defining the open sets depicted in Figure 4.3.

Since ζ is toric, ζ is uniformly recurrent (Section 4.2) and thus almost-periodic. As mentioned earlier, the suffix $w_N w_{N+1} \cdots$ of ξ is the image of ζ under a morphism, and hence is almost-periodic by [118, Section 3]. It follows that ξ is almost-periodic; **this proves Theorem 4.4.6 (a)**.

To prove the other parts, we need to show that ζ (and thus ξ) is effectively almost periodic. By Lemma 4.2.5, we only need to determine whether a given $w \in \Sigma^*$ occurs in ζ.

Lemma 4.4.9. *Let $w = w_0 \cdots w_{k-1} \in \Sigma^*$. There exists an open subset $S_w \subseteq \mathbb{T}^{d-1}$ with the following property. For all $n \in \mathbb{N}$, the pattern w occurs in ζ at the position n if and only if $g^{(n)}(\mathbf{s}) \in S_w$. Furthermore, we can compute a representation of S_w as a Boolean combination of inequalities of the form*

$$h(x_2/\delta_2, \dots, x_d/\delta_d, 1/\delta_2, \dots, 1/\delta_d) \sim 0\,, \tag{4.13}$$

where h is a $\mathbb{Q}$-affine form and $\sim$ is an inequality symbol.

Proof. Consider $b = 1b_1 \cdots b_m \in \Sigma$, where $b_i \in \{2, \dots, d\}$ for all i. Let

$$S_{b,k} = \{\mathbf{x} \in \mathbb{T}^{d-1} \colon g^{(k)}(\mathbf{x}) \in S_b\}\,.$$

Since $g\colon \mathbb{T}^{d-1} \to \mathbb{T}^{d-1}$ is a homeomorphism, the subset $S_{b,k}$ of $\mathbb{T}^{d-1}$ is open. Since $g^{(k)}(x_2, \dots, x_d) = (\{x_2 + k\delta_2\}, \dots, \{x_d + k\delta_d\})$, from (4.12) it follows that $S_{b,k}$ is the set of all $(x_2, \dots, x_d) \in \mathbb{T}^{d-1}$ satisfying

$$\frac{1 - \{x_{b_1} + k\delta_{b_1}\}}{\delta_{b_1}} < \cdots < \frac{1 - \{x_{b_m} + k\delta_{b_m}\}}{\delta_{b_m}} < 1 \wedge \bigwedge_{\substack{1 \le j \le d \\ j \neq b_1, \dots, b_m}} \frac{1 - \{x_j + k\delta_j\}}{\delta_j} > 1\,.$$

For $2 \le i \le d$, let $t_i = \lfloor k\delta_i \rfloor$. Observe that

$$1 - \{x_{b_1} + k\delta_{b_1}\} = \begin{cases} t_i + 1 - x_i - k\delta_i & \text{if } x_i + k\delta_i < t_i + 1 \\ t_i + 2 - x_i - k\delta_i & \text{if } x_i + k\delta_i \ge t_i + 1\,. \end{cases}$$

Moreover, $x_i + k\delta_i < t_i + 1$ is equivalent to $x_i/\delta_i + k < \frac{t_i+1}{\delta_i}$. Therefore, $\frac{1-\{x_i+k\delta_i\}}{\delta_i} \bowtie 1$ (where $\bowtie$ is an (in)equality symbol) is equivalent to

$$\frac{x_i}{\delta_i} + k < \frac{t_i + 1}{\delta_i} \Rightarrow \frac{t_i + 1 - x_i - k\delta_i}{\delta_i} \bowtie 1 \wedge \frac{x_i}{\delta_i} + k \ge \frac{t_i + 1}{\delta_i} \Rightarrow \frac{t_i + 2 - x_i - k\delta_i}{\delta_i} \bowtie 1\,.$$

Rearranging, this formula can be written as a Boolean combination of inequalities of the form (4.13) where h is a $\mathbb{Q}$-affine form. Similarly, $\frac{\{x_i+k\delta_i\}}{\delta_i} < \frac{\{x_j+k\delta_j\}}{\delta_j}$ can be equivalently written as a Boolean combination of inequalities of the form (4.13) by conditioning on whether $x_i/\delta_i + k < (t_i + 1)/\delta_i$ and $x_j/\delta_j + k < (t_j + 1)/\delta_j$. Finally, observe that $0 \le x_i < 1$ is equivalent to $0 \le x_i\delta_i < 1/\delta_i$. We conclude that $S_{b,k}$ can be defined by a Boolean combination of inequalities of the form (4.13).

It remains to define $S_w = \bigcap_{k=0}^{|w|-1} S_{w(k),k}$. Since each $S_{w(k),k}$ is open and defined by a Boolean combination of inequalities of the form (4.13), the same holds for S_w. □

We next prove Theorem 4.4.6 (c). Suppose $1/\log(\rho_1), \ldots, 1/\log(\rho_d)$ are linearly independent over $\mathbb{Q}$. Then for all $c_2, \ldots, c_d, k \in \mathbb{Q}$,

$$\sum_{i=2}^{d} c_i \delta_i = k \iff \frac{c_2}{\log(\rho_2)} + \cdots + \frac{c_d}{\log(\rho_d)} = \frac{k}{\log(\rho_1)}$$
$$\implies k, c_2, \ldots, c_d = 0\,.$$

Hence $G_A(\boldsymbol{\delta})$ is the trivial group, we have $\mathbb{T}_{\boldsymbol{\delta}} = \mathbb{T}^{d-1}$, and by Kronecker's theorem (Theorem 4.2.6), the set $(g^{(n)}(\mathbf{s}))_{n=0}^{\infty}$ is dense in $\mathbb{T}^{d-1}$. Therefore, a pattern $w \in \Sigma^*$ occurs in ζ if and only if $S_w \neq \emptyset$. As shown in Lemma 4.4.9, to determine whether $S_w \neq \emptyset$ we have to determine the truth of $\Psi := \exists x_2, \ldots, x_d \colon \Phi(x_2, \ldots, x_d)$, where Φ is a formula of the form

$$\bigvee_{j \in J} \bigwedge_{k \in K} h_{j,k}(x_2/\delta_2, \ldots, x_d/\delta_d, 1/\delta_2, \ldots, 1/\delta_d) \sim_{j,k} 0\,.$$

As each $h_{j,k}$ is a $\mathbb{Q}$-affine form, Ψ is equivalent to $\exists \widetilde{x}_2, \ldots, \widetilde{x}_d \colon \Phi(\widetilde{x}_2, \ldots, \widetilde{x}_d)$, where $\widetilde{x}_i = \delta_i x_i$. Applying Fourier-Motzkin elimination, we can compute finitely many $\mathbb{Q}$-affine forms $h_{\ell,m}$ and an inequality symbols $\sim_{\ell,m}$ such that Ψ is true if and only if

$$\bigvee_{\ell \in L} \bigwedge_{m \in M} h_{\ell,m}(1/\delta_2, \ldots, 1/\delta_d) \sim_{\ell,m} 0\,.$$

Recall that $\frac{1}{\delta_i} = \frac{\log(\rho_i)}{\log(\rho_1)}$. For $h(x_2, \ldots, x_d) := c_1 + c_2 x_2 + \cdots + c_d x_d$,

$$h(1/\delta_2, \ldots, 1/\delta_d) = \frac{1}{\log(\rho_1)} \big(c_1 \log(\rho_1) + \cdots + c_d \log(\rho_d) \big)\,.$$

Hence for all ℓ, m, whether $h_{\ell,m}(1/\delta_2, \ldots, 1/\delta_d) \sim_{\ell,m} 0$ can be determined using Baker's theorem (Lemma 1.1.8). Thus, when $1/\log(\rho_1), \ldots, 1/\log(\rho_d)$ are linearly independent over $\mathbb{Q}$, we can determine whether a given pattern $w \in \Sigma^*$ occurs in ζ. We conclude that ζ and hence ξ are effectively almost periodic. **This proves Theorem 4.4.6 (c).**

To prove Theorem 4.4.6 (b) and (d), we need to prove a small lemma first.

Lemma 4.4.10. *Let $d \geq 1, e \geq 2$, $c_1, \ldots, c_d \in \mathbb{R}$, and for all $1 \leq j \leq d$, $\mathbf{b}_j = (b_{1,j}, \ldots, b_{d,j}) \in \mathbb{R}^d$ are pairwise linearly independent vectors over the real numbers. Assume that f is the rational function given by*

$$f(x_1, \ldots, x_e) = \frac{c_1}{\sum_{i=1}^{e} b_{i,1} x_i} + \cdots + \frac{c_d}{\sum_{i=1}^{e} b_{i,d} x_i}$$

If f is the zero function, then $c_1 = \cdots = c_d = 0$.

Proof. Assume that f is the zero function. For $1 \leq j \leq d$, let $V_j = \{\mathbf{x} \in \mathbb{R}^e \colon \mathbf{b}_j \cdot \mathbf{x} = 0\}$. Then V_j is an $(e-1)$-dimensional subspace of $\mathbb{R}^e$. Fix some $1 \leq j \leq d$ such that $c_j \neq 0$ and let $(\mathbf{x}_n)_{n=0}^{\infty} \in (\mathbb{R}^d \setminus V_j)^{\omega}$ be a sequence converging to $\mathbf{x} \in V_j$, then $\lim_{n\to\infty} c_j / \sum_{i=1}^{e} (\mathbf{b}_j \cdot \mathbf{x}_n) = \pm\infty$. Hence, as f is the zero function, the vector $\mathbf{x}$ is in $V_{j'}$ for some $j' \in \{1, \ldots, d\} \setminus \{j\}$. Thus,

$$V_j \subset \bigcup_{j' \in \{1,\ldots,d\}\setminus\{j\}} V_{j'} \tag{4.14}$$

As all $\mathbf{b}_j$ and $\mathbf{b}_{j'}$ are linearly independent, the intersection $V_j \cap V_{j'}$ is an $(e-2)$-dimensional linear subspace of $\mathbb{R}^e$. This contradicts (4.14).

Thus, $c_j = 0$. As this holds for all $1 \leq j \leq d$, the lemma follows. □

Now Theorem 4.4.6 (d) follows from combining Theorem 4.4.6 (c) and the following lemma.

Lemma 4.4.11. *Let $d \geq 2$ and $\lambda_1, \ldots, \lambda_d \in \mathbb{R}_{>1} \cap \overline{\mathbb{Q}}$ be pairwise multiplicatively independent, and suppose* $\operatorname{rank}(G_M(\lambda_1, \ldots, \lambda_d)) \geq d-2$. *Then $1/\log(\lambda_1), \ldots, 1/\log(\lambda_d)$ are linearly independent over $\mathbb{Q}$.*

Proof. By the two assumptions, for any $1 \leq i < j < k \leq d$, there exist $b_i, b_j, b_k \in \mathbb{Z}_{\neq 0}$ such that $\lambda_i^{b_i} \lambda_j^{b_j} \lambda_k^{b_k} = 1$. Equivalently, $b_i \log(\lambda_i) + b_j \log(\lambda_j) + b_k \log(\lambda_k) = 0$. Hence, for $1 \leq j \leq d$, let $b_{1,j}, b_{2,j} \in \mathbb{Q}$ be such that $\log(\lambda_j) = b_{1,j} \log(\lambda_1) + b_{2,j} \log(\lambda_2)$. Then $b_{1,1} = b_{2,2} = 0$ and all other $b_{i,j}$ are non-zero.

To get a contradiction, let $c_1, \ldots, c_d \in \mathbb{Q}$ such that $\sum_{j=1}^{d} \frac{c_j}{\log(\lambda_j)} = 0$. Multiplying by $\prod_{j=1}^{d} \log(\lambda_j)$ gives

$$\sum_{j=1}^{d} c_j \prod_{i=1, i\neq j}^{d} \big(b_{1,j} \log(\lambda_1) + b_{2,j} \log(\lambda_2)\big) = 0\,,$$

which simplifies to

$$\sum_{i=0}^{d-1} e_i \log(\lambda_1)^i \log(\lambda_2)^{d-i} = 0$$

for some $e_i \in \mathbb{Q}$. Assume not all e_i are zero. Then dividing by $\log(\lambda_2)^{d-1}$ shows that $\log(\lambda_1)/\log(\lambda_2)$ is a root of the non-zero polynomial $\sum_{i=0}^{d-1} e_i X^i \in \mathbb{Q}[X]$. That is, $\log(\lambda_1)/\log(\lambda_2)$ is an algebraic number, say α, and so $\log(\lambda_1) - \alpha \log(\lambda_2) = 0$, contradicting Baker's theorem (Theorem 1.1.7). Hence, all e_i are zero.

Therefore, computing in $\mathbb{Q}(X_1, X_2)$

$$\sum_{j=1}^{d} \frac{c_j}{b_{1,j} X_1 + b_{2,j} x X_2} = 0\,. \tag{4.15}$$

As in Lemma 4.4.12, we deduce that as all λ_j are pairwise multiplicatively dependent and so (4.15) satisfies the hypothesis of Lemma 4.4.10. Thus, all c_i are zero. The statement follows. □

It remains to prove Theorem 4.4.6 (b), which is the trickiest. We will need the following two lemmas. The proof of the first is similar to the previous lemma.

Lemma 4.4.12. *If $\lambda_1, \ldots, \lambda_d \in \mathbb{R}_{>1} \cap \overline{\mathbb{Q}}$ are pairwise multiplicatively independent, then assuming the weak Schanuel conjecture, the numbers $1/\log(\lambda_1), \ldots, 1/\log(\lambda_d)$ are linearly independent over $\mathbb{Q}$.*

Proof. First, using Theorem 1.1.3, compute a basis for $G_M(\lambda_1, \ldots, \lambda_d)$ and select a maximum multiplicative independent subset $\{\lambda_1, \ldots, \lambda_e\}$, possibly renumbering the λ_i. For $1 \leq j \leq d$ and $1 \leq i \leq e$, let $b_{i,j}$ such that $\log(\lambda_j) = \sum_{i=1}^{e} b_{i,j} \log(\lambda_i)$. Then we have to show that the weak Schanuel conjecture implies that for $c_1, \ldots, c_d \in \mathbb{Q}$ such that

$$\frac{c_1}{\sum_{i=1}^{e} b_{i,1} \log(\lambda_i)} + \cdots + \frac{c_d}{\sum_{i=1}^{e} b_{i,d} \log(\lambda_i)} = 0\,, \tag{4.16}$$

all c_i are zero. By multiplying (4.16) by $\prod_{j=1}^{d} \sum_{i=1}^{e} b_{i,j} \log(\lambda_i)$, we obtain a polynomial $g \in \mathbb{Q}[X_1, \ldots, X_e]$ such that $g(\log(\lambda_1), \ldots, \log(\lambda_e)) = 0$.

We apply the weak Schanuel conjecture (Conjecture 1.1.5) with $\alpha_i = \log(\lambda_i)$ for $1 \leq i \leq e$. As the algebraic numbers $\lambda_1, \ldots, \lambda_e$ are multiplicatively independent, their logarithms $\log(\lambda_1), \ldots, \log(\lambda_e)$ are linearly independent over $\mathbb{Q}$. Then the weak Schanuel conjecture implies that $\log(\lambda_1), \ldots, \log(\lambda_e)$ are algebraically independent, and so $g(X_1, \ldots, X_e) = 0$. That is, the following rational function is exactly zero:

$$f(X_1, \ldots, X_e) = \frac{c_1}{\sum_{i=1}^{e} b_{i,1} X_i} + \cdots + \frac{c_d}{\sum_{i=1}^{e} b_{i,d} X_i}\,.$$

Assume that $1 \leq j, j' < d$ are distinct and that $(b_{1,j}, \ldots, b_{e,j}) = s(b_{1,j'}, \ldots, b_{e,j'})$ for some $s \in \mathbb{R}$. As $\lambda_{j'} \neq 0$, some $b_{i,j'}$ is non-zero. Then $s = b_{i,j}/b_{i,j'} \in \mathbb{Q}$. Hence,

$$\log(\lambda_j) = \sum_{i=1}^{e} b_{i,j} \log(\lambda_i) = \sum_{i=1}^{e} s b_{i,j'} \log(\lambda_i) = s \log(\lambda_{j'})\,,$$

contradicting that λ_j and $\lambda_{j'}$ are multiplicatively independent. Thus, the hypothesis of Lemma 4.4.10 is satisfied, and all c_j are 0. We conclude the statement. □

Lemma 4.4.13. *Let $\lambda_1, \ldots, \lambda_d \in \mathbb{R}_{>1} \cap \overline{\mathbb{Q}}$. Assuming the weak Schanuel conjecture, a basis for $G_A(1/\log(\lambda_1), \ldots, 1/\log(\lambda_d))$ can be computed.*

Proof. If λ_i and λ_j are multiplicatively dependent, say $\lambda_i^a = \lambda_j^b$ for some non-zero integers a and b, then $a\log(\lambda_i) = b\log(\lambda_j)$. Hence, $a/\log(\lambda_j) = b/\log(\lambda_i)$, giving a non-trivial element in $G_A(1/\log(\lambda_1), \ldots, 1/\log(\lambda_d))$.

Meanwhile, by Lemma 4.4.12, for any pairwise multiplicative independent subset of $\{\lambda_1, \ldots, \lambda_d\}$, the reciprocals of their logs are linearly independent. Together, this implies that any $\mathbb{Q}$-linear relationship among $1/\log(\lambda_1), \ldots, 1/\log(\lambda_d)$ can be reduced to a relationship generated by pairwise multiplicative relationships. One can easily find a basis among these by computing for each pair (λ_i, λ_j) whether they are multiplicatively dependent using Theorem 1.1.3. □

Assuming the weak Schanuel conjecture, we can compute a basis of the group of additive relations $G_A(1/\log(\rho_1), \ldots, 1/\log(\rho_d))$ (Lemma 4.4.13). Hence we can compute an $\mathbb{R}_{\exp}$ formula defining the compact $\mathbb{T}_{\boldsymbol{\delta},\mathbf{s}} \subseteq \mathbb{T}^{d-1}$ in which in $(f^{(n)}(\mathbf{s}))_{n=0}^{\infty}$ is dense (Section 4.2). Recall from Section 4.2 that a pattern w occurs in ζ if and only if $S_w \cap \mathbb{T}_{\boldsymbol{\delta},\mathbf{s}} \neq \emptyset$, which can be effectively verified using a decision procedure for the first-order theory of $\mathbb{R}_{\exp}$. Hence, ζ and ξ are effectively almost periodic when assuming the weak Schanuel conjecture. **This proves Theorem 4.4.6 (b) and therefore the entire theorem.**

We can now combine everything we have shown so far to prove Theorem 4.4.1. For $1 \leq i \leq d$, let $(u_n^{(i)})_{n=0}^{\infty}$ for $1 \leq i \leq d$ be as in the statement of Theorem 4.4.1 with the value set $P_i \subseteq \mathbb{N}$. Further let α be the characteristic word of $(P_1, \ldots, P_d)$, and recall that $\mathsf{MSO}_{\mathbb{N};<}(P_1, \ldots, P_d)$ is decidable if and only if Acc_α is decidable. Applying Theorem 4.4.5, we can construct $r_1, \ldots, r_d$ such that Acc_α reduces to Acc_ξ, where $\xi = \mathsf{Ord}((r_1\rho_1^n)_{n=0}^{\infty}, \ldots, (r_d\rho_d^n)_{n=0}^{\infty})$. Applying Theorem 4.4.6, we obtain conditions under which ξ is effectively almost-periodic. It remains to recall from Theorem 4.2.4 that Acc_ξ is decidable if ξ is effectively almost-periodic.

4.4.2.1 Applying the Theory of Cutting Sequences

Let

$$\xi = \mathsf{Ord}\big((r_1\rho_1^n)_{n=0}^{\infty}, \ldots, (r_d\rho_d^n)_{n=0}^{\infty}\big)$$

be as above. As mentioned earlier, it can directly be shown that ξ is the cutting sequence generated by the line $\{(t/\log(\rho_1), s_2 + t/\log(\rho_2), \ldots, s_d + t\delta_d) : t \geq 0\}$, where $s_2 \in \mathbb{R}$ for all $2 \leq i \leq d$. As in Section 1.3.1, write $p(n)$ for the number of distinct factors of ξ of length n; the function p is the *factor complexity* of ξ. The factor complexity of cutting sequences has been studied extensively, and in many cases, an

exact formula for $p(n)$ is known. We give an overview of the known results in this direction.

(i) If $d = 2$ and $\log(\rho_1)/\log(\rho_2)$ is irrational, then ξ is a Sturmian word and therefore $p(n) = n + 1$. See, e.g. [6, Chapter 10.5].

(ii) By the results of Arnoux et al. [9], if $d = 3$, and $1/\log(\rho_1), 1/\log(\rho_2), 1/\log(\rho_3)$ as well as $\log(\rho_1), \log(\rho_2), \log(\rho_3)$ are linearly independent over $\mathbb{Q}$, then $p(n) = n^2 + n + 1$.

(iii) For arbitrary $d > 0$, Bedaride [19] gives an exact formula for $p(n)$ assuming $1/\log(\rho_1), \ldots, 1/\log(\rho_d)$ as well as every triple $\log(\rho_i), \log(\rho_j), \log(\rho_k)$ for pairwise distinct i, j, k are linearly independent over $\mathbb{Q}$. This generalizes the result [17] of Baryshnikov.

Returning to our word ξ, let w be a finite pattern of length n, and suppose we know the value of $p(n)$. Then we can determine whether w occurs (as required by Lemma 4.2.5) in ξ by just reading prefixes of ξ until we have seen $p(n)$ distinct factors of length n. Using this approach, we can prove that the word ξ is effectively almost periodic under the assumption of (iii). Note, however, that this result is strictly weaker than Theorem 4.4.6 (b). Consider, for example, $\rho_1 = 2$, $\rho_2 = 3$, and $\rho_3 = 6$. We will show in Lemma 4.4.11 that $1/\log(\rho_1), \ldots, 1/\log(\rho_3)$ are linearly independent over $\mathbb{Q}$, but $\log(\rho_1), \ldots, \log(\rho_3)$ are not.

4.5 Decidability via expansions in integer bases

In this section, we discuss a second class of LRS that gives rise to interesting MSO theories; we show that these are intimately connected to base-b expansions of certain algebraic numbers. That is, we will show that the base-b expansion of $\sqrt[d]{p/q}$ is intrinsic to the pair of predicates $\{qn^d : n \in \mathbb{N}\}$ and $\{pb^{nd} : n \in \mathbb{N}\}$. For example, the binary expansion of $\sqrt[3]{1/27} = 2/3$ underlies the pair of predicates $\{27n^3 : n \in \mathbb{N}\}$ and $\{8^n : n \in \mathbb{N}\}$, while the binary expansion of $2\sqrt[3]{5}$ underlies the pair $\{n^3 : n \in \mathbb{N}\}$ and $\{5 \cdot 8^n : n \in \mathbb{N}\}$. Assuming that certain irrational numbers are *normal*, we can use these connections to give conditional decidability results for various MSO theories. The dynamical systems at play in this section differ from the ones we considered previously: they are defined by *numeration systems* [101, Chapter 7] as opposed to translations on a torus (Section 4.4.2 and 4.2).

Reductions and decidability

We begin by considering the case where $\sqrt[d]{p/q}$ is rational, which implies that its base-b expansion is ultimately periodic for any $b \geq 2$. The following is a generalisation of Corollary 4.1.3.

Theorem 4.5.1. *Let $b, d \geq 2$ and $p, q \geq 1$ be integers such that $\sqrt[d]{p/q}$ is rational. Let $P_1 = \{qn^d \colon n \in \mathbb{N}\}$, $P_2 = \{pb^{nd} \colon n \in \mathbb{N}\}$, and $\alpha \in (\{0,1\}^2)^\omega$ be the characteristic word of (P_1, P_2). Then Acc_α is decidable.*

Proof. By assumption, $p/q = A^d/B^d$ for some $A, B \in \mathbb{N}$, and so there is also a $C \in \mathbb{N}$ such that $p = CA^d$ and $q = CB^d$. We first solve the case where $A = B = C = 1$. Then, $P_1' = \{n^d \colon n \in \mathbb{N}\}$ and $P_2' = \{b^{nd} \colon n \in \mathbb{N}\}$ and so $P_2' \subset P_1'$. As P_1' is effectively sparse and $P_2' \subset P_1'$, the predicates P_1' and P_2' are pairwise effectively sparse. As P_1' and P_2' are also infinite, recursive, and effectively procyclic due to Lemma 4.2.2, Theorem 4.1.10 implies that $\mathrm{Acc}_{\alpha'}$ reduces to $\mathrm{Acc}_{\beta'}$, where β' is the order word of (P_1', P_2').

Let α' be the characteristic word of (P_1', P_2'). As $P_2' \subset P_1'$, we have $\beta'_m = \alpha'_{p_m} = \alpha'_{m^d}$, and thus $\beta'_m = (1,1)$ if m^d is a power of b^d and $(1,0)$ otherwise. Thus, β' is the characteristic word of $P' = \{b^n \colon n \in \mathbb{N}\}$ with the alphabet $\{(1,0),(1,1)\}$ instead of $(0,1)$. As P' is infinite, recursive, effectively pairwise sparse, and effectively procyclic, the problem $\mathrm{Acc}_{\beta'}$ is decidable again by Theorem 4.1.10. Thus, $\mathrm{Acc}_{\alpha'}$ is decidable as well.

We now construct a transducer $\mathcal{B}$ such that $\mathcal{B}(\alpha') = \alpha$, which due to Theorem 4.3.1 is sufficient to decide Acc_α. We do this in two steps. First, we keep track how many r_i occurrences of $(1,*)$, we have seen modulo A and change $(1,*)$ into $(0,*)$ when $r_i \not\equiv 0 \pmod{A}$. Similarly, we treat occurrences of $(*,1)$ modulo B. Making just this transduction gives the characteristic word of the predicates $\{A^d n^d \colon n \in \mathbb{N}\}$ and $\{B^d b^{nd} \colon n \in \mathbb{N}\}$. Now we add the further transduction that maps (b_1, b_2) to $(b_1, b_2)(0,0)^{C-1}$. If we combine these two steps into a single transducer $\mathcal{B}$, i.e., $\mathcal{B}(\alpha') = \alpha$, we obtain the result. □

Proof of Corollary 4.1.3. Apply the previous theorem with $p = q = 1$. □

The case where $\sqrt[d]{p/q}$ is irrational is more involved. We are not able to prove decidability, but we can show that the decidability of the MSO theory of $\langle \mathbb{N}; <, P_1, P_2 \rangle$ to Acc_β, where β is the base-b expansion of some number.

Theorem 4.5.2. *Let b, d, p, q be positive integers such that $\sqrt[d]{p/q}$ is irrational. Furthermore, let*

1. $\alpha \in (\{0,1\}^2)^\omega$ *be the characteristic word of* (P_1, P_2)*, where* $P_1 = \{qn^d \colon n \in \mathbb{N}\}$ *and* $P_2 = \{pb^{nd} \colon n \in \mathbb{N}\}$*;*

2. $\beta \in \{0, 1, \ldots, b-1\}^\omega$ *be the infinite string of digits in the base-*b *expansion of* $\eta = \sqrt[d]{p/q}$ *and*

3. γ *be the order word* (P_1, P_2)*, i.e., the word obtained by deleting all occurrences of* $(0,0)$ *from* α*.*

Then the problems Acc_α*,* Acc_β*, and* Acc_γ *are Turing-equivalent.*

Proof. We will prove the theorem by showing:

1. Acc_β reduces to Acc_α.

2. Acc_α reduces to Acc_γ.

3. Acc_γ reduces to Acc_β.

We start by breaking down the definition of β. As $\eta = \cdots \beta_{-1}\beta_0 \,.\, \beta_1\beta_2 \cdots$, we have that ηb^m is equal to $\cdots \beta_{m-1}\beta_m \,.\, \beta_{m+1}\beta_{m+2} \cdots$. That is, we shifted the divider between the integer and non-integer parts in the base-b expansion of η by m. Therefore, $\beta_m = \lfloor \eta b^m \rfloor \bmod b$, the mth digit in the base-b expansion of η. Moreover, we obtain the recursion that for all $m \geq 0$,

$$\lfloor \eta b^{m+1} \rfloor = \lfloor \eta b^m \rfloor b + \beta_{m+1} \,. \tag{4.17}$$

Also note that $(1,1)$ does not occur in α. Otherwise, there would be natural numbers m and n such that $qm^d = pb^{nd}$, which is equivalent to $m = \eta b^n$. As m and b^n are rational and η is irrational, this is impossible.

Part (1): Acc_β **reduces to** Acc_α**.** By construction and the observations above,

$$\begin{aligned} \beta_n = \lfloor \eta b^n \rfloor \bmod b &= \#\{m \in \mathbb{N}_{\geq 1} \colon m < \eta b^n\} \bmod b \\ &= \left(\#\{m \in \mathbb{N} \colon qm^d < pb^{dn}\} - 1\right) \bmod b \end{aligned}$$

Thus, we construct a transducer $\mathcal{B}$ such that $\mathcal{B}(\alpha) = \beta$ by keeping track of the number of letters $(1,0)$ modulo b (which correspond to occurrences qm^d in α), and outputting this number minus 1 each time $\mathcal{B}$ reads $(0,1)$ (which corresponds to occurrences of pb^n in α). More explicitly, $\mathcal{B}$ has the alphabet $\{0,1\}^2$, states $\{0, \ldots, b-1\}$, and

- when reading $(0,0)$ or $(1,1)$ in state q: stay in state q and output nothing;

- when reading $(1,0)$ in state q: move to state $(q+1) \bmod b$ and output nothing;
- when reading $(0,1)$ in state q: stay in state q and output $(q-1) \bmod b$.

Here, we added $(1,1)$ for completeness, but due to the observations above, it will not occur in α. Then, by construction, we indeed have that $\mathcal{B}(\alpha) = \beta$, and so Theorem 4.3.1 completes the proof in this case.

Part (2): Acc_α reduces to Acc_γ. Theorem 4.1.10 almost directly implied this reduction. As P_1 and P_2 are infinite, recursive, effectively sparse and effectively procyclic, we only have to verify that P_1 and P_2 are effectively pairwise sparse to apply Theorem 4.1.10. We do so by applying an old result of Schinzel and Tijdeman, whose proof (again) relies on Baker's theorem on linear forms in logarithms [137].

Lemma 4.5.3 (Schinzel and Tijdeman)**.** *For every $N \geq 1$, the equation $|qn^d - pb^{md}| = N$ has finitely many solutions (n,m) and they can be effectively enumerated.*

Lemma 4.5.3 implies that P_1 and P_2 are effectively pairwise sparse. Then the conditions of Theorem 4.1.10 are satisfied. This case follows.

