Decision Problems for Linear Recurrence Sequences

Joël Ouaknine and James Worrell

Department of Computer Science, Oxford University, UK
{joel,jbw}@cs.ox.ac.uk

Abstract. Linear recurrence sequences permeate a vast number of areas of mathematics and computer science. In this paper, we survey the state of the art concerning certain fundamental *decision problems* for linear recurrence sequences, namely the Skolem Problem (does the sequence have a zero?), the Positivity Problem (is the sequence always positive?), and the Ultimate Positivity Problem (is the sequence ultimately always positive?).

1 Introduction

A ***linear recurrence sequence*** is an infinite sequence $\langle u_0, u_1, u_2, \ldots \rangle$ of numbers having the following property: there exist constants $a_1, a_2, \ldots, a_k$ such that, for all n, $u_{n+k} = a_1 u_{n+k-1} + a_2 u_{n+k-2} + \ldots + a_k u_n$. If the initial values $u_0, u_1, \ldots, u_{k-1}$ of the sequence are provided, the recurrence relation defines the rest of the sequence uniquely.

The best-known example of a linear recurrence sequence was provided by Leonardo of Pisa in the 12th century: the so-called *Fibonacci sequence* $\langle 0, 1, 1, 2, 3, 5, 8, 13, \ldots \rangle$, which satisfies the recurrence relation $u_{n+2} = u_{n+1} + u_n$. Leonardo of Pisa introduced this sequence as a means to model the growth of an idealised population of rabbits. Not only has the Fibonacci sequence been extensively studied since, but linear recurrence sequences now form a vast subject in their own right, with numerous applications in mathematics and other sciences. A deep and extensive treatise on the mathematical aspects of recurrence sequences is the recent monograph of Everest *et al.* [4].

In this paper, we mainly focus on *decision problems* relating to linear recurrence sequences. In order for such problems to be well-defined, we need to place certain restrictions on the sequences under consideration. Firstly, we shall only be interested in sequences of *real* numbers, and in particular shall require that the initial values as well as all constants appearing in recurrence relations be real. We will often specialise further, by requiring for example that all initial values and constants be integral, or rational, or algebraic.

The three main decision problems which will concern us are the following:

- ***The Skolem Problem***: does a given linear recurrence sequence have a zero?

- ***The Positivity Problem***: are all the terms of a given linear recurrence sequence positive?
  (Note that this problem comes in two natural flavours, according to whether strict or non-strict positivity is required.)
- ***The Ultimate Positivity Problem***: is the given linear recurrence sequence ultimately positive, i.e., are all the terms of the sequence positive except possibly for a finite number of exceptions?
  (Note, likewise, that this problem admits two natural variants.)

The above problems (and closely related variants) have applications in many different areas, such as theoretical biology (analysis of L-systems, population dynamics), software verification (termination of linear programs), probabilistic model checking (reachability in Markov chains, stochastic logics), quantum computing (threshold problems for quantum automata), as well as combinatorics, term rewriting, and the study of generating functions. For example, a particular term of a linear recurrence sequence usually has combinatorial significance only if it is non-negative. Likewise, a linear recurrence sequence modelling population growth is biologically meaningful only if it is uniformly positive.

At the time of writing, the decidability of each of these decision problems, whether for integer, rational, or algebraic linear recurrence sequences, is open, although partial results are known. We shall review, to the best of our knowledge, the state of the art in the literature concerning these problems, and also recall a number of key facts about linear recurrence sequences.

2 Preliminaries

2.1 Linear Recurrence Sequences

We recall some fundamental properties of linear recurrence sequences. Results are stated without proofs, and we refer the reader to [4] for details.

Let $\langle u_n \rangle_{n=0}^{\infty}$ be a linear recurrence sequence satisfying the recurrence relation $u_{n+k} = a_1 u_{n+k-1} + \ldots + a_k u_n$. We say that the sequence has ***order*** k provided a_k is non-zero. (Thus the Fibonacci sequence, for example, has order 2.) The ***characteristic polynomial*** of the sequence is

$$p(x) = x^n - a_1 x^{n-1} - \ldots - a_{k-1} x - a_k \,.$$

Let $\lambda_1, \ldots, \lambda_m \in \mathbb{C}$ be a list of the distinct (possibly repeated) roots of p. Then there are complex single-variable polynomials $A_1, A_2, \ldots, A_m$ such that, for all n,

$$u_n = A_1(n)\lambda_1^n + \ldots + A_m(n)\lambda_m^n \,.$$

The A_j's are then uniquely determined by the initial values $u_0, \ldots, u_{k-1}$ of the recurrence sequence.

