Toric Certificates for Effective Skolem Tails and Subspace Orbit Reachability

Luca Blanchi

June 2026

Abstract

We develop a certificate framework for proving effective zero-free tails for scalar and vector algebraic linear recurrence sequences. The framework applies to the Skolem Problem, the Simultaneous Skolem Problem, and certified instances of the Subspace Orbit Problem.

After passing to a common Binet representation, splitting root-of-unity degeneracies, saturating residual torsion, and peeling globally cancelled dominant shells, each branch of a vector recurrence is put in the form

$$\mathbf{u}_{Mk+r} = \Lambda^k\big(\mathbf{D}(k, z(k)) + \mathbf{E}(k)\big),$$

where

$$z(k) = c\eta^k \in \mathbb{T}^\rho, \qquad \mathbf{D}(T, X) \in L[T][X_1^{\pm 1}, \dots, X_\rho^{\pm 1}]^m,$$

and

$$|\mathbf{E}(k)| \le C_E(k+1)^s\theta^k, \qquad 0 < \theta < 1.$$

The toric rank, the order of the recurrence, the number of dominant roots, and the root multiplicities are unrestricted.

The main theorem is a soundness theorem. If the dominant vector admits a verified toric lower-bound certificate, then all zeros of the branch lie below a computable bound unless the branch is identically zero. The remaining prefix is checked exactly. We give several verifiable certificate formats: joint semialgebraic certificates, Positivstellensatz and sum-of-squares certificates, Bezout–toric certificates, scalar character-polynomial certificates, and binomial-stratified certificates.

The new automatic layer is binomial-stratified. If the dominant zero locus on the compact toric closure is contained in finitely many effective binomial strata, then a semialgebraic Lojasiewicz inequality gives a lower bound by a power of the distance to those strata. Along the actual toric orbit, Baker–Matveev separation turns the binomial distances into polynomial lower bounds away from exact-hit progressions. Exact hits are computable finite unions of arithmetic progressions; on those progressions one restricts the recurrence and either peels a vanished dominant shell or lowers the toric dimension. In the automatic binomial subclass, the strata are extracted from Laurent elimination ideals and binomial decomposition.

The resulting compiler computes exact zero sets for every recurrence admitting a finite binomial-stratified toric certificate tree, and computes the strata itself in the automatic binomial-stratified class. This isolates a checkable and partially automatic mechanism for proving effective Skolem tails in high-rank, high-order families, including vector examples where scalar coordinate-wise certification fails. The geometric boundary of the method is explicit: arbitrary non-binomial toric approximation requires additional lower-bound input.

Declaration of generative AI and AI-assisted technologies in the writing process

During the preparation of this work, ChatGPT, by OpenAI, was used to assist with mathematical drafting, formalization, review, and editing. This work is shared as a preliminary AI-assisted mathematical note. The mathematical content may have been only partially reviewed and may contain errors; it should not be treated as peer-reviewed or as a fully verified manuscript.

1 Introduction

Let K be a number field. A scalar linear recurrence sequence over K is a sequence $u = (u_n)_{n\geq 0}$ satisfying

$$u_{n+d} = a_{d-1}u_{n+d-1} + \cdots + a_0 u_n$$

for some $a_0, \ldots, a_{d-1} \in K$. The Skolem Problem asks whether there exists $n \geq 0$ such that

$$u_n = 0.$$

The strong computational form asks to compute the entire zero set

$$Z(u) = \{n \geq 0 : u_n = 0\}.$$

The Skolem–Mahler–Lech theorem asserts that $Z(u)$ is a finite union of arithmetic progressions together with a finite exceptional set. Its known proofs do not provide a general algorithm for computing the finite exceptional set. The computational heart of Skolem is therefore the construction of an effective zero-free tail.

This paper studies such tails through toric certificates. The scalar case is included, but the natural level of generality is vectorial. Let

$$\mathbf{u}_n = (u_1(n), \ldots, u_m(n)) \in K^m$$

be a vector of algebraic linear recurrence sequences. The simultaneous zero set is

$$Z(\mathbf{u}) = \{n \geq 0 : \mathbf{u}_n = 0\}.$$

This formulation covers the Simultaneous Skolem Problem. It also covers affine subspace orbit reachability: if $A \in K^{d\times d}$, $x_0 \in K^d$, and

$$V = \{y \in K^d : By = b\},$$

then

$$A^n x_0 \in V \quad \Longleftrightarrow \quad BA^n x_0 - b = 0.$$

The central idea is to separate the dominant vector from zero. On each branch, after standard spectral reductions, we obtain

$$\mathbf{u}_{Mk+r} = \Lambda^k(\mathbf{D}(k, z(k)) + \mathbf{E}(k)),$$

where $z(k)$ is a toric orbit on a compact toric coset, $\mathbf{D}$ is the normalized dominant vector, and $\mathbf{E}$ is exponentially smaller. If one can certify that

$$\mathbf{D}(k, z(k)) \neq 0 \quad \Longrightarrow \quad |\mathbf{D}(k, z(k))| \geq \exp(-C(\log(k+2))^A)$$

outside a computable prefix, then the exponentially small error cannot cancel the dominant vector outside that prefix. The rest is an exact finite computation.

The vector viewpoint is essential. A scalar coordinate may contain a factor such as

$$X + Y - 1,$$

which is outside the automatic one-character Baker class. But a system such as

$$(X + Y - 1,\ X - Y)$$

has no common zero on $\mathbb{T}^2$, and the norm of the vector is bounded away from zero by a semialgebraic certificate. Thus the vector method proves effective tails in cases not visible to scalar coordinatewise methods.

The contributions are:

1. a global vector dominant-shell normal form with torsion saturation;
2. a verifier theorem for toric lower-bound certificates;
3. joint semialgebraic and Positivstellensatz/SOS certificates;
4. Bezout–toric certificates converting scalar lower bounds into vector lower bounds;
5. a scalar Baker–Matveev lemma for expressions $F(k, \delta^k)$;
6. quotient-character reduction, producing scalar atoms automatically from Laurent factors whose exponent supports become one-dimensional modulo the toric relation lattice;
7. binomial-stratified certificates combining semialgebraic Lojasiewicz bounds, Baker separation from exact binomial hits, and recursive restriction;
8. automatic extraction of binomial strata from dominant elimination ideals in the binomial subclass;
9. a partial Factor–Reduce–Certify compiler and an automatic binomial-stratified certificate-tree compiler;
10. applications to Simultaneous Skolem, Subspace Orbit, affine linear-loop reachability, algebraic target hitting, and eventual cone membership.

The framework isolates certified toric dominant shells: once a dominant lower-bound mechanism is verified, the zero-free tail and the exact finite prefix become effective. Non-binomial toric approximation remains the boundary where additional Diophantine lower bounds would be needed.

2 Effective input and certificate model

An input consists of:

1. a number field K, represented effectively;
2. either a scalar or vector linear recurrence sequence over K, or a linear dynamical system (A, x_0, B, b);

3. a complex embedding $\sigma : K \hookrightarrow \mathbb{C}$.