Part (3): Acc_γ reduces to Acc_β. As $(1,1)$ does not occur in γ, we can write γ as $\gamma = (1,0)^{m_0}(0,1)(1,0)^{m_1}(0,1)\cdots$, where $\sum_{i=0}^{n} m_i$ is the number of terms qm^d below pb^{nd}. That is, $\sum_{i=0}^{n} m_i = \lfloor \eta b^m \rfloor + 1$. Now let P be the predicate consisting of all the indices of $(0,1)$ in γ and $(p_m)_{m=0}^{\infty}$ the positive value sequence of P. Then $\mathrm{Char}(P)$ is γ with the alphabet $\{0,1\}$ instead of $\{(1,0),(0,1)\}$ and

$$p_m = m + \sum_{i=0}^{m} m_i = m + \lfloor \eta b^m \rfloor + 1\,. \tag{4.18}$$

Now, for $M \geq 1$, we have to show that $\mathrm{Acc}_{(p_m \bmod M)_{m=0}^{\infty}}$ is decidable. So fix $M \geq 1$. We claim that $p_m \bmod M$, $m \bmod M$, and β_{m+1} uniquely determine $p_{m+1} \bmod M$. Indeed, due to the first two and (4.17), uniquely determines $\lfloor \eta b^m \rfloor \bmod M$. Then, by (4.17), the number $\lfloor \eta b^{m+1} \rfloor \bmod M$ is also uniquely determined. Then, again using (4.18) and having access to $m \bmod M$ (and thus $m+1 \bmod M$), there is only one option for $p_{m+1} \bmod M$. This proves our claim. Also, note that p_0 can be computed. Thus, we can construct a finite deterministic transducer that reads β with input $(p_m \bmod M, m \bmod M)$ that outputs $(p_{m+1} \bmod M, m \bmod M)$. Thus, using Theorem 4.3.1, the problem $\mathrm{Acc}_{(p_m \bmod M)_{m=0}^{\infty}}$ reduces to Acc_β. □

Together, Lemma 4.5.1 and Theorem 4.5.2 imply Theorem 4.1.3. Conjecture 1.3.1 implies that when $\sqrt[d]{p/q}$ is irrational, its base-b expansion is disjunctive. Hence, Theorem 4.5.2, implies that $\mathsf{MSO}_{\mathbb{N};<}(\{qn^d \colon n \in \mathbb{N}\}, \{pb^{nd} \colon n \in \mathbb{N}\})$ is decidable if we

assume this conjecture. By mirroring the proof of Theorem 4.5.2 in the special case where $p = b, q = 1, d = 2$, we can show the following.

Theorem 4.5.4. *Let $b \geq 2$, $P_1 = \mathsf{N}_2$, $P_2 = b^{\mathbb{N}}$, and $\Sigma = \{0,1\}^2$. Further let $\alpha \in \Sigma^\omega$ denote the characteristic word of (P_1, P_2) and $\beta \in \{0, \ldots, b-1\}^\omega$ be the base-b expansion of $\sqrt{b}$. The problems Acc_α and Acc_β are Turing-equivalent.*

Proof. The proof mirrors that of Theorem 4.5.2 closely. To show that Acc_β reduces to Acc_α, let $\widetilde{\alpha}$ be the characteristic word of $\{n^2 \colon n \in \mathbb{N}\}$ and $\{b \cdot b^{2n} \colon n \in \mathbb{N}\}$. Then the transducer $\mathcal{B}$ that changes every $(1,1)$ into $(1,0)$ (and leaves everything else unchanged) has the property that $\mathcal{B}(\alpha) = \widetilde{\alpha}$. This follows from the fact that a power of b is not a square if and only if it is of the form $b \cdot b^{2n}$. By Lemma 4.3.1, the problem $\mathrm{Acc}_{\widetilde{\alpha}}$ reduces to Acc_α. As $\mathrm{Acc}_{\widetilde{\alpha}}$ and Acc_β are Turing-equivalent, one direction follows.

To show the other direction, let $\gamma \in \{(0,1),(1,0),(1,1)\}^\omega$ be the order word of α. By the same reasoning as in the proof of Theorem 4.5.2, we have that P_1 and P_2 are infinite, recursive, effectively sparse, and effectively procyclic, and pairwise effectively sparse where the latter is due to Schinzel's and Tijdeman's Lemma 4.5.3. Thus, P_1 and P_2 satisfy the requirements of Theorem 4.1.10 and thus Acc_α reduces to Acc_γ.

In γ, the letters $(1,1)$ and $(0,1)$ correspond to b^{2n} and $b \cdot b^{2n}$ respectively. Thus, the terms $(1,1)$ and $(0,1)$ alternate in γ, starting with $(1,1)$. We factor γ as follows:

$$\gamma = (1,0)^{r_0}(1,1)(1,0)^{s_0}(0,1)(1,0)^{r_1}(1,1)(1,0)^{s_1}(0,1)\cdots .$$

Let P' be the set of indices of $(1,1)$ and $(0,1)$, α' be the characteristic word of P', and $(p_m)_{m=0}^\infty$ the positive value sequence of P', Then, as $(1,1)$ and $(0,1)$ alternate in γ, one can construct a transducer $\mathcal{B}$ such that $\mathcal{B}(\alpha') = \gamma$. So using Theorem 4.3.1, the problem Acc_γ reduces to $\mathrm{Acc}_{\alpha'}$.

Between b^{2n} and $b^{2(n+1)}$, there are $b^{n+1} - b^n - 1 = s_n + r_{n+1}$ squares while between $b \cdot b^{2n}$ and $b \cdot b^{2(n+1)}$ there are $\lfloor \sqrt{b} b^{n+1} \rfloor - \lfloor \sqrt{b} b^n \rfloor = r_{n+1} + s_{n+1} + 1$ squares. For all $m \geq 0$, we have that $p_{2m} = r_m$ and $p_{2m+1} = s_m$. Moreover, $r_0 = 1$. By Theorem 4.1.9, to decide $\mathrm{Acc}_{\alpha'}$ we only have to be able to decide $\mathrm{Acc}_{(p_m \bmod M)_{m=0}^\infty}$ for all $M \geq 1$. Thus, it is sufficient to decide $\mathrm{Acc}_{((p_{2m} \bmod M, p_{2m+1} \bmod M))_{m=0}^\infty} = \mathrm{Acc}_{((r_m \bmod M, s_m \bmod M))_{m=0}^\infty}$. Hence, it is sufficient to decide $\mathrm{Acc}_{((s_m + r_{m+1} \bmod M, r_{n+1} + s_{n+1} + 1 \bmod M))_{m=0}^\infty}$ for all $M \geq 1$. As $(s_m + r_{m+1} \bmod M)_{m=0}^\infty = (b^{m+1} - b^m - 1 \bmod M)_{m=0}^\infty$ is ultimately periodic where the period and preperiod can be controlled. Thus, we only need to determine whether $\mathrm{Acc}_{(r_{n+1} + s_{n+1} + 1 \bmod M)_{m=0}^\infty} = \mathrm{Acc}_{(\lfloor \sqrt{b} b^{n+1} \rfloor - \lfloor \sqrt{b} b^n \bmod M)_{m=0}^\infty}$, which is Turing-interreducible with Acc_β, where β is the base-b expansion of $\sqrt{b}$ by Theorem 4.5.2. □

Applying this result with $b = 2$, the theorem $\mathsf{MSO}_{\mathbb{N};<}(\mathsf{N}_2, 2^{\mathbb{N}})$ is Turing-equivalent to Acc_β, where β is the binary expansion of $\sqrt{2}$, as stated in the introduction of this section.

4.6 One linear recurrence sequence with two dominant roots

In this section, we prove Theorem 4.1.4: the MSO theory of $\langle \mathbb{N}; <, P \rangle$ is decidable for unary predicates P comprising the set of positive values of some non-degenerate, simple, integer-valued LRS having two dominant roots.

We begin in Section 4.6.1 by untangling the definition of an LRS satisfying the above hypotheses and reduce Theorem 4.1.4 to Theorem 4.1.12: the positive value sequence of P is effectively prodisjunctive. We then provide intuition underlying the proof of the latter through an extended example in Section 4.6.2 and prove a continuous version of our problem in Section 4.6.3. Finally, we establish Theorem 4.1.12 in Section 4.6.4.

4.6.1 Reduction to prodisjunctivity

Let $(u_n)_{n=0}^{\infty}$ be an LRS satisfying the hypotheses of Theorem 4.1.4, i.e., $(u_n)_{n=0}^{\infty}$ is simple, non-degenerate LRS with two dominant roots. We first record some elementary observations.

Lemma 4.6.1. *Assume that $(u_n)_{n=0}^{\infty}$ is an LRS satisfying the hypotheses of Theorem 4.1.4 and whose two dominant roots are λ_1 and λ_2. Then $\lambda_2 = \overline{\lambda_1}$, the argument of λ_1 is not a rational multiple of π, and $|\lambda_1| > 1$.*

Proof. As the characteristic polynomial of $(u_n)_{n=0}^{\infty}$ has integer coefficients, the roots $\overline{\lambda_1}$ and $\overline{\lambda_2}$ are also dominant. Since $(u_n)_{n=0}^{\infty}$ has exactly two dominant roots, we have $\{\lambda_1, \lambda_2\} = \{\overline{\lambda_1}, \overline{\lambda_2}\}$. If $\lambda_1 = \overline{\lambda_1}$, we have that $\lambda_2 = \overline{\lambda_2}$ and so both λ_1 and λ_2 are real. Hence, $\lambda_1/\lambda_2 = \pm 1$, contradicting non-degeneracy. Thus, $\lambda_2 = \overline{\lambda_1}$.

If the argument of λ_1 is a rational multiple of π, then $\lambda_1/\overline{\lambda_1}$ also has argument a rational multiple of π. Having modulus 1, $\lambda_1/\overline{\lambda_1}$ would then be a root of unity, contradicting non-degeneracy.

Assume $|\lambda_1| \leq 1$. By the Vieta formulas, the product of the absolute values of the characteristic roots is at most 1 and also equals $|c_d| \in \mathbb{Z}_{>0}$, where c_d is the constant coefficient of the characteristic polynomial of $(u_n)_{n=0}^{\infty}$. Hence, $|\lambda_1| = 1$, and there are no non-dominant roots. Whence, by Theorem 1.1.1, we conclude that both λ_1

and λ_2 are roots of unity, and thus so is their quotient, contradicting once again non-degeneracy. Hence $|\lambda_1| > 1$ as claimed. □

When writing $(u_n)_{n=0}^{\infty}$ in its exponential-polynomial form, we split it into its dominant part $(v_n)_{n=0}^{\infty}$ and non-dominant part $(r_n)_{n=0}^{\infty}$:

$$u_n = v_n + r_n = \alpha\lambda^n + \overline{\alpha}\overline{\lambda}^n + r_n\,. \tag{4.19}$$

Here, α and λ are algebraic numbers such that $\alpha \neq 0$ (as otherwise λ and $\overline{\lambda}$ would not be characteristic roots, i.e., roots of the polynomial corresponding to the *minimal* recurrence relation that $(u_n)_{n=0}^{\infty}$ obeys), $|\lambda| > 1$, and the argument of λ is not a rational multiple of π.

Recall that $P = \{u_n \colon n \in \mathbb{N}\} \cap \mathbb{N}$ To apply Theorem 4.1.9, we need the following lemma, whose proof relies heavily on the results of Mignotte, Shorey, and Tijdeman from Section 1.2.

Lemma 4.6.2. *Let $P \subseteq \mathbb{N}$ be as above. Then, P is infinite, recursive, and effectively sparse.*

Proof. First, as discussed in Section 1.2, compute $r > 0$ and $0 < R < |\lambda|$ such that $rR^n > r_n$ for all $n \in \mathbb{N}$.

To show that P is infinite, invoke [39, Lemma 4]: for infinitely many n, we have $\alpha\lambda^n + \overline{\alpha}\overline{\lambda}^n > c|\lambda|^n$ for some real $c > 0$. As $c|\lambda|^n > rR^n$ for all but finitely many n, there is a $c' > 0$ such that $u_n \geq c'|\lambda|^n$ for infinitely many n. Hence, P is infinite. Let $(p_m)_{m=0}^{\infty}$ be the positive value sequence of P, i.e., p_{m-1} is the mth smallest element in P.

To show that $(p_m)_{m=0}^{\infty}$ is recursive, it is sufficient to find, for a given $k \in \mathbb{N}$, a number N such that $|u_n| > k$ for all $n \geq N$ as then $k \in P$ if and only if $k \in \{u_0, \ldots, u_{N-1}\}$. Using Theorem 1.2.8, we have that when $n \geq C_2$ and $u_n = k$, we have that

$$n \log|\lambda| - C_1 \log^2(n) \leq \log|u_n| < \log(k+1)\,.$$

Hence, n is bounded, giving the desired bound N.

It remains to show that $(p_m)_{m=0}^{\infty}$ is effectively sparse. Assume that $k, n_1, n_2 \in \mathbb{N}$, $n_1 > n_2$, and $|u_{n_1} - u_{n_2}| \leq k$. Then Theorem 1.2.9 asserts that whenever $n_1 \geq C_4$,

$$\begin{aligned}\log(k+1) &\geq \log|u_{n_1} - u_{n_2}| \\ &\geq |\lambda|^{n_1} - C_3 \log(n_1)^2 \log(n_2+2) \\ &\geq |\lambda|^{n_1} - C_3 \log(n_1+1)^3\,.\end{aligned}$$

Thus $|u_{n_1} - u_{n_2}| \leq k$ implies that $n_1 \leq N'$ for some computable constant N'. Hence we can write out the set $P \cap \{0, \ldots, k+1+\max_{0 \leq n \leq N'}\{u_n\}\}$ and find the two largest elements in this set having difference at most k. □

Let $(p_m)_{m=0}^{\infty}$ denote the positive value sequence of P. By Theorem 4.1.9, it now suffices to prove that for all $M \geq 1$, one can decide $\mathrm{Acc}_{(p_m \bmod M)_{m=0}^{\infty}}$ (determine whether a given deterministic Muller automaton $\mathcal{A}$ over alphabet $\{0, \ldots, M-1\}$ accepts $(p_m \bmod M)_{m=0}^{\infty}$). The next lemma shows how the effective prodisjunctivity of $(p_m)_{m=0}^{\infty}$ asserted by Theorem 4.1.12 enables this.

Lemma 4.6.3. *Theorem 4.1.12 implies Theorem 4.1.4.*

Proof. By Lemma 4.6.2, P is infinite, recursive, and effectively sparse, and so to apply Theorem 4.1.9, we only need to verify that $\mathrm{Acc}_{(p_m \bmod M)_{m=0}^{\infty}}$ is decidable for all $M \geq 1$. Let $M \geq 1$. As Theorem 4.1.12 implies that $(p_m)_{m=0}^{\infty}$ is effectively prodisjunctive, we can compute $S_M \subset \{0, \ldots, M-1\}$ and $N_M \in \mathbb{N}$ such that $(p_m \bmod M)_{m=N_M}^{\infty} \in {S_M}^{\omega}$ is disjunctive. By hard-coding the finite prefix $(p_m \bmod M)_{m=0}^{N_M-1}$ into the automaton $\mathcal{A}$ and for the remainder restricting $\mathcal{A}$ to the alphabet S_M, we can apply Theorem 4.2.8 to conclude the lemma. □

Theorem 4.1.12 asserts that, for any $M \geq 2$, $(p_m \bmod M)_{m=N_M}^{\infty} \in {S_M}^{\omega}$ is disjunctive. Unpacking this definition, we have that $(p_m \bmod M)_{m=N_M}^{\infty} \in {S_M}^{\omega}$ is disjunctive if for any $\ell \geq 1$, every pattern $(s_1, \ldots, s_\ell) \in {S_M}^{\ell}$ appears infinitely often in $(p_m \bmod M)_{m=N_M}^{\infty}$. That is, for all $\ell \in \mathbb{N}$, we have that $N \geq N_M$, and $s_1, \ldots, s_\ell \in S_M$, there are $n_1, \ldots, n_\ell \in \mathbb{N}$ such that

1. $n_1, \ldots, n_\ell \geq N$;
2. for all $1 \leq i \leq \ell$, we have that $u_{n_i} \equiv s_i \pmod{M}$;
3. $0 \leq u_{n_1} < \cdots < u_{n_\ell}$;
4. for all $m \geq 0$ such that $u_{n_1} \leq u_m \leq u_{n_\ell}$, we have $u_m \in \{u_{n_1}, \ldots, u_{n_\ell}\}$.

As the dominant part $(v_n)_{n=0}^{\infty}$ (where $v_n = \alpha\lambda^n + \overline{\alpha}\overline{\lambda}^n$, see (4.19)) only relies on two algebraic numbers, α and λ, it is easier to use $(v_n)_{n=0}^{\infty}$ than $(u_n)_{n=0}^{\infty}$.

Lemma 4.6.4. *Assume that for all natural numbers ℓ, N, $T \geq 2$, and $t_1, \ldots, t_\ell \in \{0, \ldots, T-1\}$, there are $n_1, \ldots, n_\ell \in \mathbb{N}$ such that*

(A) $n_1, \ldots, n_\ell \geq N$;

(B) for all $1 \leq j \leq \ell$, *we have* $n_j \equiv t_j \pmod{T}$;

(C) $0 < v_{n_1} < \cdots < v_{n_\ell}$;

(D) for all $m \geq 0$ *such that* $v_{n_1} \leq v_m \leq v_{n_\ell}$, *we have* $m \in \{n_1, \ldots, n_\ell\}$.

Then the conclusion of Theorem 4.1.12 holds.

Proof. We claim that for some computable number N', whenever $m_1, m_2 \geq N'$, we have that $u_{m_1} < u_{m_2}$ if and only if $v_{m_1} < v_{m_2}$. Suppose not. Then, without loss of generality, $m_1 > m_2$ and $|v_{m_1} - v_{m_2}| < |r_{m_1}| + |r_{m_2}| < 2rR^{m_1}$. But Theorem 1.2.9 implies that for $m_1 \geq C_4$,

$$\begin{aligned}\log(2r) + m_1 \log(R) &> \log|v_{m_1} - v_{m_2}| \\ &> m_1 \log|\lambda| - C_3 \log^2(m_1) \log(m_2 + 2) \\ &> m_1 \log|\lambda| - C_3 \log^3(m_1 + 1)\,,\end{aligned}$$

which cannot hold for $m_1 \geq N'$ for some computable $N' \in \mathbb{N}$ as $\log|\lambda| > \log|R|$. Our claim follows.

Assume that $(u_n \bmod M)_{n=0}^{\infty}$ has period T (which can be effectively computed). If $s_1, \ldots, s_\ell \in S$ and $1 \leq j \leq \ell$, there exists a $t_j \in \{0, \ldots, T-1\}$ such that $u_{nT+t_j} \equiv s_j \pmod{M}$ whenever n is large enough. Therefore, if $n_1, \ldots, n_\ell \in \mathbb{N}$ are at least $\max\{N, N'\}$ and satisfy the hypotheses of the lemma, it follows that $0 \leq u_{n_1} < \cdots < u_{n_\ell}$ and for all $m \in \mathbb{N}$ such that $u_{n_1} < u_m < u_{n_\ell}$, we have that $m \in \{n_1, \ldots, n_\ell\}$. In other words, if $p_m = u_{n_1}$, then for $2 \leq j \leq \ell$, we have that $p_{m+j-1} = u_{n_j}$ and $p_{m+j-1} = u_{n'T+t_j} \equiv s_j \pmod{M}$ for some $n' \in \mathbb{N}$. □

As $\alpha \neq 0$ and Lemma 4.6.4 is only concerned with inequalities $v_{n_1} < v_{n_2}$ and $v_{n_1} > 0$ for natural numbers n_1 and n_2, we can scale v_n by $1/|2\alpha|$. That is, we can assume that $|\alpha| = 1/2$. Write $\alpha = \frac{1}{2}e^{i\varphi}$ and $\lambda = |\lambda|e^{i\theta}$. Thus, $v_n = \cos(\theta n + \varphi)|\lambda|^n$.

Now, we sketch the method we will use to prove the hypothesis of Lemma 4.6.4.

Our proof of Lemma 4.6.4 relies heavily on the following notation.

Definition 4.6.5. For integers $d \neq 0$ and real numbers $0 < \gamma < \delta$, define $\mathcal{J}_d(\gamma, \delta) \subset \mathbb{R}/2\pi\mathbb{Z}$ as

$$\mathcal{J}_d(\gamma, \delta) = \left\{x \in \mathbb{R}/2\pi\mathbb{Z} \colon 0 < \gamma\cos(x) < \cos(x + d\theta)|\lambda|^d < \delta\cos(x)\right\}.$$

Moreover, we define $\mathcal{J}_d(0, \delta)$ to stand for the limit of $\mathcal{J}_d(\gamma, \delta)$ as γ tends to 0.

In Lemma 4.6.7, We will show that these open sets $\mathcal{J}_d(\gamma, \delta)$ are open intervals of a certain size that are close to a point of the form $-d\theta \pm \pi/2$.

We first want to establish a continuous version of the hypothesis of Lemma 4.6.4, where instead of using natural numbers n, we deal with real numbers. We can do this, by searching for $b_2, \dots, b_d \in \mathbb{N}$ such that $n_i = n + b_i$. Then consider that $v_m = \cos(\theta m + \varphi)|\lambda|^m$ for all $m \in \mathbb{N}$. As we only compare terms of v_n with other terms v_{n+b_i} and 0, the inequality $v_n < v_{n+b_i}$ is equivalent to $\cos(\theta n + \varphi) < \cos(\theta n + \varphi + b_i\theta)|\lambda|^{b_i}$ and $v_n > 0$ if and only if $\cos(\theta n + \varphi) > 0$. Then, we consider real numbers $x \in \mathbb{R}/2\pi\mathbb{Z}$ instead of $n\theta + \varphi$.

Thus, instead of $n_1, \dots, n_\ell$, we write $n = n_1$ and $n_i = n + b_i$ for $2 \leq i \leq \ell$. Moreover, instead of finding a natural number n, we want to find an interval I such that there are $x = \theta n + \varphi$ in I for arbitrarily large n in the correct conjugacy class. This gives us the following lemma to prove:

Lemma 4.6.6. *Let $\ell, T \geq 2$ and $t_1, \dots, t_\ell \in \{0, \dots, T-1\}$. Then there are an interval $\mathcal{I} \subset \mathbb{R}/2\pi\mathbb{Z}$, $b_2, \dots, b_\ell \geq 1$, $1 < \delta_2 < \cdots < \delta_\ell < \sqrt{|\lambda|}$, and $D \in \mathbb{N}$ such that*

(a) for all $2 \leq j \leq \ell$, we have $b_j \equiv t_j - t_1 \pmod{T}$;

(b) for all $x \in \mathcal{I}$, we have

$$\begin{aligned} 0 < \cos(x) &< \cos(x + b_2\theta)|\lambda|^{b_2} < \delta_2 \cos(x) \\ &< \cos(x + b_3\theta)|\lambda|^{b_3} < \delta_3 \cos(x) \\ &\vdots \\ &< \cos(x + b_\ell\theta)|\lambda|^{b_\ell} < \delta_\ell \cos(x)\,; \end{aligned} \tag{4.20}$$

(c) for all integers $d < D$ not in the set $\{0, b_2, \dots, b_\ell\}$, we have $\mathcal{I} \cap \mathcal{J}_d(1, \delta_\ell) = \emptyset$;

(d) $\sum_{d=D}^{\infty} \left|\mathcal{J}_d(1, \delta_\ell)\right| < |\mathcal{I}|$.

Thus, if we compare this lemma with the hypothesis of Lemma 4.6.4, we see that (a) will imply (B), (b) will imply (C), and (c) and (d) are technical conditions which ultimately will imply (A) and (B) and are used to inductively construct further values b_i. Translating (b) in the notation of intervals $\mathcal{J}$, we have that

$$\mathcal{I} \subseteq (-\pi/2, \pi/2) \cap \mathcal{J}_{b_2}(1, \delta_2) \cap \mathcal{J}_{b_3}(\delta_2, \delta_3) \cap \cdots \cap \mathcal{J}_{b_\ell}(\delta_{\ell-1}, \delta_\ell)\,.$$

We will prove Lemma 4.6.6 in Section 4.6.3 and conclude Theorem 4.1.4 in Section 4.6.4. In the following section, we begin with an example to illustrate it in a slightly less technical context.

4.6.2 An extended example

Consider an example. Let $(u_n)_{n=0}^{\infty}$ be the sequence (4.1) from Section 4.1 and assume $M = 5$. Then we study $u_n = \frac{1}{2}(2+i)^n + \frac{1}{2}(2-i)^n + 2^n$ modulo 5.

From (4.3), we get that $(u_n \bmod 5)_{n=1^{\infty}} = (4, 2, 0, 4, 4, 2, 0, 4, \dots)$ is periodic with period $T = 4$. Thus, $S_5 = \{0, 2, 4\}$. Then Theorem 4.1.12 states that every $(s_1, \dots, s_\ell) \in {S_5}^*$ appears in $(p_m \bmod 5)_{m=0}^{\infty}$ infinitely often as a factor. We will show that this indeed holds when $\ell = 3$ and $(s_1, s_2, s_3) = (0, 0, 0)$.

We now compute t_1, t_2, and t_3 (the conjugacy classes modulo $T = 4$ such that $u_{Tn+t_i} \equiv s_i \pmod 5$ for all large enough n). Thanks to (4.3), we are forced to take $t_i = 3$ for $i = 1, 2, 3$ as $s_i = 0$. Thus, to find $(0, 0, 0)$ in $(p_m)_{m=0}^{\infty}$, we want to find $n_1, n_2, n_3 \geq 1$ congruent to 3 modulo 4 such that u_{n_1}, u_{n_2}, and u_{n_3} are successive in P, i.e., $0 < u_{n_1} < u_{n_2} < u_{n_3}$, and if $m \geq 0$ and $u_{n_1} < u_m < u_{n_3}$, then $u_m \in \{u_{n_1}, u_{n_2}, u_{n_3}\}$.

By (4.2), the dominant part $(v_n)_{n=0}^{\infty}$ of $(u_n)_{n=0}^{\infty}$ is given by $v_n = \frac{1}{2}(2+i)^n + \frac{1}{2}(2-i)^n$ and the non-dominant part $(r_n)_{n=0}^{\infty}$ by $r_n = 2^n$. As shown in Lemma 4.6.4, we can use $v_n = \alpha\lambda^n + \overline{\alpha}\overline{\lambda}^n = \cos(n\theta + \varphi)|\lambda|^n$ instead of u_n for large enough n. Here, we have that $\lambda = e^{i\theta}|\lambda| = 2 + i$ and $\alpha = \frac{1}{2}e^{i\varphi} = \frac{1}{2}$. Thus, in our example, $\varphi = 0$.

Hence, we wish to find $n_1, n_2, n_3 \in \mathbb{N}$ that are congruent to 3 modulo 4 and satisfy

$$0 < \cos(n_1\theta) < \cos(n_2\theta)|\lambda|^{n_2 - n_1} < \cos(n_3\theta)|\lambda|^{n_3 - n_1},$$

and that for all m satisfying $\cos(n_1\theta) < \cos(m\theta)|\lambda|^{m-n_1} < \cos(n_3\theta)|\lambda|^{n_3-n_1}$, we have $m \in \{n_1, n_2, n_3\}$ (and hence $m = n_2$ by the strict inequalities). This is the statement of Lemma 4.6.4.

We reached this far in the previous section. Next, we aim to find $b_2, b_3 \in \mathbb{N}$ such that $n_1 = n$, $n_2 = n + b_2$, and $n_3 = n + b_3$ for some $n \in \mathbb{N}$ congruent to 3 modulo 4 that satisfies our hypotheses. To ensure that $n_2 \equiv t_2 \equiv 3 \pmod 4$ and $n_2 \equiv t_2 \equiv 3 \pmod 4$, we have to have that $b_2 \equiv b_3 \equiv 0 \pmod 4$.

To solve this discrete problem (find natural numbers n, b_2, b_3 meeting these constraints), we first want to solve a continuous variant of this problem: find an *interval* $\mathcal{I} \subset \mathbb{R}/2\pi\mathbb{Z}$ and these natural numbers b_2 and b_3 such that when $x = n\theta$ is in $\mathcal{I}$, these properties hold 'often'. We will construct an open, non-empty interval $\mathcal{I} \subset \mathbb{R}/2\pi\mathbb{Z}$ such that for all $x \in \mathcal{I}$,

$$0 < \cos(x) < \cos(x + b_2\theta)|\lambda|^{b_2} < \cos(x + b_3\theta)|\lambda|^{b_3}. \tag{4.21}$$

We cannot ensure that for all $x \in \mathcal{I}$ and $m \in \mathbb{Z}$, we have that $m = b_2$ whenever $\cos(x) < \cos(x + m\theta)|\lambda|^m < \cos(x + b_3\theta)|\lambda|^{b_3}$. However, we can ensure it happens for

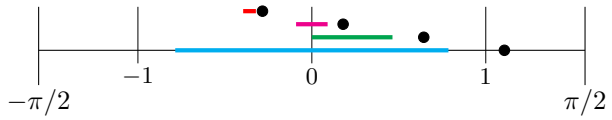

Figure 4.4: Let $\lambda = 1 + 2i$. Then, for $d = 1, 2, 3, 4$, we drew $\mathcal{J}_d(1, 3)$ in cyan, green, magenta, and red, respectively, and $|\mathcal{J}_d(1, 3)|$ are $\pi/2$, 0.464, 0.182, 0.073, respectively. As some of the intervals overlap, we have stacked them vertically for visual purposes. The points in black mark out $-d\theta \pm \pi/2$ for the corresponding value of d.

'many' $x \in \mathcal{I}$, including for infinitely many numbers of the form $x = n\theta$, where $n \equiv 3 \pmod 4$.

We translate this statement into the notation of the intervals $\mathcal{J}_d(\gamma, \delta)$. Then, arbitrarily setting $\delta_2 = 1.95$ and $\delta_3 = 2$, we strengthen the inequality (4.21) into the form of item (b) in Lemma 4.6.6: we want that for all $x \in \mathcal{I}$,

$$0 < \cos(x) < \cos(x + b_2\theta)|\lambda|^{b_2} < 1.95\cos(x) < \cos(x + b_3\theta)|\lambda|^{b_3} < 2\cos(x)\,.$$

Thus, $I \subseteq (-\pi/2, \pi/2)$ (as $\cos(x) > 0$ and $x \in \mathcal{I}$) and $I \subseteq \mathcal{J}_{b_2}(1, 1.95)$ (as $1 \cdot \cos(x) < \cos(x + b_2\theta)|\lambda|^{b_2} < 1.95\cos(x)$). Similarly, $I \subseteq \mathcal{J}_{b_3}(1.95, 2)$.