Any linear recurrence sequence $\langle u_n \rangle_{n=0}^{\infty}$ of order k can alternately be defined in matrix form, in the sense that there is a square matrix M of dimension k, together with k-dimensional vectors v and w, such that, for all n, $u_n =$

$v^T M^n w$. It suffices to take M to be the transpose of the companion matrix of the characteristic polynomial of the sequence, let v be the vector of initial values of the sequence (in reverse order), and take w to be the vector whose first $k-1$ entries are 0 and whose kth entry is 1. Conversely, any k-dimensional square matrix M and vectors v and w give rise to a linear recurrence sequence $\langle v^T M^n w\rangle_{n=0}^{\infty}$ of order at most k, thanks to the Cayley-Hamilton Theorem.

Let $\langle u_n\rangle_{n=0}^{\infty}$ and $\langle v_n\rangle_{n=0}^{\infty}$ be linear recurrence sequences of order k and l respectively. Their pointwise product $\langle u_n v_n\rangle_{n=0}^{\infty}$ and sum $\langle u_n + v_n\rangle_{n=0}^{\infty}$ are also linear recurrence sequences of order at most kl and $k+l$ respectively.

2.2 Algebraic Numbers

As is well known, the algebraic numbers form a field: given two algebraic numbers, their sum, product, and ratio (provided the divisor is non-zero) are again algebraic numbers. Moreover, algebraic numbers have canonical representations with respect to which these operations can be performed in polynomial time. Deciding whether two algebraic numbers are equal or whether an algebraic number is a root of unity can likewise be carried out in polynomial time, as can the problems of deciding membership in $\mathbb{Z}$, $\mathbb{Q}$, $\mathbb{R}$, and $\mathbb{R}_{\geq 0}$. The relevant key algorithms can be found, for example, in [3].

When discussing decision problems regarding linear recurrence sequences of algebraic numbers, we therefore assume that the relevant algebraic numbers are provided in some suitable effective canonical form, so that the underlying questions are algorithmically meaningful.

3 The Skolem Problem

Let us begin by stating one of the most fundamental results about the zeros of linear recurrence sequences, the celebrated Skolem-Mahler-Lech Theorem:

Theorem 1. *Let $\langle u_n\rangle_{n=0}^{\infty}$ be a linear recurrence sequence over the reals. Its set of zeros $\{n : u_n = 0\}$ consists of a finite set F, together with a finite number of arithmetic progressions $A_1 \cup \ldots \cup A_l$.*

This result is due to Skolem [15], and more general versions were subsequently obtained by Mahler [11, 12] and Lech [8]. All known proofs of the Skolem-Mahler-Lech Theorem make use of p-adic techniques, achieving the result in a *non-constructive* manner.

3.1 Decidability

As pointed out in [6], algorithmic decision issues in the 1930s had not yet acquired the importance that they have today. Still, it is customary to regard the Skolem Problem—*deciding* whether a given linear recurrence sequence has a zero or not—as having been open since the publication of Skolem's original paper. As alluded to earlier, it is necessary for this problem to be well-defined that linear recurrence sequences be given in effective form. To this end, we shall restrict our attention to linear recurrence sequences of integers, rationals, or algebraic numbers.

As mentioned above, the proof of the Skolem-Mahler-Lech Theorem is ineffective. Subsequently, however, Berstel and Mignotte showed how to obtain all the arithmetic progressions mentioned in the theorem effectively [1]. The critical case therefore boils down to linear recurrence sequences provably having a *finite* number of zeros, in which case one must decide whether that finite set is empty or not. As opined by Terence Tao, "[i]t is faintly outrageous that this problem is still open; it is saying that we do not know how to decide the halting problem even for 'linear' automata!" [17]. Likewise, Richard Lipton describes this state of affairs as a "mathematical embarrassment" [9].

Partial progress towards decidability of the Skolem Problem has been achieved by restricting the *order* of linear recurrence sequences. For sequences of order 1 and 2, decidability is relatively straightforward and considered to be folklore. Decidability for orders 3 and 4, however, had to wait until the 1980s before being independently settled positively by Mignotte, Shorey, and Tijdeman [13], as well as Vereshchagin [18].