All computations take place in finite extensions of K. We use standard exact algorithms in number fields: arithmetic, factorization, construction of splitting fields, root-of-unity testing, exact zero testing, comparison of algebraic absolute values under σ, and computation of multiplicative relation lattices [12]. These algorithms are not claimed to be efficient in general; they are used as exact symbolic subroutines.

A certificate is extended input. The algorithms in this paper are sound: if the certificate verifies, then the output zero set is correct. They are not complete: failure to certify means "not certified", not "no zero" and not "undecidable".

We separate three layers.

- The verifier theorem: a valid certificate implies an effective zero-free tail.
- Automatic certificate generation: for certain structural classes, the certificate is constructed by the compiler.
- Positive search: for some certificates, such as Bezout–toric identities of bounded degree, one may search positively, but failure of the search has no mathematical consequence.

3 Vector Binet form and global shells

Let

$$\mathbf{u}_n = (u_1(n), \dots, u_m(n)) \in K^m$$

be a vector of algebraic linear recurrence sequences. Passing to a common splitting field L/K, we may write

$$\mathbf{u}_n = \sum_{\lambda \in \mathcal{R}} \mathbf{P}_\lambda(n) \lambda^n,$$

where $\mathcal{R} \subset L^\times$ is finite and

$$\mathbf{P}_\lambda(T) \in L[T]^m.$$

Zero characteristic roots contribute only to a finite initial segment. That segment is checked exactly, and the sequence is shifted. Hence all roots in the remaining Binet representation are nonzero.

Fix an embedding

$$\sigma : L \hookrightarrow \mathbb{C}.$$

A global vector shell is a set

$$S_\rho = \{\lambda \in \mathcal{R} : |\sigma(\lambda)| = \rho\}.$$

The shell contribution is

$$\mathbf{U}_\rho(n) = \sum_{\lambda \in S_\rho} \mathbf{P}_\lambda(n) \lambda^n.$$

The shell is vectorially cancelled if

$$\mathbf{U}_\rho(n) \equiv 0$$

as a vector-valued exponential-polynomial function. This means every coordinate is identically zero.

After torsion splitting, distinct bases with no root-of-unity quotient give linearly independent exponential-polynomial functions. Therefore vectorial cancellation is decidable exactly.

4 Torsion splitting and torsion saturation

If λ/μ is a root of unity, let $o(\lambda,\mu)$ be its order. Let

$$M_0 = \operatorname{lcm}\{o(\lambda,\mu) : \lambda,\mu \in \mathcal{R},\ \lambda/\mu \in \mu_\infty\}.$$

For each residue class $r \pmod{M_0}$, define

$$\mathbf{v}_k^{(r)} = \mathbf{u}_{M_0 k + r}.$$

Equal bases in the resulting Binet representation are combined exactly.

A second torsion issue remains. A shell with relative multiplicative rank ρ modulo torsion need not be an exact toric coset. Its roots have the form

$$\lambda = \Lambda\, \xi_\lambda\, \gamma^a,$$

where ξ_λ is a root of unity and

$$\gamma^a = \gamma_1^{a_1} \cdots \gamma_\rho^{a_\rho}, \qquad a \in \mathbb{Z}^\rho.$$

Let

$$Q = \operatorname{lcm}_\lambda \operatorname{ord}(\xi_\lambda).$$

Passing to subbranches

$$k = Qh + s$$

makes

$$\xi_\lambda^k = \xi_\lambda^s$$

constant. The polynomial coefficients are transformed by

$$\mathbf{P}_\lambda(k) \mapsto \mathbf{P}_\lambda(Qh+s),$$

and the bases by

$$\lambda^k \mapsto \lambda^s (\lambda^Q)^h.$$

All new data are computable. We call this step torsion saturation.

5 Vector toric normal form

After torsion splitting and torsion saturation, peel globally cancelled shells until the first non-cancelled global shell is reached. Let S be that shell. Choose $\Lambda \in S$. The roots in S can be written as

$$\lambda = \Lambda\gamma^a, \qquad a \in A \subset \mathbb{Z}^\rho.$$

Since all roots in S have the same σ-absolute value, the toric parameters satisfy

$$|\sigma(\gamma_j)| = 1.$$

Define

$$z(k) = (\gamma_1^k, \ldots, \gamma_\rho^k).$$

Then the contribution of S has the form

$$\Lambda^k \mathbf{D}(k, z(k)),$$

where

$$\mathbf{D}(T, X) \in L[T][X_1^{\pm 1}, \ldots, X_\rho^{\pm 1}]^m.$$

All lower shells have strictly smaller σ-absolute value. Hence the branch has normal form

$$\mathbf{v}_k = \Lambda^k(\mathbf{D}(k, z(k)) + \mathbf{E}(k)),$$

with

$$|\mathbf{E}(k)| \leq C_E(k+1)^s\theta^k, \qquad 0 < \theta < 1.$$

Here

$$|\mathbf{w}|^2 = \sum_{i=1}^{m} |\sigma(w_i)|^2.$$

6 Toric cosets

After reindexing, it is convenient to write

$$z(k) = c\eta^k = (c_1\eta_1^k, \ldots, c_\rho\eta_\rho^k),$$

with

$$|\sigma(c_j)| = |\sigma(\eta_j)| = 1.$$

Define the exact relation lattice

$$\mathcal{L}_\eta = \{a = (a_1, \ldots, a_\rho) \in \mathbb{Z}^\rho : \eta^a = 1\},$$

where

$$\eta^a = \eta_1^{a_1} \cdots \eta_\rho^{a_\rho}.$$

This lattice is computable.

Let

$$T_\eta = \{y \in \mathbb{T}^\rho : y^a = 1 \text{ for all } a \in \mathcal{L}_\eta\}.$$

The orbit closure of $z(k) = c\eta^k$ is the compact toric coset

$$C = cT_\eta.$$

Equivalently,

$$C = \{x \in \mathbb{T}^\rho : x^a = c^a \text{ for all } a \in \mathcal{L}_\eta\}.$$

We do not assume T_η is connected. If finite torsion remains in η, T_η may have finitely many components. The equations above describe the full compact coset.

For effective real-algebraic verification, write

$$X_j = x_j + iy_j.$$

The torus is defined by

$$x_j^2 + y_j^2 = 1,$$

and the coset equations are the real and imaginary parts of

$$X^a = c^a$$

for a basis of $\mathcal{L}_\eta$.

7 The certified-tail theorem

Theorem 7.1 (Certified-tail soundness)**.** *Let a saturated branch have vector toric normal form*

$$\mathbf{v}_k = \Lambda^k\big(\mathbf{D}(k, z(k)) + \mathbf{E}(k)\big),$$

where

$$|\mathbf{E}(k)| \le C_E(k+1)^s\theta^k, \qquad 0 < \theta < 1.$$

Assume there are computable constants A, C, N_0 such that, for every $k \ge N_0$,

$$\mathbf{D}(k, z(k)) \neq 0 \quad \Longrightarrow \quad |\mathbf{D}(k, z(k))| \ge \exp(-C(\log(k+2))^A).$$

Assume also that exact zeros of

$$\mathbf{D}(k, z(k))$$

are either contained in a computable prefix or occur only on subbranches where the global dominant vector is identically cancelled and hence peeled.