Dealing with items (c) and (d) of Lemma 4.6.6 is more difficult. First, we compute that if we take $\mathcal{I}_1 := (-1.1, 1.1) \subset (-\pi/2, \pi/2)$, then for all integers $d < 0$, we have that whenever $x \in \mathcal{I}_1$, we have $\cos(x + \theta d)|\lambda|^d < \cos(x)$ and so in item (c), we can take $D = 1$.

Now we discuss the useful results on the structure of the sets $\mathcal{J}_d(\gamma, \delta)$ proven in Lemma 4.6.7. For $d \in \mathbb{Z}$, the set $\mathcal{J}_d(\gamma, \delta)$ is empty or consists of a single open interval. Furthermore, for small enough δ, we have that $|\mathcal{J}_d(\gamma, \delta))| = O((\delta - \gamma)|\lambda|^{-d})$. Thus, these intervals $\mathcal{J}_d(\gamma, \delta)$ shrink exponentially fast with respect to d. Hence, item (d) of Lemma 4.6.6 can easily be estimated using a geometric series:

$$\sum_{d=D}^{\infty} \left|\mathcal{J}_d(1, \delta_\ell)\right| \geq O\big(|\lambda|^{-D}\big)\,.$$

Moreover, we will show that every point in them is at most $O(\delta|\lambda|^{-d})$ away from $-d\theta \pm \pi/2$ in $\mathbb{R}/2\pi\mathbb{Z}$. We illustrate this in Figure 4.4.

As $\delta_2 = 1.95$ and $\delta_3 = 2$, we will construct $\mathcal{I}$ inductively as

$$\begin{aligned} \mathcal{I} := \mathcal{I}_3 &:= \mathcal{J}_{b_3}(1.95, 2) \\ &\subseteq \mathcal{I}_2 := \mathcal{J}_{b_2}(1, 1.95) \\ &\subseteq \mathcal{I}_1 := (-1.1, 1.1)\,. \end{aligned}$$

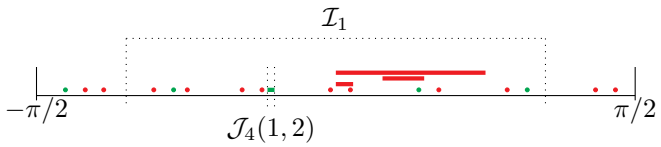

Figure 4.5: We drew the intersection of $\mathcal{I}_1$ and $\mathcal{J}_d(1,2)$ for $d = 1, \ldots, 20$ in red when $d \not\equiv 0 \pmod 4$ and in green when $d \equiv 0 \pmod 4$. For ease of visibility, the interval $\mathcal{J}_d(1,2)$ is positioned higher for $d = 1, 2, 3$, and for intervals $\mathcal{J}_d(1,2)$ that are too small to draw, their position is marked out with a dot.

In each step i, we want that for the interval $\mathcal{I}_i$, the items (a) and (b) are satisfied for $j \leq i$ while items (c) and (d) also hold.

Recall that for $x \in \mathcal{I}_1 = (-1.1, 1.1)$, we can take $D = 1$. Because $\delta_3 = 2$,

$$\Big|\bigcup_{d=1}^{\infty} \mathcal{J}_d(1,2)\Big| \leq \sum_{d=1}^{\infty} \big|\mathcal{J}_d(1,2)\big| < |\mathcal{I}_1|\,. \tag{4.22}$$

Thus, the intervals $\mathcal{J}_d(1,2)$ satisfying $d \neq 0$ do not cover $\mathcal{I}_1$ and so item (d) is satisfied.

For $\mathcal{I}_2$, we take an interval $\mathcal{J}_{b_2}(1, 1.95) \subset \mathcal{I}_1$, where $b_2 \equiv 0 \pmod 4$. We pick $b_2 = 4$ as it is the smallest possible choice for b_2. As shown in Figure 4.5, we have that $\mathcal{I}_2$ is indeed in $\mathcal{I}_1$ and for all integers $d < 20$ not equal to either 0 or 4, we have that $\mathcal{I}_2 \cap \mathcal{J}_d(1,2)$ is empty. As $\sum_{d=21}^{\infty} |\mathcal{J}_d(1,2)| < |\mathcal{J}_4(1, 1.95)|$, the interval $\mathcal{I}_2 := \mathcal{J}_4(1, 1.95)$ is not covered by the intervals $\mathcal{J}_d(1,2)$ such that $d \notin \{0, 4\}$. Hence, $D = 21$ is a valid choice in this step of our inductive procedure.

For b_3, recall that $b_3 \equiv 0 \pmod 4$. We find that $\mathcal{J}_b(1.95, 2) \cap \mathcal{I}_2$ is non-empty for $b = 0, 4, 38, 99, 160, 309, 370, \ldots$. The smallest such b that is not equal to 0 or 4 (which are already in use) and congruent to 0 modulo 4 is 160. Then,

$$\sum_{d=1, d\notin\{4,160\}}^{\infty} \big|\mathcal{J}_d(1,2) \cap \mathcal{J}_{160}(1.95, 2)\big| < \big|\mathcal{J}_{160}(1.95, 2)\big|$$

lets us take $D = 160$ and $\mathcal{I} := \mathcal{J}_{160}(1.95, 2)$. This interval $\mathcal{I}$ is tiny: it has a length of approximately $5.7 \cdot 10^{-58}$. However, $\mathcal{I} \cap \bigcup_{d=-\infty, d\notin\{0,4,160\}}^{\infty} \mathcal{J}_d(1,2)$ is an even smaller subset of $\mathcal{I}$. Thus we have found an interval $\mathcal{I}$ such that, for all $x \in \mathcal{I}$,

$$0 < \cos(x) < \cos(x + 4\theta)|\lambda|^4 < 1.95\cos(x) < \cos(x + 160\theta)|\lambda|^{160} < 2\cos(x)\,. \tag{4.23}$$

This gives a solution to the continuous version of our problem: $b_2 = 4$, $\delta_2 = 1.95$, $b_3 = 160$, $\delta_3 = 2$, and $\mathcal{I} = \mathcal{J}_{160}(1.95, 2)$ satisfy Lemma 4.6.6. It remains to show that there are infinitely many natural numbers $n \equiv 3 \pmod 4$ such that item (D) of Lemma 4.6.4 holds, as the three other conditions are already satisfied.

For a real number x, let $|x|_{2\pi}$ denote the distance from x to the nearest integer multiple of 2π. Assume that for n as above item (D) in Lemma 4.6.4 is not satisfied. Then $n\theta \in \mathcal{J}_d(1,2)$ for some integer $d \notin \{0,4,160\}$. Hence, using that $\mathcal{J}_d(1,2)$ is close to $-d\theta \pm \pi/2$ modulo 2π (we prove this in Lemma 4.6.7),

$$|d+n|^{-c_1} < \left|n\theta - (-d\theta \pm \pi/2)\right|_{2\pi} < c_2|\lambda|^{-d} \tag{4.24}$$

for two constants $c_1, c_2 > 0$. In the above, the first inequality follows from Baker's theorem, whereas the second inequality is derived from the exponential rate of shrinkage of the intervals. Thus, n is much larger than d.

As there are many $n\theta$ in $\mathcal{I}$ that are fairly 'evenly' distributed, we are able to prove that (4.24) cannot hold for all n. In other words, there are $n \in \mathbb{N}$ such that

$$0 < v_n < v_{n+4} < v_{n+160}$$

and $v_n < v_m < v_{n+160}$ implies that $m = n+4$. Translating back to u_n gives us the required result. In particular, we can calculate that when taking

$$n = 218085867698737188268427463501308698889728969450963229999559\,,$$

we have that $n\theta \in \mathcal{I}$, $n \equiv 3 \pmod 4$, and $u_n < u_m < u_{n+160}$ implies that $m = 4$. Thus, u_n, u_{n+4}, and u_{n+160} appear consecutively in $(p_m)_{m=0}^{\infty}$ and all are divisible by 5. Thus, the pattern $(0,0,0)$ does indeed appear in $(p_m \bmod 5)_{m=0}^{\infty}$.

However, this construction is far from optimal. Although our choices of $b_2 = 4$ and $b_3 = 160$ were as small as possible with our choice of δ_2 and δ_3, these choices were not optimal. Nonetheless, this would yield only a minor improvement, as one must handle inequalities such as (4.22) with greater care.

Then one can find that choosing $b_2 = 8$, $b_3 = 28$, and $n = 16958443$, we indeed get that there is no $m \in \mathbb{N} \setminus \{0,8,28\}$ such that u_m is between u_n, u_{n+b_2}, and u_{n+b_3}. Meanwhile, $\cos(n\theta) \approx 0.404$, $\cos((n+b_2)\theta) \approx 94.5$, and $\cos((n+b_3)\theta) \approx 751$, and so one should have had chosen much larger δ_2 and δ_3 for which the general bounds in Lemma 4.6.7 fail.

4.6.3 Proof of Lemma 4.6.6

In this section, we prove Lemma 4.6.6, the continuous version of Theorem 4.6.4, which will serve as our main tool to establish Theorems 4.1.12 and 4.1.4.

Before we can accomplish this, we first need to establish a few lemmas. First, we start by describing the sets $\mathcal{J}_d(\gamma,\delta)$. Then, we show that $\mathcal{J}_d(\gamma,\delta)$ is a single interval

whose size of $\mathcal{J}_d(\gamma,\delta)$ can be decently controlled. Moreover, the distance between a point in $\mathcal{J}_d(\gamma,\delta)$ and $-d\theta\pm\pi/2$ are also well-behaved, and thus these points $-d\theta\pm\pi/2$ describe the location of $\mathcal{J}_d(\gamma,\delta)$ well. Together with the size bounds, we will use this to determine whether such intervals have a non-empty intersection.

Lemma 4.6.7. *One can compute constants C_5, C_6, C_7, and C_8 such that for all $0<\gamma<\delta<\sqrt{|\lambda|}$ and all $d\geq 1$, we have that $\mathcal{J}_d(\gamma,\delta)$ consists of a single interval,*

$$C_6\frac{\delta-\gamma}{|\lambda|^d d^{C_7}} < \left|\mathcal{J}_d(\gamma,\delta)\right| < C_5\frac{\delta-\gamma}{|\lambda|^d}\,,$$

and $|x-(-d\theta\pm\pi/2)|_{2\pi}<C_8|\lambda|^{-d}$ for any $x\in\mathcal{J}_d(\gamma,\delta)$.

Proof. Identify $\mathbb{R}/2\pi\mathbb{Z}$ with $(-\pi,\pi]$. As $2\cos(x)=e^{ix}+e^{-ix}$ and $e^{id\theta}|\lambda|^d=\lambda^d$, for $\gamma\leq\eta\leq\delta$, we have that

$$\cos(x+d\theta)|\lambda|^d=\eta\cos(x) \iff \left(e^{ix}\right)^2=-\frac{\overline{\lambda}^d-\eta}{\lambda^d-\eta}\,. \tag{4.25}$$

Hence, there is a unique $x\in(-\pi/2,\pi/2]$ such that $\cos(x+d\theta)|\lambda|^d=\eta\cos(x)$. If $x=\pi/2$ and $\cos(x+d\theta)=\eta\cos(x)$, then $\cos(x)=0$ and θ/π is rational, which we excluded in Lemma 4.6.1. Thus, $\mathcal{J}_d(\gamma,\delta)$ is a single interval within $(-\pi/2,\pi/2)$.

We now tackle the size of $\mathcal{J}_d(\gamma,\delta)$. As $\mathcal{J}_d(\gamma,\delta)$ consists of a single interval, we have $|\mathcal{J}_d(\gamma,\delta)|=|x_1-x_2|_{2\pi}$, where x_1 and x_2 are solutions in $(-\pi/2,\pi/2)$ to (4.25) for $\eta=\gamma$ and $\eta=\delta$, respectively. Using the triangle inequality on the unit circle,

$$|e^{ix_1}-e^{ix_2}|\leq|x_1-x_2|_{2\pi}\leq\frac{\pi}{2}|e^{ix_1}-e^{ix_2}|\,.$$

If $\gamma=\delta$, $x_1=x_2$, and so by continuity, when $\delta-\gamma$ is small enough, $|e^{ix_1}-e^{ix_2}|<\sqrt{2}$.

Assume we are on this boundary. That is, $|e^{ix_1}-e^{ix_2}|=\sqrt{2}$. Then $e^{ix_1}=\pm ie^{ix_2}$. It follows that $e^{2ix_1}=-e^{2ix_2}$, and so

$$\begin{aligned}-1&=e^{2ix_1}e^{-2ix_2}\\&=\frac{\overline{\lambda}^d-\gamma}{\lambda^d-\gamma}\cdot\frac{\lambda^d-\delta}{\overline{\lambda}^d-\delta}\\&=\frac{\lambda^d\overline{\lambda}^d-\gamma\lambda^d-\delta\overline{\lambda}^d+\gamma\delta}{\lambda^d\overline{\lambda}^d-\delta\lambda^d-\gamma\overline{\lambda}^d+\gamma\delta}\,.\end{aligned}$$

Therefore, $2\lambda^d\overline{\lambda}^d-(\delta-\gamma)\lambda^d+(\delta-\gamma)\overline{\lambda}^d+2\gamma\delta=0$. Then $|\lambda|^{2d}\leq(\delta-\gamma)|\lambda|^d+\gamma\delta$, and so $(|\lambda|^d+\gamma)(|\lambda|^d-\delta)\leq 0$. This is impossible in our scenario as $0<\gamma<\delta<\sqrt{|\lambda|}$ and $d\geq 1$. Thus, again by continuity in γ and δ, we have that $|e^{ix_1}-e^{ix_2}|<\sqrt{2}$.

From the geometry of the unit circle it follows that $|e^{ix_1} - e^{ix_2}| < \sqrt{2}$ implies that $\sqrt{2} < |e^{ix_1} + e^{ix_2}| \leq 2$. Then, using the fact that $|e^{2ix_1} - e^{2ix_2}| = |e^{ix_1} - e^{ix_2}||e^{ix_1} + e^{ix_2}|$, we obtain

$$\frac{1}{2}|e^{i2x_1} - e^{i2x_2}| \leq |x_1 - x_2|_{2\pi} \leq \frac{\pi}{2\sqrt{2}}|e^{i2x_1} - e^{i2x_2}|\,.$$

We can rewrite $|e^{i2x_1} - e^{i2x_2}|$ as follows:

$$\begin{aligned}
&|e^{2ix_1} - e^{2ix_2}| \\
&= \left|\frac{\overline{\lambda}^d - \gamma}{\lambda^d - \gamma} - \frac{\overline{\lambda}^d - \delta}{\lambda^d - \delta}\right| \\
&= \frac{\left|\left(\overline{\lambda}^d - \gamma\right)\left(\lambda^d - \delta\right) - \left(\overline{\lambda}^d - \delta\right)\left(\lambda^d - \gamma\right)\right|}{|\lambda^d - \gamma||\lambda^d - \delta|} \\
&= \frac{1}{|\lambda^d - \gamma||\lambda^d - \delta|}\left|(\delta - \gamma)\lambda^d - (\delta - \gamma)\overline{\lambda}^d\right| \\
&= \frac{(\delta - \gamma)|\lambda|^d}{|\lambda^d - \gamma||\lambda^d - \delta|}\left|e^{i2d\theta} - 1\right|\,.
\end{aligned}$$

As $\gamma, \delta < \sqrt{|\lambda|}$, there are constants $c_1, c_2 > 0$ such that $c_1|\lambda|^d < |\lambda^d - \eta| < c_2|\lambda|^d$ for $\eta \in \{\gamma, \delta\}$ and $d \geq 1$. Moreover, using Baker's theorem (Theorem 1.1.11) on $|e^{i2d\theta} - 1|$, we obtain a constant C_7 such that $d^{-C_7} < |e^{i2d\theta} - 1| < 2$. Thus, for some constants C_5 and C_6, we have that

$$|\mathcal{J}_d(\gamma, \delta)| \leq \frac{\pi}{2\sqrt{2}}\left|e^{2ix_1} - e^{2ix_2}\right| < C_5\frac{\delta - \gamma}{|\lambda|^d} \tag{4.26}$$

and

$$|\mathcal{J}_d(\gamma, \delta)| \geq \frac{1}{2}\left|e^{2ix_1} - e^{2ix_2}\right| > C_6\frac{d^{-C_7}(\delta - \gamma)}{|\lambda|^d}\,.$$

For the last claim, we estimate $|\mathcal{J}_d(0, \sqrt{|\lambda|})|$ with (4.26). □

In the following lemma, we show that an interval in $\mathbb{R}/2\pi\mathbb{Z}$ contains intervals $\mathcal{J}_d(0, \sqrt{|\lambda|})$ and numbers $n\theta + \varphi$ for small d and n in given congruent classes.

Lemma 4.6.8. *Let $T \geq 2$. There is a number $C_9 > 0$ such that for every $t \in \{0, \ldots, T-1\}$ and small enough interval $\mathcal{I} \subset (-\pi/2, \pi/2)$, there are $|\mathcal{I}|^{-C_9} \leq n_1, n_2 \leq 2|\mathcal{I}|^{-C_9}$ such that $n_1\theta + \varphi \in \mathcal{I}$, $\mathcal{J}_{n_2}(0, |\sqrt{\lambda}|) \subset \mathcal{I}$, and $n_1, n_2 \equiv t \pmod{T}$.*

Proof. By the pigeonhole principle, there are distinct $d_1, d_2 \in \mathbb{N}$ such that $0 \leq d_1, d_2 \leq \lceil \pi|\mathcal{I}|^{-1} \rceil$ and $d_1 T\theta$ and $d_2 T\theta$ have distance at most $|\mathcal{I}|$ modulo 2π. Here, we use that θ is not a rational multiple of π (Lemma 4.6.1). The last condition implies that

$$\left|(d_1 - d_2)T\theta\right|_{2\pi} < |\mathcal{I}|\,.$$

Baker's theorem (Theorem 1.1.11) implies there is a computable number c_1 such that

$$\begin{aligned}\left|(d_1-d_2)T\theta\right|_{2\pi} &> \left|e^{iT(d_1-d_2)\theta}-1\right| \\ &> |d_1-d_2|^{-c_1} \\ &> \left\lceil \pi|\mathcal{I}|^{-1}\right\rceil^{-c_1}.\end{aligned}$$

For all $N\in\mathbb{Z}$ and $x\in\mathbb{R}/2\pi\mathbb{Z}$, there is an $N\le n\le 2\pi\lceil\pi|\mathcal{I}|^{-1}\rceil^{c_1}+N$ such that $x+nT(d_1-d_2)\theta\in\mathcal{I}$. Taking C_9 slightly larger than c_1, we have $2\pi T\lceil\pi|\mathcal{I}|^{-1}\rceil^{c_1}<|\mathcal{I}|^{-C_9}$ for small enough $\mathcal{I}$. Hence, for all $N\in\mathbb{Z}$ and $x\in\mathbb{R}/2\pi\mathbb{Z}$, there is an $n\equiv t \pmod T$ such that $x+n\theta\in\mathcal{I}$ and $N\le n\le|\mathcal{I}|^{-C_9}+N$. For n_1, let $x=\varphi$ and $N=|\mathcal{I}|^{-C_9}$.

Let $\mathcal{I}'$ be the middle half of $\mathcal{I}$. As before, there is an $n_2\equiv t \pmod T$ such that $-2|\mathcal{I}'|^{-C_9}\le -n_2\le -|\mathcal{I}'|^{-C_9}$ and $-n_2\theta\pm\pi/2\in\mathcal{I}'$. If $\mathcal{J}_{n_2}(0,\sqrt{|\lambda|})\not\subseteq\mathcal{I}$, then $\mathcal{J}_{n_2}(0,\sqrt{|\lambda|})$ intersects $\mathcal{I}'$ and the complement of $\mathcal{I}$. So, $|\mathcal{J}_{n_2}(0,\sqrt{|\lambda|})|\ge\frac14|\mathcal{I}|$ and

$$\begin{aligned}\frac14|\mathcal{I}|\le\left|\mathcal{J}_{n_2}\big(0,\sqrt{|\lambda|}\,\big)\right| &\le C_8|\lambda|^{-n_2} \\ &\le C_8|\lambda|^{-|\mathcal{I}'|^{-C_9}} \\ &= C_8|\lambda|^{-(|\mathcal{I}|/2)^{-C_9}}\end{aligned}$$

by Lemma 4.6.7. After taking logarithms, we have that

$$\log|\mathcal{I}|\le\log(4C_8)-(|\mathcal{I}|/2)^{-C_9}\log|\lambda|,$$

which cannot hold for small enough $\mathcal{I}$. Thus $\mathcal{J}_{n_2}(1,\sqrt{|\lambda|})\subseteq\mathcal{I}$. □

For the fourth condition of Lemma 4.6.6, we would like D to be large, i.e., the smallest d such that $\mathcal{J}_d(1,\delta_\ell)$ has a non-empty intersection with $\mathcal{I}$ has to be quite large. The following lemma shows this is indeed possible.

Lemma 4.6.9. *Assume $\mathcal{I}\subset\mathbb{R}/2\pi\mathbb{Z}$, $b_2,\dots,b_\ell\in\mathbb{N}$, $1<\delta_2<\cdots<\delta_\ell<\sqrt{|\lambda|}$ and $D>0$ satisfy the hypotheses of Lemma 4.6.6. Then, for every small enough $\varepsilon>0$, there is a subinterval $\mathcal{I}'$ of $\mathcal{I}$ of length ε for which the hypotheses of Lemma 4.6.6 hold for these $b_2,\dots,b_\ell$ and $\delta_1,\dots,\delta_\ell$ and some $D'>\varepsilon^{-1/2}$.*

Proof. For an interval $\mathcal{I}'\subset\mathbb{R}/2\pi\mathbb{Z}$, let $D(\mathcal{I}')$ denote the smallest natural number $d\ge1$ such that $d\ne b_2,\dots,b_\ell$ and $\mathcal{J}_d(1,\delta_\ell)\cap\mathcal{I}'$ is non-empty. Let

$$c_1=|\mathcal{I}|-\sum_{\substack{d=D\\ d\notin\{b_2,\dots,b_\ell\}}}^{\infty}|\mathcal{J}_d(1,\delta_\ell)|.$$

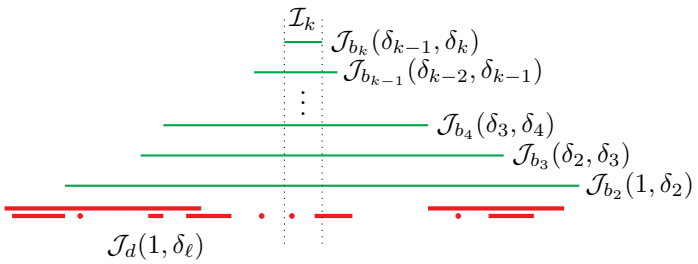

Figure 4.6: The inductive construction for $\mathcal{I}$. In green, we see that at each step, the interval $\mathcal{I}_k$ is taken to be $\mathcal{J}_{b_k}(\delta_{k-1},\delta_k)$. This gives (4.20) and by construction, we have the inclusion $\mathcal{I}_{k+1}\subset\mathcal{I}_k$. The red intervals are dense in $\mathcal{I}_k$, but by being careful, do not cover $\mathcal{I}_k$ at any step. The picture is not to scale as both the red and green intervals decrease in size exponentially fast.

By construction, $c_1>0$. Let $k\in\mathbb{N}$ such that $\frac{c_1}{k+2}<\varepsilon\leq\frac{c_1}{k+1}$ (which always exists for small enough ε). We will study the set

$$X=\mathcal{I}\setminus\bigcup_{\substack{d=D\\ d\notin\{b_2,\dots,b_\ell\}}}^{k}\mathcal{J}_d(1,\delta_\ell)\,,$$

which is an interval $\mathcal{I}$ from which at most k intervals are removed. Thus, X consists out of at most $k+1$ intervals By assumption, $|X|>c_1$, and so $|X|$ contains an interval $\mathcal{I}'$ of length $\varepsilon\leq\frac{c_1}{k+1}$ such that $D(\mathcal{I}')\geq k+1$. Then, for large enough k, we have that

$$D(I')\geq k+1>\sqrt{\frac{k+2}{c_1}}>\varepsilon^{-1/2}\,.$$

When ε is small enough, the result follows when setting $D'=D(\mathcal{I}')$. □

Now we are in a position to prove Lemma 4.6.6.

Proof of Lemma 4.6.6. We apply induction on k and require that the conditions hold for $1,\dots,k-1$ for an interval $\mathcal{I}_{k-1}$, and construct an interval $\mathcal{I}_k$ that satisfies the theorem when the first three conditions are restricted to $1,\dots,k$. Then we take $\mathcal{I}=\mathcal{I}_\ell$. The iterative process (not to scale) is depicted in Figure 4.6.

For the base case, let $\mathcal{I}_1=\left\{x\in\mathbb{R}/2\pi\mathbb{Z}\colon\ \cos(x)>|\lambda|^{-1}\right\}$, $D=1$, and $1<\delta_\ell<\sqrt{|\lambda|}$ small enough such that that the last inequality of

$$\sum_{d=D}^{\infty}\left|\mathcal{J}_d(1,\delta_\ell)\right|\leq\sum_{d=1}^{\infty}C_5\frac{\delta_\ell-1}{|\lambda|^d}\leq C_5\frac{\delta_\ell-1}{|\lambda|-1}<|\mathcal{I}_1|$$

holds, where the first inequality follows from Lemma 4.6.7 and the second by a geometric series. Hence, Condition (D) holds, and the first three conditions follow by construction. The base case follows.

For the other cases, choose $\delta_2, \dots, \delta_{\ell-1}$ such that $1 < \delta_2 < \cdots < \delta_{\ell-1} < \delta_\ell$. Furthermore, for simplicity, set $\delta_1 = 1$.

Now assume $k \geq 1$ and let $\varepsilon > 0$. For ε small enough, applying Lemma 4.6.9 on $\mathcal{I}_{k-1}$ gives intervals $\mathcal{I}_\varepsilon \subset \mathcal{I}_{k-1}$ of length ε where the smallest d such that $\mathcal{J}_d(1, \delta_\ell) \cap \mathcal{I}_\varepsilon \neq \emptyset$ is at least $\varepsilon^{-1/2}$.

Lemma 4.6.8 implies that when ε is again small enough, there is an $\varepsilon^{-C_9} < b_k < 2\varepsilon^{-C_9}$ such that $b_k \equiv t_k - t_1 \pmod{T}$ and $\mathcal{J}_{b_k}(\delta_{k-1}, \delta_k) \subset \mathcal{I}_\varepsilon$. If $d > 0$ and $\mathcal{J}_{b_k}(\delta_{k-1}, \delta_k) \cap \mathcal{J}_d(1, \delta_\ell) \neq \emptyset$, then $d > \varepsilon^{-1/2}$ by Lemma 4.6.9. By Lemma 4.6.7,

$$\left|(b_k - d)\theta \pm_1 \pi/2 \pm_2 \pi/2\right|_{2\pi} \leq \frac{C_8}{|\lambda|^{b_k}} + \frac{C_8}{|\lambda|^d} \leq \frac{2C_8}{|\lambda|^{\min\{b_k, d\}}}\,.$$

We put $j\pi = \pm_1 \pi/2 \pm_2 \pi/2$ for $j \in \mathbb{Z}$.[2] Then by Baker's theorem (Theorem 1.1.11), there is a constant $c_1 > 0$,

$$\begin{aligned} |b_k - d|^{-c_1} &< \left|e^{i((b_k - d)\theta + j\pi)} - 1\right| \\ &\leq \left|(b_k - d)\theta + j\pi\right|_{2\pi} \\ &\leq \frac{2C_8}{|\lambda|^{\min\{b_k, d\}}}\,. \end{aligned}$$

Taking logarithms, we obtain

$$\min\{b_k, d\} \log|\lambda| < \log(2C_8) + c_1 \log|b_k - d|\,. \tag{4.27}$$

Hence, if $d < b_k$, we have $|b_k - d| < b_k$ (as $d > 0$) and

$$d \log|\lambda| < \log(2C_8) + c_1 \log(b_k)\,.$$

Using that $d > \varepsilon^{-1/2}$ and $b_k \leq 2(\varepsilon/2)^{-C_9}$, we obtain

$$\varepsilon^{-1/2} \log|\lambda| < \log(2C_8) + c_1 \log(2) - c_1 C_9 \log(\varepsilon/2)\,.$$

This is impossible for sufficiently small ε. We can therefore assume that $d > b_k$. We take $\mathcal{I}_k := \mathcal{J}_{b_k}(\delta_{k-1}, \delta_k)$ such that Conditions (A) and (B) are satisfied. Now assume $d > b_k$ and let $D = d$ (and so Condition (C) automatically holds). For a contradiction, assume Condition (D) is violated. Lemma 4.6.7 and the geometric series give that for some $c_2 > 0$,

$$\frac{C_6(\delta_k - \delta_{k-1})}{|\lambda|^{b_k} b_k^{C_7}} \leq |\mathcal{I}_k| \leq \sum_{d'=d}^{\infty} |\mathcal{J}_{d'}(1, \delta_\ell)| < c_2 |\lambda|^{-d}\,.$$

[2] We have adorned the $\pm$ operator with subscripts (1 and 2) to indicate how the particular choice of signs should be preserved.

Hence, setting $c_3 = \log(\frac{c_2}{C_6(\delta_k - \delta_{k-1})})$ and taking logarithms, we get

$$(d - b_k)\log|\lambda| \leq c_3 + C_7 \log(b_k)\,. \tag{4.28}$$

Inserting the latter in (4.27) gives

$$b_k \log|\lambda| \leq \log(2C_8) + c_1 \log\left(\frac{c_3 + C_7\log(b_k)}{\log|\lambda|}\right),$$

which upper bounds b_k (independently of ε). As taking ε small gives arbitrarily large b_k, the result follows for k. Induction completes the proof. □

4.6.4 Proof of Theorems 4.1.4 and 4.1.12

Now that we have solved the continuous version of our problem (Lemma 4.6.6), we can solve the discrete version and conclude the proofs of Theorems 4.1.4 and 4.1.12.

Proof of Theorems 4.1.4 and 4.1.12. As Theorem 4.1.12 implies Theorem 4.1.4 due to Lemma 4.6.3, it is sufficient to prove Theorem 4.1.12. In turn, Lemma 4.6.4 states that Theorem 4.1.12 is implied by the following statement: for all $N \in \mathbb{N}$, $T \geq 2$, and $t_1, \ldots, t_\ell \in \{0, \ldots, T-1\}$, we can find $n_1, \ldots, n_\ell \in \mathbb{N}$ such that

1. $n_1, \ldots, n_\ell \geq N$;
2. $n_j \equiv t_j \pmod{T}$ for $1 \leq j \leq \ell$;
3. $0 < v_{n_1} < \cdots < v_{n_\ell}$;
4. for all $m \in \mathbb{N}$ such that $v_{n_1} \leq v_m \leq v_{n_\ell}$, we have $m \in \{n_1, \ldots, n_\ell\}$.

