Algorithmic Applications of Schanuel's Conjecture

Toghrul Karimov[1], Joris Nieuwveld[1], Joël Ouaknine[1], Mihir Vahanwala[1], and James Worrell[2]

[1] Max Planck Institute for Software Systems, Saarland Informatics Campus, Saarbrücken, Germany {toghs,jnieuwve,joel,mvahanwa}@mpi-sws.org
[2] University of Oxford, Department of Computer Science, Oxford, United Kingdom jbw@cs.ox.ac.uk

Abstract. We present a survey of algorithms, mostly drawn from the broad field of logic in computer science, that rely on Schanuel's conjecture for termination and/or correctness. Schanuel's conjecture is a central hypothesis in transcendental number theory that generalises many existing classical results such as the Lindemann-Weierstrass theorem and Baker's theorem on linear independence of logarithms of algebraic numbers. The algorithmic use of Schanuel's conjecture was spearheaded by computer algebraists in the 1970s, as well as by Macintyre and Wilkie in the 1990s, most notably to establish the decidability of real arithmetic expanded with the exponential function. Since then, many further applications have been recorded in the literature. We present and discuss several of these algorithms, with a particular focus on the precise role played by Schanuel's conjecture.

1 Introduction

It is not uncommon for correctness properties of algorithms to be conditional upon hypotheses that are unproven, but plausibly true. The most emblematic examples come from cryptography, where the security of protocols, i.e., the computational infeasibility of breaking a scheme, is predicated on conjectures such as the impossibility to factor large numbers efficiently, or more generally $\mathsf{P} \neq \mathsf{NP}$.

In the fields of computer algebra, automata theory, and dynamical systems, various decision problems hinge on variants of the following subroutine:

Problem 1. Given a polynomial $p \in \mathbb{Q}[x_1, \ldots, x_k]$, and $q_1, \ldots, q_k \in \mathbb{Q}_{>0}$, decide whether $p(\log q_1, \ldots, \log q_k) > 0$.

The number-theoretic hurdle in following the obvious approach of using increasingly precise approximations of $\log q_1, \ldots, \log q_k$ is that these numbers might "unexpectedly" be *algebraically dependent*[3] with p as a witness, i.e., $p(\log q_1, \ldots, \log q_k) = 0$. The purported decision procedure could not effectively detect this case because it would never terminate if it were to arise.

[3] Throughout this paper, algebraic (in)dependence and transcendence are meant to be over the field $\mathbb{Q}$ of rational numbers, unless otherwise specified.

However, if we are promised that $\log q_1, \ldots, \log q_k$ are *algebraically independent*, i.e., $f(\log q_1, \ldots, \log q_k) \neq 0$ for all nonzero $f \in \mathbb{Q}[x_1, \ldots, x_k]$, then instead of the usual trichotomy, we are only faced with a dichotomy between strict inequalities, which the strategy above is guaranteed to resolve. Crucially, only *termination* is conditional upon the promise: *correctness* of the output is unconditional, since there are well-known techniques for approximating logarithms within any required error bound.

Unfortunately, unconditional promises of algebraic independence often stretch beyond the capabilities of contemporary number-theoretic techniques. A classical example is the widely expected algebraic independence of e and π, which has yet to be established. This is where *Schanuel's conjecture* [19, Pages 30-31], a central hypothesis in transcendental number theory going back to the 1960s, enters the picture.

Conjecture 1 (Schanuel's conjecture). If the complex numbers $\beta_1, \ldots, \beta_k$ are linearly independent over $\mathbb{Q}$, then the set $\{\beta_1, \ldots, \beta_k, \exp(\beta_1), \ldots, \exp(\beta_k)\}$ contains a subset of k algebraically independent numbers.

As immediate examples, consider:

- The numbers 1 and $i\pi$, where i is the imaginary unit, are linearly independent over $\mathbb{Q}$. If Schanuel's conjecture holds, the set $\{1, i\pi, e, -1\}$ contains a subset of two algebraically independent numbers: by inspection, these must be $i\pi$ and e. From this assertion we can then prove the algebraic independence of π and e through elementary algebra.
- The transcendental numbers $\log 2, \log 3, \log 5$ are linearly independent over $\mathbb{Q}$, thanks to the fundamental theorem of arithmetic. Indeed, if they weren't, there would be an integer linear relationship between $\log 2, \log 3, \log 5$, e.g., $a \log 2 = b \log 3 + c \log 5$ with a, b, c positive (the other possibilities are analogous). Upon taking exponents, it would imply that 2^a can also be factorised as $3^b 5^c$, a contradiction. If Schanuel's conjecture holds, these linearly independent numbers must also be the three *algebraically* independent numbers in the set $\{\log 2, \log 3, \log 5, 2, 3, 5\}$. There being no nonzero polynomial $p \in \mathbb{Q}[x_1, x_2, x_3]$ such that $p(\log 2, \log 3, \log 5) = 0$ guarantees that the obvious approach discussed earlier would resolve this specific class of instances of Prob. 1.

Schanuel's conjecture is in fact a powerful and far-reaching generalisation of a number of classical results in transcendental number theory, three of which we state below. Recall that a complex number $\alpha \in \mathbb{C}$ is *algebraic* if $p(\alpha) = 0$ for some polynomial $p \in \mathbb{Q}[x]$, and is *transcendental* otherwise. The collection of algebraic numbers forms an algebraically closed field, which we denote by $\overline{\mathbb{Q}}$.

Theorem 1 (Lindemann-Weierstrass, 1885). *If $\alpha_1, \ldots, \alpha_k$ are algebraic numbers that are linearly independent over $\mathbb{Q}$, then $\exp(\alpha_1), \ldots, \exp(\alpha_k)$ are algebraically independent over $\overline{\mathbb{Q}}$.*

Theorem 2 (Gelfond-Schneider, 1934). *If α and β are algebraic numbers such that $\alpha \neq 0, 1$ and β is irrational, then α^β is transcendental.*

Theorem 3 (Baker, 1966). *If* $\alpha_1, \ldots, \alpha_k$ *are algebraic numbers such that* $\log(\alpha_1), \ldots, \log(\alpha_k)$ *are linearly independent over* $\mathbb{Q}$*, then* $1, \log(\alpha_1), \ldots, \log(\alpha_k)$ *are linearly independent over* $\overline{\mathbb{Q}}$*.*

As we observed in the example of $\log 2, \log 3, \log 5$ above, Schanuel's conjecture would strengthen the consequence of Baker's theorem to assert that $\log(\alpha_1), \ldots, \log(\alpha_k)$ are in fact *algebraically* independent.

In this paper, we survey the algorithmic applications of Schanuel's conjecture. We start with the role it plays in the foundations of computer algebra, and consequently, the logical reasoning about real numbers. The latter proves to be an especially convenient interface for further applications in logic, as well as in both discrete and continuous dynamical systems, which include recurrence sequences, automata, and stochastic processes.

Schanuel's conjecture was naturally of immediate interest to computer algebraists, who sought to expand their domain of operations to encompass elementary functions such as $\exp x, \log x, \sin x, \arccos x$, without compromising on the critical ability to recognise when an expression evaluates to 0. There were significant strides in this endeavour as early as 1970 [7]. Throughout the last three decades of the 20th century, Richardson developed an influential line of work culminating in [28], which achieved the above goal: this body of research showed, through both theory and practice, that assuming Schanuel's conjecture, the exponential field of *elementary* numbers (Sec. 2.2), which is countable, algebraically closed, and also closed under the elementary functions, is computable, i.e., there is an effective representation scheme that is amenable to arithmetic operations, elementary functions such as $\exp, \log, \sin, \arccos$, etc., and zero testing. The specific role of Schanuel's conjecture is to guarantee the termination of the fundamental zero-testing algorithm for elementary numbers.

As [28, Introduction] notes, the consensus [2,3,30] among the eminent number theorists of the 1970s was that Schanuel's conjecture is very likely correct, but would be extremely hard to prove. Little has changed since. In order to approach the resolution of Schanuel's conjecture, Zilber [38] showed in 2005 that there is a unique exponential field $\mathbb{B}$ with cardinality $|\mathbb{C}|$ that axiomatically "imitates" the complex numbers with exponentiation, and satisfies Schanuel's conjecture along with a property called *strong exponential-algebraic closure*. If $\mathbb{B}$ and $\mathbb{C}$ are isomorphic, then Schanuel's conjecture for $\mathbb{C}$ follows. Conversely, if $\mathbb{B}$ and $\mathbb{C}$ are not isomorphic, then at least one among Schanuel's conjecture and strong exponential-algebraic closure fails to hold for the complex numbers. Unfortunately, neither proof nor refutation of the isomorphism seems accessible.

The connection between model theory and transcendence theory, however, had come to the fore nearly a decade prior to Zilber's work when in 1996, Macintyre and Wilkie [22] used Schanuel's conjecture to prove the termination of their now celebrated algorithm to decide the first-order theory $T_{\exp}$ of the real numbers with the exponential function. As we survey decision procedures that rely upon Schanuel's conjecture, we observe that several do so solely by virtue of queries to $T_{\exp}$: such is the influence of [22] on making the algorithmic consequences of Schanuel's conjecture accessible. Although Schanuel's conjecture is

only used to prove termination of the algorithm deciding $T_{\exp}$, [22, Sec. 5] nevertheless hedges against its failure by identifying that the decidability of $T_{\exp}$ is equivalent to an ostensibly weaker hypothesis (Conj. 2), whose resolution was still expected to be inaccessible.