The proofs of Mignotte *et al.* and Vereshchagin are complex and deep. In addition to p-adic techniques and Galois theory, these proofs rely in a fundamental way on versions of Baker's Theorem, discovered in the late 1960s and for which Alan Baker was awarded the Fields Medal in 1970. To date, all known proofs of decidability of the Skolem Problem at orders 3 and 4 make essential use of versions of Baker's Theorem. An excellent reference for Baker's Theorem and variants is Waldschmidt's book [19].

An instructive and accessible paper, available as a technical report [6], paints a much more detailed history of the Skolem Problem than we have sketched above. It also includes self-contained primers on the relevant mathematical tools, including algebraic numbers, Galois theory, rings and ideals, p-adic techniques, and Baker's Theorem. The paper also claims to prove that the Skolem Problem is decidable for integer linear recurrence sequences of order 5. Unfortunately, it appears that the proof is incorrect, with no immediately apparent way to repair it. The critical case which is not adequately handled by the authors is that of a linear recurrence sequence $\langle u_n \rangle_{n=0}^{\infty}$ whose characteristic polynomial has five distinct roots, four of which (λ_1, λ_2, $\overline{\lambda_1}$, $\overline{\lambda_2}$) are complex and of the same magnitude, and one of which, r, is real and of strictly smaller magnitude.

In this case, the terms of the linear recurrence sequence are of the form

$$u_n = |\lambda_1|^n \left(a\cos(n\theta_1 + \varphi_1) + b\cos(n\theta_2 + \varphi_2)\right) + cr^n \,,$$

where θ_1 and θ_2 are the arguments of λ_1 and λ_2 respectively, a, b, and c are real algebraic numbers, and φ_1 and φ_2 are the arguments of two algebraic numbers, all of which can be effectively calculated. If $|a|$ and $|b|$ differ, there does not appear to be a general mechanism to decide whether $u_n = 0$ for some n. More precisely, in cases where none of the u_n's are 0, there does not seem to be a way to substantiate this fact.

Finally, let us mention an earlier paper claiming decidability of the Skolem Problem for rational linear recurrence sequences of all orders [10]. Unfortunately, this paper is also incorrect. In it, the author correctly expresses the terms of a given linear recurrence sequence $\langle u_n \rangle_{n=0}^{\infty}$ as $u_n = v^T M^n w$, for M a matrix comprising only non-negative entries. (We reproduce a similar construction in Subsection 3.3 below.) The author then invokes the Perron-Frobenius Theorem to argue that the dominant terms in the explicit solution for u_n are 'well-behaved' in that they all involve eigenvalues λ of M such that $\lambda/|\lambda|$ is a root of unity. Unfortunately the argument breaks down because the explicit solution for u_n is an expression which need not involve all the eigenvalues of M, and therefore the dominant terms need not be functions of the eigenvalues of M of maximum modulus.

3.2 Complexity

To the best of our knowledge, no upper complexity bounds have been published in relation to the decidability of the Skolem Problem at orders 4 and below. In [2], Blondel and Portier showed that the Skolem Problem for integer linear recurrence sequences is NP-hard. We are not aware of other lower bounds, whether for linear recurrence sequences of arbitrary order (as in [2]) or for restricted classes of sequences.

PSPACE-hardness in the case of rational linear recurrence sequences of arbitrary order is claimed in [10]. Unfortunately, this is also incorrect. The purported proof attempts to reduce non-universality for two-letter automata to the Skolem Problem. The author effectively defines a linear recurrence sequence $\langle c_n \rangle_{n=0}^{\infty}$ such that c_n is the sum over all words w of length n of the number of accepting computations along w. But non-universality of the automaton clearly does *not* imply that one of the c_n's must be 0. The reduction is therefore incorrect.

3.3 Variants and Reductions

It will undoubtedly come as no surprise to the reader that many computational problems are equivalent, or reducible, to the Skolem Problem. We record below a small sample of such observations.

We begin by noting that the integral and rational versions of the Skolem Problem are equivalent: given a rational linear recurrence sequence $\langle u_n \rangle_{n=0}^{\infty}$, there exists an integer linear recurrence sequence $\langle v_n \rangle_{n=0}^{\infty}$, of the same order, such that, for all n, $u_n = 0$ iff $v_n = 0$. The sequence $\langle v_n \rangle_{n=0}^{\infty}$ is straightforwardly obtained from $\langle u_n \rangle_{n=0}^{\infty}$ by multiplying the recurrence relation as well as the initial values of the sequence by a suitable integer so as to clear the denominators of all rational numbers.