Then the branch has a computable zero-free tail. Consequently the zeros on the branch are computable exactly as a finite set plus any identically zero subprogressions.

Proof. Choose $B \ge N_0$ such that

$$C_E(k+1)^s\theta^k < \frac{1}{2}\exp(-C(\log(k+2))^A)$$

for all $k \ge B$. Such B is computable. For $k \ge B$, if $\mathbf{D}(k, z(k)) \neq 0$, then

$$|\mathbf{D}(k, z(k)) + \mathbf{E}(k)| \ge |\mathbf{D}(k, z(k))| - |\mathbf{E}(k)| > 0.$$

Hence $\mathbf{v}_k \neq 0$. The remaining possible zeros are in a computable prefix or in identically zero subbranches produced by peeling. The finite prefix is checked exactly in L. □

8 Joint semialgebraic certificates

Definition 8.1 (Joint semialgebraic certificate)**.** *Let*

$$\mathbf{D}(T, X) = (D_1(T, X), \dots, D_m(T, X)).$$

A joint semialgebraic certificate, or joint S-certificate, on a toric coset $C = cT_\eta$ consists of explicit data

$$B_0 \in \mathbb{N}, \qquad L_0 \in \mathbb{N}, \qquad q \in \mathbb{Q}_{>0}$$

such that

$$\sum_{i=1}^{m} |\sigma(D_i(t, x))|^2 \ge q(t+1)^{-L_0}$$

for every

$$t \ge B_0, \qquad x \in C.$$

The assertion is a universal formula over a real closed field. Laurent monomials are represented on the torus by

$$X_j = x_j + iy_j, \qquad X_j^{-1} = x_j - iy_j.$$

The coset C is represented by polynomial equations

$$x_j^2 + y_j^2 = 1$$

and the real and imaginary parts of

$$X^a = c^a$$

for a basis of $\mathcal{L}_\eta$.

The certificate condition becomes

$$\forall t, x, y : \ (t \geq B_0 \wedge (x+iy) \in C) \Rightarrow (t+1)^{L_0} \sum_i |\sigma(D_i(t, x+iy))|^2 - q \geq 0.$$

This formula is effectively decidable by quantifier elimination over real closed fields.

Proposition 8.2. *A joint S-certificate implies the hypothesis of the certified-tail theorem with a polynomial lower bound:*

$$|\mathbf{D}(k, z(k))| \geq q^{1/2}(k+1)^{-L_0/2}.$$

Proof. Since $z(k) \in C$, substituting $t = k$ and $x = z(k)$ into the certificate gives the result. □

9 Binomial-stratified certificates

Uniform separation on the whole compact toric coset is a useful certificate, but many natural dominant vectors vanish on structured subtori. The certificate can still be effective when the entire dominant zero locus is forced into finitely many binomial strata.

Definition 9.1 (Binomial stratum)**.** *Let $C \subseteq \mathbb{T}^\rho$ be an effective compact toric coset. A binomial stratum in C is a subset*

$$H = \{x \in C : \chi_{a_1}(x) = \zeta_1, \ldots, \chi_{a_\ell}(x) = \zeta_\ell\},$$

where

$$a_j \in \mathbb{Z}^\rho, \qquad \chi_{a_j}(X) = X^{a_j},$$

and each ζ_j is algebraic of complex absolute value 1. The stratum is proper if $H \neq C$.

For such a stratum set

$$\Delta_H(x) = \sum_{j=1}^{\ell} |\chi_{a_j}(x) - \zeta_j|^2.$$

Then

$$\Delta_H(x) = 0 \quad \Longleftrightarrow \quad x \in H.$$

For a finite family

$$\mathcal{H} = \{H_1, \ldots, H_N\},$$

define

$$\Delta_{\mathcal{H}}(x) = \prod_{\nu=1}^{N} \Delta_{H_\nu}(x).$$

Thus

$$\Delta_{\mathcal{H}}(x) = 0 \quad \Longleftrightarrow \quad x \in \bigcup_{\nu=1}^{N} H_\nu.$$

Definition 9.2 (Binomial stratification certificate)**.** *Let*

$$\mathbf{D}(T, X) \in L[T][X_1^{\pm 1}, \dots, X_\rho^{\pm 1}]^m$$

be a dominant vector on C. A binomial stratification certificate consists of

$$B \in \mathbb{N}$$

and a finite family $\mathcal{H}$ of proper effective binomial strata in C such that

$$\mathbf{D}(t, x) = 0 \quad \Longrightarrow \quad x \in \bigcup_{H \in \mathcal{H}} H$$

for all real $t \geq B$ and all $x \in C$.

Once the candidate strata are supplied, the implication is a universal formula over a real closed field after writing the torus and the coset in real coordinates. Hence it is effectively checkable by quantifier elimination.

Lemma 9.3 (Effective semialgebraic Lojasiewicz inequality)**.** *Let K_0 be a compact effective semialgebraic set, and let*

$$f, g : K_0 \to \mathbb{R}_{\geq 0}$$

be effective continuous semialgebraic functions. If

$$Z(f) \subseteq Z(g),$$

then there exist effectively searchable constants

$$N \in \mathbb{N}, \qquad c > 0,$$

with c algebraic, such that

$$f(y) \geq c\, g(y)^N$$

for every $y \in K_0$.

Proof. The existence statement is the standard semialgebraic Lojasiewicz inequality on compact semialgebraic sets. For effectivity, fix N. The assertion

$$\exists c > 0 \ \forall y \in K_0 : \ f(y) \geq c\, g(y)^N$$

is a first-order sentence over real closed fields. Quantifier elimination decides whether such a c exists and, when it does, gives an effective semialgebraic description of the admissible c's. Enumerating $N = 0, 1, 2, \dots$ therefore finds a valid exponent and a positive algebraic c. □

Theorem 9.4 (Stratified semialgebraic lower bound)**.** *Assume that $\mathbf{D}$ admits a binomial stratification certificate $(B, \mathcal{H})$ on C. Put*

$$P(t, x) = \sum_{i=1}^{m} |D_i(t, x)|^2, \qquad \Delta(x) = \Delta_{\mathcal{H}}(x).$$

Then there exist effectively computable constants

$$B' \in \mathbb{N}, \qquad q > 0, \qquad L_0 \in \mathbb{N}, \qquad R \in \mathbb{N}$$

such that

$$P(t, x) \geq q(t+1)^{-L_0} \Delta(x)^R$$

for every $t \geq B'$ and every $x \in C$.