We will prove this statement for given N, T, ℓ, and $t_1, \ldots, t_\ell$.

Lemma 4.6.6 gives numbers $b_2, \ldots, b_\ell \in \mathbb{N}$ and $1 < \delta_\ell < \sqrt{|\lambda|}$ and an interval $\mathcal{I}$. We take $n_1 = n$ and $n_j = n + b_j$ for $2 \leq j \leq \ell$ and claim that $n \geq N$, $n \equiv t_1 \pmod{T}$, and $n\theta + \varphi \in \mathcal{I}$, imply the first three conditions. Indeed, $n_1 \geq N$ and $n_1 \equiv t_1 \pmod{T}$ hold. For $2 \leq j \leq \ell$, we have $n_j = n + b_j \geq N$ and

$$n_j \equiv n + b_j \equiv n + (t_j - t_1) \equiv t_j \pmod{T}\,.$$

For the third condition, we have that as $n\theta + \varphi \in \mathcal{I}$,

$$0 < \cos(n\theta + \varphi) < \cos\big((n+b_2)\theta + \varphi\big)|\lambda|^{b_2} < \cdots < \cos\big((n+b_\ell)\theta + \varphi\big)|\lambda|^{b_\ell}$$

and multiplying the last inequalities by $2|\lambda|^n$ and noting that $2\cos(m\theta+\varphi)|\lambda|^m = v_m$ for all $m \in \mathbb{N}$, we obtain our claim.

Let $0 < \varepsilon < |\mathcal{I}|$ be small enough. Lemma 4.6.9 implies that there is an interval $\mathcal{I}_\varepsilon \subset \mathcal{I}$ of length ε such that for all $d < \varepsilon^{-1/2}$, we have that $\mathcal{I}_\varepsilon \cap \mathcal{J}_d(1, \delta_\ell) = \emptyset$. By Lemma 4.6.8, there is an n such that $\varepsilon^{-C_9} < n < 2\varepsilon^{-C_9}$, $n\theta + \varphi \in \mathcal{I}_\varepsilon$, and $n \equiv t_1 \pmod{T}$. Thus, for small enough ε, we also have that $n \geq N$.

Now assume that $m \notin \{n_1, \ldots, n_\ell\}$ and that $v_{n_1} < v_m < v_{n_\ell}$. Then, setting $d = n - m$, we have that $d \notin \{0, b_2, \ldots, b_\ell\}$ and

$$0 < \cos(n\theta + \varphi)|\lambda|^n < \cos\big((n+d)\theta + \varphi\big)|\lambda|^{n+d} < \cos\big((n+b_\ell)\theta + \varphi\big)|\lambda|^{n+b_\ell}\,.$$

As $\cos((n+b_\ell)\theta + \varphi)|\lambda|^{b_\ell} < \delta_\ell \cos(n\theta + \varphi)$ because $n\theta + \varphi \in \mathcal{I}$, it follows that $n\theta + \varphi$ is in $\mathcal{J}_d(1, \delta_\ell)$. Thus, $d \geq \varepsilon^{-1/2}$. By Lemma 4.6.7, $n\theta + \varphi \in \mathcal{J}_d(1, \delta_\ell)$ and $-d\theta \pm \pi/2$ are at most $C_8|\lambda|^{-d}$ apart. Thus, for a constant c_1 derived from Theorem 1.1.11,

$$\begin{aligned} C_8|\lambda|^{-d} &\geq \big|(n\theta + \varphi) - (-d\theta \pm \pi/2)\big|_{\pi/2} \\ &\geq \big|e^{i(n\theta+\varphi)-(-d\theta\pm\pi/2)}\big| \\ &> \frac{\pi}{2}|n+d|^{-c_1}\,. \end{aligned}$$

Taking logarithms, we obtain

$$\log(C_8) + c_1 \log(n+d) > d \log|\lambda|\,.$$

As $d \geq \varepsilon^{-1/2}$ and $n \geq \varepsilon^{-C_9}$, it follows that $n, d \geq 2$ for small enough ε. In that case, $n + d \leq dn$ and so

$$\log(C_8) + c_1 \log(n) + c_1 \log(d) > d \log|\lambda|\,.$$

Hence, either

$$2\log(C_8) + 2c_1 \log(d) > d \log|\lambda|$$

or

$$2c_1 \log(n) > d \log|\lambda|\,.$$

The former is impossible for large enough d (and thus small enough ε), while for the latter, the upper and lower bounds for n and d give that

$$2c_1 \log(2\varepsilon^{-C_9}) > \log|\lambda|\varepsilon^{-1/2}\,,$$

which again is impossible for small enough ε. Hence, the fourth condition also follows. □

4.7 Concluding remarks

Our two main positive results of this chapter, Theorems 4.1.1 and 4.1.4, significantly expand the decidability landscape of the MSO theory of $\langle \mathbb{N}; <, P_1, \dots, P_d \rangle$, where the predicates P_i are the positive values of an LRS $(u_n^{(i)})_{n=0}^{\infty}$. However, the techniques are diametrically opposed to each other. In the case of Theorem 4.1.1, we obtain a toric word, restricting which kind of factors appear in the characteristic word. In the case of Theorem 4.1.4, we obtain a disjunctive word, showing that the characteristic word contains every finite word as a factor. Recall from Proposition 4.2.10, that being toric and disjunctive are almost orthogonal properties.

How far can these methods be pushed? In this section, we make a conjecture.

Further generalizing these properties for more general LRS forces us to prove that predicates derived from such LRS are infinite, recursive, effectively sparse, and effectively pairwise sparse. Proving such properties is at least Skolem- and Positivity-hard, so for simplicity, assume we can avoid these problems. Then we use the Elgot-Rabin contraction method in the form of Theorem 4.3.2: We need to prove that $\mathrm{Acc}_{(\beta_m, p_m \bmod M)_{m=0}^{\infty}}$ is decidable for all $M \geq 1$, where β and $(p_m)_{m=0}^{\infty}$ denote the order word and the positive value sequence, respectively.

As explored in Section 4.5, dealing with polynomials is awkward as their behaviour wildly differs from exponentially growing LRS. For non-simple LRS, we do not know how to decide ultimate positivity, let alone what kind of structure their positive value set can have. Thus, let us assume that we are dealing with simple LRS whose spectral radius is larger than 1. Lastly, for non-degenerate LRS, cancellations within the exponential-polynomial form can occur regularly. We want to avoid such problems here.

Hence, we focus on simple, non-degenerate LRS whose spectral radius is larger than 1. Then, for an LRS $(u_n)_{n=0}^{\infty}$ with a dominant root λ, we can compute the closure of the normalised sequence $(u_n/|\lambda|^n)_{n=0}^{\infty}$ as an interval as in Section 3.2. That is, we compute $\mu_-, \mu_+ \in \overline{\mathbb{Q}}$ such that

$$[\mu_-, \mu_+] = \overline{\{u_n/|\lambda|^n \colon n \in \mathbb{N}\}}\,. \tag{4.29}$$

Call $[\mu_-, \mu_+]$ the *normalised closure* of $(u_n)_{n=0}^{\infty}$. If $\mu_+ \leq 0$, the predicate $P = \{u_n \colon n \in \mathbb{N}\} \cap \mathbb{N}$ is finite, so such predicates are not very interesting. We conjecture the following.

Conjecture 4.7.1. *Let $(u_n^{(1)})_{n=0}^{\infty}, \dots, (u_n^{(d)})_{n=0}^{\infty}$ be pairwise sparse LRS of spectral radius larger than 1 whose dominant parts have finite overlap. For $1 \leq i \leq d$, let*

$[\mu_-^{(i)}, \mu_+^{(i)}]$ *be the normalised closure of* $(u_n^{(i)})_{n=0}^{\infty}$ *and* $P_i = \{u_n^{(i)} : n \in \mathbb{N}\} \cap \mathbb{N}$ *be infinite. Let* $M \geq 1$, β *be the order word of* $P_1, \ldots, P_d$, *and* $(p_m)_{m=0}^{\infty}$ *be the positive value sequence of* $\bigcup_{i=1}^{d} P_i$. *Then, there are* $N \in \mathbb{N}$ *and* $\Sigma \subseteq \{0,1\}^d \times \{0, \ldots, M-1\}$ *such that*

1. $\big((\beta_m, p_m \bmod M)\big)_{m=N}^{\infty} \in \Sigma^{\omega}$ *is toric if* $\mu_-^{(i)} > 0$ *for all* $1 \leq i \leq d$;
2. $\big((\beta_m, p_m \bmod M)\big)_{m=N}^{\infty} \in \Sigma^{\omega}$ *is disjunctive if* $\mu_-^{(i)} \leq 0$ *for all* $1 \leq i \leq d$.

Thus, we split the set of non-degenerate, simple LRS $(u_n)_{n=0}^{\infty}$ such that $\{u_n : n \in \mathbb{N}\}$ into two different classes that we expect to have a very distinct structure while combining multiple such LRS preserves these structures. However, we remain silent on what we expect to occur when some of $(u_n^{(1)})_{n=0}^{\infty}, \ldots, (u_n^{(d)})_{n=0}^{\infty}$ satisfy Condition (1) and others satisfy Condition (2) of Conjecture 4.7.1.

Chapter 5

Presburger arithmetic expanded with multiple powers

5.1 Introduction and main results

In this chapter, we study the decidability of Presburger arithmetic with predicates derived from multiple linear recurrence sequences.

Recall that Presburger arithmetic is the first-order theory of the integers with addition and order, i.e., the first-order theory of the structure $\langle \mathbb{N}; 0, 1, <, + \rangle$. Presburger first established its decidability in 1929 via a quantifier-elimination procedure [128]; yet Presburger arithmetic remains to this day a topic of active research owing, among others, to its deep connections to automata theory and formal languages (see, e.g., the survey [74]) as well as symbolic dynamics and combinatorics on words (see, e.g., [141]).

Another rich line of inquiry has consisted in investigating *expansions* of Presburger arithmetic, i.e., theories obtained by augmenting Presburger arithmetic with particular predicates or functions. Here one must proceed with care: adding, for example, the multiplication function $\times : \mathbb{Z}^2 \to \mathbb{Z}$ (or even simply the 'squaring' function, from which multiplication is easily recovered) to Presburger arithmetic immediately results in undecidability, thanks to Gödel's incompleteness theorem [70]. Even the existential fragment of the first-order theory of $\langle \mathbb{Z}; 0, 1, <, +, \times \rangle$ is undecidable, as shown by Matiyasevich in his negative solution of Hilbert's 10th problem (see [114]). Nevertheless, many decidable expansions of Presburger arithmetic have been discovered and studied (see, for instance, the survey [29]). Decidability is usually established in one of two ways: either via quantifier elimination, following Presburger's original approach, or through automata-theoretic means, where integers are encoded in a given base as strings of digits and are processed by automata.

Before giving examples of such expansions, we introduce some notation. For a fixed integer $\alpha \geq 2$, as in the previous chapter, we let $\alpha^{\mathbb{N}}$ denote the set $\{\alpha^n : n \in \mathbb{N}\}$ and let α^x denote the $\mathbb{N} \to \mathbb{N}$-function $n \mapsto \alpha^n$. We also write $V_\alpha(n)$ to represent the function taking n to the largest power of α that divides n (thus, for example, $V_2(24) = 8$). Moreover, if $P_1, \ldots, P_\ell$ are predicates and $f_1, \ldots, f_m$ functions, we denote the first-order theory of $\langle \mathbb{Z}; 0, 1, <, +, P_1, \ldots, P_\ell, f_1, \ldots, f_m \rangle$ as $\mathcal{PA}(P_1, \ldots, P_\ell, f_1, \ldots, f_m)$.

Using an automata-theoretic construction, Büchi showed that, for any α, the theory $\mathcal{PA}(V_\alpha)$ is decidable [42].[1] Semenov used quantifier elimination to show that, for any 'effectively sparse' predicate $P \subset \mathbb{Z}$, the first-order theory of $\langle \mathbb{Z}; 0, 1, <, +, P \rangle$ is decidable. Examples of sparse predicates include the sets of powers $\alpha^{\mathbb{N}}$ as well as the set of factorial numbers $\{n! : n \in \mathbb{N}\}$.[2]

As in Chapter 4, we are interested in adding multiple such predicates simultaneously. Unfortunately, the results are far less positive. Villemaire [157] proved that, for multiplicatively independent α and β, the theory $\mathcal{PA}(V_\alpha, V_\beta)$ is undecidable by encoding multiplication, and in 1997, Bès [28] achieved the same for $\mathcal{PA}(V_\alpha, \beta^{\mathbb{N}})$. Decidability was open for the even further restricted case $\mathcal{PA}(\alpha^{\mathbb{N}}, \beta^{\mathbb{N}})$ for 25 years, for which Hieronymi and Schulz recently established undecidability [78].

Automata-theoretic techniques perform well when all numbers in play can be represented over a common base. But unfortunately, for multiplicatively independent α and β (such as 2 and 3), this is not the case: powers of 2, for example, can be easily be described using their base-2 expansion but not in their base-3 expansion. For example, Erdős [63] conjectured that every large enough power of 2 contains a 2 in its base-3 expansion, and this conjecture is still open. Multiplicatively independent power predicates enable one to formulate non-trivial number-theoretic assertions about integers, such as that a fixed number $N \geq 0$, the equation $|2^n - 3^m| < N$ has only finitely many solutions $(m, n) \in \mathbb{N}^2$. Such an assertion can already be formulated in $\mathcal{PA}(2^{\mathbb{N}}, 3^{\mathbb{N}})$.

Hieronymi and Schulz's undecidability result is quite intricate. Often, when proving undecidability in such theories, one encodes multiplication. Unfortunately, multiplication cannot be defined in this theory [138]. The undecidability construction in [78] makes use of three quantifier alternations (i.e., four blocks of quantifiers of alternating polarity), and so the $\forall\exists\forall\exists$-fragment of $\mathcal{PA}(2^{\mathbb{N}}, 3^{\mathbb{N}})$ is already undecidable. This naturally raises the question of whether weaker fragments might be decidable.

[1] Thus, the first-order theory of $\mathcal{PA}(\alpha^{\mathbb{N}}\rangle)$ is decidable as well as the set of powers of α can be defined in $\mathcal{PA}(0, 1, <, +, V_\alpha)$ using that x is a power of α if and only if $V_\alpha(x) = x$

[2] The complexity of expansions of Presburger arithmetic by a power predicate $\alpha^{\mathbb{N}}$ or a powering function α^x was very recently investigated [22, 49].

In [78, Section 5], Hieronymi and Schulz conjecture that the *existential* fragment of $\mathcal{PA}(\alpha^{\mathbb{N}}, \beta^{\mathbb{N}})$ is decidable subject to certain number-theoretic effectiveness assumptions.

Main results

The main result of this chapter is that the existential fragment of $\mathcal{PA}(\alpha^{\mathbb{N}}, \beta^{\mathbb{N}})$ is indeed decidable.

Theorem 5.1.1. *There is an algorithm that, given natural numbers $\alpha, \beta > 1$ together with an existential formula φ of $\mathcal{PA}(\alpha^{\mathbb{N}}, \beta^{\mathbb{N}})$, determines whether φ is true or not.*

As automata-theoretic methods are seemingly insufficient, our methods depend on number-theoretical tools like Baker's theorem on linear forms in logarithms (Theorem 1.1.11) and Kronecker's theorem on Diophantine approximation (Theorem 1.1.14) in a manner similar to [27, 148].

Secondly, we provide a shorter proof of Hieronymi and Schulz's undecidability result, requiring only two quantifier alternations (rather than three). We use the same approach but reduce from the Halting problem for 2-counter machines as opposed to the Halting problem for Turing machines, which results in a simpler construction.

Theorem 5.1.2. *$\exists\forall\exists$-fragment of $\mathcal{PA}(\alpha^{\mathbb{N}}, \beta^{\mathbb{N}})$ is undecidable for multiplicatively independent α and β.*

For any multiplicatively independent α, β, the decidability of $\mathcal{PA}(\alpha^{\mathbb{N}}, \beta^{\mathbb{N}})$ remains open only for formulas containing exactly two alternating blocks of quantifiers.

Finally, we also investigate the existential fragment of $\mathcal{PA}(\alpha^x, \beta^x)$, in which the power predicates have been replaced by powering functions.[3] We have not been able to establish either decidability or undecidability; however, we prove the following by way of hardness.

Theorem 5.1.3. *Let $\alpha, \beta > 1$ be multiplicatively independent integers. Write $(A_n)_{n=0}^{\infty}$ for the base-β expansion of $\log_\beta(\alpha)$ and $(B_n)_{n=0}^{\infty}$ for the base-α expansion of $\log_\alpha(\beta)$. Suppose that the existential fragment of $\mathcal{PA}(\alpha^x, \beta^x)$ is decidable. Then the following are in turn decidable:*

(A) Whether a given pattern appears in $(A_n)_{n=0}^{\infty}$.

[3] To remain within the realm of integers, we set $\alpha^n = \beta^n = 0$ for all $n < 0$

(B) Whether a given pattern appears at some index simultaneously in $(A_n)_{n=0}^{\infty}$ *and* $(B_n)_{n=0}^{\infty}$.

(C) Whether a given pattern appears in $(A_{\alpha^n})_{n=0}^{\infty}$.

To place Theorem 5.1.3 in context, consider the case of $\alpha = 2$ and $\beta = 3$. The constant $\log_3(2)$ is a transcendental number that is widely conjectured to be normal (and thus in base 3, every length-ℓ pattern should appear within $(A_n)_{n=0}^{\infty}$ with density $3^{-\ell}$). A fortiori, this would entail that the answer to the first query is always positive. However, normality on its own is not sufficient to settle either of the other two queries.

Organization of the chapter

We will use the majority of this chapter to prove Theorem 5.1.1 in Sections 5.2–5.4. The other two results, Theorems 5.1.2 and 5.1.3 have relatively short and straightforward proofs which can be found in Sections 5.6 and 5.5, respectively. Therefore, we only give an overview of our proof of Theorem 5.1.1.

Recall that our central problem is to determine, given $\alpha, \beta \in \mathbb{N}_{\geq 2}$, whether an existential formula φ in the language of $\mathcal{PA}(\alpha^{\mathbb{N}}, \beta^{\mathbb{N}})$ holds. As the first step in our decidability proof (Theorem 5.1.1), we will prove we can assume α and β are multiplicatively independent in Lemma 5.2.1. Then, in Section 5.2, we will reduce Theorem 5.1.1 to the following problem.

Problem 5.1.4. Given multiplicatively independent $\alpha, \beta \in \mathbb{N}_{\geq 2}$, $z_1, \ldots, z_\ell \in \{\alpha, \beta\}$, $r, s \geq 0$, $A \in \mathbb{Z}^{r \times \ell}$, $\mathbf{b} \in \mathbb{Z}^r$, $C \in \mathbb{Z}^{s \times \ell}$, and $\mathbf{d} \in \mathbb{Z}^s$, determine whether there exists $\mathbf{z} = (z_1^{n_1}, \ldots, z_\ell^{n_\ell})$ such that $A\mathbf{z} > \mathbf{b}$ and $C\mathbf{z} = \mathbf{d}$.

Lemma 5.1.5. *Let* $\alpha, \beta \in \mathbb{N}_{\geq 2}$ *be multiplicatively independent. Then deciding the existential fragment of* $\mathcal{PA}(\alpha^{\mathbb{N}}, \beta^{\mathbb{N}})$ *reduces to Problem 5.1.4.*

We approach Problem 5.1.4 by first studying how to solve systems of the form $C\mathbf{z} = \mathbf{d}$, i.e., the case where there are no inequalities. The following definition captures the structure of solutions of such systems.

Definition 5.1.6. A set $X \subseteq \mathbb{N}^\ell$ belongs to the class $\mathfrak{A}$ if it can be written in the form

$$X = \bigcup_{i \in I} \bigcap_{j \in J_i} X_j \tag{5.1}$$

where I and J_i for every $i \in I$ are finite, and each X_j is either of the form

$$X_j = \{(n_1, \ldots, n_\ell) \in \mathbb{N}^\ell \colon n_{\mu(j)} = n_{\sigma(j)} + c_j\} \tag{5.2}$$

or of the form

$$X_j = \{(n_1, \ldots, n_\ell) \in \mathbb{N}^\ell \colon n_{\xi(j)} = b_j\} \tag{5.3}$$

where $1 \leq \xi(j), \mu(j), \sigma(j) \leq \ell$ and $b_j, c_j \in \mathbb{N}$.

The sets belonging to $\mathfrak{A}$ are semilinear. Observe that every finite subset of $\mathbb{N}^\ell$ belongs to $\mathfrak{A}$, and the class $\mathfrak{A}$ is closed under finite unions and intersections. In Section 5.3, we will prove the following structure and effectiveness result about the system $C\mathbf{z} = \mathbf{d}$. Our main tool is Baker's theorem on linear forms in logarithms, which we employ to solve Diophantine equations where the unknowns appear in the exponent position.

Theorem 5.1.7. *Let $\alpha, \beta \in \mathbb{N}_{\geq 2}$ be multiplicatively independent and $z_1, \ldots, z_\ell \in \{\alpha, \beta\}$ for some $\ell \geq 1$. Further let $s \geq 1$, $C \in \mathbb{Z}^{s \times \ell}$, $\mathbf{d} \in \mathbb{Z}^s$, and $\mathcal{S} \subseteq \mathbb{N}^\ell$ be the set of solutions of $C\mathbf{z} = \mathbf{d}$, where $\mathbf{z} = (z_1^{n_1}, \ldots, z_\ell^{n_\ell})$. Then $\mathcal{S} \in \mathfrak{A}$. Moreover, a representation of $\mathcal{S}$ in the form (5.1) can be effectively computed, with the additional property that $z_{\mu(j)} = z_{\sigma(j)}$ for every X_j of the form (5.2).*

When proving Theorem 5.1.7, because the class $\mathfrak{A}$ is closed under intersections, it suffices to consider a single equality

$$c_1 z_1^{n_1} + \cdots + c_\ell z_\ell^{n_\ell} = d \tag{5.4}$$

where $c_1, \ldots, c_\ell \in \mathbb{Z}_{\neq 0}$, $d \in \mathbb{Z}$, $\alpha, \beta \geq 2$ are multiplicatively independent, and $z_1, \ldots, z_\ell \in \{\alpha, \beta\}$. We will show that the set $\mathcal{S}$ of solutions of (5.4) belongs to $\mathfrak{A}$ and has an effectively computable representation. We further stipulate that $z_i = \alpha$ and $z_j = \beta$ for some i, j, and that no proper subsum of the left-hand side of (5.4) is zero. In this case, we will show that the set of solutions is finite and can be effectively computed; see the proof of Theorem 5.3.2.[4] The idea is to apply Baker's theorem on linear forms in logarithms iteratively to bound the gaps between $n_1, \ldots, n_\ell$, which, in case $d \neq 0$, will yield an upper bound on all of $n_1, \ldots, n_\ell$. If $d = 0$, then we need an additional argument involving p-adic valuations. On the other hand, if $\bigcap_{j \in J_i} X_j$ is infinite for some i in the representation of $\mathcal{S}$ in the form (5.1), then some subsum of $c_1 z_1^{n_1} + \cdots + c_\ell z_\ell^{n_\ell}$ must be zero at infinitely many points $(n_1, \ldots, n_\ell)$.

[4]For convenience, we make additional assumptions on $z_1^{n_1}, \ldots, z_\ell^{n_\ell}$ in the statement of Theorem 5.3.2. A slightly modified proof can be given to show finiteness of solutions of (5.4) only assuming that both α, β appear among $z_1, \ldots, z_\ell$ and requiring that all proper subsums of $c_1 z_1^{n_1} + \cdots + c_\ell z_\ell^{n_\ell}$ be non-zero.

Example 5.1.8. Consider the equation

$$15 \cdot 3^{n_1} - 5 \cdot 3^{n_2} + 2^{n_3} = 8\,. \tag{5.5}$$

The only proper subsum of the left-hand side that can be zero is $15 \cdot 3^{n_1} - 5 \cdot 3^{n_2}$. Therefore, we have infinitely many solutions:

$$X := \left\{(n_1, n_2, n_3) \in \mathbb{N}^3 \colon n_2 = n_1 + 1 \wedge n_3 = 3\right\}.$$

Suppose no proper subsum is zero. We additionally stipulate that $3^{n_2} \geq 3^{n_1} \geq 2^{n_3}$. In this case, if $n_2 > n_1 + 1$, then the summand $5 \cdot 3^{n_2}$ becomes too large in magnitude: we have that $5 \cdot 3^{n_2} \geq 45 \cdot 3^{n_1}, 45 \cdot 2^{n_3}$ and hence (5.5) cannot hold. Therefore, only the possibilities $n_2 = n_1$ and $n_2 = n_1 + 1$ remain. If we substitute $n_2 = n_1$ into (5.5), we obtain $10 \cdot 3^{n_2} + 2^{n_3} = 8$, which does not have a solution. The substitution $n_2 = n_1 + 1$, meanwhile, is not permitted as $15 \cdot 3^{n_1} - 5 \cdot 3^{n_2}$ becomes zero.

By the same argument, we can handle the case where $3^{n_1} \geq 3^{n_2} \geq 2^{n_3}$. The four remaining cases (e.g., $3^{n_1} \geq 2^{n_3} \geq 3^{n_2}$), however, require an iterated application of Baker's theorem on linear forms in logarithms as in Theorem 5.3.2 to bound the solutions. Checking all possible (n_1, n_2, n_3) up to this bound, we obtain that the set of all solutions of (5.5) is $\{(0, 3, 7), (1, 8, 15)\} \cup X$. □

Once we know how to solve systems of linear equations in powers of α and β, we discuss how we deal with inequalities. In Section 5.4, we argue as follows. Consider a system $A\mathbf{z} > \mathbf{b}$ and $C\mathbf{z} = \mathbf{d}$ as in the statement of Problem 5.1.4. Observing that $x > -c$ is equivalent to $x = -c + 1 \vee \cdots \vee x = 0 \vee x > 0$ for any variable x and positive integer c, we rewrite our system in the form

$$\bigvee_{k \in K} A_k\mathbf{z} > \mathbf{b}_k \wedge C_k\mathbf{z} = \mathbf{d}_k$$

where $A_k \in \mathbb{Z}^{r \times \ell}, C_k \in \mathbb{Z}^{s \times \ell}, \mathbf{b}_k \in \mathbb{Z}^r, \mathbf{d}_k \in \mathbb{Z}^s$ and, importantly, $\mathbf{b}_k \geq \mathbf{0}$ for all k. We can now solve each $A_k\mathbf{z} > \mathbf{b}_k \wedge C_k\mathbf{z} = \mathbf{d}_k$ separately. Let $\mathcal{S}_k$ denote the set of all $(n_1, \ldots, n_\ell) \in \mathbb{N}^\ell$ that satisfy $C_k\mathbf{z} = \mathbf{d}_k$. By Theorem 5.1.7, the set $\mathcal{S}_k$ is defined by equations of the form either $n_a = n_b + c$ or $n_a = c$, where $1 \leq a, b \leq \ell$, $c \in \mathbb{N}$, and $z_a = z_b$ in the former case. We can use each such equation as a substitution rule to eliminate the variable z_a; see the proof of Theorem 5.1.1 on Page 177 for the exact procedure. In the end, we construct $\widetilde{A}_k \in \mathbb{Z}^{r \times \ell}$ such that $A_k\mathbf{z} > \mathbf{b}_k \wedge C_k\mathbf{z} = \mathbf{d}_k$ has a solution if and only if $\widetilde{A}_k\mathbf{z} > \mathbf{b}_k$ has a solution.

It remains to show how to solve the system $\widetilde{A}_k\mathbf{z} > \mathbf{b}_k$. To do this, we first argue that $\widetilde{A}_k\mathbf{z} > \mathbf{b}_k$ has a solution if and only if $\widetilde{A}_k\mathbf{z} > \mathbf{0}$ has a solution. Next, using a

form of Fourier-Motzkin elimination, we reduce solving the latter system to solving systems of the form

$$\begin{cases} \frac{h_i(z_3^{n_3},\ldots,z_\ell^{n_\ell})}{z_2^{n_2}} < \frac{z_1^{n_1}}{z_2^{n_2}} - a < \frac{h_j(z_3^{n_3},\ldots,z_\ell^{n_\ell})}{z_2^{n_2}} \quad \text{for all } (i,j) \in I_- \times I_+ \\ z_1^{n_1}, z_2^{n_2} > z_3^{n_3} > \cdots > z_\ell^{n_\ell} \\ h_i(z_3^{n_3},\ldots,z_\ell^{n_\ell}) > 0 \quad \text{for all } i \in I \\ n_2 - n_3 > N \end{cases} \tag{5.6}$$

where I_-, I_+, I are finite sets of indices, each h_i is a $\mathbb{Q}$-linear form, $a \in \mathbb{Q}_{>0}$, and $N \in \mathbb{N}$. Our algorithm for solving the system (5.6) proceeds by first inductively solving the subsystem consisting of the inequalities $h_i(z_3^{n_3},\ldots,z_\ell^{n_\ell}) < h_j(z_3^{n_3},\ldots,z_\ell^{n_\ell})$ for all $(i,j) \in I_- \times I_+$, $z_3^{n_3} > \cdots > z_\ell^{n_\ell}$, and $h_i(z_3^{n_3},\ldots,z_\ell^{n_\ell}) > 0$ for $i \in I$. If no such solution exists, then (5.6) does not have a solution either. Otherwise, let $(m_3,\ldots,m_\ell)$ be a solution to the subsystem. In Section 5.4, we use arguments from Diophantine approximation to prove that in the latter case, the system (5.6) does always have a solution.

5.2 From formulas to systems of inequalities

To start, we can assume that α and β are multiplicatively independent.

Lemma 5.2.1. *Theorem 5.1.1 holds when α and β are multiplicatively dependent.*

Proof. As α and β are multiplicatively dependent, one can compute natural numbers D_α, D_β such that $\alpha^{D_\alpha} = \beta^{D_\beta}$. Let $\gamma = \alpha^{D_\alpha}$. Then, $\gamma^{\mathbb{N}}$ can be defined in $\mathcal{PA}(V_\gamma)$ (by $x \in \gamma^{\mathbb{N}}$ if and only if $V_\gamma(x) = x$), $\alpha^{\mathbb{N}}$ by $\bigcup_{i=0}^{D_\alpha - 1} \alpha^i \gamma^{\mathbb{N}}$, and $\beta^{\mathbb{N}}$ by $\bigcup_{i=0}^{D_\beta - 1} \beta^i \gamma^{\mathbb{N}}$. Thus, $\alpha^{\mathbb{N}}$ and $\beta^{\mathbb{N}}$ can be defined in $\mathcal{PA}(V_\gamma)$, which is decidable due to the results of Büchi [42]. Hence, the result follows as one can encode any formula in $\mathcal{PA}(\alpha^{\mathbb{N}}, \beta^{\mathbb{N}})$ thus in $\mathcal{PA}(V_\gamma)$. □

From now on, we assume that α and β are multiplicatively independent and continue by proving Theorem 5.1.5. Our main tool is the fact that semilinear sets have quantifier-free representations constructed from linear inequalities and divisibility constraints. We note that our reduction from the decision problem for the existential fragment of $\mathcal{PA}(\alpha^{\mathbb{N}}, \beta^{\mathbb{N}})$ to Problem 5.1.4 does not preserve α and β.