The synergy between transcendence theory, computer algebra, and model theory continues to be explored to this day. In 2016, Macintyre [21] exhibited remarkable evidence to suggest that Schanuel's conjecture is fundamental to computability: more precisely, he showed that any countable exponential field that obeys Zilber's axioms (which, in particular, entail Schanuel's conjecture) is computable.

In this survey, we begin by explaining the foundational computer-algebraic applications in Sec. 2. In Sec. 3.1 and 3.2, we then explain how these advances, when combined with model theory, established the decidability of first-order logical theories. Queries formulated in first-order logic are typically (but not exclusively) the means through which Schanuel's conjecture is invoked by algorithms that decide problems in monadic second-order logic (Sec. 3.3), discrete dynamical systems (Sec. 4), quantitative verification (Sec. 5), and continuous dynamical systems and MDPs (Sec. 6).

2 Computer Algebra

As discussed earlier, Schanuel's conjecture is useful for zero-testing in computer algebra because it can be intuited as a promise of algebraic dependence among numbers being formally certifiable. We shall next make this intuition concrete by summarising [29], and subsequently survey the more general case of recognising zero among the elementary numbers [27,28]. These procedures are unconditionally correct, and require Schanuel's conjecture only to guarantee termination.

2.1 Zero-testing elementary expressions

The problem considered in [29] is that of testing whether expressions involving addition, subtraction, multiplication, division, taking n-th roots for a positive integer n, division, exp, and log[4] evaluate to zero. Formally, an expression E is given as a parse tree whose leaves represent rational numbers, and internal nodes signify elementary operations, and we have to decide whether it is the case that the value $V(E) = 0$. Observe that this entails recursively solving subproblems: in order for $V(E)$ to be well defined, for every subexpression[5] E_i that evaluates a divisor, or the argument of log, it is necessary that $V(E_i) \neq 0$.

[4] We take the principal branch of the logarithm. A complex number $z = p + iq$ can also be expressed as $z = e^{x+iy}$, choosing $y \in (\pi, \pi]$. This choice defines $\log z = x + iy$. Similarly, we assume that n-th roots of numbers except 0 are interpreted as $\exp((\log z)/n)$. Note that with this set of operations, we can also implement trigonometric, inverse trigonometric functions, and their hyperbolic analogues.

[5] Technically, we take the parse tree to be a directed acyclic graph in order to avoid duplicate subexpressions. We assume an ordering such that if $i < j$, then E_i cannot have E_j as a subexpression.

For example, i can be expressed as either $(-1)^{1/2}$ or $\exp(\log(-1)/2)$, and thus the difference of these two expressions is 0. Similarly, π can be expressed as $\log(-1)/(-1)^{1/2}$. We shall adopt $E = (-1)^{1/2} - \exp(\log(-1)/2)$ as our running (toy) example.

We begin by classifying subexpressions. We shall use the letter A to denote subexpressions that are either arguments to exp or outputs of log, and the letter B to denote subexpressions that are either outputs of exp or arguments to log. In our example, we have $A_1 = \log(-1)$, $A_2 = \log(-1)/2$, $B_1 = -1$, $B_2 = \exp(\log(-1)/2)$. In this way, we shall maintain that either $A_j = \log(B_j)$, or $B_j = \exp(A_j)$. We note that an expression can be both an A_j and a B_k, e.g., if $E' = \exp(\exp(1))$, then $B'_1 = A'_2 = \exp(1)$.

We shall use α_j, β_j to respectively denote $V(A_j), V(B_j)$. In our running example, $A_1 = i\pi, A_2 = i\pi/2, B_1 = -1, B_2 = i$. By definition, we always have that $\beta_j = \exp(\alpha_j)$, and for every pair (α_j, β_j), there is an input to the operand, and an output to the operand, e.g., if $A_j = \log(B_j)$ then α_j is the output number. We use γ_j to denote the output component of the pair. By the structure of E, we have that $V(E)$, and in fact each $\alpha_1, \beta_1, \ldots, \alpha_k, \beta_k$ is an algebraic expression in $\gamma_1, \ldots, \gamma_k$ (where k is the number of subexpressions rooted at exp or log).

Now, if we assume that $\alpha_1, \ldots, \alpha_k$ are linearly independent over $\mathbb{Q}$, then Schanuel's conjecture promises us that there are k algebraically independent numbers among $\alpha_1, \ldots, \alpha_k, \beta_1, \ldots, \beta_k$. By the preceding observation, the only way to ensure this is for $\gamma_1, \ldots, \gamma_k$ to be algebraically independent. Recall that we have $V(E) = f(\gamma_1, \ldots, \gamma_k)$ for some algebraic function f. If $f \equiv 0$, then we have proof that $V(E) = 0$ regardless of our assumption of linear independence above. If not, then we are promised in particular that $V(E) \neq 0$. We thus run a semi-algorithm to compute $V(E)$ to arbitrary precision to cover this case, Schanuel's conjecture assuring us that this will eventually certify that $V(E) \neq 0$.

If the linear independence assumption fails, then there exist $c_1, \ldots, c_j \in \mathbb{Z}$ with $c_j \neq 0$ such that $c_1\alpha_1 + \cdots + c_j\alpha_j = 0$. If we find these integers, we can use them to obtain an equivalent expression E' of a "reduced order", i.e., E' has fewer subexpressions using $\exp, \log$, and $V(E') = V(E)$. In our example, we have that $2\alpha_2 = \alpha_1$, and $B_2 = \exp(A_2)$. We can replace B_2 with $(B_1)^{1/2}$, to get the reduced order $E' = (-1)^{1/2} - (-1)^{1/2}$. More generally, if $B_j = \exp(A_j)$, then we replace B_j with $(B_1^{-c_1} \cdots B_{j-1}^{-c_{j-1}})^{1/c_j}$, and if $A_j = \log(B_j)$, then we replace A_j with $-(c_1A_1 + \cdots + c_{j-1}A_{j-1})/c_j$. Clearly, the order can be reduced only finitely often.

However, in order to soundly obtain such an equivalent E', one needs to *prove* that a purported[6] linear dependency holds. Towards such a formal proof, we shall recursively associate with each subexpression E_i of E, an algebraic function $\eta(E_i)$ as follows. If E_i is $\exp(A_j)$ for some j, then $\eta(E_i) = x_j$, likewise if E_i is $\log(B_j)$ for some j, then $\eta(E_i) = y_j$. Otherwise, if $E_i = \mathsf{op}(E_{i_1}, \ldots)$ for some algebraic operation op, then we set $\eta(E_i) = \mathsf{op}(\eta(E_{i_1}), \ldots)$. Finally, if E_i is a rational constant, then $\eta(E_i)$ returns the same constant, e.g., $\eta(-1) = -1$.

[6] Algorithms such as LLL (Lenstra-Lenstra-Lovász lattice basis reduction) or PSLQ [13] can greatly optimise the enumeration of potential dependencies.

Returning to our running example, we have that $\eta(A_1) = x_1, \eta(B_1) = -1, \eta(A_2) = x_1/2, \eta(B_2) = y_2$. By easy symbolic reasoning, we see that if $c_1\eta(A_1)+\cdots+c_j\eta(A_j) \equiv 0$ (is identically 0), then $c_1\alpha_1+\cdots+c_j\alpha_j = 0$. Similarly, if $\eta(B_1)^{c_1}\cdots\eta(B_j)^{c_j}-1 \equiv 0$ and $|c_1\alpha_1+\cdots+c_j\alpha_j| < 1$, then $c_1\alpha_1+\cdots+c_j\alpha_j = 0$ because it must be an integer multiple of $2\pi i$. In this manner, associated algebraic functions serve as certificates; however, if a linear dependency exists, is there guaranteed to be a certificate?

We make a thematic invocation of Schanuel's conjecture to answer affirmatively. Suppose that we have $c_1\alpha_1 + \cdots + c_j\alpha_j = 0$ with j minimal, i.e., $c_j \neq 0$ and $\alpha_1, \ldots, \alpha_{j-1}$ are linearly independent. Let λ_j be the input component among α_j, β_j. We have by structure that $\alpha_1, \beta_1, \ldots, \alpha_{j-1}, \beta_{j-1}, \lambda_j$ are algebraic in $\gamma_1, \ldots, \gamma_{j-1}$. In particular, either $\eta(c_1A_1 + \cdots + c_jA_j)$, or $\eta(B_1^{c_1}\cdots B_j^{c_j} - 1)$ is of the form $f(z_1, \ldots, z_{j-1})$, where $f(\gamma_1, \ldots, \gamma_{k-1})$ would return $c_1\alpha_1 + \cdots + c_j\alpha_j$ or $\beta_1^{c_1}\cdots\beta_j^{c_j} - 1$ respectively. The arguments $\gamma_1, \ldots, \gamma_{k-1}$, by Schanuel's conjecture, are algebraically independent. The only way to realise $f(\gamma_1, \ldots, \gamma_{k-1}) = 0$ therefore, would be $f \equiv 0$, implying that the minimal dependency, when enumerated, would constitute a desired symbolic proof.