Proof. After clearing Laurent monomials, P is represented by an effective real semialgebraic function on $[B,\infty)\times C$. Choose d at least the degree in t of the cleared expression and compactify

$$u=\frac{1}{t+1},\qquad t=\frac{1-u}{u}.$$

Set

$$Q(u,x)=u^dP\left(\frac{1-u}{u},x\right)$$

on the compact semialgebraic set

$$K_0=\left[0,\frac{1}{B+1}\right]\times C.$$

For $u>0$, the certificate gives

$$Q(u,x)=0\quad\Longrightarrow\quad\Delta(x)=0.$$

At $u=0$, compactification may create additional zeros. Hence

$$Z(Q)\cap K_0\subseteq Z(u\Delta).$$

Apply the effective semialgebraic Lojasiewicz inequality to

$$f=Q,\qquad g=u\Delta.$$

There are N and $c>0$ such that

$$Q(u,x)\geq c\,u^N\Delta(x)^N.$$

Since $Q=u^dP$, this gives

$$P(t,x)\geq c\,u^{N-d}\Delta(x)^N.$$

If $N\geq d$, take $L_0=N-d$. If $N<d$, weaken the estimate on $0<u\leq 1$ to $P(t,x)\geq c\Delta(x)^N$. This gives the stated form with effective constants. □

10 SOS and Positivstellensatz certificates

Quantifier elimination proves decidability, but one may also use algebraic certificates that are easier to verify.

Let C be represented by equalities

$$f_1=\cdots=f_r=0$$

and inequalities

$$g_1\geq 0,\dots,g_s\geq 0.$$

These include the torus equations, the coset equations, and the tail condition $t-B_0\geq 0$. A Positivstellensatz/SOS certificate for the joint lower bound consists of an identity

$$(t+1)^L\sum_i|D_i(t,X)|^2-q=\sigma_0+\sum_j\sigma_jg_j+\sum_\ell h_\ell f_\ell,$$

where the σ_j are sums of squares and the h_ℓ are arbitrary polynomials, all with algebraic coefficients.

Such an identity is checked by exact polynomial arithmetic and by verifying the supplied sum-of-squares decompositions. When present, it gives a joint S-certificate.

This SOS formulation is not required for decidability, but it gives a more explicit certificate format than quantifier elimination.

11 Automatic joint certificates from leading vectors

Write

$$\mathbf{D}(T, X) = T^d \mathbf{D}_d(X) + T^{d-1}\mathbf{D}_{d-1}(X) + \cdots + \mathbf{D}_0(X),$$

with $\mathbf{D}_d \neq 0$.

Proposition 11.1 (Leading-vector criterion)**.** *If*

$$\mathbf{D}_d(x) \neq 0 \qquad \text{for every } x \in C,$$

then a joint S-certificate exists and can be effectively found.

Proof. The function

$$x \mapsto |\mathbf{D}_d(x)|^2$$

is continuous and positive on the compact coset C. Its positivity can be verified by quantifier elimination. Once verified, one can find a rational $q_0 > 0$ such that

$$|\mathbf{D}_d(x)|^2 \geq q_0$$

for all $x \in C$.

For $j < d$, compute an upper bound M_j for $|\mathbf{D}_j(x)|$ on C. Since monomials have modulus 1 on the torus, summing absolute values of algebraic coefficients gives an effective upper bound.

For t large enough,

$$\left| \sum_{j<d} t^j \mathbf{D}_j(x) \right| \leq \frac{1}{2} q_0^{1/2} t^d$$

uniformly on C. Hence

$$|\mathbf{D}(t, x)| \geq \frac{1}{2} q_0^{1/2} t^d.$$

This is a joint S-certificate. □

12 Scalar character-polynomial atoms

A scalar atom is an expression

$$F(T, \chi(X)),$$

where

$$F(T, Y) \in L[T, Y] \setminus \{0\}$$

and

$$\chi(X) = X_1^{a_1} \cdots X_\rho^{a_\rho}$$

is a toric character. Along the branch

$$z(k) = c\eta^k,$$

we have

$$\chi(z(k)) = \chi(c)\chi(\eta)^k.$$

Absorbing the constant $\chi(c)$ into F, one reduces to

$$F(k, \delta^k), \qquad \delta = \chi(\eta).$$

13 The Baker–Matveev lemma

We use the following standard consequence of effective lower bounds for linear forms in logarithms.

Lemma 13.1 (Scalar character lower bound)**.** *Let*

$$F(T,Y) \in L[T,Y] \setminus \{0\},$$

let

$$\delta \in L^{\times}, \qquad |\sigma(\delta)| = 1,$$

and suppose δ is not a root of unity. Then there exist computable constants

$$C > 0, \qquad N_0 \in \mathbb{N}$$

such that, for every $k \geq N_0$,

$$F(k,\delta^k) \neq 0 \quad \Longrightarrow \quad |\sigma(F(k,\delta^k))| \geq \exp(-C(\log(k+2))^2).$$

Moreover, exact zeros of $F(k,\delta^k)$ are contained in a computable finite prefix.

Proof. Write

$$F(T,Y) = \sum_{j=0}^{r} f_j(T) Y^j.$$

For each k, let

$$F_k(Y) = F(k,Y).$$

The set of k such that $F_k \equiv 0$ is finite and computable unless all f_j vanish identically, which is excluded. The case in which F_k is constant nonzero is harmless and gives a Liouville lower bound.

Assume F_k is nonconstant. Factor

$$F_k(Y) = a_k \prod_{\nu} (Y - \eta_{\nu,k})$$

over $\overline{\mathbb{Q}}$. The degree of F_k is at most r, though it may drop for special k. This only decreases the number of factors.

The coefficients $f_j(k)$ have logarithmic height $O(\log(k+2))$. Standard height estimates for roots of a polynomial give

$$h(\eta_{\nu,k}) \leq C_1 \log(k+2)$$

with C_1 effective. Liouville's inequality gives

$$|\sigma(a_k)| \geq \exp(-C_2 \log(k+2))$$

whenever $a_k \neq 0$.

Let $\eta = \eta_{\nu,k}$. If $\eta = 0$, then

$$|\sigma(\delta^k - \eta)| = 1.$$

Assume $\eta \neq 0$ and $\delta^k \neq \eta$. Apply an effective lower bound for linear forms in logarithms to

$$\eta^{-1}\delta^k - 1.$$

The degree of the field $\mathbb{Q}(\delta, \eta)$ is uniformly bounded in terms of L and r. The height of η is $O(\log k)$, and the exponent of δ is k. A standard Baker–Wuestholz or Matveev estimate yields

$$|\sigma(\eta^{-1}\delta^k) - 1| \geq \exp(-C_3(\log(k+2))^2).$$

Liouville gives

$$|\sigma(\eta)| \geq \exp(-C_4 \log(k+2)).$$

Therefore

$$|\sigma(\delta^k - \eta)| \geq \exp(-C_5(\log(k+2))^2).$$

If $F(k, \delta^k) \neq 0$, no factor vanishes. Multiplying the lower bounds for the leading coefficient and all nonzero factors gives

$$|\sigma(F(k, \delta^k))| \geq \exp(-C(\log(k+2))^2).$$

Now suppose

$$F(k, \delta^k) = 0.$$

Then δ^k is a root of F_k, so

$$h(\delta^k) \leq C_1 \log(k+2).$$

Since δ is not a root of unity, $h(\delta) > 0$, and

$$h(\delta^k) = kh(\delta).$$

Thus

$$kh(\delta) \leq C_1 \log(k+2),$$

which fails beyond a computable bound. □

14 Baker separation from binomial strata

The preceding scalar lemma controls one-character atoms. Binomial strata reduce to finitely many such one-character conditions along the toric orbit.