Proof of Theorem 5.1.5. Suppose we are given multiplicative independent integers $\alpha, \beta \geq 2$ and an existential formula $\exists \mathbf{z} \colon \varphi(\mathbf{z})$ in the language of $\mathcal{PA}(\alpha^{\mathbb{N}}, \beta^{\mathbb{N}})$, where φ is quantifier-free and $\mathbf{z}$ is a collection of variables. We start by applying a sequence

of transformations to $\exists \mathbf{z}\colon \varphi(\mathbf{z})$. Therefore, every atomic formula in the language of $\mathcal{PA}(\alpha^{\mathbb{N}}, \beta^{\mathbb{N}})$ is equivalent to either $t > 0$, $t = 0$, $t \in \alpha^{\mathbb{N}}$, or $t \in \beta^{\mathbb{N}}$ for an integer linear combination of integer constants and variables t. For a term t and $\gamma \in \{\alpha, \beta\}$, we can rewrite the formula $\neg(t \in \gamma^{\mathbb{N}})$ as

$$t < 1 \vee \exists x\colon x \in \gamma^{\mathbb{N}} \wedge x < t < \underbrace{x + \cdots + x}_{\gamma \text{ times}} .$$

Since $\neg(t > 0)$ and $\neg(t = 0)$ are equivalent to $t < 0 \vee t = 0$ and $t > 0 \vee t < 0$ respectively, we can construct a formula $\exists \mathbf{x}\colon \widehat{\varphi}(\mathbf{x})$ equivalent to $\exists \mathbf{z}\colon \varphi(\mathbf{z})$ in which the negation symbol does not occur. We can also rewrite $t \in \gamma^{\mathbb{N}}$ as $y = t \wedge y \in \gamma^{\mathbb{N}}$, where y is a fresh variable. Therefore, we can construct a formula

$$\widetilde{\varphi}(\mathbf{y}, \mathbf{x}) := \bigvee_{e \in E} \bigwedge_{j \in J_e} \mu_j(\mathbf{y}, \mathbf{x})$$

not containing the negation symbol, where $\mathbf{y}$ denotes a collection $y_1, \ldots, y_\ell$ of fresh variables, with the following properties.

- $\exists \mathbf{z} \in \mathbb{Z}^k\colon \varphi(\mathbf{z}) \iff \exists \mathbf{y} \in \mathbb{Z}^\ell, \mathbf{x} \in \mathbb{Z}^k\colon \widetilde{\varphi}(\mathbf{y}, \mathbf{x})$.
- For each y_i, there exists unique $\gamma_i \in \{\alpha, \beta\}$ such that $y_i \in \gamma_i^{\mathbb{N}}$ is a sub-formula of $\widetilde{\varphi}$.
- Each $\mu_j(\mathbf{y}, \mathbf{x})$ is an atomic formula either of the form $t(\mathbf{y}, \mathbf{x}) \sim 0$ for a term $t(\mathbf{y}, \mathbf{x})$ and $\sim \in \{>, =\}$, or of the form $y_i \in \gamma_i^{\mathbb{N}}$ for some i.

Next, write each $\bigwedge_{j \in J_e} \mu_j(\mathbf{y}, \mathbf{x})$ in the form

$$\bigwedge_{j \in A_e} y_{\sigma(j)} \in \gamma_{\sigma(j)}^{\mathbb{N}} \wedge \bigwedge_{j \in B_e} t_j(\mathbf{y}, \mathbf{x}) \sim_j 0$$

where $\sigma(j) \in \{1, \ldots, \ell\}$ and $\sim_j \in \{>, =\}$. We can then write $\exists \mathbf{y} \in \mathbb{Z}^\ell, \mathbf{x} \in \mathbb{Z}^k\colon \widetilde{\varphi}(\mathbf{y}, \mathbf{x})$ equivalently as

$$\bigvee_{e \in E} \left(\exists \mathbf{y} \in \mathbb{Z}^\ell\colon \bigwedge_{j \in A_e} y_{\sigma(j)} \in \gamma_{\sigma(j)}^{\mathbb{N}} \wedge \exists \mathbf{x} \in \mathbb{Z}^k\colon \bigwedge_{j \in B_e} t_j(\mathbf{y}, \mathbf{x}) \sim_j 0 \right). \tag{5.7}$$

For $e \in E$, let S_e be the semilinear set consisting of all $\mathbf{y} \in \mathbb{Z}^\ell$ such that $\exists \mathbf{x} \in \mathbb{Z}^k\colon \bigwedge_{j \in B_e} t_j(\mathbf{y}, \mathbf{x}) \sim_j 0$ holds. Observe that each S_e is computable as it is defined in Presburger arithmetic. Setting $z_i = \gamma_i$ and $y_i = z_i^{n_i}$ for $1 \leq i \leq \ell$, we rewrite (5.7) as

$$\bigvee_{e \in E} \exists n_1, \ldots, n_\ell \in \mathbb{N}\colon (z_1^{n_1}, \ldots, z_\ell^{n_\ell}) \in S_e \,,$$

which is equivalent to

$$\exists n_1, \ldots, n_\ell \in \mathbb{N}\colon (z_1^{n_1}, \ldots, z_\ell^{n_\ell}) \in S$$

for the semilinear set $S = \bigcup_{e \in E} S_e$.

Recall from Section 1.3.2 that each semilinear set has a representation in the form (1.5). For $x, y, r \geq 0$ and $\lambda, D \geq 1$, the equivalence $x + y \equiv r \pmod{D}$ can be rewritten as

$$\bigvee_{\substack{0 \leq r_1, r_2 < D \\ r_1 + r_2 \equiv r \pmod{D}}} x \equiv r_1 \pmod{D} \wedge y \equiv r_2 \pmod{D}$$

and $x \equiv r \pmod{D}$ is equivalent to $\bigvee_{k=0}^{\lambda-1} x \equiv r + kD \pmod{\lambda D}$. Hence, we can construct $D \geq 1$ and a representation of S of the form

$$\bigvee_{p \in P} \Bigg(\bigwedge_{i=1}^{\ell} x_i \equiv r_{i,p} \pmod{D} \wedge \bigwedge_{s \in S_p} h_s(x_1, \ldots, x_d) \sim_s b_s \Bigg) \tag{5.8}$$

where each $r_{i,p} \geq 0$, h_s is a $\mathbb{Z}$-linear form, $b_s \in \mathbb{Z}$, and $\sim_s \in \{>, =\}$. Write $\widetilde{S}_p$ for the set defined by $p \in P$ in (5.8), so that $S = \bigcup_{p \in P} \widetilde{S}_p$. It suffices to reduce deciding $\exists n_1, \ldots, n_\ell \in \mathbb{N}\colon (z_1^{n_1}, \ldots, z_\ell^{n_\ell}) \in \widetilde{S}_p$ to Problem 5.1.4. To do this, apply Lemma 1.2.1 on the LRS $(z_i^n)_{n=0}^{\infty}$ to deduce that the sequence $(z_i^n \bmod D)_{n=0}^{\infty}$ is ultimately periodic with a computable period D_{z_i}. Thus, for all $1 \leq i \leq \ell$ and $p \in P$, the set $\{n_i \in \mathbb{N}\colon z_i^{n_i} \equiv r_{i,p} \pmod{D}\}$ is the union of finitely many points and finitely many arithmetic progressions with period D_{z_i} that can be computed explicitly. Therefore, $\exists n_1, \ldots, n_\ell\colon (z_1^{n_1}, \ldots, z_\ell^{n_\ell}) \in \widetilde{S}_p$ can be equivalently expressed as a disjunction of formulas of the form

$$\exists m_1, \ldots, m_\ell \in \mathbb{N}\colon \bigwedge_{s \in S_p} h_s\big(z_1^{t_{s,1}}, \ldots, z_\ell^{t_{s,\ell}}\big) \sim_s b_s$$

where for all s and ℓ, we have $t_{s,\ell} = a$ or $t_{s,\ell} = a + m_i \cdot D_{z_i}$ for a constant $a \in \mathbb{N}$. It remains to observe that $z_i^{a_i + m_i \cdot D_{z_i}} = z_i^{a_i} \left(z_i^{D_{z_i}}\right)^{m_i}$. Therefore, we have reduced deciding the truth value of $\exists \mathbf{x}\colon \varphi(\mathbf{x})$ to solving systems of (in)equalities in powers of $\gamma_\alpha = \alpha^{D_\alpha}$ and $\gamma_\beta = \beta^{D_\beta}$. Note that $\gamma_\alpha, \gamma_\beta$ are also multiplicatively independent. The statement follows. □

5.3 Solving Diophantine equations

We now discuss solutions of systems of affine Diophantine equations in powers of α and β. This is the first step towards showing the decidability of Problem 5.1.4. Our goal in this section is to prove the following theorem.

Theorem 5.1.7. *Let $\alpha, \beta \in \mathbb{N}_{\geq 2}$ be multiplicatively independent and $z_1, \ldots, z_\ell \in \{\alpha, \beta\}$ for some $\ell \geq 1$. Further let $s \geq 1$, $C \in \mathbb{Z}^{s \times \ell}$, $\mathbf{d} \in \mathbb{Z}^s$, and $\mathcal{S} \subseteq \mathbb{N}^\ell$ be the set of solutions of $C\mathbf{z} = \mathbf{d}$, where $\mathbf{z} = (z_1^{n_1}, \ldots, z_\ell^{n_\ell})$. Then $\mathcal{S} \in \mathfrak{A}$. Moreover, a representation of $\mathcal{S}$ in the form (5.1) can be effectively computed, with the additional property that $z_{\mu(j)} = z_{\sigma(j)}$ for every X_j of the form (5.2).*

Proof. As the class $\mathfrak{A}$ is closed under intersections, it is sufficient to show that the solution set of a single equation $c_1 z_1^{n_1} + \cdots + c_\ell z_\ell^{n_\ell}$ belongs to $\mathfrak{A}$ and can be effectively computed. Moreover, we can assume that all c_i are non-zero. Then the theorem holds when $\ell = 0$, and if $\ell = 1$, then $c_1 z_1^{n_1} = d$ has at most one solution that can be computed. Thus, we can assume that $\ell \geq 2$.

To prove the theorem above, we want to assume that no proper subsum of $c_1 z_1^{n_1} + \cdots + c_\ell z_\ell^{n_\ell}$ vanishes. We can make this assumption by partitioning $\{1, \ldots, \ell\}$ into non-empty subsets $A_1, \ldots, A_r$ such that for a single subset A_j we have that

$$\sum_{i \in A_j} c_i z_i^{n_i} = d \tag{5.9}$$

and for all other subsets A_j, we have that

$$\sum_{i \in A_j} c_i z_i^{n_i} = 0\,, \tag{5.10}$$

where in all cases, we can assume that no proper subsum vanishes. Note that any solution to $c_1 z_1^{n_1} + \cdots + c_\ell z_\ell^{n_\ell} = d$ occurs in such a partition. As the class $\mathfrak{A}$ is closed under intersections and unions, it follows that finding the solutions to equations without vanishing proper subsums is sufficient. Thus, without loss of generality, we can assume that we need to find the solutions to $c_1 z_1^{n_1} + \cdots + c_\ell z_\ell^{n_\ell} = d$ for integers d where no proper subsum vanishes.

Next, we will consider the $\ell!$ ways in which we can order $z_1^{n_1}, \ldots, z_\ell^{n_\ell}$ using the permutations $\sigma : \{1, \ldots, \ell\} \to \{1, \ldots, \ell\}$. We will show that the set $(n_1, \ldots, n_\ell)$ such that

$$\begin{cases} c_1 z_1^{n_1} + \cdots + c_\ell z_\ell^{n_\ell} = d \\ \forall\, \emptyset \subsetneq I \subsetneq \{1, \ldots, \ell\} \colon \sum_{i \in I} c_i z_i^{n_i} \neq 0 \\ z_{\sigma(1)}^{n_{\sigma(1)}} \geq \cdots \geq z_{\sigma(\ell)}^{n_{\sigma(\ell)}} \end{cases}$$

all belong to the class $\mathfrak{A}$ and can be effectively computed. To prove this by reordering the variables, we can assume that σ is the identity. That is, the set $\mathcal{S}$ of tuples $(n_1, \ldots, n_\ell)$ such that

$$\begin{cases} c_1 z_1^{n_1} + \cdots + c_\ell z_\ell^{n_\ell} = d \\ \forall\, \emptyset \subsetneq I \subsetneq \{1, \ldots, \ell\} \colon \sum_{i \in I} c_i z_i^{n_i} \neq 0 \\ z_1^{n_1} \geq \cdots \geq z_\ell^{n_\ell} \end{cases} \tag{5.11}$$

is in $\mathfrak{A}$ and can be represented as such.

Claim 5.3.1. *Without loss of generality, $z_1 \neq z_2$.*

Proof. After renaming the variables α and β when needed, assume that $\alpha = z_1 = z_2$ and that (5.11) holds. Then,

$$|c_1|\alpha^{n_1} \leq |d| + \sum_{i=2}^{\ell} |c_i|\alpha^{n_i} \leq \Big(|d| + \sum_{i=2}^{\ell} |c_i|\Big)\alpha^{n_2}\,,$$

set $N = \log_\alpha\left(\frac{|d|+\sum_{i=2}^{\ell}|c_i|)}{|c_1|}\right)$ such that we have to have that $0 \leq n_1 - n_2 \leq N$ (we use that $c_1 \neq 0$). Let $\widetilde{\mathcal{S}} = \bigcup_{k=0}^{N} \widetilde{\mathcal{S}}_k$, where each $\widetilde{\mathcal{S}}_k$ is the set of all solutions of

$$\begin{cases} (c_1\alpha^k + c_2)\alpha^{n_2} + c_3\alpha^{n_3} + \cdots + c_\ell\alpha^{n_\ell} = d \\ \alpha^{n_2} \geq \cdots \geq \alpha^{n_\ell}\,, \end{cases}$$

where we have eliminated n_1 by setting $n_1 = n_2 + k$. If $\ell > 2$ and $c_1\alpha^k + c_2 = 0$, a proper subsum vanishes, which is not allowed, and we set $\widetilde{\mathcal{S}}_k = \emptyset$. Otherwise, we inductively construct a representation of $\widetilde{\mathcal{S}}_k$, by finding a representation of the system above and adding the condition $n_1 = n_2 + k$. □

The remainder of the proof follows from the next theorem, whose proof is below.

Theorem 5.3.2. *Let $\alpha, \beta \in \mathbb{N}_{\geq 2}$ be multiplicatively independent, $\ell \geq 2$, $z_1, \ldots, z_\ell \in \{\alpha, \beta\}$ with $z_1 = \alpha$ and $z_2 = \beta$, $c_1, \ldots, c_\ell \in \mathbb{Z}_{\neq 0}$, and $d \in \mathbb{Z}$. Let $\mathcal{S}$ denote the set of all $(n_1, \ldots, n_\ell) \in \mathbb{N}^\ell$ satisfying all of the following.*

(a) $c_1 z_1^{n_1} + \cdots + c_\ell z_\ell^{n_\ell} = d$;

(b) $z_1^{n_1}, z_2^{n_2} \geq z_3^{n_3} \geq \cdots \geq z_\ell^{n_\ell}$;

(c) $\sum_{i\in I} c_i z_i^{n_i} \neq 0$ *for every non-empty proper subset I of $\{1, \ldots, \ell\}$.*

Define $\mu(j)$ to be 1 if $z_j = \alpha$ and $\mu(j) = 2$ if $z_j = \beta$. We have the following.

(i) We can compute $\xi_1, \xi_2 \in \mathbb{Q}$ such that $n_1 \geq \frac{\log(\beta)}{\log(\alpha)} n_2 - \xi_1$ and $n_2 \geq \frac{\log(\alpha)}{\log(\beta)} n_1 - \xi_2$.

(ii) There exist effectively computable polynomials $p_1, \ldots, p_\ell \in \mathbb{Q}[x, y]$ such that

$$n_{\mu(j)} - n_j < p_j\big(\log(1 + n_1), \log(1 + n_2)\big)$$

for all $(n_1, \ldots, n_\ell) \in \mathcal{S}$ and $1 \leq j \leq \ell$.

(iii) The set $\mathcal{S}$ is finite and can be effectively computed.

As finite sets are in $\mathfrak{A}$, Theorem 5.3.2 completes the proof of Theorem 5.1.7. $\square$

Next, we show how to prove Theorem 5.1.7: under certain assumptions, equations involving powers of both α and β have only finitely many solutions that can be computed using Baker's theorem on linear forms in logarithms in an iterative fashion.

Proof of Theorem 5.3.2. Observe that $n_j \leq n_{\mu(j)}$ for all $j \geq 1$ by (b).

Proof of (i). Together (a) and (b) imply that, for all $(n_1, \ldots, n_\ell) \in \mathcal{S}$,

$$|c_1| z_1^{n_1} \leq \left(|d| + \sum_{i=2}^{\ell} |c_i| \right) z_2^{n_2} .$$

Taking logarithms and dividing by $\log(\beta) = \log(z_2)$ gives

$$n_2 \geq \frac{\log(\alpha)}{\log(\beta)} n_1 - \frac{\log\left(|d| + \sum_{i=2}^{\ell} |c_i|\right) - \log|c_1|}{\log(\beta)} .$$

This lets us find ξ_2. To compute ξ_1, observe that by (a) and (b),

$$|c_1| z_1^{n_1} \leq \left(|d| + |c_1| + \sum_{i=2}^{\ell} |c_i| \right) z_2^{n_2}$$

and proceed similarly.

Proof of (ii). By finite induction. Note that we can choose $p_1(x, y), p_2(x, y) = 1$. Suppose, therefore, that $p_1, \ldots, p_j$ have already been computed for some $j \geq 2$. By swapping z_1 and z_2 if necessary, we can assume $z_{j+1} = z_1 = \alpha$ as the roles z_1 and z_2 in the statement of our theorem are completely symmetrical.

For $(n_1, \ldots, n_\ell) \in \mathcal{S}$ define

$$\begin{aligned} X &:= - \sum_{\substack{1 \leq i \leq j \\ z_i = \alpha}} c_i \alpha^{n_i - n_1}, \\ Y &:= \sum_{\substack{1 \leq i \leq j \\ z_i = \beta}} c_i \beta^{n_i - n_2}, \\ \Lambda &:= \alpha^{-n_1} \beta^{n_2} X^{-1} Y - 1 = \left(- \sum_{\substack{1 \leq i \leq j \\ z_i = \alpha}} c_i \alpha^{n_i} \right)^{-1} \cdot \sum_{\substack{1 \leq i \leq j \\ z_i = \beta}} c_i \beta^{n_i} - 1 . \end{aligned}$$

By (c), X is non-zero and hence X^{-1} is well-defined, and similarly, Y and Λ are non-zero. Next, observe that (a) can be written as

$$\Lambda = \left(\sum_{i=j+1}^{\ell} c_i z_i^{n_i} - d \right) \cdot \left(\sum_{\substack{1 \leq i \leq j \\ z_i = \alpha}} c_i \alpha^{n_i} \right)^{-1} . \tag{5.12}$$

We will estimate the magnitude of terms on both sides of (5.12), starting with the left-hand side. Recall the definition and the properties of the absolute logarithmic Weil height $h(\cdot)$. We have that $h(X^{-1}Y) \leq h(X) + h(Y)$ and

$$h(X) \leq \log(j) + \sum_{\substack{1 \leq i \leq j \\ z_i = \alpha}} \log|c_i| + (n_1 - n_i)\log|\alpha|,$$
$$h(Y) \leq \log(j) + \sum_{\substack{1 \leq i \leq j \\ z_i = \beta}} \log|c_i| + (n_2 - n_i)\log|\beta|.$$

Therefore, using Corollary 1.1.12 we can compute $\kappa_1 \in \mathbb{Q}_{>0}$ such that

$$\log|\Lambda| > -\kappa_1 \cdot \big(1 + \log(1 + \max(1, n_1, n_2))\big) \cdot \max_{1 \leq i \leq j}(n_{\mu(i)} - n_i). \tag{5.13}$$

Applying the induction hypothesis, there exists computable $q \in \mathbb{Q}[x, y]$ such that

$$\log|\Lambda| > -\kappa_1 \cdot q\big(\log(1 + n_1), \log(1 + n_2)\big).$$

Next, consider the right-hand side of (5.12). Let a be the largest integer $1 \leq i \leq j$ such that $z_i = \alpha$. We have that

$$\left|\sum_{i=j+1}^{\ell} c_i z_i^{n_i} - d\right| \leq \left|\sum_{i=j+1}^{\ell} c_i - d\right| z_{j+1}^{n_{j+1}} = \kappa_2 \alpha^{n_{j+1}}$$

for some computable $\kappa_2 \in \mathbb{Z}_{>0}$ as $\alpha = z_{j+1}$. By (c), $\sum_{1 \leq i \leq j, z_i = \alpha} c_i \alpha^{n_i}$ is a non-zero integer multiple of α^{n_a} and so by the induction hypothesis,

$$\left|\sum_{\substack{1 \leq i \leq j \\ z_i = \alpha}} c_i \alpha^{n_i}\right| \geq \alpha^{n_a} > \alpha^{n_1 - r(\log(1+n_1), \log(1+n_2))}$$

for a computable polynomial $r \in \mathbb{Q}[x, y]$. Hence the magnitude of the right-hand side of (5.12) is bounded by $\kappa_2 \alpha^{r(\log(1+n_1), \log(1+n_2)) - n_1 + n_{j+1}}$, and a necessary condition for (5.12) to hold is that

$$-\kappa_1 \cdot q\big(\log(1 + n_1), \log(1 + n_2)\big) < \log\big(\kappa_2 \cdot \alpha^{r(\log(1+n_1), \log(1+n_2)) - n_1 + n_{j+1}}\big)$$

which is equivalent to

$$n_1 - n_{j+1} < \frac{\kappa_1 \cdot q\big(\log(1 + n_1), \log(1 + n_2)\big) - \log(\kappa_2)}{\log(\alpha)} + r\big(\log(1 + n_1), \log(1 + n_2)\big). \tag{5.14}$$

It remains to choose $p_{j+1} \in \mathbb{Q}[x, y]$ such that $p_{j+1}(\log(1 + n_1), \log(1 + n_2))$ is at least as large as the right-hand side of (5.14).

Proof of (iii). Since $\alpha \neq \beta$ by the multiplicative independence assumption, without loss of generality, we can assume that $\alpha < \beta$. (The roles of $\alpha = z_1$ and $\beta = z_2$ are symmetric, and we can swap them if necessary.) If $n_2 \geq n_1$, (i) gives that $n_2 \geq n_1 \geq \frac{\log(\beta)}{\log(\alpha)} n_2 - \xi_1$, effectively bounding n_2 (and thus n_1) as $\frac{\log(\beta)}{\log(\alpha)} > 1$. Thus, we can assume that $n_1 > n_2$. Call $\mathcal{S}_1$ the set of elements in $\mathcal{S}$ satisfying $n_1 > n_2$.

Case 1: Suppose $d \neq 0$. This case is very similar to the proof of (ii). Let

$$\begin{aligned} X &= -\sum_{\substack{1 \leq i \leq \ell \\ z_i = \alpha}} c_i \alpha^{n_i - n_1}, \\ Y &= \sum_{\substack{1 \leq i \leq \ell \\ z_i = \beta}} c_i \beta^{n_i - n_2}, \\ \Lambda &= \alpha^{-n_1} \beta^{n_2} \cdot X^{-1} Y - 1 \,. \end{aligned}$$

Then again X, Y, and Λ are non-zero, and we rewrite (a) in the form

$$\Lambda = -d \Big(\sum_{\substack{1 \leq i \leq \ell \\ z_i = \alpha}} c_i \alpha^{n_i} \Big)^{-1} \tag{5.15}$$

and bound the magnitude on both sides. Because $n_1 > n_2 \geq 0$ for all solutions in $\mathcal{S}$, application of Corollary 1.1.12 and (ii) yields,

$$\log |\Lambda| > -\kappa_2 p\big(\log(n_1)\big) \,,$$

where $\kappa_2 > 0$ and $p \in \mathbb{Q}[x]$ are computed effectively. It remains to compute an upper bound for the right-hand side. Let a be the largest integer $1 \leq i \leq \ell$ such that $z_i = \alpha$. Then,

$$\Big| \sum_{\substack{1 \leq i \leq \ell \\ z_i = \alpha}} c_i \alpha^{n_i} \Big| > \alpha^{n_a} > \alpha^{n_1 - f(\log(n_1))} \,,$$

where $f \in \mathbb{Q}[x]$ is computable. Hence, a necessary condition for $(n_1, \ldots, n_\ell) \in \mathcal{S}$ is

$$\kappa_2 p\big(\log(n_1)\big) > \big(n_1 - f(\log(n_1))\big) \log(\alpha) - \log |d| \,,$$

from which we can compute a bound on n_1. Once we bound n_1, a bound on the remaining variables can be computed using (b) and (i).

Case 2: Suppose $d = 0$. This case is trickier. We will need a lemma.

Lemma 5.3.3. *There exists a prime number $p \in \mathbb{N}$ such that $\nu_p(\beta) > 0$ and*

$$\frac{\log(\alpha)}{\log(\beta)} > \frac{\nu_p(\alpha)}{\nu_p(\beta)} \,.$$

Proof. If some prime p divides β but not α, the statement is immediate. Suppose, therefore, that α, β have exactly the same prime divisors $p_1, \ldots, p_k$. We have

$$\frac{\log(\alpha)}{\log(\beta)} = \frac{\nu_{p_1}(\alpha)\log(p_1) + \cdots + \nu_{p_k}(\alpha)\log(p_k)}{\nu_{p_1}(\beta)\log(p_1) + \cdots + \nu_{p_k}(\beta)\log(p_k)}\,.$$

By the multiplicative independence of α and β, we have that $\log(\alpha)/\log(\beta) \notin \mathbb{Q}$ and hence we have that $\nu_{p_i}(\alpha)/\nu_{p_i}(\beta) \neq \log(\alpha)/\log(\beta)$ for all $1 \leq i \leq k$. It follows that $\log(\alpha)/\log(\beta) > \nu_{p_i}(\alpha)/\nu_{p_i}(\beta)$ for some $1 \leq i \leq k$. □

To bound the elements of $\mathcal{S}_1$, let $a = \max\{i\colon z_i = \alpha\}$, $b = \max\{i\colon z_i = \beta\}$,

$$A = \sum_{\substack{1\leq i\leq \ell \\ z_i=\alpha}} c_i\alpha^{n_i}\,, \quad \text{and}$$
$$B = -\sum_{\substack{1\leq i\leq \ell \\ z_i=\beta}} c_i\beta^{n_i}\,.$$

Let p be a prime number as in Theorem 5.3.3. A necessary condition for (a) to hold is that $\nu_p(A) \geq \nu_p(B)$. Using the properties of the p-adic valuation, we have

$$\nu_p(B) \geq \min_{z_i=\beta}\big(\nu_p(c_i\beta^{n_i})\big) \geq \nu_p(\beta^{n_b}) = n_b \cdot \nu_p(\beta)\,.$$

Under the assumption $n_1 > n_2$, by (ii), there exists effectively computable $q_1 \in \mathbb{Q}[x]$ such that $n_b > n_2 - q_1(\log(n_1))/\nu_p(\beta)$. Hence,

$$\nu_p(B) > n_2\nu_p(\beta) - q_1\big(\log(n_1)\big)\,.$$

Meanwhile,

$$\begin{aligned}\nu_p(A) &= \nu_p(\alpha^{n_a}) + \nu_p\bigg(\sum_{\substack{1\leq i\leq \ell \\ z_i=\alpha}} c_i\alpha^{n_i-n_a}\bigg) \\ &\leq n_1\nu_p(\alpha) + \log_p\bigg|\sum_{\substack{1\leq i\leq \ell \\ z_i=\alpha}} c_i\alpha^{n_i-n_a}\bigg|\,.\end{aligned}$$

Applying (ii), we obtain that

$$\nu_p(A) \leq n_1\nu_p(\alpha) + q_2\big(\log(n_1)\big)$$

for an effectively computable $q_2 \in \mathbb{Q}[x]$. Thus, a necessary condition for (a) to hold is that

$$n_2\nu_p(\beta) - q_1\big(\log(n_1)\big) \leq n_1\nu_p(\alpha) + q_2\big(\log(n_1)\big)\,,$$

which is equivalent to

$$n_2 - n_1 \frac{\nu_p(\alpha)}{\nu_p(\beta)} \leq \frac{q_1\big(\log(n_1)\big) + q_2\big(\log(n_1)\big)}{\nu_p(\beta)}\,.$$

Applying (i), we obtain that

$$-\xi_2 + n_1 \left(\frac{\log(\alpha)}{\log(\beta)} - \frac{\nu_p(\alpha)}{\nu_p(\beta)}\right) \leq \frac{q_1\big(\log(n_1)\big) + q_2\big(\log(n_1)\big)}{\nu_p(\beta)}\,.$$

By construction of p, the left-hand side of the inequality above grows linearly in n_1 while the right-hand side grows poly-logarithmically. Hence, we can compute a bound on n_1, from which bounds on every n_i can be derived. □

5.4 Handling inequalities

In this section, we prove the decidability of Problem 5.1.4. This, in conjunction with Theorem 5.1.5, completes the proof of our main decidability result (Theorem 5.1.1). The following lemma is one of our main technical tools. In particular, it says that if $A\mathbf{z} > \mathbf{0}$ has a solution, then it has infinitely many solutions.