Note that in the worst case, if Schanuel's conjecture is false, then no proof would be observed, and we would continue our search for a dependency, but never make unsound progress towards a decision in the algorithm.

In summary, we run a semi-algorithm that seeks to prove that $V(E) \neq 0$ by brute approximation, and a semi-algorithm that seeks to reduce the (finite) order of E. If Schanuel's conjecture holds, at least one of them will terminate. Having terminated, the algorithm gives a certificate (an approximate, or an identically zero function) of whether $V(E) = 0$.

2.2 Zero-testing elementary numbers

Recall that elementary numbers constitute an algebraically closed subfield of the complex numbers, which is closed under applications of elementary functions and computable (i.e., there is an effective representation scheme that is amenable to arithmetic operations, elementary functions such as $\exp, \log, \sin, \arccos$, etc., and zero testing) assuming Schanuel's conjecture. We now follow [28, Sec. 2] and record the concepts required to define elementary numbers.

The first building block is an *exponential system*, which consists of polynomials $p_1, \ldots, p_r \in \mathbb{Q}[x_1, \ldots, x_n]$ and expressions $w_1 - \exp(z_1), \ldots, w_k - \exp(z_k)$ where $\{w_1, \ldots, w_k, z_1, \ldots, z_k\} \subseteq \{x_1, \ldots, x_n\}$. We use $F = \langle F_1, \ldots, F_{r+k}\rangle$ to collectively denote the entire exponential system.

An *elementary point* $\overline{\gamma} = (\gamma_1, \ldots, \gamma_n) \in \mathbb{C}^n$ is a nonsingular root of an exponential system $F = \langle F_1, \ldots, F_n\rangle$, i.e., $\overline{\gamma}$ satisfies $F(\overline{\gamma}) = 0$, and the Jacobian determinant $J_F(\overline{\gamma}) = \det\left(\frac{\partial F_i}{\partial x_j}\right)_{1\leq i,j\leq n}(\overline{\gamma}) \neq 0$. A number $\zeta \in \mathbb{C}$ is *elementary* if there is an elementary point $\overline{\gamma} \in \mathbb{C}^n$ and a polynomial $p \in \mathbb{Q}[x_1, \ldots, x_n]$ such that $\zeta = p(\overline{\gamma})$. Naturally, an elementary number ζ is represented by $(\overline{\gamma}, p)$. It is clear how to represent p; we discuss how we represent an elementary point $\overline{\gamma}$.

1. We are given an exponential system $F = \langle F_1, \ldots, F_n \rangle$, which is straightforward to represent.
2. We are given an approximate $\overline{\alpha} \in (\mathbb{Q}[i])^n$, i.e., all coordinates of $\overline{\alpha}$ have rational real and imaginary parts.
3. We are given a rational precision $\varepsilon > 0$, which defines an ε-neighbourhood $B(\overline{\alpha}, \varepsilon)$ around $\overline{\alpha}$ such that:
   - The infimum over $\overline{\beta} \in B(\overline{\alpha}, \varepsilon)$ of $|J_F(\overline{\beta})|$ is greater than 0, and is effectively computable.
   - The neighbourhood passes a fixed standard test (e.g., [27, Sec. 2], see also [28, Sec. 2.1] for a survey of alternatives) that Newton's iteration[7] converges to the unique root of F in the neighbourhood, thereby proving that $\overline{\alpha}$ approximates $\overline{\gamma}$ to precision ε.

The algorithm of [28] decides whether a representation as described above encodes 0, and relies on Schanuel's conjecture to guarantee termination.

3 Logic

In this section, we discuss how decidability results for expansions of classical structures follow when the computability enabled by Schanuel's conjecture is combined with model theory in the case of first-order logic, and with automata theory, dynamical systems, and infinite-word combinatorics in the case of monadic second-order logic.

3.1 First-order theory of the reals with exponentiation

Tarski [31] famously showed that the first-order theory T_0 of the structure $\langle \mathbb{R}; +, -, \cdot, <, 0, 1 \rangle$ is decidable by virtue of admitting quantifier elimination. In other words, there is an algorithm to translate any first-order formula that is interpreted over the real numbers, and uses the functions $+, -, \cdot$, the predicate $<$, and the constants $0, 1$, into an equivalent formula (over the same signature) without quantified variables. In particular, sentences (formulae without free variables) are translated into Boolean combinations of inequalities involving integers. The question naturally arose: could we expand the signature with, e.g., the exponentiation function and yet obtain a structure with a decidable theory?

Macintyre and Wilkie [22] proved that the first-order theory $T_{\exp}$ of the structure $\langle \mathbb{R}; +, -, \cdot, <, 0, 1, \exp \rangle$, i.e., the real numbers expanded with the exponential function is decidable subject to Schanuel's conjecture (for $\mathbb{R}$). To be specific, the decidability of $T_{\exp}$ is *equivalent* to a weaker version of Schanuel's conjecture [22, Sec. 5], which is stated in computational terms (Conj. 2). In view of the fact that this work built an accessible interface for the algorithmic applications of Schanuel's conjecture, we thus make the important observation: results that are conditional because they invoke the decidability of the first-order theory of the reals with exponentiation are technically reliant upon a weaker hypothesis.

[7] This approximates a root as the limit of the recurrence $\overline{\alpha_{i+1}} = \overline{\alpha_i} - J_F^{-1}(\overline{\alpha_i})F(\overline{\alpha_i})$.

The decision procedure is outlined by the proposition that the complete theory $T_{\exp}$ can be axiomatised by $T \cup \mathcal{E}_{\exp}$, where T is a recursive set of sentences, and $\mathcal{E}_{\exp}$ is the existential fragment of $T_{\exp}$ [22, Thm. 2.6]. In order to decide whether a given sentence φ is in $T_{\exp}$, we enumerate formulae in $T_{\exp}$ until arriving at either φ or $\neg\varphi$. This enumeration performs the following tasks in parallel: (i) enumerate sentences from the recursive set T; (ii) enumerate sentences that can be deduced from the ones enumerated thus far; (iii) iterate in parallel over existential sentences φ_e, running the subroutine described below on each φ_e until either: (a) it certifies that $\varphi_e \in \mathcal{E}_{\exp}$, (b) $\neg\varphi_e$ has been enumerated as a member of $T_{\exp}$. Schanuel's conjecture is needed to guarantee that each $\varphi_e \in \mathcal{E}_{\exp}$ will indeed be certified.

We now explain the ingredients to devise a subroutine whose termination certifies that $\varphi_e \in \mathcal{E}_{\exp}$. An arbitrary existential sentence can effectively be translated[8] into the form

$$\exists x_1 \cdots \exists x_n.\ p(x_1, \ldots, x_n, \exp(x_1), \ldots, \exp(x_n)) = 0, \tag{1}$$

where $p \in \mathbb{Z}[x_1, \ldots, x_{2n}]$. We shall use $F_p : \mathbb{R}^n \to \mathbb{R}$ to denote the function that maps $\overline{x} = (x_1, \ldots, x_n)$ to $p(x_1, \ldots, x_n, \exp(x_1), \ldots, \exp(x_n))$. The task is thus to certify that F_p has a root. Due to the following lemma [36, Lem. 6], the existence of a root of F_p implies the existence of one that can be "isolated".

Lemma 1. *Suppose F_p has a root. Then there exist $q_1, \ldots, q_n \in \mathbb{Z}[x_1, \ldots, x_{2n}]$ and $\overline{\gamma} = (\gamma_1, \ldots, \gamma_n) \in \mathbb{R}^n$ such that:*

1. *$\overline{\gamma}$ is a root of F_p, i.e., $F_p(\gamma_1, \ldots, \gamma_n) = 0$;*
2. *$\overline{\gamma}$ is a nonsingular root of the function $\langle F_{q_1}, \ldots, F_{q_n} \rangle : \mathbb{R}^n \to \mathbb{R}^n$.*

For example, $(\log 2, \log 3, \log 5)$ is a nonsingular root of $\langle \exp(x) - 2, \exp(y) - 3, \exp(z) - 5 \rangle$. Nonsingularity ensures that such roots can be approximated via (an appropriate version of) Newton's method [36, Lem. 5], and as discussed in Sec. 2.2, given a sufficiently precise neighbourhood of the root, we can use standard tests to prove that Newton's method will converge.

Our task is thus to enumerate $(q_1, \ldots, q_n, \overline{\alpha}, \varepsilon)$ until we find a tuple such that $\alpha \in \mathbb{Q}^n$ approximates to precision $\varepsilon \in \mathbb{Q}$ a point $\overline{\gamma}$ which is a nonsingular root of $\langle F_{q_1}, \ldots F_{q_n} \rangle$, and moreover satisfies $F_p(\overline{\gamma}) = 0$. The verification at each iteration is reminiscent (albeit with slight technical differences) of the zero-testing in Sec. 2.2, and requires Schanuel's conjecture for termination.