Lemma 14.1 (One-character separation)**.** *Let*

$$\alpha, \zeta \in \overline{\mathbb{Q}}^\times, \qquad |\sigma(\alpha)| = |\sigma(\zeta)| = 1.$$

Then the solution set

$$S(\alpha, \zeta) = \{k \geq 0 : \alpha^k = \zeta\}$$

is effectively computable as a finite union of arithmetic progressions and a finite set. Moreover, there exist effective constants

$$c > 0, \qquad A \in \mathbb{N},$$

such that

$$|\sigma(\alpha^k - \zeta)| \geq c(k+1)^{-A}$$

for every $k \notin S(\alpha, \zeta)$.

Proof. If α is a root of unity, the sequence α^k is periodic. Thus $S(\alpha, \zeta)$ is empty or a finite union of residue classes modulo the order of α, and the nonzero distances outside $S(\alpha, \zeta)$ have a positive periodic minimum.

Assume α is not a root of unity. Then $\alpha^k = \zeta$ has at most one solution, since two distinct solutions would make α torsion. The possible exact solution is decidable by the standard algorithms for multiplicative relations among algebraic numbers. For non-solutions, apply an effective lower bound for linear forms in logarithms to

$$\alpha^k \zeta^{-1} - 1.$$

Since α and ζ are fixed, Baker–Wuestholz or Matveev estimates give a lower bound of the form

$$|\sigma(\alpha^k \zeta^{-1} - 1)| \geq c_1 (k+1)^{-A}.$$

Multiplying by the fixed number $|\sigma(\zeta)| = 1$ gives the stated estimate. □

Lemma 14.2 (Separation from one binomial stratum)**.** *Let $z(k) = c\eta^k \in C$, and let*

$$H = \{x \in C : \chi_{a_1}(x) = \zeta_1, \ldots, \chi_{a_\ell}(x) = \zeta_\ell\}$$

be an effective binomial stratum. Set

$$S_H = \{k \geq 0 : z(k) \in H\}.$$

Then S_H is effectively computable as a finite union of arithmetic progressions and a finite set. Moreover, there exist effective constants

$$c_H > 0, \qquad A_H \in \mathbb{N},$$

such that

$$\Delta_H(z(k)) \geq c_H (k+1)^{-A_H}$$

for every $k \notin S_H$.

Proof. The condition $z(k) \in H$ is the finite system

$$\chi_{a_j}(\eta)^k = \zeta_j \, \chi_{a_j}(c)^{-1}, \qquad j = 1, \ldots, \ell.$$

Each equation has an effectively computable solution set by the one-character lemma, and finite intersections of finite unions of progressions and finite sets are effectively computable.

If $k \notin S_H$, then at least one equation fails. For each j, the one-character lemma gives a polynomial lower bound away from the exact solution set of the j-th equation. Taking the minimum over the finitely many failed possibilities gives

$$\Delta_H(z(k)) = \sum_j |\chi_{a_j}(z(k)) - \zeta_j|^2 \geq c_H (k+1)^{-A_H}.$$

□

Lemma 14.3 (Separation from a finite family of binomial strata)**.** *Let $\mathcal{H} = \{H_1, \ldots, H_N\}$ be a finite family of effective binomial strata and set*

$$S_{\mathcal{H}} = \bigcup_{\nu=1}^{N} S_{H_\nu}.$$

Then $S_{\mathcal{H}}$ is effectively computable as a finite union of arithmetic progressions and a finite set. Moreover, there exist effective constants

$$c > 0, \qquad A \in \mathbb{N},$$

such that

$$\Delta_{\mathcal{H}}(z(k)) \geq c(k+1)^{-A}$$

for every $k \notin S_{\mathcal{H}}$.

Proof. The exact-hit set is a finite union of effectively computable sets. If $k \notin S_{\mathcal{H}}$, then $z(k)$ lies in none of the strata. Hence every $\Delta_{H_\nu}(z(k))$ is bounded below by a negative power of $k+1$. Multiplying the finitely many inequalities gives the stated lower bound for

$$\Delta_{\mathcal{H}}(z(k)) = \prod_{\nu=1}^{N} \Delta_{H_\nu}(z(k)).$$

□

Theorem 14.4 (Stratified zero-free tail outside exact hits)**.** *Let a saturated branch have normal form*

$$\mathbf{v}_k = \Lambda^k(\mathbf{D}(k, z(k)) + \mathbf{E}(k)), \qquad z(k) = c\eta^k,$$

with

$$|\mathbf{E}(k)| \leq C_E(k+1)^s\theta^k, \qquad 0 < \theta < 1.$$

Assume that $\mathbf{D}$ admits a binomial stratification certificate $\mathcal{H}$ on C, and let

$$S_{\mathcal{H}} = \{k \geq 0 : z(k) \in \bigcup_{H \in \mathcal{H}} H\}.$$

Then there is an effectively computable B such that

$$\mathbf{v}_k \neq 0$$

for every $k \geq B$ with $k \notin S_{\mathcal{H}}$.

Proof. The stratified semialgebraic lower bound gives

$$\sum_i |D_i(k, z(k))|^2 \geq q(k+1)^{-L_0}\Delta_{\mathcal{H}}(z(k))^R.$$

For $k \notin S_{\mathcal{H}}$, Baker separation from the finite family of strata gives

$$\Delta_{\mathcal{H}}(z(k)) \geq c(k+1)^{-A}.$$

Therefore

$$|\mathbf{D}(k, z(k))| \geq c'(k+1)^{-A'}$$

for effective constants $c' > 0$ and A'. The exponentially small error is eventually less than half of this lower bound, and the conclusion follows from the same comparison used in the certified-tail theorem. □

15 Torsion scalar atoms

If δ is a root of unity of order Q, pass to subbranches

$$k = Qh + s.$$

Then

$$F(k, \delta^k) = F(Qh + s, \delta^s) = G_s(h),$$

where $G_s \in L[h]$.

If $G_s \not\equiv 0$, its integer zeros are finite and computable, and Liouville gives a polynomial lower bound for nonzero values.

If $G_s \equiv 0$, the scalar atom vanishes identically on that subbranch. In a scalar product certificate this may force the dominant scalar to vanish. In a vector or Bezout certificate it means only that the particular scalar test has failed. The algorithm either peels the branch if the full dominant vector is identically zero, or requires a new certificate on that subbranch.

Termination follows because every peeling removes one global shell and all splitting operations are finite.

16 Bezout–toric certificates

Let

$$\mathbf{D} = (D_1, \dots, D_m).$$

Let

$$a^{(1)}, \dots, a^{(\ell)}$$

be a basis of $\mathcal{L}_\eta$, and define the coset binomials

$$G_j(X) = X^{a^{(j)}} - c^{a^{(j)}}.$$

Then

$$G_j(z(k)) = 0.$$

Definition 16.1 (Bezout–toric certificate)**.** *A Bezout–toric certificate consists of Laurent polynomials*

$$A_i(T, X), \quad B_j(T, X), \quad M(T, X)$$

such that

$$M(T, X) = \sum_{i=1}^{m} A_i(T, X) D_i(T, X) + \sum_{j=1}^{\ell} B_j(T, X) G_j(X),$$

together with a scalar certificate for $M(k, z(k))$ giving

$$|M(k, z(k))| \geq \exp(-C(\log(k+2))^A)$$

outside a computable prefix.