Lemma 5.4.1 (Pumping Lemma)**.** *Suppose we are given*

(a) $\mathbb{Q}$*-linear forms* $h_1, \ldots, h_r$ *in* $\ell \geq 1$ *variables,*

(b) multiplicatively independent $\alpha, \beta \in \mathbb{N}_{\geq 2}$*,*

(c) $z_1, \ldots, z_\ell$ *satisfying* $z_i \in \{\alpha, \beta\}$ *for all* i *and* $z_1 = \beta$*,*

(d) $m_1, \ldots, m_\ell \in \mathbb{N}$*, and*

(e) $\varepsilon \in \mathbb{Q}_{>0}$*.*

Write $J = \{j \in \{1, \ldots, r\} \colon h_j(z_1^{m_1}, \ldots, z_\ell^{m_\ell}) > 0\}$*. We can compute* $\mu, \delta \in \mathbb{Q}_{>0}$ *with the following property. Suppose* $n_1 > m_1$ *is such that there exists* $k \in \mathbb{N}$ *for which* $|\alpha^k/\beta^{n_1} - \mu| < \delta$*. Then there exist* $n_2, \ldots, n_\ell$ *such that for all* $1 \leq j \leq r$*,*

(i) if $j \in J$*, then* $h_j(z_1^{n_1}, \ldots, z_\ell^{n_\ell}) > 0$*, and*

(ii) $\left| \frac{h_j\left(z_1^{n_1}, \ldots, z_\ell^{n_\ell}\right)}{\beta^{n_1}} - \frac{h_j\left(z_1^{m_1}, \ldots, z_\ell^{m_\ell}\right)}{\beta^{m_1}} \right| < \varepsilon.$

In particular, there exist infinitely many n_1 *that can be extended to* $(n_1, \ldots, n_\ell)$ *satisfying (i) and (ii) for all* $1 \leq j \leq r$*.*

Proof. By re-ordering the numbers $z_2, \dots, z_\ell$, we can without loss of generality assume that $z_1, \dots, z_b = \beta$ and $z_{b+1}, \dots, z_\ell = \alpha$ for some $1 \le b \le \ell + 1$. For $1 \le j \le r$, write

$$h_j(x_1, \dots, x_\ell) = t_j(x_1, \dots, x_b) + s_j(x_{b+1}, \dots, x_\ell)\,,$$

where s_j, t_j are $\mathbb{Q}$-linear forms. Compute $\nu \in \mathbb{Q}_{>0}$ be such that

(A) $t_j(\beta^{m_1}, \dots, \beta^{m_b}) + c \cdot s_j(\alpha^{m_{b+1}}, \dots, \alpha^{m_\ell}) > 0$ whenever $c \in (1-\nu, 1+\nu)$ and $j \in J$, and

(B) $\nu \left| s_j(\alpha^{m_{b+1}}, \dots, \alpha^{m_\ell})/\beta^{m_1} \right| < \varepsilon$ for all $1 \le j \le r$.

Choose $\mu = \beta^{-m_1}$ and $\delta \in (0, \nu\beta^{-m_1}) \cap \mathbb{Q}$. It remains to argue the correctness of our choice of μ, δ.

Suppose $n_1 > m_1$ satisfies $|\alpha^k/\beta^{n_1} - \mu| < \delta$ for some $k \in \mathbb{N}$. Due to Theorem 1.1.14, there are infinitely many such n_1. Write $m_\alpha = k$ and $m_\beta = n_1 - m_1$. We have that

$$\left| \frac{\alpha^{m_\alpha}}{\beta^{m_\beta}} - 1 \right| = \beta^{m_1} \left| \frac{\alpha^k}{\beta^{n_1}} - \mu \right| < \beta^{m_1}\delta < \nu\,.$$

For $2 \le i \le b$ define $n_i = m_i + m_\beta$ and for $b+1 \le i \le \ell$, define $n_i = m_i + m_\alpha$. Then, for all $j \in J$,

$$\begin{aligned} \frac{1}{\beta^{m_\beta}} h_j(z_1^{n_1}, \dots, z_\ell^{m_\ell}) &= t_j(\beta^{m_1}, \dots, \beta^{m_b}) + \frac{\alpha^{m_\alpha}}{\beta^{m_\beta}} s_j(\alpha^{m_{b+1}}, \dots, \alpha^{m_\ell}) \\ &> 0\,, \end{aligned}$$

where the inequality follows from (A). This proves (i). To prove (ii), first observe that for $1 \le i \le b$, we have $z_i^{n_i}/z_1^{n_1} = 1$ and for $b+1 \le i \le \ell$, we have $z_i^{n_i}/z_1^{n_1} = \alpha^{m_\alpha}/\beta^{m_\beta}$. Hence,

$$\frac{t_j(\beta^{n_1}, \dots, \beta^{n_b})}{z_1^{n_1}} = \frac{t_j(\beta^{m_1}, \dots, \beta^{m_b})}{z_1^{m_1}}$$

for all $1 \le j \le r$. Therefore, for all j,

$$\begin{aligned} &\frac{h_j(z_1^{n_1}, \dots, z_\ell^{n_\ell})}{z_1^{n_1}} - \frac{h_j(z_1^{m_1}, \dots, z_\ell^{m_\ell})}{z_1^{m_1}} \\ &= \frac{s_j(\alpha^{n_{b+1}}, \dots, \alpha^{n_\ell})}{\beta^{n_1}} - \frac{s_j(\alpha^{m_{b+1}}, \dots, \alpha^{m_\ell})}{\beta^{m_1}} \\ &= \frac{s_j(\alpha^{m_{b+1}+m_\alpha}, \dots, \alpha^{m_\ell+m_\alpha})}{\beta^{m_1+m_\beta}} - \frac{s_j(\alpha^{m_{b+1}}, \dots, \alpha^{m_\ell})}{\beta^{m_1}} \\ &= \frac{s_j(\alpha^{m_{b+1}}, \dots, \alpha^{m_\ell})}{\beta^{m_1}} \left(\frac{\alpha^{m_\alpha}}{\beta^{m_\beta}} - 1 \right). \end{aligned}$$

It remains to invoke (B). Finally, that there exist infinitely many n_1 that can be extended to $(n_1, \dots, n_\ell)$ satisfying (i) and (ii) for all j follows from Theorem 1.1.14. □

Corollary 5.4.2. *Let $\alpha, \beta \in \mathbb{N}_{\geq 2}$ be multiplicatively independent, $z_1, \ldots, z_\ell \in \{\alpha, \beta\}$, $A \in \mathbb{Z}^{r \times \ell}$, and $\mathbf{b} \in \mathbb{Z}^r_{>0}$. There exists $\mathbf{z} = (z_1^{m_1}, \ldots, z_\ell^{m_\ell})$ with $m_1, \ldots, m_\ell \geq 0$ satisfying $A\mathbf{z} > \mathbf{0}$ if and only if there exists $\widetilde{\mathbf{z}} = (z_1^{n_1}, \ldots, z_\ell^{n_\ell})$ with $n_1, \ldots, n_\ell \geq 0$ satisfying $A\widetilde{\mathbf{z}} > \mathbf{b}$.*

Proof. The 'if'-direction is trivial as one can take $\mathbf{z} := \widetilde{\mathbf{z}}$. Thus, we focus on the other direction. Let $\mathbf{z} = (z_1^{m_1}, \ldots, z_\ell^{m_\ell})$ be as above. For $1 \leq j \leq r$, define the form

$$h_j(x_1, \ldots, x_\ell) = e_j^\top A(x_1, \ldots, x_\ell)\,.$$

Let $\varepsilon \in \mathbb{Q}_{>0}$ be such that $h_j(z_1^{m_1}, \ldots, z_\ell^{m_\ell})/z_1^{m_1} > 2\varepsilon$ for all j. Invoke Theorem 5.4.1 with the forms $h_1, \ldots, h_r$ and the values $m_1, \ldots, m_\ell, \varepsilon$. Therefore, there exist infinitely many $(\widetilde{m}_1, \ldots, \widetilde{m}_\ell)$ (where $\widetilde{m}_1$ can be arbitrarily large) such that, for all j,$h_j(z_1^{\widetilde{m}_1}, \ldots, z_\ell^{\widetilde{m}_\ell}) > 0$ and

$$\left| \frac{h_j\big(z_1^{\widetilde{m}_1}, \ldots, z_\ell^{\widetilde{m}_\ell}\big)}{z_1^{\widetilde{m}_1}} - \frac{h_j\big(z_1^{m_1}, \ldots, z_\ell^{m_\ell}\big)}{z_1^{m_1}} \right| < \varepsilon\,. \tag{5.16}$$

Since $h_j(z_1^{m_1}, \ldots, z_\ell^{m_\ell})/z_1^{m_1} > 2\varepsilon$, (5.16) implies that $h_j\big(z_1^{\widetilde{m}_1}, \ldots, z_\ell^{\widetilde{m}_\ell}\big) > \varepsilon z_1^{\widetilde{m}_1}$. It remains to choose $(\widetilde{m}_1, \ldots, \widetilde{m}_\ell)$ with $z_1^{\widetilde{m}_1}$ sufficiently large. □

The following useful lemma shows how to eliminate a variable n_a if we can bound the gap between n_a and some other (suitable) variable n_b.

Lemma 5.4.3. *Let $\alpha, \beta \in \mathbb{N}_{\geq 2}$, $z_1, \ldots, z_\ell \in \{\alpha, \beta\}$ for $\ell \geq 2$, and $1 \leq a, b \leq \ell$ be distinct with $z_a = z_b$. Suppose we are given the system*

$$\begin{cases} A\mathbf{z} > \mathbf{0} \\ N_1 \leq n_a - n_b \leq N_2 \end{cases} \tag{5.17}$$

where $A \in \mathbb{Z}^{r \times \ell}$ for $r \geq 1$, $\mathbf{z} = (z_1^{n_1}, \ldots, z_\ell^{n_\ell})$, and $N_1, N_2 \in \mathbb{Z}$. Then we can construct matrices $\widetilde{A}_k \in \mathbb{Z}^{r \times (\ell-1)}$ for $N_1 \leq k \leq N_2$ and $y_1, \ldots, y_{\ell-1} \in \{\alpha, \beta\}$ with the following property. There exists $\mathbf{y} = \big(y_1^{n_1}, \ldots, y_{\ell-1}^{n_{\ell-1}}\big)$ satisfying $n_1, \ldots, n_{\ell-1} \geq 0$ and $\bigvee_{k=N_1}^{N_2} \widetilde{A}_k\mathbf{y} > \mathbf{0}$ if and only if the system (5.17) has a solution.

Proof. Choose $(y_1, \ldots, y_{\ell-1})$ to be any ordering of $\{z_1, \ldots, z_\ell\} \setminus \{z_a\}$. It suffices to construct $\widetilde{A}_k$ for $N_1 \leq k \leq N_2$ such that $\widetilde{A}_k \cdot \mathbf{y} > \mathbf{0}$ has a solution if and only if

$$\begin{cases} A\mathbf{z} > \mathbf{0} \\ n_a = n_b + k \end{cases} \tag{5.18}$$

has a solution. The system (5.18) has a solution if and only if there exist natural numbers $n_1, \dots, n_{a-1}, n_{a+1}, \dots, n_\ell$ such that

$$\left(A_{j,a} z_b^k + A_{j,b}\right) z_b^{n_b} + \sum_{\substack{i=1 \\ i \neq a,b}}^{\ell} A_{j,i} z_i^{n_i} > 0 \tag{5.19}$$

for all $1 \leq j \leq r$. Thus, we have eliminated the variable n_a and can construct $\widetilde{A}_k$ by writing (5.19) for $1 \leq j \leq r$ in a matrix form. □

By Theorem 5.4.2, to solve the inequality $A\mathbf{z} > \mathbf{b}$ for $\mathbf{b} \geq 0$ it suffices to solve $A\mathbf{z} > \mathbf{0}$. Next, we show how to do the latter.

Theorem 5.4.4. *Suppose we are given multiplicatively independent $\alpha, \beta \in \mathbb{N}_{\geq 2}$, $z_1, \dots, z_\ell \in \{\alpha, \beta\}$ for some $\ell \geq 1$, and $A \in \mathbb{Z}^{r \times \ell}$ with $r > 0$. It is decidable whether there exist $n_1, \dots, n_\ell \in \mathbb{N}$ such that $A\mathbf{z} > \mathbf{0}$, where $\mathbf{z} = (z_1^{n_1}, \dots, z_\ell^{n_\ell})$.*

Proof. The proof is by induction on ℓ. For $\ell = 1$, the statement is immediate. Suppose $\ell = 2$. Then $A\mathbf{z} > \mathbf{0}$ is equivalent to $z_1^{n_1}/z_2^{n_2} \in (c, d)$ for some computable $c, d \in \mathbb{Q} \cup \{+\infty\}$. If $z_1 = z_2$, then a solution exists if and only $z_1^k \in (c, d)$ for some $k \in \mathbb{Z}$, which is trivial to determine. If $z_1 \neq z_2$, then applying Theorem 1.1.14, a solution exists if and only if $d > 0$ and (c, d) is non-empty.

Suppose $\ell > 2$. If we additionally assume that $z_a^{n_a} = z_b^{n_b}$ for some $a \neq b$, then we can eliminate at least one variable and solve the resulting system inductively, as follows. If $z_a = z_b$, then $n_a = n_b$, and we can invoke Theorem 5.4.3 with $N_1 = N_2 = 0$. If $z_a \neq z_b$, then by multiplicative independence, we have $n_a = n_b = 0$, and we can eliminate two variables. Hence, we have reduced our problem to solving $\ell(\ell - 1)/2$ systems in at most $\ell - 1$ variables (which can be solved inductively), and the system

$$\begin{cases} A\mathbf{z} > \mathbf{0} \\ z_a^{n_a} \neq z_b^{n_b} \text{ for all } a \neq b\,. \end{cases} \tag{5.20}$$

Claim 5.4.5. *We can assume that there is no $1 \leq j \leq \ell$ such that $A_{i,j} \geq 0$ for all $1 \leq i \leq r$.*

Proof. Let j be as in the hypothesis, which we can assume to be $j = 1$ by reordering. Set $K = \{1 \leq i \leq r \colon A_{i,1} = 0$. If there is a solution $(n_2, \dots, n_\ell)$ to the subsystem defined by $\forall k \in K \colon h_k(z_2^{n_2}, \dots, z_\ell^{n_\ell}) > 0$, where $h_k(x_2, \dots, x_\ell) = \sum_{i=2}^{\ell} A_{k,i} x_i$, then the original system has a solution when taking $z_1^{n_1}$ large enough. If this subsystem does not have a solution, the original system has no solution. □

Next, by case analysis on the largest two terms among $z_1^{n_1}, \ldots, z_\ell^{n_\ell}$ and the order of the remaining terms, we reduce solving (5.20) to solving systems of the form

$$\begin{cases} A\mathbf{z} > \mathbf{0} \\ z_{\sigma(1)}^{n_{\sigma(1)}}, z_{\sigma(2)}^{n_{\sigma(2)}} > z_{\sigma(3)}^{n_{\sigma(3)}} > \cdots > z_{\sigma(\ell)}^{n_{\sigma(\ell)}} \\ z_{\sigma(1)}^{n_{\sigma(1)}} \neq z_{\sigma(2)}^{n_{\sigma(2)}} \end{cases}$$

where σ is a permutation of $\{1, \ldots, \ell\}$. By renaming variables and rearranging the rows of A, we reduce to solving systems of the form

$$\begin{cases} A\mathbf{z} > 0 \\ z_1^{n_1}, z_2^{n_2} > z_3^{n_3} > \cdots > z_\ell^{n_\ell} \\ z_1^{n_1} \neq z_2^{n_2}\,. \end{cases} \tag{5.21}$$

Suppose $z_1 = z_2$. In this case, we consider the two possibilities $n_1 > n_2$ and $n_1 < n_2$. We will only show how to solve the system

$$\begin{cases} A\mathbf{z} > 0 \\ z_1^{n_1} > z_2^{n_2} > z_3^{n_3} > \cdots > z_\ell^{n_\ell} \end{cases} \tag{5.22}$$

as the same argument applies to the case of $n_1 < n_2$. By Claim 5.4.5, $A_{j,1} < 0$ for some j. Then we can compute N such that $1 \leq n_1 - n_2 \leq N$ in every solution of (5.22). Using Theorem 5.4.3, we then eliminate the variable n_1 and solve the resulting system in $\ell - 1$ variables inductively.

Suppose $z_1 \neq z_2$; this is the more difficult case. As $\ell > 2$, there is a $n_3 \geq 0$. As $z_1^{n_1}, z_2^{n_2} > z_3^{n_3}$, both n_1 and n_2 are non-zero and so by multiplicative independence, we have $z_1^{n_1} \neq z_2^{n_2}$. By exchanging z_1 and z_2 if necessary, and maybe α and β as well, we can assume that $\alpha = z_1 \neq z_2 = z_3 = \beta$. Note that this implies $n_2 > n_3$.

By multiplying inequalities with different rational constants if necessary, write the system (5.21) in the form

$$\begin{cases} p_i(z_2^{n_2}, \ldots, z_\ell^{n_\ell}) < z_1^{n_1} & \text{for } i \in I_- \\ p_i(z_2^{n_2}, \ldots, z_\ell^{n_\ell}) > z_1^{n_1} & \text{for } i \in I_+ \\ p_i(z_2^{n_2}, \ldots, z_\ell^{n_\ell}) > 0 & \text{for } i \in J \\ z_1^{n_1}, z_2^{n_2} > z_3^{n_3} > \cdots > z_\ell^{n_\ell} \end{cases} \tag{5.23}$$

where I_-, I_+, J are disjoint finite sets and each p_i is a $\mathbb{Q}$-linear form. We can assume that I_- is non-empty by adding the identically zero $\mathbb{Q}$-linear form over $\ell - 1$ variables if necessary. Using Claim 5.4.5, I_+ is also non-empty.

For $i \in I_-, I_+, J$, let a_i be the coefficient of $z_2^{n_2}$ in $p_i(z_2^{n_2}, \ldots, z_\ell^{n_\ell})$. Then set $a_- = \max_{i \in I_-}(a_i)$ and $a_+ = \min_{i \in I_+}(a_i)$. Moreover, let $\widetilde{I}_- = \{i \in I_- \colon a_i = a_-\}$,

$\widetilde{I}_+ = \{i \in I_+ \colon a_i = a_+\}$, and $\widetilde{J} = \{j \in J \colon a_i = 0\}$. Further, write $p_i(z_2^{n_2}, \dots, z_\ell^{n_\ell}) = a_- z_2^{n_2} + h_i(z_3^{n_3}, \dots, z_\ell^{n_\ell})$ for $i \in \widetilde{I}_-$, write $p_i(z_2^{n_2}, \dots, z_\ell^{n_\ell}) = a_+ z_2^{n_2} + h_i(z_3^{n_3}, \dots, z_\ell^{n_\ell})$ for $i \in \widetilde{I}_+$, and write $h_i(z_3^{n_3}, \dots, z_\ell^{n_\ell}) = p_i(z_2^{n_2}, \dots, z_\ell^{n_\ell})$ for $i \in \widetilde{J}$.

Then one can effectively compute a number N such that whenever $n_2 - n_3 > N$,

- $p_i(z_2^{n_2}, \dots, z_\ell^{n_\ell}) \le p_j(z_2^{n_2}, \dots, z_\ell^{n_\ell})$ for all $i \in I_- \setminus \widetilde{I}_-$ and $j \in \widetilde{I}_-$;
- $p_i(z_2^{n_2}, \dots, z_\ell^{n_\ell}) \le p_j(z_2^{n_2}, \dots, z_\ell^{n_\ell})$ for all $i \in I_+ \setminus \widetilde{I}_+$ and $j \in \widetilde{I}_+$; and
- $\operatorname{sign}(c_i) = \operatorname{sign}\left(p_i(z_2^{n_2}, \dots, z_\ell^{n_\ell})\right)$ for all $i \in J \setminus \widetilde{J}$.

If we add $0 \le n_2 - n_3 \le N$ to (5.23), we can solve the resulting system by eliminating n_2 using Theorem 5.4.3. Meanwhile, if $n_2 - n_3 > N$, we have reduced (5.23) (and hence our original decision problem) to solving systems of the following form:

$$\begin{cases} a_- + \dfrac{h_i(z_3^{n_3}, \dots, z_\ell^{n_\ell})}{z_2^{n_2}} < \dfrac{z_1^{n_1}}{z_2^{n_2}} < a_+ + \dfrac{h_j(z_3^{n_3}, \dots, z_\ell^{n_\ell})}{z_2^{n_2}} & \text{for all } i \in \widetilde{I}_-,\ j \in \widetilde{I}_+ \quad (5.24) \\ z_1^{n_1} > z_3^{n_3} & (5.25) \\ z_3^{n_3} > \cdots > z_\ell^{n_\ell} & (5.26) \\ h_i(z_3^{n_3}, \dots, z_\ell^{n_\ell}) > 0 \quad \text{for all } i \in \widetilde{J} & (5.27) \\ n_2 - n_3 > N. & (5.28) \end{cases}$$

Note that as $N \ge 0$ and $z_2 = z_3$, the condition (5.28) implies $z_2^{n_2} > z_3^{n_3}$. It remains to show how to solve the system (5.24–5.28).

Case 1. Suppose $a_- = a_+ = a > 0$ for some $a \in \mathbb{Q}$. This is the only difficult case. Recalling that $z_2 = z_3 = \beta$, (5.24) is equivalent to

$$\frac{h_i(z_3^{n_3}, \dots, z_\ell^{n_\ell})}{z_3^{n_3}} \cdot \frac{1}{\beta^{n_2 - n_3}} < \frac{\alpha^{n_1}}{\beta^{n_2}} - a < \frac{h_j(z_3^{n_3}, \dots, z_\ell^{n_\ell})}{z_3^{n_3}} \cdot \frac{1}{\beta^{n_2 - n_3}} \quad \text{for all } i \in \widetilde{I}_-,\ j \in \widetilde{I}_+ .$$

Observe that $h_i(z_3^{n_3}, \dots, z_\ell^{n_\ell}) < h_j(z_3^{n_3}, \dots, z_\ell^{n_\ell})$ is implied by (5.24). Inductively solve the system consisting of the inequalities $h_i(z_3^{n_3}, \dots, z_\ell^{n_\ell}) < h_j(z_3^{n_3}, \dots, z_\ell^{n_\ell})$ for $i \in \widetilde{I}_-$ and $j \in \widetilde{I}_+$, (5.26), and (5.27). If no solution exists, then the system (5.24–5.28) does not have a solution either. Otherwise, let $(m_3, \dots, m_\ell)$ be a solution to the smaller system. We will argue that the system (5.24–5.28) also has a solution.

Let $x_- = \max_{i \in \widetilde{I}_-} \left(h_i(z_3^{m_3}, \dots, z_\ell^{m_\ell})/z_3^{m_3}\right)$, $x_+ = \min_{i \in \widetilde{I}_+} \left(h_i(z_3^{m_3}, \dots, z_\ell^{m_\ell})/z_3^{m_3}\right)$, and $\varepsilon = (x_+ - x_-)/4$. From the construction of $m_3, \dots, m_\ell$, it follows that $\varepsilon > 0$. We construct a solution $(k_1, \dots, k_\ell) \in \mathbb{N}^\ell$ to the system (5.24–5.28). To do this, it suffices to construct $(k_1, \dots, k_\ell)$ satisfying (5.25–5.28) with the following additional properties:

(a) $\frac{x_- + \varepsilon}{\beta^{k_2 - k_3}} < \frac{\alpha^{k_1}}{\beta^{k_2}} - a < \frac{x_+ - \varepsilon}{\beta^{k_2 - k_3}}$;

(b) $\frac{h_i\left(z_3^{k_3}, \ldots, z_\ell^{k_\ell}\right)}{z_3^{k_3}} < x_- + \varepsilon$ for all $i \in \widetilde{I}_-$;

(c) $\frac{h_i\left(z_3^{k_3}, \ldots, z_\ell^{k_\ell}\right)}{z_3^{k_3}} > x_+ - \varepsilon$ for all $i \in \widetilde{I}_+$.

Observe that (a), (b), and (c) indeed imply (5.24) as for all $i \in \widetilde{I}_-$ and $j \in \widetilde{I}_+$,

$$\frac{h_i\left(z_3^{k_3}, \ldots, z_\ell^{k_\ell}\right)}{\beta^{k_3}} < \frac{x_- + \varepsilon}{\beta^{k_2 - k_3}} < \frac{\alpha^{k_1}}{\beta^{k_2}} - a < \frac{x_+ - \varepsilon}{\beta^{k_2 - k_3}} < \frac{h_j\left(z_3^{k_3}, \ldots, z_\ell^{k_\ell}\right)}{\beta^{k_3}}\,.$$

Next, invoke the Pumping Lemma with $m_1, \ldots, m_\ell$, ε as above and the linear forms

- $-h_i(z_3^{n_3}, \ldots, z_\ell^{n_\ell}) + (x_- + \varepsilon) z_3^{k_3}$ for all $i \in \widetilde{I}_-$,
- $h_i(z_3^{n_3}, \ldots, z_\ell^{n_\ell}) - (x_+ - \varepsilon) z_3^{k_3}$ for all $i \in \widetilde{I}_+$,
- $h_i(z_3^{n_3}, \ldots, z_\ell^{n_\ell})$ for all $i \in \widetilde{J}$, and
- $z_i^{n_i} - z_{i+1}^{n_{i+1}}$ for $3 \leq i \leq \ell - 1$

to compute $\mu, \delta > 0$. We have that any $n_3 > m_3$ satisfying $|\alpha^{\widetilde{n}}/\beta^{n_3} - \mu| < \delta$ for a $\widetilde{n} \in \mathbb{N}$ can be extended to $(n_3, \ldots, n_\ell) \in \mathbb{N}^{\ell-2}$ satisfying (5.26–5.27) and (b–c). Let $0 < \Delta < \min(a, \frac{a\delta}{2\mu})$. Then $\mu\Delta/a < \delta/2$. We will need the following lemma that, intuitively, shows we can simultaneously satisfy the Diophantine approximation conditions arising from the above application of the Pumping Lemma and item (a).

Lemma 5.4.6. *Let a, μ, δ, Δ be as above. Given $M \in \mathbb{N}$, we can compute $d > M$ and $m \in \mathbb{N}$ with the following property. For all $k \geq m$, if there exists $\widetilde{n} \in \mathbb{N}$ such that*

$$\left|\alpha^{\widetilde{n}}/\beta^k - a\right| < \Delta\,,$$

then there exists $\widehat{n} \in \mathbb{N}$ such that

$$\left|\alpha^{\widehat{n}}/\beta^{k-d} - \mu\right| < \delta\,.$$

Proof. Let $\xi = \delta/(4a)$. Using Theorem 1.1.14, choose $d, m \in \mathbb{N}$ such that $d > M$ and

$$\left|\frac{\beta^d}{\alpha^m} - \frac{\mu}{a}\right| < \xi\,.$$

Suppose $|\alpha^{\widetilde{n}}/\beta^k - a| < \Delta$ for some $\widetilde{n} \geq m$. Let $\widehat{n} = \widetilde{n} - m$. Then

$$\begin{aligned}
\left|\frac{\alpha^{\widehat{n}}}{\beta^{k-d}} - \mu\right| &= \left|\frac{\alpha^{\widetilde{n}}}{\beta^k} \cdot \frac{\beta^d}{\alpha^m} - \mu\right| \\
&= \left|\left(\frac{\alpha^{\widetilde{n}}}{\beta^k} - a\right)\frac{\beta^d}{\alpha^m} + a\left(\frac{\beta^d}{\alpha^m} - \frac{\mu}{a}\right)\right| \\
&< \Delta(\xi + \mu/a) + a\xi \\
&< 2a\xi + \frac{\mu\Delta}{a} \\
&\leq \delta\,,
\end{aligned}$$

where the last two inequalities follow from $\Delta < a$, $\mu\Delta/a < \delta/2$, and $a\xi = \delta/4$. □

Choose $M > N$ such that $\Delta\beta^M > |x_- + \varepsilon|, |x_+ - \varepsilon|$, and $x_- + \varepsilon + a\beta^d > 1$. Then apply Lemma 5.4.6 with this value of M to construct d and m. We will next construct $(k_1, \ldots, k_\ell) \in \mathbb{N}^\ell$ satisfying (5.25–5.28) and (a-c); recall that such $(k_1, \ldots, k_\ell)$ will also be a solution to (5.24–5.28). First, choose k_1, k_2 such that $k_2 > \max(d, m)$, and

$$\frac{x_- + \varepsilon}{\beta^d} < \frac{\alpha^{k_1}}{\beta^{k_2}} - a < \frac{x_+ - \varepsilon}{\beta^d}\,.$$

As $d > M$, we have $|\alpha^{k_1}/\beta^{k_2} - a| < \Delta$. Then set $k_3 = k_2 - d$. By the construction of d and m via Lemma 5.4.6, and the fact that $k_2 > m$, there exists $\widehat{n}$ such that

$$\left|\frac{\alpha^{\widehat{n}}}{\beta^{k_2-d}} - \mu\right| = \left|\frac{\alpha^{\widehat{n}}}{\beta^{k_3}} - \mu\right| < \delta\,.$$

Hence, by construction of μ, δ via the Pumping Lemma, we can extend k_3 to a tuple $(k_3, \ldots, k_\ell)$ that satisfies (5.26–5.27) as well as (b–c). Inequality (5.28) and property (a) are satisfied by construction. It remains to show that (5.25) is satisfied. By (a), $\alpha^{k_1} - a\beta^{k_2} > (x_- + \varepsilon)\beta^{k_3}$. Hence,

$$\alpha^{k_1} > (x_- + \varepsilon)\beta^{k_3} + a\beta^{k_2} = \beta^{k_3}(x_- + \varepsilon + a\beta^d) > \beta^{k_3}\,.$$

Case 2. Suppose $a_+ > 0$ and $a_+ > a_-$. Let $\varepsilon = \frac{a_+ - \max(a_-, 0)}{4}$. Compute $M \geq N$ such that for all $n_2, \ldots, n_\ell \in \mathbb{N}$ satisfying $z_2^{n_2} > z_3^{n_3} > \cdots > z_\ell^{n_\ell}$ and $n_2 - n_3 > M$, we have

$$\left|\frac{h_i(z_3^{n_3}, \ldots, z_\ell^{n_\ell})}{z_2^{n_2}}\right| < \varepsilon \quad \text{for all } i \in \widetilde{I}_- \cup \widetilde{I}_+\,.$$

Next, inductively solve the subsystem comprising inequalities (5.26) and (5.27). If there is no solution, then (5.24–5.28) does not have a solution, and we are done. Otherwise, let $(k_3, \ldots, k_\ell)$ be a solution of the subsystem. Applying Theorem 1.1.14,

construct $k_1, k_2 \in \mathbb{N}$ such that $z_1^{k_1} > z_3^{k_3}$, $k_2 - k_3 > M$, and $z_1^{k_1}/z_2^{k_2} \in (a_- + \varepsilon, a_+ - \varepsilon)$. Then $(k_1, \ldots, k_\ell)$ is a solution of (5.24–5.28).

Case 3. Suppose $a_+ < a_-$. Let ε, M, and $(k_3, \ldots, k_\ell)$ be as in Case 2; If no $(k_3, \ldots, k_\ell)$ exist, once again we are done. Observe that any $(n_1, \ldots, n_\ell) \in \mathbb{N}^\ell$ such that $n_2 - n_3 > M$ is not a solution of (5.24–5.28). Hence the system (5.24–5.28) has a solution if and only if the system comprising (5.24–5.27) and $N < n_2 - n_3 \leq M$ has a solution, which can be checked using Theorem 5.4.3.