Intuitively, we use Schanuel's conjecture to argue that if some $\overline{\gamma}$ satisfies algebraic and transcendental dependencies from $\langle F_{q_1}, \ldots F_{q_n} \rangle$ as well as F_p, then it is no coincidence; rather, $p, q_1, \ldots, q_n$ are related in a manner that we can elicit, and furthermore use as a formal proof of F_p having a root [36, Cor. of SC]. In fact, [22, Sec. 5] identifies that the validity of the following weaker hypothesis is sufficient as well as *necessary* for the decidability of $T_{\exp}$:

[8] We translate $\exists \overline{x}.\ F_1(\overline{x}) < 0 \wedge \cdots \wedge F_k(\overline{x}) < 0 \wedge F_{k+1}(\overline{x}) = 0 \wedge \cdots \wedge F_m(\overline{x}) = 0$ to $\exists \overline{x}. \exists \overline{t}.\ (F_1(\overline{x}) + \exp(t_1))^2 + \cdots + (F_k(\overline{x}) + \exp(t_k))^2 + F_{k+1}(\overline{x})^2 + \cdots + F_m(\overline{x})^2 = 0$.

Conjecture 2 (Weak Schanuel's conjecture). Given $p, q_1, \ldots, q_n \in \mathbb{Z}[x_1, \ldots, x_{2n}]$, we can compute a positive $\eta \in \mathbb{N}$ such that if $\overline{\gamma} \in \mathbb{R}^n$ is a nonsingular root of $\langle F_{q_1}, \ldots F_{q_n} \rangle$ and $|F_p(\overline{\gamma})| \leq 1/\eta$, then $F_p(\overline{\gamma}) = 0$.

We note that this η certainly exists: [22] argues this via Khovanskii's theorem, which implies that $\langle F_{q_1}, \ldots, F_{q_n} \rangle$ only has finitely many nonsingular roots. It is only the effectiveness of η that is conditional.

Before proceeding, we illustrate the difference between Schanuel's conjecture and the above weak variant with the example problem of determining, given nonzero $p \in \mathbb{Z}[x, y, z]$, whether $p_0 = p(\log 2, \log 3, \log 5) = 0$. Schanuel's conjecture immediately declares that this cannot be the case. The weak variant is more circumspect: it computes $\eta(e^x - 2, e^y - 3, e^z - 5, p(x, y, z))$ such that if $p_0 \neq 0$, then $|p_0| > 1/\eta$. It remains to approximate p_0 to precision $1/3\eta$.

We now return our attention to showing how to verify a purported certificate $(q_1, \ldots, q_n, \overline{\alpha}, \varepsilon)$ that F_p has a root. We first check (by a standard Newton's method-based test) that it is well-formed, i.e., $\overline{\alpha}$ indeed approximates to precision ε a nonsingular root $\overline{\gamma}$ of $\langle F_{q_1}, \ldots F_{q_n} \rangle$. We then evaluate $F_p(\overline{\alpha})$ and use continuity to check that this value implies $F_p(\overline{\gamma}) < 1/\eta$.

Conversely, it is elementary to argue by continuity of F_p that if $F_p(\overline{\gamma}) = 0$, then there is a well-formed certificate with sufficiently high precision which is guaranteed to be accepted.

3.2 Extensions and related first-order decidability results

The result above has an analogue of a more technical number-theoretic flavour: the 2013 PhD thesis of Mariaule [24] shows the decidability of the first-order theory of the ring $\mathbb{Z}_p$ of p-adic integers with the (p-adic) exponential function.

For the reals, however, Macintyre and Wilkie have actually adapted the above techniques to prove a stronger (unpublished) result (see [21, Thm. 3.1(4)]): assuming Schanuel's conjecture, the first-order theory T_{el} of the reals with exponentiation and *restricted* trigonometric functions, i.e., the theory of the structure $\langle \mathbb{R}; +, \cdot, <, \exp, \sin \restriction [0, n], \cos \restriction [0, n] \rangle$ is decidable. Here, a restricted function $f \restriction [0, n]$ (where $n \in \mathbb{N}$ returns $f(x)$ for $x \in [0, n]$, and 0 otherwise. We believe that the proof of decidability of T_{el} requires Schanuel's conjecture for $\mathbb{C}$ only for termination of the algorithm, and proceeds by adapting the model-theoretic machinery of [35] to prove a "combined" (and effective) version of the two main results therein,[9] and hence deduce that T_{el} is effectively model-complete and axiomatised analogously to T_{exp}. Furthermore, [35, Thm. 5.1] in particular can be seen as the required analogue of Lem. 1 to enable the detection of roots of the ensuing elementary functions by enumerating guesses for an elementary-point root, and using the techniques surveyed in Sec. 2.2 to verify the guesses.

We observe that the introduction of trigonometric functions takes us to the frontiers of decidability: if we were to allow *unrestricted* trigonometric functions,

[9] In fact, for the applications we survey in Sec. 6, it suffices to consider the restriction of *all* three functions to bounded intervals. In this case, one only needs carefully assess the proof of the first main result of of [35] for effectiveness.

we would get undecidable theories: indeed the expanded first-order theory of $\langle \mathbb{R}; +, \cdot, <, \sin \rangle$ is undecidable because one can express the predicate "x is an integer" as $\forall w.\ \sin(w) = 0 \Rightarrow \sin(xw) = 0$. With access to both addition and multiplication, can encode any given instance of Hilbert's 10th problem[10] as a sentence. However, decidability subject to Schanuel's conjecture is recovered for the theory $T_{+,\sin}$ of $\langle \mathbb{R}; +, <, \sin \rangle$ [6, Thm. 2.10].

The question of whether we can get decidable theories if we relinquish multiplication in favour of direct access to the integers was very recently considered in [6], which studied Presburger arithmetic with sine. In this setting, we have access to only the addition and sine function, the variables are interpreted over the integers, and the terms are interpreted over the reals. The first main result of [6] is that the resulting theory is undecidable, already with four alternating blocks of quantifiers. The second main result, however, is that the existential fragment is decidable subject to Schanuel's conjecture. Rather atypically, the decision procedure relies on Schanuel's conjecture for correctness (of the step [6, Thm. 4.5]) as well as termination (deciding whether a sentence obtained by [6, Thm. 4.18] is in $T_{+,\sin}$). In particular, if Schanuel's conjecture does not hold, the step of [6, Thm. 4.5] fails and the algorithm might incorrectly decide sentences of the form $\exists x_1, \ldots x_n.\ \bigwedge_{i=1}^{k} t_i(x_1, \ldots, x_n) = 0$, i.e., a conjunction of equalities.

3.3 Monadic second-order (MSO) theories of the natural numbers with integer-power predicates

In 1966, Elgot and Rabin [12] showed that monadic second-order (MSO) theories of structures of the form $\langle \mathbb{N}; <, a^{\mathbb{N}} \rangle$ are decidable, where $a \geq 2$ is a natural number and $a^{\mathbb{N}}$ denotes the predicate $\{a^n : n \in \mathbb{N}\}$. It remained open whether expanding the structure with multiple such predicates results in decidable theories until recently, when [4] provided a positive resolution: for any a, b, the MSO theory of $\langle \mathbb{N}; <, a^{\mathbb{N}}, b^{\mathbb{N}} \rangle$ is decidable. However, the decidability of the MSO theory of $\langle \mathbb{N}; <, a_1^{\mathbb{N}}, \ldots, a_d^{\mathbb{N}} \rangle$ (e.g., the MSO theory of $\langle \mathbb{N}; <, 2^{\mathbb{N}}, 3^{\mathbb{N}}, 5^{\mathbb{N}} \rangle$) was shown subject to Schanuel's conjecture. In this subsection, we explain the role of Schanuel's conjecture, and show a slightly improved result.

Theorem 4. *Let $2 \leq a_1 < \cdots < a_d$ be natural numbers. The MSO theory of $\langle \mathbb{N}; <, a_1^{\mathbb{N}}, \ldots, a_d^{\mathbb{N}} \rangle$ is decidable subject to the decidability of $T_{\exp}$ (the first-order theory of the reals with exponentiation).*[11]

We can assume, without loss of generality, that $a_1, \ldots, a_d$ are pairwise multiplicatively independent, i.e., for all $i, j \in \{1, \ldots, d\}$ and integers $p, q \geq 1$, $a_i^p \neq a_j^q$. We then define the *order word* $\alpha \in \{a_1, \ldots, a_d\}^\omega$, which records the

[10] Hilbert's 10th problem gives as input a polynomial equation and asks whether it has an integer solution. It was famously shown undecidable in 1970 by Matiyasevich.

[11] Technically, we only work with the structure $\langle \mathbb{R}; <, +, \cdot, \log a_1, \ldots, \log a_d \rangle$ (the usual first-order structure of the reals, expanded with constants), which, due to Tarski [32], admits quantifier elimination. As a result, the technical hinge is the ability to decide polynomial (in)equalities in $\log a_1, \ldots, \log a_d$.

order in which the powers of $a_1, \ldots, a_d$ (excluding 1) appear. For example, the order word of the powers $2, 3, 4, 5, 8, 9, 16, 25, 27, 32, 64, 81, 125, 128, \ldots$ of $2, 3, 5$ is $23252325322352\cdots$. In [4], it was shown that deciding whether a sentence is in the MSO theory of $\langle \mathbb{N}; <, a_1^{\mathbb{N}}, \ldots, a_d^{\mathbb{N}} \rangle$ is equivalent to deciding whether a given automaton $\mathcal{A}$ accepts the order word α.