A second observation regards a variant of the Skolem Problem in which we are given a linear recurrence sequence $\langle u_n \rangle_{n=0}^{\infty}$ of order k, together with a constant c, and are asked whether there exists n such that $u_n = c$. This easily reduces to an instance of the standard Skolem Problem for a linear recurrence sequence $\langle v_n \rangle_{n=0}^{\infty}$ *of order* $k + 1$. Indeed, by letting $v_n = u_n - c$, and noting that the constant sequence $\langle -c \rangle_{n=0}^{\infty}$ has order 1, we can invoke results on the sum of linear recurrence sequences to establish the desired claim.

Finally, let us turn to the matricial representation of linear recurrence sequences. Given a k-dimensional square matrix M, along with k-dimensional vectors v and w, we noted earlier that setting $u_n = v^T M^n w$ yields a linear recurrence sequence. As we now show, it is also possible to express this sequence as $u_n = \widetilde{v}^T \widetilde{M}^n \widetilde{w}$ for a square matrix $\widetilde{M}$ comprising exclusively *strictly positive* entries. In performing this reduction, however, we end up with a matrix $\widetilde{M}$ of dimension $2k$.

The reduction rests on the fact that any real number can always be written as the difference of two strictly positive real numbers. Thus one replaces each entry c of M by a 2×2 submatrix $\begin{bmatrix} a & b \\ b & a \end{bmatrix}$ where $a, b > 0$ and $a - b = c$.

One easily observes that this representation commutes with matrix addition and multiplication, since

$$\begin{bmatrix} a & b \\ b & a \end{bmatrix} + \begin{bmatrix} a' & b' \\ b' & a' \end{bmatrix} = \begin{bmatrix} a+a' & b+b' \\ b+b' & a+a' \end{bmatrix} \text{ and } \begin{bmatrix} a & b \\ b & a \end{bmatrix} \begin{bmatrix} a' & b' \\ b' & a' \end{bmatrix} = \begin{bmatrix} aa'+bb' & ab'+ba' \\ ab'+ba' & aa'+bb' \end{bmatrix}$$

whereas indeed $(a-b)+(a'-b') = (a+a')-(b+b')$ and $(a-b)(a'-b') = (aa'+bb')-(ab'+ba')$.

Having so obtained $\widetilde{M}$, one accordingly adjusts the vectors v and w into $\widetilde{v}$ and $\widetilde{w}$ to obtain the desired result.

We can go slightly further and achieve the following: given a linear recurrence sequence $\langle u_n \rangle_{n=0}^{\infty}$ of order k, one can construct a *stochastic* matrix[1] $\widehat{M}$, together with vectors $\widehat{v}$ and $\widehat{w}$, such that, for all n, $u_n = 0$ iff $\widehat{v}^T \widehat{M}^n \widehat{w} = 0$. The general Skolem Problem is therefore equivalent to that specialised to linear recurrence sequences arising from stochastic matrices. Let us however point out that in this reduction, the dimension of $\widehat{M}$ will be $2k+1$.

We briefly sketch the construction below. Begin by writing $u_n = v^T M^n w$ for a k-dimensional square matrix M. Next, write $u_n = \widetilde{v}^T \widetilde{M}^n \widetilde{w}$ for a $2k$-dimensional square matrix $\widetilde{M}$ comprising only non-negative entries. Divide each entry of $\widetilde{M}$ by a sufficiently large number, so that the resulting matrix $\overline{M}$ is substochastic, and observe that, for all n, $u_n = 0$ iff $\widetilde{v}^T \overline{M}^n \widetilde{w} = 0$. Finally, add a dummy 'padding' $(2k+1)$th column to $\overline{M}$, as well as a dummy bottom row, so that all rows of the resulting square matrix $\widehat{M}$ sum to 1. It remains to augment both vectors $\widetilde{v}$ and $\widetilde{w}$ by an extra 0 entry, thereby obtaining $\widehat{v}$ and $\widehat{w}$, which completes the construction.

4 Positivity and Ultimate Positivity

A perhaps surprising observation concerning the Positivity Problem is that its decidability would immediately entail the decidability of the Skolem Problem,

[1] Recall that a stochastic matrix is a square matrix each of whose rows consists of nonnegative real numbers, with each row summing to 1.

since $u_n = 0$ iff $u_n^2 \leq 0$, and we know that the pointwise square of a linear recurrence sequence is again a linear recurrence sequence.[2] In the worst case, however, this trick reduces an instance of the Skolem Problem for a linear recurrence sequence of order k to an instance of the Positivity Problem for a linear recurrence sequence of order $k^2 + 1$. An immediate consequence of the reduction is the NP-hardness of all versions of the Positivity Problem. Let us note, however, that no such reduction is known for the Ultimate Positivity Problem, and therefore that no non-trivial lower bounds are known for this problem.