The identity is verified exactly in the Laurent polynomial ring.

Lemma 16.2. *There exist computable constants C_A, R such that*

$$\left(\sum_i |A_i(k, z(k))|^2 \right)^{1/2} \leq C_A (k+1)^R.$$

Proof. Write

$$A_i(T,X) = \sum_{\nu} a_{i,\nu}(T) X^{b_{i,\nu}}.$$

On the torus,

$$|X^{b_{i,\nu}}| = 1.$$

Each polynomial $a_{i,\nu}(k)$ grows at most polynomially in k. Summing finitely many terms gives the stated bound. □

Proposition 16.3. *A Bezout–toric certificate implies the lower-bound hypothesis of the certified-tail theorem.*

Proof. Along the orbit,

$$M(k,z(k)) = \sum_i A_i(k,z(k)) D_i(k,z(k)).$$

By Cauchy–Schwarz,

$$|M(k,z(k))| \le \left(\sum_i |A_i(k,z(k))|^2 \right)^{1/2} |\mathbf{D}(k,z(k))|.$$

Using the preceding lemma,

$$|M(k,z(k))| \le C_A (k+1)^R |\mathbf{D}(k,z(k))|.$$

A subexponential lower bound for M therefore gives a subexponential lower bound for $|\mathbf{D}|$. □

One may search for Bezout–toric certificates by fixing degree bounds for the unknown A_i, B_j, M, imposing the identity as a finite linear system over L, and checking whether M admits a scalar S/B/T certificate. If the search succeeds, the certificate is valid. If it fails at a chosen bound, no conclusion follows. Thus this is a positive semidecision procedure, not a complete algorithm.

17 Quotient-character reduction

Let

$$H(T,X) = \sum_{a \in A} h_a(T) X^a$$

be a Laurent factor. Consider the image of A in

$$\mathbb{Z}^\rho / \mathcal{L}_\eta.$$

We say that H has quotient affine rank at most one if the differences

$$a - a_0, \qquad a \in A,$$

span a vector space of dimension at most one in

$$(\mathbb{Z}^\rho / \mathcal{L}_\eta) \otimes_{\mathbb{Z}} \mathbb{Q}.$$

Proposition 17.1. *If H has quotient affine rank at most one, then after finite torsion splitting,*

$$H(k,z(k)) = \alpha_0^k F(k, \delta^k)$$

for some algebraic α_0, δ of σ-absolute value 1 and some

$$F(T,Y) \in L[T,Y].$$

Proof. Choose $a_0 \in A$. Since the quotient affine rank is at most one, there exists $v \in \mathbb{Z}^\rho$ such that every $a - a_0$ maps to a rational multiple of v in the quotient. Clearing denominators and splitting finite torsion in the quotient, we may assume that for each $a \in A$,

$$a = a_0 + r_a v + \ell_a$$

with $r_a \in \mathbb{Z}$ and $\ell_a \in \mathcal{L}_\eta$.

On the coset $z(k) = c\eta^k$,

$$z(k)^a = c^a \eta^{ak} = c^a \eta^{a_0 k} (\eta^v)^{r_a k},$$

because $\eta^{\ell_a} = 1$. Thus

$$H(k, z(k)) = \eta^{a_0 k} \sum_{a \in A} h_a(k) c^a (\eta^v)^{r_a k}.$$

After multiplying by a power of Y if necessary, the sum becomes $F(k, \delta^k)$ with

$$\delta = \eta^v.$$

The factor $\eta^{a_0 k}$ is α_0^k and has modulus 1. □

Corollary 17.2. *A quotient-affine-rank-one factor is automatically converted into a scalar Baker or torsion atom.*

18 Automatic binomial extraction

The binomial-stratified theorem is a verifier: it proves a tail once the strata are known. In an important subclass the strata can be extracted from the dominant equations.

18.1 Dominant elimination ideals

Let $C \subseteq \mathbb{T}^\rho$ be the compact toric coset attached to a branch, and let

$$I_C \subseteq L[X_1^{\pm 1}, \dots, X_\rho^{\pm 1}]$$

be its Laurent ideal, generated by the binomials

$$X^a - c^a, \qquad a \in \mathcal{L}_\eta.$$

For

$$\mathbf{D}(T, X) = (D_1, \dots, D_m)$$

define

$$I_D = I_C + \langle D_1, \dots, D_m \rangle \subseteq L[T, X_1^{\pm 1}, \dots, X_\rho^{\pm 1}]$$

and

$$J_D = \operatorname{rad}\left(I_D \cap L[X_1^{\pm 1}, \dots, X_\rho^{\pm 1}] \right).$$

The variety $V(J_D) \cap C$ contains every point $x \in C$ for which there exists a complex value $T = t$ with

$$\mathbf{D}(t, x) = 0.$$

It is therefore an algebraic over-approximation of the eventual real projection

$$\{x \in C : \exists t \gg 0,\ \mathbf{D}(t, x) = 0\}.$$

The over-approximation is useful: if it is already covered by proper binomial strata, then it gives a valid certificate without guessing the strata.

Definition 18.1 (Automatically binomial branch)**.** *A dominant branch is automatically binomial if the computed set*

$$V(J_D) \cap C$$

is effectively verified to be contained in a finite union of proper binomial strata

$$H_1 \cup \cdots \cup H_N$$

inside C. In the strict binomial case, these strata are obtained from binomial radical and decomposition algorithms [11], followed by lattice basis reduction and Smith normal form.

Theorem 18.2 (Automatic binomial certificate extraction)**.** *Let a nonzero dominant branch be automatically binomial. Then one can effectively compute a finite family of proper binomial strata*

$$\mathcal{H} = \{H_1, \ldots, H_N\}$$

such that

$$\mathbf{D}(t, x) = 0 \quad \Longrightarrow \quad x \in \bigcup_{\nu=1}^{N} H_\nu$$

for every t and every $x \in C$. Hence the branch admits a binomial stratification certificate with $B = 0$.

Proof. Compute I_C, form I_D, eliminate T, and take the radical J_D. By the automatic binomial hypothesis, $V(J_D) \cap C$ is covered by effectively listed proper binomial strata. If $\mathbf{D}(t, x) = 0$ for some $x \in C$, then $(t, x) \in V(I_D)$. Every Laurent polynomial in $I_D \cap L[X_1^{\pm 1}, \ldots, X_\rho^{\pm 1}]$ therefore vanishes at x, and so $x \in V(J_D) \cap C$. The verified binomial cover gives the desired implication. □

19 Binomial-stratified certificate trees

Exact hits of binomial strata are not exceptional failures. They are arithmetic progressions on which the recurrence can be restricted and renormalized.