The key insight is that the order word enjoys a *toric* structure. Consider the example of $2, 3, 5$, and the factorisation of α based on occurrences of powers of $a_1 = 2$. We get $2 \cdot 32 \cdot 52 \cdot 32 \cdot 532 \cdot 2 \cdot 352 \cdots$. The k-th factor conveys whether there were powers of 3 and 5 between 2^{k-1} and 2^k, and in what order they occurred. Equivalently, it conveys whether there were integer multiples of $\log 3, \log 5$ (there can be at most one of each kind) between $(k-1)\log 2$ and $k \cdot \log 2$, and in what order. We capture the above through the torus $\mathbb{T}$ in Fig. 1. Due to [25, Thm. 3] (stated below), it would suffice to elicit *effective almost-periodicity* from the toric order word in order to prove our decidability result.

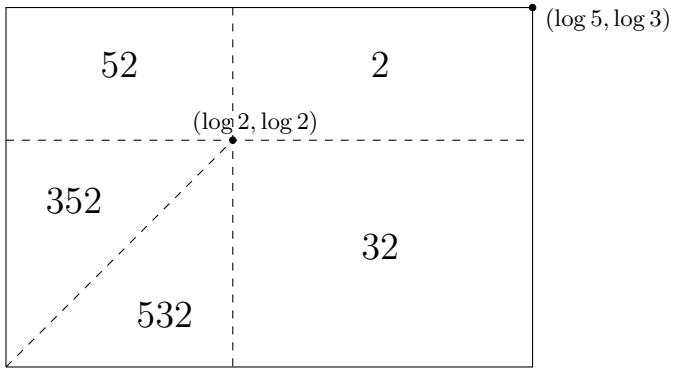

Fig. 1. The above torus $\mathbb{T}$ helps construct the order word through the orbit of a point that starts at $(\log 2, \log 2)$ and travels in discrete steps of $(\log 2, \log 2)$. The label of the (open) region it lands in determines the next letters of the order word; the starting point prints 2 and is the only point that lands on a boundary.

Theorem 5. *Define a word $\alpha \in \Sigma^\omega$ to be effectively almost-periodic if, given any $u = u(0)\cdots u(\ell-1) \in \Sigma^+$, we can compute a return time $R \in \mathbb{N}$ such that either:*

- *For all $n \geq R$, $\alpha(n)\cdots\alpha(n+\ell-1) \neq u$, or*
- *For all n, there exists an $m \in [n, n+R)$ such that $\alpha(m)\cdots\alpha(m+\ell-1) = u$.*

If α is effectively almost-periodic, then given any Büchi automaton $\mathcal{A}$, we can decide whether $\mathcal{A}$ accepts α.

It remains to show how, given a word $u \in \{a_1, \ldots, a_d\}^+$, the toric structure helps us to compute a return time. We work with our running example of powers of $2, 3, 5$ for ease of exposition; it is straightforward to generalise the arguments.

Let us take the example of $u = 5232532$. As the first step, we factorise[12] it based on occurrences of $a_1 = 2$, and obtain $52 \cdot 32 \cdot 532 = b_0 b_1 b_2$. In order for $u = b_0 \cdots b_\ell$ to be observed, there must be a sequence of points $(x_0, y_0), \ldots, (x_\ell, y_\ell)$ in the torus $\mathbb{T}$ such that:

- For each i, (x_{i+1}, y_{i+1}) is obtained by taking a single step of $(\log 2, \log 2)$ from (x_i, y_i), i.e., $x_{i+1} = x_i + \log 2$ or $x_i + \log 2 - \log 5$ as appropriate to ensure that $0 \leq x_i < \log 5$ (and likewise for y). Thus, the sequence is uniquely determined by (x_0, y_0).
- The point (x_0, y_0) is in a region that has b_0 as its suffix. In our example, (x_0, y_0) must be in the region of 52 or 352. Likewise (x_ℓ, y_ℓ) must be in a region that has b_ℓ as its prefix.
- All other points (x_i, y_i) must be in the region that prints b_i.

Points (x_0, y_0) that satisfy the above constraints, i.e., define a sequence of points that land in the appropriate regions, are easily (e.g., by linear programming) seen to comprise a bounded union $\mathcal{I}_u$ of open sets of the torus $\mathbb{T}$. By construction, there is a direct correspondence between occurrences of u and visits to $\mathcal{I}_u$: in particular, u has gaps of at most R between consecutive occurrences in the order word if and only if the orbit visits $\mathcal{I}_u$ at least once in every R steps.

The orbit is dense in some (not necessarily proper) subtorus of $\mathbb{T}$. More specifically, by Kronecker's theorem in Diophantine approximation [15], the orbit is dense in the subtorus[13] $\mathbb{T}_{\mathsf{orbit}}$ given by

$$\left\{(x, y) \in \mathbb{T} : \text{for all } b, c \in \mathbb{Z}, \text{ if } \frac{b \cdot \log 2}{\log 5} + \frac{c \cdot \log 2}{\log 3} \in \mathbb{Z}, \text{ then } \frac{bx}{\log 5} + \frac{cy}{\log 3} \in \mathbb{Z}\right\}.$$

By invoking the compactness of the (sub)torus $\mathbb{T}_{\mathsf{orbit}}$, we can prove that if $\mathcal{J}_u = \mathcal{I}_u \cap \mathbb{T}_{\mathsf{orbit}}$ is nonempty, then there exists an R such that the orbit of any point in $\mathbb{T}_{\mathsf{orbit}}$ visits $\mathcal{J}_u$, and hence $\mathcal{I}_u$, within R steps. On the other hand, if $\mathcal{J}_u$ is empty, then u can never occur in the order word.

We have thus far proven that the order word α is not merely almost-periodic, but in fact enjoys the stronger property of *uniform recurrence*: every u occurs either with bounded gaps $R(u)$, or never at all. It remains to effectively compute a return time R.[14] The problem with the strategy of brute enumeration to compute R is that we cannot unconditionally obtain $\mathbb{T}_{\mathsf{orbit}}$: it can so happen that we enumerate R forever because $\mathcal{I}_u$ is nonempty but $\mathcal{J}_u$ is empty as a consequence of $\mathbb{T}_{\mathsf{orbit}} \subset \mathbb{T}$.

The original paper [4] invoked Schanuel's conjecture to argue that $\mathbb{T}_{\mathsf{orbit}} = \mathbb{T}$, since indeed in our example, the reciprocals of $\log 2, \log 3, \log 5$ would have to be linearly independent over the rationals. This guaranteed that the enumeration

[12] The case where there is just a single factor (because, e.g., there is no a_1) is trivial

[13] We continue with the example of powers of $2, 3, 5$ for ease of exposition, the general statement for $a_1, \ldots, a_d$ is completely analogous.

[14] One implementation can use the open-cover characterisation above. Alternately, we can also use the property that all words v of length $R + |u|$ for which $\mathcal{J}_v$ is nonempty must contain u as a contiguous subword.

of potential R would terminate. In this case, Baker's theorem (Thm. 3) allows us to determine whether $\mathcal{J}_u = \mathcal{I}_u$ is nonempty.

However, we show that relying upon the decidability of $T_{\exp}$ suffices. The key observation is that it is not always necessary to know $\mathbb{T}_{\mathsf{orbit}}$ exactly in order to compute R. We justify this as follows. Consider a "restricted" torus $\mathbb{T}_{\mathsf{res}}$ such that $\mathbb{T} \supseteq \mathbb{T}_{\mathsf{res}} \supseteq \mathbb{T}_{\mathsf{orbit}}$. For an arbitrary word v, we define $\mathcal{K}_v$ as $\mathcal{I}_v \cap \mathbb{T}_{\mathsf{res}}$, and observe that $\mathcal{I}_v \supseteq \mathcal{K}_v \supseteq \mathcal{J}_v$. If $\mathcal{K}_u$ is empty, it implies the emptiness of $\mathcal{J}_u$, and thus the non-occurrence of u. Conversely, if for some R, if the orbit of every point in $\mathbb{T}_{\mathsf{res}}$ visits $\mathcal{K}_u$ within R steps, it implies that the actual orbit visits $\mathcal{I}_u$ every R steps, thereby certifying the occurrence of u with gaps bounded by R.

Each candidate $\mathbb{T}_{\mathsf{res}}$ thus outlines a correct semi-algorithm to compute a return time R given u. In case $\mathbb{T}_{\mathsf{res}} = \mathbb{T}_{\mathsf{orbit}}$, then termination is guaranteed (provided the attendant linear programming is effective) and we have an algorithm. For effectiveness of R, it suffices to show that there is an enumerable set Tori of candidate $\mathbb{T}_{\mathsf{res}}$ such that we are guaranteed $\mathbb{T}_{\mathsf{orbit}} \in \mathsf{Tori}$.

Recall that $\mathbb{T}_{\mathsf{orbit}}$ was defined in terms of the integer linear dependencies of $\frac{1}{\log 2}, \frac{1}{\log 3}, \frac{1}{\log 5}$. More generally, the integer linear dependencies of $\frac{1}{\log a_1}, \ldots, \frac{1}{\log a_d}$ will have a basis[15] $\overline{b_1}, \ldots, \overline{b_p}$, where for all i, $\overline{b_i} \in \mathbb{Z}^d$, $\sum_j \frac{b_{ij}}{\log a_j} = 0$, and $p < d - 1$. Our candidate $\mathbb{T}_{\mathsf{res}}$ will be enumerated by linearly independent sets $\{\overline{c_1}, \ldots, \overline{c_q}\} \subset \mathbb{Z}^d$, and have the form

$$\left\{(x_2, \ldots, x_d) \in \mathbb{T} : \text{for all } i, \frac{c_{i2}x_2}{\log a_2} + \cdots + \frac{c_{id}x_d}{\log a_d} \in \mathbb{Z}\right\},$$

where for each i, we have $\frac{c_{i1}}{\log a_1} + \cdots + \frac{c_{id}}{\log a_d} = 0$.