Case 4. $a_+ = a_- = 0$. In this case, (5.24) is equivalent to

$$h_i(z_3^{n_3}, \ldots, z_\ell^{n_\ell}) < z_1^{n_1} < h_j(z_3^{n_3}, \ldots, z_\ell^{n_\ell}) \quad \text{for all } i \in \widetilde{I}_- \text{ and } j \in \widetilde{I}_+ \tag{5.29}$$

in which the variable n_2 does not appear. Hence, we can first inductively solve the subsystem comprising (5.24–5.27). If a solution exists, then choose n_2 to be sufficiently large to satisfy (5.28). Otherwise, conclude that (5.24–5.28) does not have a solution either.

Case 5. Suppose $a_+ = 0 > a_-$. This case is similar to Case 4. Let M be such that for all $(n_1, \ldots, n_\ell)$, if $n_2 - n_3 > M$ then

$$a_- + \frac{h_i(z_3^{n_3}, \ldots, z_\ell^{n_\ell})}{z_2^{n_2}} < 0 \quad \text{for all } i \in \widetilde{I}_- \,.$$

Hence for such $(n_1, \ldots, n_\ell)$, (5.24) is equivalent to

$$z_1^{n_1} < h_i(z_3^{n_3}, \ldots, z_\ell^{n_\ell}) \quad \text{for all } i \in \widetilde{I}_+ \,. \tag{5.30}$$

Therefore, we solve two systems. First, inductively check if the system comprising (5.25–5.27) and (5.30) has a solution $(k_1, k_3, \ldots, k_\ell)$. If yes, then choose k_2 to be sufficiently large so that (5.28) is satisfied. Thereafter, solve (5.24–5.28) together with $N < n_2 - n_3 \leq M$ using Theorem 5.4.3. The system (5.24–5.28) has a solution if and only if at least one of the two systems has a solution.

Case 6. Finally, suppose, $a_+ < 0$. Let M be such that for all $n_2 - n_3 > M$,

$$a_+ + \frac{h_i(z_3^{n_3}, \ldots, z_\ell^{n_\ell})}{z_2^{n_2}} < 0 \quad \text{for all } i \in \widetilde{I}_+ \,.$$

It remains to solve (5.24–5.28) together with $N < n_2 - n_3 \leq M$ using Theorem 5.4.3. □

We can finally prove that Problem 5.1.4 is decidable.

Proof of Theorem 5.1.1. We use induction on ℓ to solve Problem 5.1.4. If $\ell = 1$, the result is immediate. Suppose $\ell \geq 2$. Write the system $A\mathbf{z} > \mathbf{b} \wedge C\mathbf{z} = \mathbf{d}$ in the form

$$\bigvee_{k \in K} A_k\mathbf{z} > \mathbf{b}_k \wedge C_k\mathbf{z} = \mathbf{d}_k$$

where each $\mathbf{b}_k \geq \mathbf{0}$. By Theorem 5.4.2, this is equivalent to the system

$$\bigvee_{k \in K} A_k\mathbf{z} > \mathbf{0} \wedge C_k\mathbf{z} = \mathbf{d}_k\,.$$

It suffices to solve each disjunct separately. Fix $k \in K$. If C_k is empty, then we can solve $A_k\mathbf{z} > 0$ using Theorem 5.4.4. Suppose C_k is non-empty. Then first solve $C_k\mathbf{z} = \mathbf{d}_k$ and write the set of solutions $\mathcal{S}$ in the form

$$\mathcal{S} = \bigcup_{i \in I} \bigcap_{j \in J_i} X_j$$

as in Theorem 5.1.7. It suffices to check, for every $i \in I$, whether $A_k\mathbf{z} > \mathbf{0}$ has a solution belonging to $\bigcap_{j \in J_i} X_j$. Fix $1 \leq i \leq I$. If J_i is empty, then we simply solve $A_k\mathbf{z} > \mathbf{0}$ using Theorem 5.4.4. When J_i is non-empty, we will carry out a variable elimination as follows.

Lemma 5.4.7. *Let $\alpha, \beta \in \mathbb{N}_{\geq 2}$, $z_1, \ldots, z_\ell \in \{\alpha, \beta\}$, $E \in \mathbb{Z}^{r \times \ell}$, $\mathbf{u} \in \mathbb{Z}^r$, and $X_1, \ldots, X_M \subseteq \mathbb{N}^\ell$ where each X_j is defined by either*

$$n_{\mu(j)} = n_{\sigma(j)} + c_j \tag{5.31}$$

or

$$n_{\xi(j)} = b_j \tag{5.32}$$

for some $b_j, c_j \in \mathbb{Z}$ and $1 \leq \xi(j), \mu(j), \sigma(j) \leq \ell$ satisfying $z_{\mu(j)} = z_{\sigma(j)}$. We can construct $\lambda < \ell$, $F \in \mathbb{Z}^{r \times \lambda}$, $\mathbf{v} \in \mathbb{Z}^r$, and $y_1, \ldots, y_\lambda \in \{\alpha, \beta\}$ such that

$$E \cdot (z_1^{n_1}, \ldots, z_\ell^{n_\ell}) > \mathbf{u} \wedge (n_1, \ldots, n_\ell) \in \bigcap_{1 \leq j \leq M} X_j \tag{5.33}$$

has a solution if and only if $F \cdot (y_1^{n_1}, \ldots, y_\lambda^{n_\lambda}) > \mathbf{v}$ has a solution.

Proof. We use induction on M. Write $\mathbf{y} = (y_1^{n_1}, \ldots, y_\lambda^{n_\lambda})$ and $\mathbf{z} = (z_1^{n_1}, \ldots, z_\ell^{n_\ell})$. If $M = 0$, we are done. So let $M \geq 1$ and consider X_1. If X_1 is of the form (5.31), we split into two cases. First, consider $\mu(j) = \sigma(j)$. If $c_j = 0$, (5.31) holds trivially true, and we can remove it to return to $M - 1$ equations while if $_j \neq 0$, (5.31) never holds, and we simply return $\lambda = 1$, $F = (-1)$, $\mathbf{v} = 0$, and $y_1 = \alpha$ such that $F\mathbf{y} > 0$ never

holds. Secondly, if $\mu(j) \neq \sigma(j)$, let $(y_1, \ldots, y_{\ell-1}) = (z_1, \ldots, z_{\mu(j)-1}, z_{\mu(j)+1}, \ldots, z_\ell)$ and update $X_2, \ldots, X_M$ by substituting each occurrence of $\mu(j)$ with $\sigma_j + c_j$ to obtain a system $\widetilde{E}\mathbf{y} > \widetilde{\mathbf{u}}$ equivalent to $E\mathbf{z} > \mathbf{u}$.

Similarly, assume that X_1 is of the form (5.32). Then, if $b_j < 0$, the system has no solution, and we return some F and $\mathbf{v}$ that do not have a solution. If $b_j \geq 0$, substitute n_{ξ_j} with b_j in $X_2, \ldots, X_M$, let $(y_1, \ldots, y_{\ell-1}) = (z_1, \ldots, z_{\xi(j)-1}, z_{\xi(j)+1}, \ldots, z_\ell)$ and obtain a system $\widetilde{E}\mathbf{y} > \widetilde{\mathbf{u}}$ equivalent to $E\mathbf{z} > \mathbf{u}$ by setting $z_{\xi_j}^{n_{\xi_j}}$ to by $z_{\xi_j}^{b_j}$. This again reduces M. Thus, in each case, we reduce M by at least 1, proving the result. □

Using the lemma above, we can construct $\lambda < \ell$, $F \in \mathbb{Z}^{r \times \lambda}$, $\mathbf{v} \in \mathbb{Z}^r$, and $y_1, \ldots, y_\lambda \in \{\alpha, \beta\}$ such that

$$A_k\mathbf{z} > 0 \wedge (n_1, \ldots, n_\ell) \in \bigcap_{j \in J_i} X_j$$

has a solution if and only if

$$F \cdot (y_1^{n_1}, \ldots, y_\lambda^{n_\lambda}) > \mathbf{v} \tag{5.34}$$

has a solution. Since $\lambda < \ell$, we can use the induction hypothesis to solve (5.34). Thus, Problem 5.1.4 is decidable and so using Lemma 5.1.5, Theorem 5.1.1 holds. □

5.5 Presburger arithmetic expanded with powering functions

We now consider the existential fragment of $\mathcal{PA}(\alpha^x, \beta^x)$ for multiplicatively independent α and β, where γ^x denotes the function $x \mapsto \gamma^x$ for $\gamma \in \{\alpha, \beta\}$. Unlike the case for $\mathcal{PA}(\alpha^{\mathbb{N}}, \beta^{\mathbb{N}})$, we show that decidability of the existential fragment of $\mathcal{PA}(\alpha^x, \beta^x)$ would give us algorithms for deciding various properties of base-α and base-β expansions of a large class of numbers, captured by the following definition.

Definition 5.5.1. A sequence $(u_n)_{n=0}^\infty$ over $\mathbb{N}$ is existentially definable if for every $k \geq 1$ there exists an existential formula φ with $k+1$ free variables in the language of $\mathcal{PA}(\alpha^x, \beta^x)$ such that for all $n, y_0, \ldots, y_{k-1} \in \mathbb{N}$, the sentence $\varphi(n, y_0, \ldots, y_{k-1})$ holds if and only if $u_{n+i} = y_i$ for all $0 \leq i < k$.

The set of definable sequences is closed under many operations. Let $(u_n)_{n=0}^\infty$ and $(v_n)_{n=0}^\infty$ be definable, and $c \in \mathbb{N}$. Then $(u_n + v_n)_{n=0}^\infty$, $(c + u_n)_{n=0}^\infty$, $(c \cdot u_n)_{n=0}^\infty$, $(\alpha^{u_n})_{n=0}^\infty$, $(\beta^{u_n})_{n=0}^\infty$, and $(u_{v_n})_{n=0}^\infty$ are also definable. Write $\{x\}$ for the fractional part of x. Let $\alpha, \beta \in \mathbb{N}_{\geq 2}$ be multiplicatively independent, $(A_n)_{n=0}^\infty$ be the base-α

expansion of $\{\log_\beta(\alpha)\}$, and $(B_n)_{n=0}^\infty$ be the base-β expansion of $\{\log_\alpha(\beta)\}$. Note that $\log_\alpha(\beta), \log_\beta(\alpha)$ are both irrational, and for any $\gamma \in \mathbb{N}_{>0}$ and $x \in \mathbb{R}_{\geq 0}$, the base-$\gamma$ expansions of x and $\{x\}$ differ only by a finite prefix.

Proposition 5.5.2. *The sequences $(A_n)_{n=0}^\infty$ and $(B_n)_{n=0}^\infty$ are definable.*

Proof. By symmetry, it is sufficient to prove the proposition for $(A_n)_{n=0}^\infty$. For $x \geq 1$, let $f(x)$ denote the integer α^m such that $\alpha^m \leq x < \alpha^{m+1}$, noting that $f(x) = \alpha^{\lfloor \log_\alpha(x) \rfloor}$. Fix $k \geq 1$, and let $w \in \{0, \ldots, \alpha - 1\}^k$. Let $\lambda(w)$ denote the natural number whose base-α expansion equals w. That w occurs at position n in $(A_n)_{n=0}^\infty$ can be expressed as

$$\lambda(w) < \alpha^{n+k}\{\log_\alpha(\beta)\} < \lambda(w) + 1\,,$$

which is equivalent to

$$\alpha^{\lambda(w)} < \left(\frac{\beta^{\alpha^n}}{\alpha^{\lfloor \alpha^n \log_\alpha(\beta) \rfloor}}\right)^{\alpha^k} < \alpha^{\lambda(w)+1}\,.$$

For any constant c and a term t, we can express $c \cdot t$ as $\underbrace{t + \cdots + t}_{c \text{ times}}$. As $\alpha^{\lfloor \alpha^n \log_\alpha(\beta) \rfloor} = f(\alpha^n)$, we thus have that

$$\begin{aligned}\varphi(n, y_0, \ldots, y_{k-1}) := \exists m\colon\ & \alpha^m \leq \beta^{\alpha^n} < \alpha^{m+1} \\ & \wedge\ \alpha^{\lambda(y_0 \cdots y_{k-1}) + m\alpha^k} < \beta^{\alpha^{n+k}} < \alpha^{\lambda(y_0 \cdots y_{k-1})+1+m\alpha^k}\end{aligned}$$

for $k \geq 1$ define $(A_n)_{n=0}^\infty$ as required. □

Observe that we can express whether a pattern $w_0 \cdots w_{k-1}$ occurs in an existentially definable sequence $(u_n)_{n=0}^\infty$ using the existential formula $\exists n\colon \varphi(n, w_0, \ldots, w_{k-1})$, where φ is the formula described in Definition 5.5.1. Therefore, decidability of the existential fragment of $\mathcal{PA}(\alpha^x, \beta^x)$ would entail the existence of oracles, among others, for deciding the following problems.

(A) Whether a given pattern w appears in the base-β expansion of $\log_\beta(\alpha)$.

(B) Whether a given pattern w appears at some index simultaneously in the base-β expansions of $\log_\beta(\alpha)$ and $\log_\alpha(\beta)$.

(C) Whether a given pattern w appears in $(A_{\alpha^n})_{n=0}^\infty$.

This proves Theorem 5.1.3.

To the best of our knowledge, for no base $\gamma \in \mathbb{N}_{\geq 2}$ and multiplicatively independent $\alpha, \beta \in \mathbb{N}_{\geq 2}$, an algorithm is known that determines whether a given pattern appears in the base-γ expansion of $\log_\alpha(\beta)$. Proof of normality for the sequences $(A_n)_{n=0}^{\infty}$ and $(B_n)_{n=0}^{\infty}$ would make Problem (A) above trivially decidable. However, normality alone is not strong enough to deal with Problems (B) and (C): Deciding the latter problems in the same way as Problem (A) would require a far stronger 'randomness' property. Even if such properties are proven, we might still be unable to prove the decidability of the full existential fragment of $\mathcal{PA}(\alpha^x, \beta^x)$.

5.6 Undecidability of expansions with two sets of powers

In this section, let $\alpha, \beta \in \mathbb{N}_{\geq 2}$ be multiplicatively independent. In [78], Hieronymi and Schulz reduce from the Halting problem for Turing machines to show that $\mathcal{PA}(\alpha^{\mathbb{N}}, \beta^{\mathbb{N}})$ is undecidable. We provide an alternative (and shorter) undecidability proof by reducing from the Halting problem for 2-counter Minksy machines, which is also undecidable [117, Chapter 14]. Our proof shows that already for formulas containing three alternating blocks of quantifiers, membership in $\mathcal{PA}(\alpha^{\mathbb{N}}, \beta^{\mathbb{N}})$ is undecidable.

A *2-counter Minsky machine* is given by $R > 0$ instructions, numbered $1, \dots, R$, and two counters $c^{(1)}, c^{(2)}$ that take values in $\mathbb{N}$. Each instruction except the Rth one is either of the form

$$c^{(i)} = c^{(i)} + 1;\ \mathsf{GOTO}\, r \quad or$$
$$c^{(i)} = 0;\ \mathsf{GOTO}\, r;\ \mathsf{ELSE}\, \mathsf{c^{(i)} = c_i^{(i)} - 1};\ \mathsf{GOTO}\, \widetilde{r}\,.$$

The execution starts at line $r = 1$ with both counters set to zero, and the machine halts if the line $r = R$ is reached. Let $c_i^{(n)}$ denote the value of the counter c_i and by r_n the current instruction number after n steps. We refer to $(c_n^{(1)}, c_n^{(2)}, r_n)$ as the *configuration* of the machine at time n. The *transition function* $f\colon \mathbb{N} \times \mathbb{N} \times \{1, \dots, R\} \to \mathbb{N} \times \mathbb{N} \times \{1, \dots, R\}$ of the machine describes how the configuration is updated. By definition, we have that $c_0^{(1)} = c_0^{(2)} = 0$ and $r_0 = 1$.

We will represent the trace of the machine by the sequence

$$\left(\alpha^{R+c_0^{(1)}}, \alpha^{R+c_0^{(2)}}, \alpha^{r_0-1}, \alpha^{R+c_1^{(1)}}, \alpha^{R+c_1^{(2)}}, \alpha^{r_1-1}, \dots\right).$$

Here, $\alpha^{R+c_n^{(1)}}$ and $\alpha^{R+c_n^{(2)}}$ are at least α^R while $\alpha^{r_n-1} < \alpha^R$ for every $n \geq 0$. Note that every entry in the sequence is a power of α, and the nth entry is smaller than α^R if

and only if $n \equiv 2 \pmod 3$. It remains to represent such sequences using arithmetic of powers of α and β.

For $x \in \mathbb{N}$, let $\mu(x)$ denote the most significant digit in the base-α expansion of x, and by $\delta(x)$ the number α^n (whenever it exists) such that the digit corresponding to α^n in the base-α expansion of x is the second most significant digit that is non-zero. For example, if $\alpha = 10$, then $\mu(3078) = 3$ and $\delta(3078) = 10^1$. Next, consider $\mathcal{A}_\ell, \mathcal{A}_u \in \alpha^{\mathbb{N}}, \mathcal{B}_\ell, \mathcal{B}_u \in \beta^{\mathbb{N}}$ with $\mathcal{A}_\ell < \mathcal{A}_u$ and $\mathcal{B}_\ell < \mathcal{B}_u$. Let $\mathcal{P}$ be the set of all $b \in \beta^{\mathbb{N}} \cap [\mathcal{B}_\ell, \mathcal{B}_u]$ such that $\mu(b) = 1$ and $\delta(b) \in [\mathcal{A}_\ell, \mathcal{A}_u]$. Write $N = |\mathcal{P}| - 1$ and order the elements of $\mathcal{P}$ as $B_0 < \cdots < B_N$. We say that the tuple $(\mathcal{A}_\ell, \mathcal{A}_u, \mathcal{B}_\ell, \mathcal{B}_u)$ *defines* the finite sequence $(u_n)_{n=0}^N$ over $\alpha^{\mathbb{N}}$ given by $u_n = \delta(B_n)/\mathcal{A}_\ell$. The following result, Lemma 3.4 in [78], serves a crucial role in their and our undecidability proofs.

Theorem 5.6.1. *Every finite sequence $(u_n)_{n=0}^N$ over $\alpha^{\mathbb{N}}$ is defined by some tuple $(\mathcal{A}_\ell, \mathcal{A}_u, \mathcal{B}_\ell, \mathcal{B}_u)$.*

By choosing $\mathcal{B}_\ell$ to be the smallest element of $\mathcal{P}$ and $\mathcal{B}_u$ to be the largest element of $\mathcal{P}$ if necessary, we can always assume that $\mathcal{B}_\ell, \mathcal{B}_u \in \mathcal{P}$. We will encode the Halting problem for 2-counter machines by constructing a formula that expresses the existence of a tuple $(\mathcal{A}_1, \mathcal{A}_2, \mathcal{B}_1, \mathcal{B}_2)$ that defines a sequence corresponding to a finite trace of the machine ending with the halting instruction. Let $\mathcal{A}_\ell, \mathcal{A}_u \in \alpha^{\mathbb{N}}, \mathcal{B}_\ell, \mathcal{B}_u \in \beta^{\mathbb{N}}$ define the sequence $(u_n)_{n=0}^N$, and $\mathcal{P} = \{B_0, \ldots, B_N\}$ be as above. Define

$$\begin{aligned}\varphi_{\mathcal{A}_\ell,\mathcal{A}_u,\mathcal{B}_\ell,\mathcal{B}_u}(C, A, B) := {} & C \in \alpha^{\mathbb{N}} \wedge A \in \alpha^{\mathbb{N}} \cap [\mathcal{A}_\ell, \mathcal{A}_u] \wedge B \in \beta^{\mathbb{N}} \cap [\mathcal{B}_\ell, \mathcal{B}_u] \wedge \\ & \wedge C \leq B < 2C \wedge A \leq B - C < \alpha \cdot A\,.\end{aligned}$$

This formula states that $B \in \mathcal{P}$, which is witnessed by C and A. Here, C is the largest power of α not exceeding B, the atomic formula $C \leq B < 2C$ ensures that $\mu(B) = 1$, and $A \leq B - C < \alpha \cdot A$ ensures that $A = \delta(B)$. If $\varphi_{\mathcal{A}_\ell,\mathcal{A}_u,\mathcal{B}_\ell,\mathcal{B}_u}(C, A, B)$ holds, then $u_n = A/\mathcal{A}_\ell$ where n is the position of B in $\mathcal{P}$. The next formula, on input B_1, B_2 that belong to $\mathcal{P}$, returns whether B_1 is immediately followed by B_2 in the ordering of $\mathcal{P}$.

$$\psi_{\mathcal{A}_\ell,\mathcal{A}_u,\mathcal{B}_\ell,\mathcal{B}_u}(B_1, B_2) := \forall C, A, B_1 < B < B_2 \colon \neg\varphi_{\mathcal{A}_\ell,\mathcal{A}_u,\mathcal{B}_\ell,\mathcal{B}_u}(C, A, B)\,.$$

We omit the subscript from φ and ψ when $\mathcal{A}_\ell, \mathcal{A}_u, \mathcal{B}_\ell, \mathcal{B}_u$ are clear from the context. We can now construct a formula in the language of $\mathcal{PA}(\alpha^{\mathbb{N}}, \beta^{\mathbb{N}})$ that is true if and only if the given 2-counter machine halts. Write $\mathbf{X}$ for the collection of variables $\mathcal{A}_\ell, \mathcal{A}_u, \mathcal{B}_\ell, \mathcal{B}_u, \widehat{B}_1, \widehat{B}_2, \widehat{C}_0, \widehat{C}_1, \widehat{C}_2, C_{\text{last}}$, and $\mathbf{Y}$ for the collection of variables $C_0, A_0, B_0, \ldots, C_5, A_5, B_5$. The variables

- $\mathcal{A}_\ell, \mathcal{A}_u, \mathcal{B}_\ell, \mathcal{B}_u$ serve to define a finite sequence over $\alpha^{\mathbb{N}}$,
- $\mathcal{B}_\ell, \widehat{B}_1, \widehat{B}_2$ denote the first three elements of $\mathcal{P}$ with witnesses $(\widehat{C}_0, \alpha^R \cdot \mathcal{A}_\ell)$, $(\widehat{C}_1, \alpha^R \cdot \mathcal{A}_\ell)$, and $(\widehat{C}_2, \mathcal{A}_\ell)$, respectively,
- $\mathcal{B}_u$ is the final element of $\mathcal{P}$ with the witness $(\widehat{C}_{\text{last}}, \alpha^{R-1} \cdot \mathcal{A}_\ell)$, and
- $C_0, A_0, B_0, \ldots, C_5, A_5, B_5$ represent arbitrary 6 consecutive terms of the sequence defined by $(\mathcal{A}_\ell, \mathcal{A}_u, \mathcal{B}_\ell, \mathcal{B}_u)$, which correspond to two consecutive configurations of the machine. (Recall that each configuration of the machine consists of three numbers.)

The required formula is then

$$\begin{aligned}\exists \mathbf{X}\colon\ & \psi(\mathcal{B}_\ell, \widehat{B}_1) \wedge \psi(\widehat{B}_1, \widehat{B}_2) \wedge \varphi(\widehat{C}_0, \alpha^R \cdot \mathcal{A}_\ell, \mathcal{B}_\ell) \wedge \varphi(\widehat{C}_1, \alpha^R \cdot \mathcal{A}_\ell, \widehat{B}_1) \\ & \wedge \varphi(\widehat{C}_2, \mathcal{A}_\ell, \widehat{B}_2) \wedge \varphi(C_{\text{last}}, \mathcal{B}_u, \alpha^{R-1} \cdot \mathcal{A}_\ell) \\ & \wedge \forall \mathbf{Y}\colon \Big(\bigwedge_{i=0}^{4} \psi(B_i, B_{i+1}) \wedge \bigwedge_{i=0}^{5} \varphi(C_i, A_i, B_i) \wedge A_2 < \alpha^R \cdot \mathcal{A}_\ell \Big) \\ & \qquad \Longrightarrow \Phi(C_0, A_0, B_0, \ldots, C_5, A_5, B_5)\end{aligned}$$

where Φ implements the transition function of the machine. Note that $\mathcal{A}_\ell, \mathcal{A}_u, \mathcal{B}_\ell, \mathcal{B}_u$ also appear in the definitions of φ and ψ. The first row in the formula above fixes the initial configuration of the machine to $(0, 0, 1)$ by requiring that the first three elements of the sequence defined by $(\mathcal{A}_\ell, \mathcal{A}_u, \mathcal{B}_\ell, \mathcal{B}_u)$ must be $\alpha^R, \alpha^R, 1$, respectively. The second row says that the last term in the sequence must be α^{R-1}, which represents the halting instruction. The condition $A_2 < \alpha^R \cdot \mathcal{A}_\ell$ in the third row, in conjunction with $\varphi(C_2, A_2, B_2)$, ensures that the term of the sequence at the position defined by B_2 represents an instruction number, as opposed to a counter value. Thus $(C_0, A_0, B_0), \ldots, (C_5, A_5, B_5)$ represent two consecutive configurations of the machine. Regarding Φ, observe that we can define a function mapping α^n to α^{n+1} (which corresponds to incrementing a counter) by the formula $\chi(x, y) := y = \underbrace{x + \cdots + x}_{\alpha \text{ times}}$, and a function mapping α^{n+1} to α^n (corresponding to decrementing a counter) by $\widetilde{\chi}(x, y) := \chi(y, x)$. Finally, to see that the formula above has quantifier alternation depth 2 (i.e., three alternating blocks of quantifiers), recall that $\chi_1 \Rightarrow \chi_2$ is equivalent to $\neg\chi_1 \vee \chi_2$ and the definition of ψ involves a single universal quantifier. give Therefore, we can conclude that Theorem 5.1.2 holds: The $\exists\forall\exists$-fragment of $\mathcal{PA}(2^{\mathbb{N}}, 3^{\mathbb{N}})$ is undecidable.

Chapter 6

Conclusion

In this thesis, we investigated a range of decision problems for linear recurrence sequences with a particular emphasis on the Skolem and Positivity problems. We also explored expansions of various logics, in particular monadic second-order logic and Presburger arithmetic, in combination with predicates derived from linear recurrence sequences.

We start by discussing the Positivity problem, where we are pessimistic about the prospects of any further progress. Recall from Section 3.2 that the Positivity and Ultimate Positivity problems are 'Diophantine-hard', implying that their decidability would entail that certain mathematical constants are computable for which no known method currently exists. Moreover, these issues already arise for order-6 LRS, which is the lowest order for which we cannot resolve all instances of the Positivity and Ultimate Positivity problems. Consequently, making any substantial advances on the (Ultimate) Positivity problem seems out of reach. Even when restricting ourselves to simple LRS where this Diophantine hardness does not occur, the tools to approach this case have not yet been developed.

Compared to the Positivity problem, we are more optimistic about potential progress towards the Skolem problem. To show the decidability of fragments of the Positivity problem, one has to rely on the growth of the LRS, where Baker's theorem is the strongest (and only) existing method. For the Skolem problem, one can also leverage *local* methods. The results of Mignotte, Shorey, Tijdeman and Vereshchagin [116, 156] rely on Baker's theorem for complex logarithms and a variant for p-adic logarithms. We presented another local method in Chapter 2, where we examined LRS modulo integers. Although we relied on the Skolem and p-adic Schanuel's conjectures to guarantee termination, we expanded the repertoire of techniques for the Skolem problem. While the Skolem conjecture does not apply to non-simple LRS,

similar arguments sometimes still apply to non-simple LRS. Hence, we are far more optimistic about the possibilities for progress towards a positive solution of the Skolem problem than of the Positivity problem.

As discussed in Section 2.7, instead of focusing on natural numbers n for which $u_n = 0$, one can extend the definition of a zero of an LRS to the p-adic numbers. A non-degenerate, non-zero LRS has only finitely many p-adic zeros that can be explicitly approximated [12]. Empirically, contrary to classical Diophantine equations, for most non-degenerate, non-zero LRS $(u_n)_{n=0}^{\infty}$, the largest non-trivial index $n \in \mathbb{N}$ such that $u_n = 0$ tends to be rather small. Hence, these tools allow one to rapidly identify all zeros of an LRS while being reasonably confident that the remaining terms are non-zero. Otherwise, any additional $n \in \mathbb{N}$ such that u_n has to correspond with a p-adic zero and thus has to be exceedingly large.

Thus, for the Skolem problem, compared to the Positivity problem, we have a large toolbox and can reasonably estimate the probable outcome of a specific instance.

Other classes of fundamental sequences beyond LRS represent an additional, largely unexplored direction. Let Fac denote $\{n! \colon n \in \mathbb{N}\}$. Then it is currently unknown whether theories like $\mathsf{MSO}_{\mathbb{N},<}(2^{\mathbb{N}}, \mathsf{Fac})$ and $\mathcal{PA}(2^{\mathbb{N}}, \mathsf{Fac})$ are decidable. The factorials are in a more general class of sequences, the holonomic or P-finite sequences, for which one can also formulate the Skolem and Positivity problems (see, for example, [81, 90]).

In addition to the Skolem and Positivity problems, we have studied combinations of logic and linear recurrence sequences, which can be placed within the broader field of arithmetical theories. This field explores the interplay of logic and arithmetic. Since the dawn of computer science, arithmetical theories have been studied, leading to the undecidability of Hilbert's tenth problem and the development of automata theory to understand Presburger arithmetic and monadic second-order logic. Despite the field's long history, notable progress has been made in recent years. For example, Hieronymi and Schulz established the undecidability of $\mathcal{PA}(2^{\mathbb{N}}, 3^{\mathbb{N}})$ and Chistikov, Mansutti, and Starchak [49] showed that the existential fragment of $\mathcal{PA}(V_2, n \mapsto 2^n)$ (so-called Büchi-Semënov arithmetic) is in NP.

In Chapter 5, we examined Presburger arithmetic expanded with the sets of powers of 2 and 3, which are examples of 2- and 3-*recognizable sets*, respectively.[1] As a consequence of Hieronymi and Schulz's result that $\mathcal{PA}(2^{\mathbb{N}}, 3^{\mathbb{N}})$ is undecidable, the

[1] A set $X \subset \mathbb{N}$ is k-recognizable if there is a deterministic finite automaton over the alphabet $\{0, \ldots, k-1\}$ whose language corresponds with the base-k representations of the elements of X. Equivalently, X is k-recognizable if and only if X be defined in base-k Büchi arithmetic.

theory $\mathcal{PA}(X_2, X_3)$ is undecidable when X_2 is 2-recognizable, X_3 is 3-recognizable, and both are not definable in Presburger arithmetic alone. Barely anything is known about the existential fragments of such expansions. It is even unknown if one can determine whether the intersection of X_2 and X_3 is empty. In this problem, one can encode a conjecture of Erdős [63]: Does the digit 2 appear in the base-3 expansion of 2^n for all large enough n? Here, the set of powers of 2 is 2-recognizable and the set of numbers not having a 2 in its base-3 expansion is also 3-recognizable. More generally, one could ask to decide certain fragments of $\mathcal{PA}(V_2, V_3)$, where one combines Büchi arithmetic in certain bases. For example although the full theory is undecidable, the decidability of the existential fragment is still open.