Definition 19.1 (Binomial-stratified toric certificate tree)**.** *A binomial-stratified toric certificate tree for a vector recurrence is a finite rooted tree whose nodes carry restricted branches*

$$\mathbf{v}_k = \mathbf{u}_{ak+b}$$

with saturated dominant normal forms

$$\mathbf{v}_k = \Lambda^k(\mathbf{D}(k, z(k)) + \mathbf{E}(k)).$$

A terminal node is one of the following:

1. *an identically zero branch;*
2. *a branch with a joint S-certificate, an SOS/Positivstellensatz certificate, a Bezout–toric certificate, or scalar atoms proving a dominant lower bound;*
3. *a closure-avoiding branch for which the leading-vector or semialgebraic compactness test gives a lower bound.*

A nonterminal node carries a binomial stratification certificate $\mathcal{H}$. Its children are the infinite arithmetic progressions contained in

$$S_{\mathcal{H}} = \{k \geq 0 : z(k) \in \bigcup_{H \in \mathcal{H}} H\}.$$

On each child, the recurrence is restricted to that progression and the dominant normal form is recomputed. Along every edge, either a dominant spectral shell is peeled or the dimension of the compact toric orbit closure strictly decreases.

Definition 19.2 (Automatic binomial-stratified tree)**.** *A binomial-stratified toric certificate tree is automatic if every nonterminal node obtains its binomial stratification certificate from automatic binomial extraction.*

Theorem 19.3 (Exact zero set from a certificate tree)**.** *If a vector recurrence admits a finite binomial-stratified toric certificate tree, then its zero set is computable exactly as a finite union of arithmetic progressions and a finite set. If the tree is automatic, the binomial strata at all nonterminal nodes are computed from the dominant branch equations.*

Proof. Proceed by induction on the finite tree. An identically zero terminal node contributes its whole arithmetic progression. Any other terminal node has a verified dominant lower bound, hence a computable zero-free tail by the certified-tail theorem; its finite prefix is checked exactly.

At a nonterminal node, the stratified zero-free tail theorem gives a computable bound beyond which no zeros occur outside $S_{\mathcal{H}}$. The finitely many remaining indices outside $S_{\mathcal{H}}$, together with the finite exceptional part of $S_{\mathcal{H}}$, are checked exactly. The infinite progressions in $S_{\mathcal{H}}$ are precisely the children, whose zero sets are computed by induction. Taking the union over all nodes gives the exact zero set. □

20 The Factor–Reduce–Certify compiler

The partial compiler operates branchwise.

1. Compute the vector toric normal form.
2. Compute the toric coset $C = cT_\eta$.
3. Attempt the leading-vector criterion.
4. If it fails, attempt to construct a Bezout–toric certificate by bounded-degree linear algebra.
5. For scalar factors or scalar tests, factor in the Laurent polynomial ring.
6. For each irreducible factor, compute its quotient affine rank modulo $\mathcal{L}_\eta$.
7. If the rank is at most one, reduce it to a scalar atom.
8. Verify remaining S-certificates by quantifier elimination, SOS identities, or Positivstellensatz identities.
9. If the branch is not terminally certified, attempt automatic binomial extraction from J_D.
10. If binomial extraction succeeds, compute the exact-hit progressions, restrict to the infinite progressions, and recompute the dominant normal form on each child.

11. Accept the recursive step only when it peels a dominant shell or strictly lowers the toric dimension.

12. If every terminal branch is certified and every recursive branch belongs to a finite certificate tree, compute the zero-free tails and exact prefixes.

13. If some branch is not certified, return "not certified".

Theorem 20.1 (Compiler soundness)**.** *If Factor–Reduce–Certify succeeds on all branches of a vector recurrence* $\mathbf{u}$*, then it computes*

$$Z(\mathbf{u}) = \{n \geq 0 : \mathbf{u}_n = 0\}$$

exactly as a finite union of arithmetic progressions and a finite set.

Proof. Every terminal branch has one of the terminal certificates listed above, so it satisfies the certified-tail theorem. Identically zero branches give progressions. Nonzero terminal branches have zero-free tails and exact finite prefix checks.

Every accepted nonterminal step is justified by a binomial stratification certificate. The stratified zero-free tail theorem removes all sufficiently large indices outside the exact-hit progressions. The children are precisely the restrictions to those infinite progressions. Since a successful run produces a finite certificate tree, the exact zero set is computed by the theorem on certificate trees. □

21 Certified and automatic binomial-stratified classes

Definition 21.1. *A vector recurrence belongs to the certified binomial-stratified toric class if its dominant branch decomposition admits a finite binomial-stratified toric certificate tree.*

It belongs to the automatic binomial-stratified toric class if it admits such a tree and every nonterminal stratum is produced by automatic binomial extraction.

Corollary 21.2 (Certified binomial-stratified Skolem)**.** *For every algebraic vector recurrence* $\mathbf{u}$ *in the certified binomial-stratified toric class, the set*

$$Z(\mathbf{u}) = \{n \geq 0 : \mathbf{u}_n = 0\}$$

is computable exactly.

Proof. Apply the certificate-tree theorem to every residue branch in the dominant toric normal form and take the finite union of the branch zero sets. □

Corollary 21.3 (Automatic binomial-stratified Skolem)**.** *For every algebraic vector recurrence* $\mathbf{u}$ *in the automatic binomial-stratified toric class, the zero set* $Z(\mathbf{u})$ *is computable exactly. In this class the binomial strata required at nonterminal nodes are computed from the dominant elimination ideals.*

Remark 21.4. *The restriction is geometric rather than dimensional. There is no a priori bound here on recurrence order, root multiplicity, number of dominant roots, toric rank, vector dimension, ambient matrix dimension, or target dimension. The required structure is that the dominant toric zero loci are either terminally separated from zero or recursively controlled by binomial strata.*

22 Applications

Theorem 22.1 (Simultaneous Skolem)**.** *Let*

$$u^{(1)}, \dots, u^{(m)}$$

be algebraic linear recurrence sequences and set

$$\mathbf{u}_n = (u_n^{(1)}, \dots, u_n^{(m)}).$$

If Factor–Reduce–Certify succeeds on $\mathbf{u}$*, then*

$$\{n \geq 0 : u_n^{(1)} = \cdots = u_n^{(m)} = 0\}$$

is computable exactly.

This is stronger than separately solving the scalar Skolem instances and intersecting their zero sets, because a joint vector certificate may exist even when a scalar coordinate is outside the automatic scalar certificate class.

Theorem 22.2 (Subspace Orbit and affine linear loops)**.** *Let*

$$A \in K^{d \times d}, \qquad x_0 \in K^d,$$

and let

$$V = \{y \in K^d : By = b\}$$

be an affine subspace. Define

$$\mathbf{u}_n = BA^n x_0 - b.$$

If Factor–Reduce–Certify succeeds on $\mathbf{u}_n = BA^n x_0 - b$*, then the hitting-time set*

$$\{n \geq 0 : A^n x_0 \in V\}$$

is computable exactly.