The conditions necessitate that $\overline{c_1}, \ldots, \overline{c_q}$ lie in the free module generated by $\overline{b_1}, \ldots, \overline{b_p}$, by linear independence it suffices to consider $q < d - 1$, and by construction $\mathbb{T}_{\mathsf{res}} \supseteq \mathbb{T}_{\mathsf{orbit}}$. By definition, $\mathbb{T}_{\mathsf{orbit}}$ will necessarily be enumerated when we guess $\overline{b_1}, \ldots, \overline{b_p}$ as the basis.

We finally address the outstanding effectiveness concerns. In order to make the above enumeration effective, we need to resolve the last condition consisting of equalities, and hence invoke the decidability of $T_{\exp}$ (the first-order theory of the reals with exponentiation). Finally, we argue that having assumed the decidability of $T_{\exp}$, one can effectively implement the linear programming required to check that $\mathcal{K}_u \subseteq \mathbb{T}_{\mathsf{res}}$ is nonempty. This completes the proof of Thm. 4.

4 Applications to Discrete Recurrence Sequences

By the turn of the century, Schanuel's conjecture had been used to prove the termination of fundamental algorithms in computer algebra and first-order logic.

[15] These dependencies constitute a submodule L of the free module $\mathbb{Z}^d$; hence L is also free and generated by a basis B with cardinality at most d (see the text [20, App. 2, page 880] for a proof). If $|B| = d$, then by elementary linear algebra, the reciprocals of logs are all 0, a contradiction. If $|B| = d - 1$, linear algebra would imply pairwise multiplicative dependencies between $a_1, \ldots, a_d$, again a contradiction.

While choosing an instruction set to devise algorithms, the computability of elementary numbers is an immensely powerful tool; at a higher level of abstraction, the decidability of the first-order theory $T_{\exp}$ of the real numbers with the exponential function serves as a conveniently accessible interface.

In this section, we discuss how these subroutines help in solving problems about discrete recurrence sequences. We start with the Skolem problem for linear recurrence sequences (LRS). An integer LRS of order k satisfies an integer recurrence relation $u_{n+k} = a_0 u_n + \cdots + a_{k-1} u_{n+k-1}$, and is given by the coefficients $a_0, \ldots, a_{k-1} \in \mathbb{Z}$ as well as the initial terms $u_0, \ldots, u_{n-1} \in \mathbb{Z}$. The Skolem problem asks to decide whether a given LRS has a term that is equal to 0. A practical algorithm to solve the Skolem problem for *simple* LRS was given in [5]: this algorithm is unconditionally correct, but relies on two number-theoretic conjectures for termination, one of which is a p-adic version of Schanuel's conjecture. The role of Schanuel's conjecture is to test whether expressions [5, Prop. 8] of the form $f(\log \lambda_1, \ldots, \log \lambda_k)$ are equal to 0, where $\lambda_1, \ldots, \lambda_d$ are *characteristic roots* of the LRS, and f is a polynomial whose coefficients are *p-adic integers*.

A more sophisticated application is that to hypergeometric sequences [18]. These sequences are given by an initial term $u_0 \in \mathbb{Q}$, and satisfy the recurrence $q(n)u_{n+1} = p(n)u_n$, where $p, q \in \mathbb{Q}[x]$ are polynomials without roots in $\mathbb{N}$. The membership problem asks whether some term equals a given $t \in \mathbb{Q}$. Difficulties arise when p, q are *harmonious*, implying that the sequence converges to a finite nonzero limit ℓ. This limit is expressed in terms of the gamma function, and is not known to be elementary, unless we restrict the spectra of p, q [18, Sec. 4, Property S]. Because the convergence to ℓ is effectively eventually monotone, the critical task is to decide whether $\ell = t$: the result would determine an upper bound on the iterate by which t can occur in the sequence. We thus use Schanuel's conjecture for testing the resulting elementary expression for 0 [18, Sec. 4.2].

For a setting where the above kind of "invariant synthesis" is broader in scope, we return to the realm of linear algebra, and consider linear dynamical systems (LDS). An LDS is specified by a starting point $s \in \mathbb{Q}^d$ and an (invertible) update matrix $A \in \mathbb{Q}^{d \times d}$, and defines a trajectory $(s, As, A^2 s, \ldots)$. The halting problem for LDS considered in [1] additionally takes a set $F \subseteq \mathbb{R}^d$ represented by a first-order formula over the structure $\mathfrak{R}_0 = \langle \mathbb{R}; +, \cdot, < \rangle$ or $\mathfrak{R}_{\exp}$ ($\mathfrak{R}_0$ expanded with the exponential function, as discussed in Sec. 3.1), and asks to decide whether the trajectory intersects F. A "yes" answer is certified by n such that $A^n s \in F$; a "no" answer is certified by an *invariant*, i.e., a set $\mathcal{I} \subseteq \mathbb{R}^d$ such that for all $x \in \mathcal{I}$, $Ax \in \mathcal{I}$, $s \in \mathcal{I}$, and $\mathcal{I} \cap F$ is empty. The paper [1] considers the task of synthesising invariants that consist of finitely many connected components by virtue of being defined in $\mathfrak{R}_{\exp}$, a structure that is *o-minimal*.

We now outline the techniques. It is first shown [1, Thm. 5] that there is a family $\mathcal{J}$ of $\mathfrak{R}_{\exp}$-definable sets parametrised by t_0, such that the set $\mathcal{J}(t_0)$ contains the trajectory for every $t_0 \geq 1$. Furthermore, [1, Lem. 11] asserts that an invariant of the desired form must belong to this family. Critically, the set $\mathcal{T}$ of t_0 for which the set $\mathcal{J}(t_0)$ is indeed an invariant is $\mathfrak{R}_{\exp}$-definable. Finally, using the decidability of $T_{\exp}$ (whose termination is conditional on Schanuel's

conjecture), and the fact that $\mathcal{T}$ consists of finitely many connected components, we can test whether $\mathcal{T}$ is non-empty, and if yes, effectively return $\mathcal{I} = \mathcal{J}(t_0)$ for some $t_0 \in \mathcal{T}$ [1, Thm. 12]. We remark that if the set F is defined in $\mathfrak{R}_0$, then the $\mathfrak{R}_{\exp}$-definable invariant synthesis is unconditional [1, Thm. 13].

5 Applications to Quantitative Verification

A fundamental object in quantitative verification is the *weighted automaton*, which can be intuited as "implementing (linear) recurrence with branching." In this section, we survey two instances of algorithms for weighted automata invoking Schanuel's conjecture through queries to $T_{\exp}$.

Formally, a weighted automaton computes over a (commutative) semiring $\mathcal{R}$, i.e., a set that is equipped with distinguished elements $0 \neq 1$, a commutative and associative addition operation $+$ that has 0 as its identity, and a commutative and associative multiplication operation $\cdot$ that has 1 as its identity, distributes over $+$, and has 0 as its absorbing element (for all x, $0 \cdot x = 0$). For example, the usual finite-word automata compute over the Boolean semiring $\{0, 1\}$ where $+$ is disjunction and $\cdot$ is conjunction. In this section, we shall survey results for weighted automata that compute over the semiring of nonnegative rational numbers with the usual addition and multiplication.

In general, a weighted automaton $\mathcal{A}$ over the semiring $\mathcal{R}$ is given by the tuple $(Q, \Sigma, \Delta, I, F)$, where Q is the finite set of states, Σ is the alphabet, $\Delta \subseteq Q \times \Sigma \times \mathcal{R} \times Q$ is the finite set of transitions, $I : Q \to \mathcal{R}$ is the initial weight function, and $F : Q \to \mathcal{R}$ is the final weight function.[16] Given an input $u = u(0) \cdots u(\ell - 1) \in \Sigma^*$, a path of $\mathcal{A}$ on u is given as $w_0, q_0, u(0), w_1, q_1, \ldots, q_{\ell-1}, u(\ell-1), w_\ell, q_\ell, w_{\ell+1}$, where $w_0 = I(q_0)$, for all $j \in \{0, \ldots, \ell - 1\}$, we have $(q_j, u(j), w_{j+1}, q_{j+1}) \in \Delta$, and $w_{\ell+1} = F(q_\ell)$. A path on the empty word is simply w_0, q_0, w_1, where $w_0 = I(q_0), w_1 = F(q_0)$. The weight of the path is the product $w_0 \cdot w_1 \cdots w_{\ell+1}$. The weight of the word u is the sum of weights of all paths of $\mathcal{A}$ on u (hence the weight is 0 if there is no path). The automaton thus defines a function $[\![\mathcal{A}]\!] : \Sigma^* \to \mathcal{R}$.