Many other open problems for arithmetic theories remain, and many such instances combine number-theoretic, combinatoric, and automata-theoretic properties, which creates an exciting, interdisciplinary field of research.

Bibliography

[1] Boris Adamczewski and Yann Bugeaud. On the complexity of algebraic numbers I. Expansions in integer bases. *Annals of Mathematics*, 165(2):547–565, 2007.

[2] Manindra Agrawal, Sundararaman Akshay, Blaise Genest, and PS Thiagarajan. Approximate verification of the symbolic dynamics of markov chains. *Journal of the ACM (JACM)*, 62(1):1–34, 2015.

[3] S. Akshay, Timos Antonopoulos, Joël Ouaknine, and James Worrell. Reachability problems for Markov chains. *Information Processing Letters*, 115(2):155–158, 2015.

[4] S. Akshay, Nikhil Balaji, and Nikhil Vyas. Complexity of Restricted Variants of Skolem and Related Problems. In *42nd International Symposium on Mathematical Foundations of Computer Science (MFCS 2017)*, volume 83, pages 78:1–78:14. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2017.

[5] S. Akshay, Hugo Bazille, Blaise Genest, and Mihir Vahanwala. On Robustness for the Skolem and Positivity Problems. In *STACS 2022-39th International Symposium on Theoretical Aspects of Computer Science*, 2022.

[6] Jean-Paul Allouche and Jeffrey Shallit. *Automatic sequences: Theory, Applications, Generalizations*. Cambridge University Press, 2003.

[7] Shaull Almagor, Dmitry Chistikov, Joël Ouaknine, and James Worrell. O-Minimal Invariants for Linear Loops. In *45th International Colloquium on Automata, Languages, and Programming (ICALP 2018)*, volume 107, pages 114:1–114:14. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2018.

[8] Shaull Almagor, Toghrul Karimov, Edon Kelmendi, Joël Ouaknine, and James Worrell. Deciding ω-regular properties on linear recurrence sequences. *Proceedings of the ACM on Programming Languages*, 5(POPL):1–24, 2021.

[9] Pierre Arnoux, Christian Mauduit, Iekata Shiokawa, and Jun-Ichi Tamura. Complexity of sequences defined by billiard in the cube. *Bulletin de la Société Mathématique de France*, 122(1):1–12, 1994.

[10] Piotr Bacik. Completing the Picture for the Skolem Problem on Order-4 Linear Recurrence Sequences. *arXiv preprint arXiv:2409.01221*, 2024.

[11] Piotr Bacik, Yuri Bilu, Florian Luca, Joris Nieuwveld, Joël Ouaknine, David Purser, and James Worrell. The SKOLEM-tool, 2025. https://skolem.mpi-sws.org/.

[12] Piotr Bacik, Joël Ouaknine, David Purser, and James Worrell. On the p-adic Skolem Problem. *arXiv preprint arXiv:2504.14413*, 2025.

[13] C. Baier, F. Funke, S. Jantsch, T. Karimov, E. Lefaucheux, F. Luca, J. Ouaknine, D. Purser, M. A. Whiteland, and J. Worrell. The Orbit Problem for parametric linear dynamical systems. In *32nd International Conference on Concurrency Theory, CONCUR 2021*, volume 203 of *LIPIcs*, pages 28:1–28:17. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2021.

[14] David H. Bailey, Jonathan M. Borwein, Cristian S. Calude, Michael J. Dinneen, Monica Dumitrescu, and Alex Yee. An Empirical Approach to the Normality of π. *Experimental Mathematics*, 21(4):375–384, 2012.

[15] Gilles Barthe, Charlie Jacomme, and Steve Kremer. Universal equivalence and majority of probabilistic programs over finite fields. *ACM Transactions on Computational Logic (TOCL)*, 23(1):1–42, 2021.

[16] B. Bartolome, Y. Bilu, and F. Luca. On the exponential local-global principle. *Acta Arithmetica*, 159(2):101–111, 2013.

[17] Yu. Baryshnikov. Complexity of Trajectories in Rectangular Billiards. *Communications in mathematical physics*, 174:43–56, 1995.

[18] Paul T. Bateman, Carl G. Jockusch, and Alan R. Woods. Decidability and Undecidability of Theories with a Predicate for the Primes. *The Journal of Symbolic Logic*, 58(2):672–687, 1993.

[19] Nicolas Bedaride. Directional complexity of the hypercubic billiard. *Discrete mathematics*, 309(8):2053–2066, 2009.

[20] Paul C. Bell, Jean-Charles Delvenne, Raphaël M. Jungers, and Vincent D. Blondel. The continuous Skolem-Pisot problem. *Theoretical Computer Science*, 411(40-42):3625–3634, 2010.

[21] Amir M. Ben-Amram, Samir Genaim, and Abu Naser Masud. On the Termination of Integer Loops. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 34(4):1–24, 2012.

[22] Michael Benedikt, Dmitry Chistikov, and Alessio Mansutti. The complexity of presburger arithmetic with power or powers. In *50th International Colloquium*

on Automata, Languages, and Programming (ICALP 2023), volume 261, page 112. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2023.

[23] Jean Berstel and Christophe Reutenauer. *Noncommutative Rational Series with Applications*. Encyclopedia of Mathematics and its Applications. Cambridge University Press, 2011.

[24] Valérie Berthé, Toghrul Karimov, Joris Nieuwveld, Joël Ouaknine, Mihir Vahanwala, and James Worrell. On the Decidability of Monadic Second-Order Logic with Arithmetic Predicates. In *2024 39th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 11:1–11:14. ACM, 2024.

[25] Valérie Berthé, Toghrul Karimov, Joris Nieuwveld, Joël Ouaknine, Mihir Vahanwala, and James Worrell. The Monadic Theory of Toric Words. *Theoretical Computer Science*, 1025:114959, 2025.

[26] Csanád Bertók and Lajos Hajdu. A Hasse-type principle for exponential diophantine equations over number fields and its applications. *Monatshefte für Mathematik*, 187(3):425–436, 2018.

[27] Csanád Bertók, Lajos Hajdu, Florian Luca, and Divyum Sharma. On the number of non-zero digits of integers in multi-base representations. *Publicationes Mathematicae Debrecen*, 90:181–194, 01 2017.

[28] Alexis Bès. Undecidable Extensions of Büchi Arithmetic and Cobham-Semënov Theorem. *The Journal of Symbolic Logic*, 62(4):1280–1296, 1997.

[29] Alexis Bès. A Survey of Arithmetical Definability. *Bulletin de la Société Mathématique de Belgique*, pages 1–54, 2002. A Tribute to Maurice Boffa.

[30] Yuri Bilu. Skolem Problem for Linear Recurrence Sequences with 4 Dominant Roots (after Mignotte, Shorey, Tijdeman, Vereshchagin and Bacik). *arXiv preprint arXiv:2501.16290*, 2025.

[31] Yuri Bilu, Florian Luca, Joris Nieuwveld, Joël Ouaknine, David Purser, and James Worrell. Skolem Meets Schanuel. In *47th International Symposium on Mathematical Foundations of Computer Science*, volume 241, pages 20:1–20:15. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2022.

[32] Yuri Bilu, Florian Luca, Joris Nieuwveld, Joël Ouaknine, and James Worrell. On the p-adic zeros of the Tribonacci sequence. *Mathematics of Computation*, 93(347):1333–1353, 2024.

[33] Yuri Bilu, Florian Luca, Joris Nieuwveld, Joël Ouaknine, and James Worrell. Twisted rational zeros of linear recurrence sequences. *Journal of the London Mathematical Society*, 111(3):e70126, 2025.

[34] V. Blondel and J. Tsitsiklis. A survey of computational complexity results in systems and control. *Automatica*, 36(9):1249–1274, 2000.

[35] Vincent D. Blondel and Natacha Portier. The presence of a zero in an integer linear recurrent sequence is NP-hard to decide. *Linear algebra and its Applications*, 351:91–98, 2002.

[36] Achim Blumensath. Monadic Second-Order Model Theory. `http://www.fi.muni.cz/~blumens/`, 2024. [Online; accessed on 24 January 2025].

[37] Émile M. Borel. Les probabilités dénombrables et leurs applications arithmétiques. *Rendiconti del Circolo Matematico di Palermo (1884-1940)*, 27(1):247–271, 1909.

[38] Marius Bozga, Radu Iosif, and Filip Konečnỳ. Deciding Conditional Termination. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, pages 252–266. Springer, 2012.

[39] Mark Braverman. Termination of Integer Linear Programs. In *CAV*, volume 4144 of *Lecture Notes in Computer Science*, pages 372–385. Springer, 2006.

[40] Bradley W. Brock, Noam D. Elkies, and Bruce W. Jordan. Periodic continued fractions over s-integers in number fields and Skolem's p-adic method. *Acta Arithmetica*, 197:379–420, 2021.

[41] Armand Brumer. On the units of algebraic number fields. *Mathematika*, 14(2):121–124, 1967.

[42] J. Richard Büchi. Weak second-order arithmetic and finite automata. *Mathematical Logic Quarterly*, 6(1–6), 1960.

[43] J. Richard Büchi. On a Decision Method in Restricted Second Order Arithmetic. In *The Collected Works of J. Richard Büchi*, pages 425–435. Springer New York, 1990.

[44] J. Richard Büchi and Lawrence H. Landweber. Definability in the Monadic Second-Order Theory of Successor. *The Journal of Symbolic Logic*, 34(2):166–170, 1969.

[45] Yann Bugeaud. *Distribution modulo one and Diophantine approximation*, volume 193 of *Cambridge Tracts in Mathematics*. Cambridge University Press, 2012.

[46] F. Calegari and B. Mazur. Nearly ordinary Galois deformations over arbitrary number fields. *J. Inst. Math. Jussieu*, 8(1):99–177, 2009.

[47] Olivier Carton and Wolfgang Thomas. The Monadic Theory of Morphic Infinite Words and Generalizations. *Information and Computation*, 176(1):51–65, 2002.

[48] Krishnendu Chatterjee and Laurent Doyen. Stochastic processes with expected stopping time. *Logical Methods in Computer Science*, 20, 2024.

[49] Dmitry Chistikov, Alessio Mansutti, and Mikhail R. Starchak. Integer Linear-Exponential Programming in NP by Quantifier Elimination. In *51st International Colloquium on Automata, Languages, and Programming (ICALP 2024)*, volume 297, pages 132:1–132:20. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2024.

[50] V. Chonev, J. Ouaknine, and J. Worrell. The Polyhedron-Hitting Problem. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015*, pages 940–956. SIAM, 2015.

[51] Ventsislav Chonev, Joël Ouaknine, and James Worrell. On the Skolem Problem for Continuous Linear Dynamical Systems. In *43rd International Colloquium on Automata, Languages, and Programming (ICALP 2016)*, volume 55, pages 100:1–100:13, 2016.

[52] Benoit Cloitre and Jeffrey Shallit. Some fibonacci-related sequences. *arXiv preprint arXiv:2312.11706*, 2023.

[53] Henri Cohen. *A course in computational algebraic number theory*, volume 138 of *Graduate Texts in Mathematics*. Springer-Verlag, Berlin, 1993.

[54] George B. Dantzig and B. Curtis Eaves. Fourier-Motzkin elimination and its dual. *J. Comb. Theory, Ser. A*, 14(3):288–297, 1973.

[55] Julian D'Costa, Toghrul Karimov, Rupak Majumdar, Joël Ouaknine, Mahmoud Salamati, and James Worrell. The Pseudo-Reachability Problem for Diagonalisable Linear Dynamical Systems. *arXiv preprint arXiv:2204.12253*, 2022.

[56] Harm Derksen. A Skolem–Mahler–Lech theorem in positive characteristic and finite automata. *Inventiones mathematicae*, 168(1):175–224, 2007.

[57] Harm Derksen and David Masser. Linear equations over multiplicative groups, recurrences, and mixing II. *Indagationes Mathematicae*, 26(1):113–136, 2015.

[58] The SageMath Developers. *SageMath, the Sage Mathematics Software System (Version 9.7)*, 2022. `https://www.sagemath.org`.

[59] A. Dubickas and C. J. Smyth. On the Remak height, the Mahler measure and conjugate sets of algebraic numbers lying on two circles. *Proc. Edinb. Math. Soc. (2)*, 44(1):1–17, 2001.

[60] Artūras Dubickas and Min Sha. Positive density of integer polynomials with some prescribed properties. *Journal of Number Theory*, 159:27–44, 2016.

[61] Andrej Dujella and Attila Pethő. A Generalization of a Theorem of Baker and Davenport. *The Quarterly Journal of Mathematics*, 49(195):291–306, 1998.

[62] Calvin C. Elgot and Michael O. Rabin. Decidability and Undecidability of Extensions of Second (First) Order Theory of (Generalized) Successor. *The Journal of Symbolic Logic*, 31(2):169–181, 1966.

[63] Paul Erdős. Some unconventional problems in number theory. *Mathematics Magazine*, 52(2):67–70, 1979.

[64] Graham Everest, Alf van der Poorten, Igor Shparlinski, and Thomas Ward. *Recurrence Sequences*, volume 104. American Mathematical Society Providence, RI, 2003.

[65] Jan-Hendrik Evertse. On sums of S-units and linear recurrences. *Compositio Mathematica*, 53(2):225–244, 1984.

[66] P. Fatou. Sur les séries entières à coefficients entiers. *Comptes Rendus Acad. Sci. Paris*, 138(130):342–344, 1904.

[67] Ronald Ferguson. Irreducible polynomials with many roots of equal modulus. *Acta Arithmetica*, 78(3):221–225, 1997.

[68] N. Fijalkow, J. Ouaknine, A. Pouly, J. Sousa Pinto, and J. Worrell. On the decidability of reachability in linear time-invariant systems. In *Proceedings of the 22nd ACM International Conference on Hybrid Systems: Computation and Control, HSCC 2019*, pages 77–86. ACM, 2019.

[69] N. Pytheas Fogg, Valérie Berthé, Sébastien Ferenczi, Christian Mauduit, and Anne Siegel. *Substitutions in dynamics, arithmetics and combinatorics*. Springer, 2002.

[70] Kurt Gödel. Über formal unentscheidbare Sätze der Principia Mathematica und verwandter Systeme I. *Monatshefte für Mathematik und Physik*, 38:173–198, 1931.

[71] Steven M. Gonek and Hugh L. Montgomery. Kronecker's approximation theorem. *Indagationes Mathematicae*, 27(2):506–523, 2016. In Memoriam J.G. Van der Corput (1890–1975) Part 2.

[72] Fernando Q. Gouvêa. *p-adic Numbers*. Springer Berlin Heidelberg, 1997.

[73] Quentin Guilmant, Engel Lefaucheux, Joël Ouaknine, and James Worrell. The 2-Dimensional Constraint Loop Problem Is Decidable. In *51st International Colloquium on Automata, Languages, and Programming (ICALP 2024)*, volume 297, pages 140:1–140:21. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2024.

[74] Christoph Haase. A Survival Guide to Presburger Arithmetic. *ACM SIGLOG News*, 5(3):67–82, July 2018.

[75] Vesa Halava and Tero Harju. Mortality in Matrix Semigroups. *The American Mathematical Monthly*, 108(7):649–653, 2001.

[76] Vesa Halava, Tero Harju, and Mika Hirvensalo. Positivity of second order linear recurrent sequences. *Discrete Applied Mathematics*, 154(3):447–451, 2006.

[77] Glyn Harman. One Hundred Years of Normal Numbers. In *Surveys in Number Theory*, pages 57–74. A K Peters/CRC Press, 2002.

[78] Philipp Hieronymi and Christian Schulz. A Strong Version of Cobham's Theorem. In *STOC*, pages 1172–1179. ACM, 2022.

[79] Tony Hoare. The verifying compiler: A grand challenge for computing research. *Journal of the ACM (JACM)*, 50(1):63–69, 2003.

[80] Mehran Hosseini, Joël Ouaknine, and James Worrell. Termination of Linear Loops over the Integers. In *46th International Colloquium on Automata, Languages, and Programming*, volume 132, pages 118:1–118:13. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2019.

[81] Alaa Ibrahim and Bruno Salvy. Positivity Proofs for Linear Recurrences through Contracted Cones. *Journal of Symbolic Computation*, page 102463, 2025.

[82] Ravindran Kannan and Richard J. Lipton. The orbit problem is decidable. In *Proceedings of the twelfth annual ACM symposium on Theory of computing*, STOC '80, pages 252–261. Association for Computing Machinery, 1980.

[83] Toghrul Karimov. *Algorithmic Verification of Linear Dynamical Systems*. PhD thesis, Saarland University, 2023.

[84] Toghrul Karimov, Edon Kelmendi, Joris Nieuwveld, Joël Ouaknine, and James Worrell. The Power of Positivity. In *2023 38th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 1–11. IEEE, 2023.

[85] Toghrul Karimov, Engel Lefaucheux, Joël Ouaknine, David Purser, Anton Varonka, Markus A. Whiteland, and James Worrell. What's decidable about linear loops? *Proc. ACM Program. Lang.*, 6(POPL):1–25, 2022.

[86] Toghrul Karimov, Florian Luca, Joris Nieuwveld, Joël Ouaknine, and James Worrell. On the Decidability of Presburger Arithmetic Expanded with Powers. In *Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2755–2778. SIAM, 2025.

[87] Toghrul Karimov, Joris Nieuwveld, Joël Ouaknine, Mihir Vahanwala, and James Worrell. Algorithmic Applications of Schanuel's Conjecture. preprint.

[88] George Kenison. On the Skolem Problem for Reversible Sequences. In *47th International Symposium on Mathematical Foundations of Computer Sci-*

ence (MFCS 2022), volume 241, pages 61:1–61:15. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2022.

[89] George Kenison, Joris Nieuwveld, Joël Ouaknine, and James Worrell. Positivity Problems for Reversible Linear Recurrence equences. In *50th International Colloquium on Automata, Languages, and Programming (ICALP 2023)*, volume 261, pages 130:1–130:17. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2023.

[90] George Kenison, Klara Nosan, Mahsa Shirmohammadi, and James Worrell. The Membership Problem for Hypergeometric Sequences with Quadratic Parameters. In *Proceedings of the 2023 International Symposium on Symbolic and Algebraic Computation*, pages 407–416. Association for Computing Machinery, 2023.

[91] L. Kronecker. Zwei Sätze über Gleichungen mit ganzzahligen Coefficienten. *Journal für die reine und angewandte Mathematik*, 53:173–175, 1857.

[92] Serge Lang. *Introduction to Transcendental Numbers*. Addison-Wesley, 1966.

[93] Vichian Laohakosol and Pinthira Tangsupphathawat. Positivity of third order linear recurrence sequences. *Discrete Applied Mathematics*, 157(15):3239–3248, 2009.

[94] Christer Lech. A note on recurring series. *Arkiv för Matematik*, 2(5):417–421, 1953.

[95] Engel Lefaucheux, Joël Ouaknine, David Purser, and Mohammadamin Sharifi. Model checking linear dynamical systems under floating-point rounding. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, pages 47–65. Springer, 2023.

[96] Tamás Lengyel. The order of the Fibonacci and Lucas numbers. *The Fibonacci Quarterly*, 33(3):234–239, 1995.

[97] Arjen K. Lenstra, Hendrik W. Lenstra, and László Lovász. Factoring Polynomials with Rational Coefficients. *Mathematische Annalen*, 261(4):515–534, 1982.

[98] Aristid Lindenmayer and Grzegorz Rozenberg, editors. *Automata, Languages, Development*. North-Holland Publishing Company, 1976.

[99] Richard Lipton, Florian Luca, Joris Nieuwveld, Joël Ouaknine, David Purser, and James Worrell. On the Skolem Problem and the Skolem Conjecture. In *37th Annual ACM/IEEE Symposium on Logic in Computer Science*, pages 1–9. ACM, 2022.

[100] Richard J. Lipton. Mathematical embarrassments. In *The* P = NP *Question and Gödel's Lost Letter*, pages 209–213. Springer, 2010.

[101] M. Lothaire. *Algebraic Combinatorics on Words*, volume 90. Cambridge University Press, 2002.

[102] Florian Luca. Exponential Diophantine equations. *Notes from the International Autumn School on Computational Number Theory*, pages 267–309, 2019.

[103] Florian Luca. Personal communication, 2022.

[104] Florian Luca, James Maynard, Armand Noubissie, Joël Ouaknine, and James Worrell. Skolem Meets Bateman-Horn. *arXiv preprint arXiv:2308.01152*, 2023.

[105] Florian Luca, Joël Ouaknine, and James Worrell. On Large Zeros of Linear Recurrence Sequences. preprint.

[106] Florian Luca, Joël Ouaknine, and James Worrell. Universal Skolem Sets. In *2021 36th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 1–6. IEEE, 2021.

[107] Florian Luca, Joël Ouaknine, and James Worrell. Algebraic Model Checking for Discrete Linear Dynamical Systems. In *International Conference on Formal Modeling and Analysis of Timed Systems*, pages 3–15. Springer, 2022.

[108] Florian Luca, Joël Ouaknine, and James Worrell. A Universal Skolem Set of Positive Lower Density. In *47th International Symposium on Mathematical Foundations of Computer Science*, volume 241, pages 73:1–73:12. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2022.

[109] A. Macintyre and A. J. Wilkie. On the Decidability of the Real Exponential Field. In *Kreiseliana. About and Around Georg Kreisel*, pages 441–467. A K Peters, 1996.

[110] Kurt Mahler. Eine arithmetische Eigenschaft der Taylor-koeffizienten rationaler Funktionen. *Proceedings of the Koninklijke Nederlandse Akademie van Wetenschappen*, 38:56–60, 1935.

[111] Diego Marques and Tamás Lengyel. The 2-adic Order of the Tribonacci Numbers and the Equation $T_n = m!$. *Journal of Integer Sequences*, 17(2):3, 2014.

[112] David Masser. Linear relations on algebraic groups. *New Advances in Transcendence Theory*, pages 248–262, 1988.

[113] David Masser. Alan Baker, FRS, 1939–2018. *Bulletin of the London Mathematical Society*, 53(6):1916–1949, 2021.

[114] Yuri V. Matiyasevich. *Hilbert's Tenth Problem*. Foundations of Computing. MIT Press, Cambridge, MA, 1993.

[115] Eugene M. Matveev. An explicit lower bound for a homogeneous rational linear form in the logarithms of algebraic numbers. II. *Izvestiya: Mathematics*, 64(6):1217, 2000.

[116] Maurice Mignotte, Tarlok Shorey, and Robert Tijdeman. The distance between terms of an algebraic recurrence sequence. *Journal für die Reine und Angewandte Mathematik*, 349:63–76, 1984.

[117] Marvin Lee Minsky. *Computation: Finite and Infinite Machines*. Prentice-Hall Englewood Cliffs, 1967.

[118] An. Muchnik, A. Semenov, and M. Ushakov. Almost periodic sequences. *Theoretical Computer Science*, 304(1):1–33, 2003.

[119] Joris Nieuwveld and Joël Ouaknine. On Expansions of Monadic Second-Order Logic with Dynamical Predicates. preprint.

[120] Joël Ouaknine and James Worrell. On the Positivity Problem for Simple Linear Recurrence Sequences. In *International Colloquium on Automata, Languages, and Programming (ICALP 2014)*, pages 318–329. Springer, 2014.

[121] Joël Ouaknine and James Worrell. Positivity Problems for Low-Order Linear Recurrence Sequences. In *Proceedings of the twenty-fifth annual ACM-SIAM Symposium on Discrete Algorithms*, pages 366–379. SIAM, 2014.

[122] Joël Ouaknine and James Worrell. Ultimate Positivity is Decidable for Simple Linear Recurrence Sequences. In *International Colloquium on Automata, Languages, and Programming (ICALP 2014)*, pages 330–341. Springer, 2014.

[123] Joël Ouaknine and James Worrell. On Linear Recurrence Sequences and Loop Termination. *ACM SIGLOG News*, 2(2):4–13, April 2015.

[124] Jakob Piribauer and Christel Baier. On Skolem-Hardness and Saturation Points in Markov Decision Processes. In *47th International Colloquium on Automata, Languages, and Programming (ICALP 2020)*, volume 168, pages 138:1–138:17. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2020.

[125] Jakob Piribauer and Christel Baier. Positivity-hardness results on Markov decision processes. *TheoretiCS*, 3, 2024.

[126] Andreas Podelski and Andrey Rybalchenko. A Complete Method for the Synthesis of Linear Ranking Functions. In *International Workshop on Verification, Model Checking, and Abstract Interpretation*, pages 239–251. Springer, 2004.

[127] Alf J. van der Poorten. Some problems of recurrent interest. Technical report, Technical Report 81-0037, School of Mathematics and Physics, Macquarie, 1981.

[128] Mojzesz Presburger. Über die Vollständigkeit eines gewissen Systems der Arithmetik ganzer Zahlen, in welchen die Addition als einzige Operation hervortritt. In *Comptes Rendus du ler Congrès de Mathématiciens des Pays Slaves*, 1929.

[129] Martine Queffélec. Old and new results on normality. *Dynamics & Stochastics: Festschrift in Honour of MS Keane*, 48:225, 2006.

[130] Alexander Rabinovich. On decidability of monadic logic of order over the naturals extended by monadic predicates. *Information and Computation*, 205(6):870–889, 2007.

[131] Alexander Rabinovich and Wolfgang Thomas. Decidable Theories of the Ordering of Natural Numbers with Unary Predicates. In *International Workshop on Computer Science Logic*, pages 562–574. Springer, 2006.

[132] J. J. Rotman. *An Introduction to the Theory of Groups*. Graduate Texts in Mathematics. Springer New York, 2012.

[133] Grzegorz Rozenberg and Arto Salomaa. *Cornerstones of undecidability*. Prentice-Hall, Inc., 1995.

[134] Arto Salomaa. Formal power series and growth functions of Lindenmayer systems. In *International Symposium on Mathematical Foundations of Computer Science*, pages 101–113. Springer, 1975.

[135] Andrzej Schinzel. Abelian binomials, power residues and exponential congruences. *Acta Arithmetica*, 32(3):245–274, 1977.

[136] Andrzej Schinzel. On the congruence $u_n \equiv c \pmod{\mathfrak{p}}$ where u_n is a recurring sequence of the second order. *Acta Acad. Paedagog. Agriensis Sect. Math.*, 30:147–165, 2003.

[137] Andrzej Schinzel and Robert Tijdeman. On the equation $y^m = P(x)$. *Acta Arithmetica*, 31:199–204, 1976.

[138] Chris Schulz. Undefinability of multiplication in Presburger arithmetic with sets of powers. *The Journal of Symbolic Logic*, pages 1–15, 2023.

[139] A. L. Semenov. On certain extensions of the arithmetic of addition of natural numbers. *Mathematics of The USSR-Izvestiya*, 15:401–418, 1980.

[140] A. L. Semenov. Logical theories of one-place functions on the set of natural numbers. *Mathematics of the USSR-Izvestiya*, 22(3):587, 1984.

[141] Jeffrey Shallit. *The Logical Approach to Automatic Sequences: Exploring Combinatorics on Words with Walnut*, volume 482. Cambridge University Press, 2022.

[142] Saharon Shelah. The monadic theory of order. *Annals of Mathematics*, 102(3):379–419, 1975.

[143] Thoralf Skolem. Einige Sätze über gewisse Reihenentwicklungen und exponentiale Beziehungen mit Anwendung auf Diophantische Gleichungen. *Norsk videnskaps-akademi i Oslo. Skrifter. I*, 6, 1933.

[144] Thoralf Skolem. Einige Sätze über π-adische Potenzreihen mit Anwendung auf gewisse exponentielle Gleichungen. *Mathematische Annalen*, 111(1):399–424, 1935.

[145] Thoralf Skolem. Anwendung exponentieller Kongruenzen zum Beweis der Unlösbarkeit gewisser diophantischer Gleichungen. *Avhandlinger fra Det Norske Videnskaps-Akademi Oslo I*, 12:1–16, 1937.

[146] C. J. Smyth. Conjugate algebraic numbers on conics. *Acta Arithmetica*, 40(4):333–346, 1982.

[147] M. Soittola. On D0L synthesis problem. In A. Lindenmayer and G. Rozenberg, editors, *Automata, Languages, Development*. North-Holland, 1976.

[148] C. L. Stewart. On the representation of an integer in two different bases. *Journal für die reine und angewandte Mathematik*, 319:63–72, 1980.

[149] R. J. Stroeker and N. Tzanakis. On the Application of Skolem's p-adic Method to the Solution of Thue Equations. *Journal of Number Theory*, 29(2):166–195, 1988.

[150] Terence Tao. Structure and randomness: pages from year one of a mathematical blog, 2008.

[151] Alfred Tarski. A Decision Method for Elementary Algebra and Geometry. In *Quantifier elimination and cylindrical algebraic decomposition*, pages 24–84. Springer, 1998.

[152] Wolfgang Thomas. Ehrenfeucht Games, the Composition Method, and the Monadic Theory of Ordinal Words. In *Structures in Logic and Computer Science*, volume 1261 of *Lecture Notes in Computer Science*, pages 118–143. Springer, 1997.

[153] Wolfgang Thomas. Languages, Automata, and Logic. In *Handbook of Formal Languages: Volume 3 Beyond Words*, pages 389–455. Springer, 1997.

[154] Ashish Tiwari. Termination of Linear Programs. In *Computer Aided Verification: 16th International Conference, CAV 2004, July 13-17, 2004. Proceedings 16*, pages 70–82. Springer, 2004.

[155] Mihir Vahanwala. Skolem and Positivity Completeness of Ergodic Markov Chains. *Information Processing Letters*, 186:106481, 2024.

[156] Nikolai Vereshchagin. Occurrence of zero in a linear recursive sequence. *Mathematical notes of the Academy of Sciences of the USSR*, 38(2):609–615, 1985.

[157] Roger Villemaire. The Theory of ($\mathbb{N}$, $+$, V_k, V_l) is Undecidable. *Theoretical Computer Science*, 106(2):337–349, 1992.

[158] Michel Waldschmidt. *Diophantine Approximation on Linear Algebraic Groups: Transcendence Properties of the Exponential Function in Several Variables.* Springer Berlin Heidelberg, Berlin, Heidelberg, 2000.

[159] Kunrui Yu. Report on p-adic logarithmic forms. *A Panorama of Number Theory or the View from Baker's Garden (Zürich, 1999), Cambridge University Press, Cambridge*, pages 11–25, 2002.

[160] Richard Zach. Hilbert's Program Then and Now. In *Philosophy of logic*, pages 411–447. Elsevier, 2007.