In view of the above reduction, it might be natural to consider that the Positivity Problem has been open since the advent of the Skolem Problem. The earliest explicit reference that we have found in the literature, however, goes back to the 1970s: the problem is mentioned (in equivalent formulation) in papers of Salomaa [14] and Soittola [16]; in the latter, the author opines that "in our estimation, [Skolem and Positivity] will be very difficult problems."

Indeed, the Positivity Problem has proven rather resilient over time. Decidability for integer linear recurrence sequences of order two was only established 6 years ago [5] whereas order three held out a little longer [7]. The latter paper points out that the Ultimate Positivity Problem for linear recurrence sequences of order two and three can be handled via similar techniques. Neither of these papers however offer any upper bounds on the complexity of their algorithms.

Finally, we note that the variants and reductions which we discussed for the Skolem Problem apply almost verbatim to the Positivity and Ultimate Positivity Problem. We leave the precise formulation of these facts to the reader.

References

[1] J. Berstel and M. Mignotte. Deux propriétés désirables des suites récurrentes linéaires. *Bull. Soc. Math. France*, 104, 1976.

[2] V. D. Blondel and N. Portier. The presence of a zero in an integer linear recurrent sequence is NP-hard to decide. *Linear Algebra and Its Applications*, 2002.

[3] H. Cohen. *A Course in Computational Algebraic Number Theory*. Springer-Verlag, 1993.

[4] G. Everest, A. van der Poorten, I. Shparlinski, and T. Ward. *Recurrence Sequences*. American Mathematical Society, 2003.

[5] V. Halava, T. Harju, and M. Hirvensalo. Positivity of second order linear recurrent sequences. *Discrete Applied Mathematics*, 154(3), 2006.

[6] V. Halava, T. Harju, M. Hirvensalo, and J. Karhumäki. Skolem's problem — on the border between decidability and undecidability. Technical Report 683, Turku Centre for Computer Science, 2005.

[7] V. Laohakosol and P. Tangsupphathawat. Positivity of third order linear recurrence sequences. *Discrete Applied Mathematics*, 157(15), 2009.

[8] C. Lech. A note on recurring series. *Ark. Mat.*, 2, 1953.

[2] This reduction can be made to work for both strict and non-strict Positivity, and for integral, rational, or algebraic linear recurrence sequences, with one exception: the Skolem Problem for algebraic linear recurrence sequences is not known to reduce to the non-strict Positivity Problem for algebraic linear recurrence sequences.

[9] R. J. Lipton. Mathematical embarrassments. Blog entry, December 2009. http://rjlipton.wordpress.com/2009/12/26/mathematical-embarrassments/.

[10] B. Litow. A decision method for the rational sequence problem. *Electronic Colloquium on Computational Complexity (ECCC)*, 4(55), 1997.

[11] K. Mahler. Eine arithmetische Eigenschaft der Taylor Koeffizienten rationaler Funktionen. *Proc. Akad. Wet. Amsterdam*, 38, 1935.

[12] K. Mahler. On the Taylor coefficients of rational functions. *Proc. Cambridge Philos. Soc.*, 52, 1956.

[13] M. Mignotte, T. N. Shorey, and R. Tijdeman. The distance between terms of an algebraic recurrence sequence. *Journal für die reine und angewandte Mathematik*, 349, 1984.

[14] A. Salomaa. Growth functions of Lindenmayer systems: Some new approaches. In A. Lindenmayer and G. Rozenberg, editors, *Automata, Languages, Development*. North-Holland, 1976.

[15] T. Skolem. Ein Verfahren zur Behandlung gewisser exponentialer Gleichungen. In *Comptes rendus du congrès des mathématiciens scandinaves*, 1934.

[16] M. Soittola. On D0L synthesis problem. In A. Lindenmayer and G. Rozenberg, editors, *Automata, Languages, Development*. North-Holland, 1976.

[17] T. Tao. Open question: effective Skolem-Mahler-Lech theorem. Blog entry, May 2007. http://terrytao.wordpress.com/2007/05/25/open-question-effective-skolem-mahler-lech-theorem/.

[18] N. K. Vereshchagin. The problem of appearance of a zero in a linear recurrence sequence (in Russian). *Mat. Zametki*, 38(2), 1985.

[19] M. Waldschmidt. *Diophantine approximation on linear algebraic groups*. Springer-Verlag, 2000.