For example, a conventional automaton $(Q, \Sigma, \Delta, I, F)$ (where $\Delta \subseteq Q \times \Sigma \times Q$, $I, F \subseteq Q$, as is familiar) can be interpreted as a weighted automaton over the Boolean semiring as follows: we replace each $(q, a, q') \in \Delta$ with $(q, a, 1, q')$, construct the initial weight function to assign 1 to elements of I and 0 to others, and similarly the final weight function assigns 1 to elements of F and 0 to others. The weight of a path is 1 if it starts in an initial state and ends in an accepting state, and 0 otherwise; a word is assigned weight 1 if it is accepted by the automaton, and weight 0 otherwise.

As mentioned before, we shall consider weighted automata over nonnegative rational numbers. If, in addition, we have that all weights are at most 1, the sum of all initial weights is 1, all final weights are 0 or 1, and for each state q, the sum of weights of all outgoing transitions is 1, the automaton is said to be *probabilistic*. Probabilistic automata are closely related to other stochastic models such as Markov chains, and Markov decision processes.

[16] In this paper, we only survey works that consider $F : Q \to \{0, 1\}$.

Weighted, and indeed even probabilistic automata constitute a powerfully expressive model of computation: it is folklore [26] (see also [14] for a modern proof) that the emptiness problem for probabilistic automata, which gives $\mathcal{A}$ and $t \in [0,1]$ and asks whether there is some word $u \in \Sigma^*$ such that $[\![\mathcal{A}]\!](u) > t$, is undecidable. Even over a unary alphabet, which makes the automaton a Markov chain, the problem is as hard as the yet unresolved positivity problem for linear recurrence sequences [33]. In order to obtain decidability results for weighted automata, works have thus restricted either the structure of the input automaton, or the language $\mathcal{L}(\mathcal{A})$ of words that are assigned nonzero weight.

The former is done by restricting *ambiguity*: an automaton $\mathcal{A}$ is k-ambiguous if there are at most k nonzero paths for every $u \in \Sigma^*$. We say that $\mathcal{A}$ is finitely ambiguous if it is k-ambiguous for some k, and unambiguous if $k = 1$. We can further parametrise ambiguity, and say that an automaton is polynomially (respectively, linearly) ambiguous if for all u, there are at most $p(|u|)$ nonzero paths, where p is a polynomial (respectively, a polynomial of degree 1). This distinction is technically relevant, because if an automaton fails to be finitely ambiguous, then it is at least linearly ambiguous [34, Sec. 3]. Furthermore, it has been shown [11, Proof of Thm. 1] that the emptiness problem is undecidable even for linearly ambiguous probabilistic automata.

The only hope to recover decidability with this restriction, therefore, is to assume finite ambiguity. Indeed, [11, Thm. 2] shows that assuming Schanuel's conjecture, the containment problem, which gives probabilistic automata $\mathcal{A}, \mathcal{B}$ over the alphabet Σ, and asks whether for all $u \in \Sigma^*$, $[\![\mathcal{A}]\!](u) \leq [\![\mathcal{B}]\!](u)$, is decidable for the class of finitely ambiguous probabilistic automata, provided at least one of the input automata is unambiguous. In fact, the decidability is unconditional if $\mathcal{A}$ is finitely ambiguous and $\mathcal{B}$ is unambiguous [11, Prop. 5].

Otherwise, the task is equivalent to the integer program with exponentiation problem that asks whether the system $M\overline{x} < \overline{c}$, $\sum_{i=1}^{\ell} r_i s_{i,1}^{x_1} \cdots s_{i,n}^{x_n} < 1$ has a solution $\overline{x} \in \mathbb{Z}^n$, where M is an integer matrix, $\overline{c}$ is an integer vector, and $r_i, s_{i,j}$ are all positive rationals. Obviously, if there is a solution, a semi-algorithm that simply enumerates integer vectors will find it. It remains to describe a semi-algorithm that would certify the absence of integral feasible points.

The key technical observation [11, Lem. 10] is that if there is no integer solution, then the feasible region X of *real* solutions must be contained in a "tube", i.e., a set $\{\overline{y} \in \mathbb{R}^n : \overline{d}^\top \overline{y} \in [a,b]\}$, where $\overline{d} \in \mathbb{Z}^n, a, b \in \mathbb{Z}$. If we find such an encompassing tube, we know that any integer feasible point must satisfy $\overline{d}^\top \overline{x} = i$ for some integer $i \in [a,b]$, and hence can reduce the search for a certificate of absence to finitely many lower-dimensional systems [11, Lem. 11].

To make the proposed semi-algorithm effective, we need to check whether a tube purported by $\overline{d}, a, b$ indeed contains the feasible set X. This is the step that uses Schanuel's conjecture for termination, because it is implemented as a query to $T_{\exp}$. However, given the lack of explicit transcendence in the integer programming with exponentiation problem, we cannot yet rule out the circumvention of Schanuel's conjecture.

The general undecidability of the containment problem motivated the consideration of an approximate variant, namely, the big-O problem for weighted automata [8], which gives as input a weighted automaton $\mathcal{A}$, two states q, q', and asks whether there exists $c > 0$ such that for all $u \in \Sigma^*$, we have $[\![\mathcal{A}_q]\!](u) \leq c \cdot [\![\mathcal{A}_{q'}]\!](u)$, where $\mathcal{A}_q$ is the automaton obtained by modifying the initial weight function of $\mathcal{A}$ to return 1 for q and 0 for all other states. Unfortunately, as [8] establishes, the big-O problem is also undecidable in general, and relatively tractable in very special cases (in P for unambiguous automata, in coNP if the alphabet is unary).

The interesting restriction is that of requiring the languages $\mathcal{L}(\mathcal{A}_q)$ and $\mathcal{L}(\mathcal{A}_{q'})$ be *bounded*: a language $L \subseteq \Sigma^*$ is bounded if it is contained in $w_1^* \cdots w_k^*$ for some words $w_1, \ldots, w_k \in \Sigma^*$. The big-O problem for these instances is decidable subject to Schanuel's conjecture [8, Thm. 28].

The technical algorithm requires a subroutine to test whether a first-order sentence is in $T_{\exp}$ [8, Lem. 37]. The expressions therein involve the logarithms of variables, as well as the logarithms of real positive algebraic constants as coefficients of linear terms, and hence, in this case, Schanuel's conjecture appears to play a necessary role.

6 Applications to Continuous-Time Systems

In this final technical section, we survey how queries about elementary functions naturally arise when reasoning about continuous-time (linear) dynamical systems. These queries often involve trigonometric functions: recall from Sec. 3.2 that in order for the first-order theory of the attendant structure to be decidable, these functions must be restricted to bounded intervals. Consequently, the literature predominantly considers decision problems whose inputs specify a bounded time interval $[0, N]$ of interest. As discussed in Sec. 3.2, Schanuel's conjecture then assures[17] the decidability of the theory $T_{\mathrm{el},N}$ of the structure $\langle \mathbb{R}; <, +, \cdot, \exp \restriction [0,N], \sin \restriction [0,N], \cos \restriction [0,N] \rangle$.

A fundamental continuous-time decision problem is the bounded continuous Skolem problem, which asks whether the solution $f : \mathbb{R} \to \mathbb{R}$ of the ordinary differential equation $f^{(n)} + a_{n-1} f^{(n-1)} + \cdots + a_0 f \equiv 0$, with the coefficients $a_0, \ldots, a_{n-1}$, and initial conditions $f(0), \ldots, f^{(n-1)}(0)$ being real algebraic numbers, has a zero in the interval $[c, d]$. We have that the function f is the first component of the vector $\exp(tA) \cdot \mu_0$, where A is the *companion matrix*[18] derived from $a_0, \ldots, a_{n-1}$, and μ_0 is the vector of initial conditions. We can therefore express $f(t) = \sum_{j=1}^{m} \exp(\rho_j t)(p_j(t) \sin(\omega_j t) + q_j(t) \cos(\omega_j t))$, where ρ_j, ω_j, and the coefficients of the polynomials p_j, q_j are real algebraic numbers. As in [9, Introduction], we can choose $N = (d+1) \cdot \max(\rho_1, \omega_1, \ldots, \rho_m, \omega_m)$, and query $T_{\mathrm{el},N}$ to arrive at a decision.

[17] We recall that this claim is based on unpublished results; however, we outlined why we believe that Schanuel's conjecture is needed only for the termination of the deciding algorithm.

[18] The matrix exponential $\exp(M)$ is defined as the limit of $I + M + M^2/2! + M^3/3! + \cdots$.

However, since the decidability proof of $T_{\mathrm{el},N}$ is not published, [9] gives a decision procedure for the bounded continuous Skolem problem by invoking Schanuel's conjecture directly (via Prop. 5 therein) in technical algebraic arguments. The technical ingredient that is the difference between the raw Schanuel's conjecture and the off-the-shelf decidability of $T_{\mathrm{el},N}$ is the model-completeness of the theory. Bereft of this refinement, the direct algorithm of [9] relies on Schanuel's conjecture for correctness in Case (i) of the proof of its Thm. 7 (the main result, which states that the bounded-time continuous Skolem problem is decidable subject to Schanuel's conjecture, and is proven via a technical case distinction): the algorithm trusts Schanuel's conjecture's assertion that a function f has no root without a certificate; on the other hand, in case the claim is true, the (effectively) model-complete $T_{\mathrm{el},N}$ would provide a proof in the form of the claim being a formal consequence of the theory. As usual, the termination is also conditional as Case (ii) of the proof of Thm. 7 indicates. See also [10] for an alternative presentation of this argument.

The choice of whether to trust Schanuel's conjecture for correctness and forego the cumbersome queries to logical theories that give unconditional correctness has practical consequences. The work [17] is an applied example of prioritising efficiency: they give an algorithm to isolate roots of an exponential-polynomial, and rely on Schanuel's conjecture to argue its completeness, i.e., that the algorithm does not discount any roots, and termination [17, Discussion of Example 16].

Such root isolation algorithms find immediate applications to verify the evolution of the probability distributions in continuous time Markov chains. A continuous time Markov chain with k states is specified by a $k \times k$ matrix M, where the entry $M(i,j) = m$ indicates that a transition from state j to state i is possible, and its timing is exponentially distributed with rate m. Succinctly, the dynamics of the probability distribution μ are given by $\mu^{(1)}(t) = M\mu(t)$, and the distribution at time t is given by $\exp(Mt) \cdot \mu_0$, where μ_0 is the initial distribution at $t = 0$.

In order to verify the evolution of the distribution, [16] considers CLL, a continuous (bounded-time) counterpart of LTL, whose atomic propositions are of the form "the probability of being in state s is at least 0.8"; an example specification $\mathit{true}\ \mathcal{U}^{[3,7]}\langle s, \geq\ 0.8\rangle$ asserts "the probability of being in state s is at least 0.8 at some moment $t \in [3,7]$." Upon observing the expansion of $\exp(Mt) \cdot \mu_0$ and the semantics of CLL, we remark that the satisfaction of a formula can, in theory, be verified with unconditional correctness by querying $T_{\mathrm{el},N}$ for sufficiently large N. However, as [16] demonstrates, trusting Schanuel's conjecture for correctness makes the problem more tractable in practice.

As in the discrete case, having techniques to reason about continuous-time Markov chains equips us to reason about certain *policies* in continuous-time Markov *decision processes* (MDPs), which augment Markov chains by determining transition dynamics at any given time by an agent's resolution of a finite choice of *actions*. A policy is a function that prescribes these choices; it is called

stationary if the prescription does not change with time. We observe that following a stationary policy induces a Markov chain.

MDPs form the building blocks of reinforcement learning, and [23] considers the problem of deciding whether there is a policy that achieves the objective of visiting a designated good state within given time bound is achieved with probability above a given threshold. The fact that the optimal policy for this task is piecewise stationary allows the problem to be formulated as one about Markov chains in spirit, whence the decidability of $T_{\mathrm{el},N}$ can be called upon for a solution.

Finally, we note [37] uses a slight augmentation of MDPs to model the runs of probabilistic programs while accounting for scheduling delays, and considers the problem of quantitatively verifying termination within a time bound. The number-theoretic engine [37, Lem. 6.2] for their results is the root-isolation subroutine, which is efficient in practice provided that the implementation trusts Schanuel's conjecture for correctness, but in theory can also be implemented using queries to $T_{\mathrm{el},N}$ too.

References

1. 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)*. Schloss-Dagstuhl-Leibniz Zentrum für Informatik, 2018.
2. James Ax. On Schanuel's Conjectures. *Annals of Mathematics*, 93(2):252–268, 1971.
3. Alan Baker. *Transcendental Number Theory*. Cambridge Mathematical Library. Cambridge University Press, 1975.
4. 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 *Proceedings of the 39th Annual ACM/IEEE Symposium on Logic in Computer Science*, pages 1–14, 2024.
5. 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*, 2022.
6. Eion Blanchard and Philipp Hieronymi. Decidability bounds for Presburger arithmetic extended by sine. *Annals of Pure and Applied Logic*, page 103487, 2024.
7. B. F. Caviness. On Canonical Forms and Simplification. *J. ACM*, 17(2):385–396, April 1970.
8. Dmitry Chistikov, Stefan Kiefer, Andrzej S. Murawski, and David Purser. The big-O problem for labelled Markov chains and weighted automata. *Leibniz International Proceedings in Informatics (LIPIcs)*, 171:41, 2020.
9. 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)*. Schloss-Dagstuhl-Leibniz Zentrum für Informatik, 2016.
10. Ventsislav Chonev, Joël Ouaknine, and James Worrell. On the zeros of exponential polynomials. *J. ACM*, 70(4):26:1–26:26, 2023.

11. Laure Daviaud, Marcin Jurdziński, Ranko Lazić, Filip Mazowiecki, Guillermo A Pérez, and James Worrell. When is containment decidable for probabilistic automata? In *45th International Colloquium on Automata, Languages, and Programming, ICALP 2018*, page 121. Schloss Dagstuhl-Leibniz-Zentrum fur Informatik GmbH, Dagstuhl Publishing, 2018.
12. 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.
13. Helaman R. P. Ferguson, David H. Bailey, and Steve Arno. Analysis of pslq, an integer relation finding algorithm. *Math. Comput.*, 68(225):351–369, January 1999.
14. Hugo Gimbert and Youssouf Oualhadj. Probabilistic Automata on Finite Words: Decidable and Undecidable Problems. In Samson Abramsky, Cyril Gavoille, Claude Kirchner, Friedhelm Meyer auf der Heide, and Paul G. Spirakis, editors, *Automata, Languages and Programming*, pages 527–538, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg.
15. 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.
16. Ji Guan and Nengkun Yu. A probabilistic logic for verifying continuous-time Markov chains. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, pages 3–21. Springer, 2022.
17. Cheng-Chao Huang, Jing-Cao Li, Ming Xu, and Zhi-Bin Li. Positive root isolation for poly-powers by exclusion and differentiation. *Journal of Symbolic Computation*, 85:148–169, 2018.
18. George Kenison. The Threshold Problem for Hypergeometric Sequences with Quadratic Parameters. In *51st International Colloquium on Automata, Languages, and Programming (ICALP 2024)*. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2024.
19. Serge Lang. *Introduction to Transcendental Numbers*. Addison-Wesley series in mathematics. Addison-Wesley Publishing Company, 1966.
20. Serge Lang. *Algebra*. Graduate Texts in Mathematics. Springer New York, NY, 3rd edition, 2002.
21. Angus Macintyre. Turing meets Schanuel. *Annals of Pure and Applied Logic*, 167(10):901–938, 2016. Logic Colloquium 2012.
22. Angus Macintyre, A Wilkie, and P Odifreddi. On the decidability of the real exponential field. *Kreisel's Mathematics*, 115:451, 1996.
23. Rupak Majumdar, Mahmoud Salamati, and Sadegh Soudjani. On decidability of time-bounded reachability in CTMDPs. In *47th International Colloquium on Automata, Languages, and Programming (ICALP 2020)*. Schloss-Dagstuhl-Leibniz Zentrum für Informatik, 2020.
24. Nathanaël Mariaule. *On the decidability of the p-adic exponential ring*. The University of Manchester (United Kingdom), 2013.
25. A. Muchnik, A. Semenov, and M. Ushakov. Almost periodic sequences. *Theoretical Computer Science*, 304(1):1–33, 2003.
26. Azaria Paz. *Introduction to probabilistic automata (Computer science and applied mathematics)*. Academic Press, Inc., USA, 1971.
27. Daniel Richardson. A simplified method of recognizing zero among elementary constants. In *Proceedings of the 1995 international symposium on Symbolic and algebraic computation*, pages 104–109, 1995.
28. Daniel Richardson. How to Recognize Zero. *Journal of Symbolic Computation*, 24(6):627–645, 1997.

29. Daniel Richardson. Zero tests for constants in simple scientific computation. *Mathematics in Computer Science*, 1:21–37, 2007.
30. Maxwell Rosenlicht. On Liouville's theory of elementary functions. *Pacific Journal of Mathematics*, 65:485–492, 1976.
31. Alfred Tarski. A decision method for elementary algebra and geometry. In Bob F. Caviness and Jeremy R. Johnson, editors, *Quantifier Elimination and Cylindrical Algebraic Decomposition*, pages 24–84, Vienna, 1998. Springer Vienna.
32. Alfred Tarski. A Decision Method for Elementary Algebra and Geometry. In Bob F. Caviness and Jeremy R. Johnson, editors, *Quantifier Elimination and Cylindrical Algebraic Decomposition*, pages 24–84, Vienna, 1998. Springer Vienna.
33. Mihir Vahanwala. Skolem and positivity completeness of ergodic Markov chains. *Information Processing Letters*, 186:106481, 2024.
34. Andreas Weber and Helmut Seidl. On the degree of ambiguity of finite automata. *Theoretical Computer Science*, 88(2):325–349, 1991.
35. A. J. Wilkie. Model Completeness Results for Expansions of the Ordered Field of Real Numbers by Restricted Pfaffian Functions and the Exponential Function. *Journal of the American Mathematical Society*, 9(4):1051–1094, 1996.
36. A. J. Wilkie. *Schanuel's Conjecture and the Decidability of the Real Exponential Field*, pages 223–230. Springer Netherlands, Dordrecht, 1997.
37. Ming Xu and Yuxin Deng. Time-bounded termination analysis for probabilistic programs with delays. *Information and Computation*, 275:104634, 2020.
38. B. Zilber. Pseudo-exponentiation on algebraically closed fields of characteristic zero. *Annals of Pure and Applied Logic*, 132(1):67–95, 2005.