Proof. We have

$$A^n x_0 \in V \quad \Longleftrightarrow \quad BA^n x_0 - b = 0.$$

Apply compiler soundness. □

Corollary 22.3 (Binomial-stratified Subspace Orbit)**.** *With notation as above, suppose*

$$\mathbf{u}_n = BA^n x_0 - b$$

belongs to the certified binomial-stratified toric class. Then

$$\mathrm{Hit}(A, x_0, V) = \{n \geq 0 : A^n x_0 \in V\}$$

is computable exactly. If $\mathbf{u}_n$ *belongs to the automatic binomial-stratified toric class, the required binomial strata are computed from the dominant branch equations.*

Proof. The sequence $\mathbf{u}_n$ is a vector recurrence, and

$$A^n x_0 \in V \quad \Longleftrightarrow \quad \mathbf{u}_n = 0.$$

Apply the certified or automatic binomial-stratified Skolem corollary. □

This gives a certified reachability theorem for affine linear loops

$$x := Ax$$

with target condition $Bx = b$.

Corollary 22.4 (Algebraic target hitting)**.** *Let*

$$W = \{y : P_1(y) = \cdots = P_m(y) = 0\}$$

be an affine algebraic set over K. Since coordinates of $A^n x_0$ are linear recurrence sequences and such sequences are closed under addition and pointwise multiplication, each sequence

$$P_i(A^n x_0)$$

is a linear recurrence sequence. If Factor–Reduce–Certify succeeds on

$$(P_1(A^n x_0), \ldots, P_m(A^n x_0)),$$

then

$$\{n \geq 0 : A^n x_0 \in W\}$$

is computable exactly.

23 Eventual cone membership and ultimate positivity

Let $C \subseteq \mathbb{R}^m$ be a semialgebraic cone or a closed semialgebraic set. Suppose a real vector recurrence has normal form

$$\mathbf{u}_k = \Lambda^k(\mathbf{D}(k, z(k)) + \mathbf{E}(k))$$

with $\Lambda > 0$, and suppose a certificate proves that $\mathbf{D}(k, z(k))$ has distance at least

$$c(k+1)^{-L}$$

from the boundary of C for $k \geq B$, with a fixed side of the boundary. Then $\mathbf{u}_k \in C$ eventually if and only if the dominant vector lies on that side. The finite prefix is checked exactly.

For $m = 1$ and $C = [0, \infty)$, this gives certified ultimate positivity.

24 A high-rank vector family

Let $\rho \geq 2$, and let

$$\gamma_1, \ldots, \gamma_\rho \in L^\times$$

satisfy

$$|\sigma(\gamma_i)| = 1$$

and be multiplicatively independent modulo torsion.

Define

$$\mathbf{D}_\rho(X_1, \ldots, X_\rho) = (X_2 - X_1,\, X_3 - X_1,\, \ldots,\, X_\rho - X_1,\, X_1 + \cdots + X_\rho - 1)\,.$$

Proposition 24.1. *The vector $\mathbf{D}_\rho$ has no zero on $\mathbb{T}^\rho$.*

Proof. If $\mathbf{D}_\rho(X) = 0$, then

$$X_2 = X_1, \quad \dots, \quad X_\rho = X_1.$$

Thus

$$X_1 = \cdots = X_\rho.$$

The last component gives

$$\rho X_1 = 1.$$

But $|X_1| = 1$, whereas $|1/\rho| < 1$ for $\rho \geq 2$. This is impossible. □

Therefore

$$Q_\rho(X) = |\mathbf{D}_\rho(X)|^2$$

has a strictly positive minimum on $\mathbb{T}^\rho$, and a joint S-certificate can be verified.

Let

$$\mathbf{u}_n = \Lambda^n \left(\mathbf{D}_\rho(\gamma_1^n, \dots, \gamma_\rho^n) + \mathbf{E}(n) \right),$$

where

$$|\mathbf{E}(n)| \leq C\theta^n, \qquad 0 < \theta < 1.$$

Then Factor–Reduce–Certify succeeds and computes the simultaneous zero set.

The scalar coordinate

$$X_1 + \cdots + X_\rho - 1$$

has exponent support of affine rank ρ, so it is not an automatic quotient-character scalar factor. The vector system is nevertheless certified by joint separation. This demonstrates the genuine advantage of the vector method.

25 Realization as Subspace Orbit

Let

$$A = \operatorname{diag}(\gamma_1, \dots, \gamma_\rho), \qquad x_0 = (1, \dots, 1).$$

Then

$$A^n x_0 = (\gamma_1^n, \dots, \gamma_\rho^n).$$

Let $V_\rho \subseteq K^\rho$ be the affine subspace defined by

$$x_2 - x_1 = 0, \quad x_3 - x_1 = 0, \quad \dots, \quad x_\rho - x_1 = 0,$$

and

$$x_1 + \cdots + x_\rho = 1.$$

Then

$$A^n x_0 \in V_\rho \quad \Longleftrightarrow \quad \mathbf{D}_\rho(\gamma_1^n, \dots, \gamma_\rho^n) = 0.$$

By the preceding proposition, the pure dominant model has no hitting time. With exponentially smaller subdominant perturbations, the subspace-orbit certificate theorem gives a computable zero-free tail and exact prefix check.

This gives a high-rank family of certified Subspace Orbit instances in which one scalar equation contains high-rank toric support, but the vector system is jointly separated.

26 Relation to previous work

The Skolem Problem is open in general. The Skolem–Mahler–Lech theorem gives the eventual structure of the zero set but is not a general algorithm for computing the finite exceptional set.

Low-order and few-dominant-root results decide important special cases. The present framework is different: it is not based on low order or few dominant roots, but on a verifiable toric separation certificate.

Modern algorithms for broader recurrence classes often produce correct certificates when they terminate, with termination proved under strong Diophantine conjectures. The present method is unconditional inside the certified class.

The Evertse–Schlickewei–Schmidt theory gives explicit upper bounds for the number of non-degenerate solutions of linear equations in variables from a finite-rank multiplicative group [10]. Such results are close in spirit, but a bound on the number of solutions does not by itself give a computable bound on the positions of the indices n. The present method produces an explicit zero-free tail once a certificate is verified.

Bacik and Varonka study Subspace Orbit and Simultaneous Skolem from a dimension-theoretic perspective [5]. Their results are orthogonal to the present work. They prove decidability in certain target-dimension regimes and hardness in others. The present framework imposes no a priori target-dimension threshold; it requires a toric separation certificate for the dominant vector.

The vector feature is essential. The family above shows that scalar coordinatewise certification can fail while joint vector certification succeeds.

27 Boundary of the certified class

The certified class is governed by the available lower-bound mechanisms. Joint semialgebraic certificates handle dominant vectors that avoid zero on the compact toric closure. Binomial-stratified certificates handle dominant zero loci contained in finitely many character strata, because exact character hits are computable and Baker separation gives polynomial distance from the strata away from those hits.

A scalar expression such as

$$X + Y - 1$$

on a genuinely rank-two orbit is different. Its zero locus is not a finite union of binomial strata. Along

$$(X, Y) = (\alpha^n, \beta^n)$$

one would need effective lower bounds for

$$|\alpha^n + \beta^n - 1|$$

outside exact zeros. Such non-binomial toric approximation is not supplied by the binomial-stratified layer.

Thus the compiler is sound and incomplete: success produces a verified exact zero set, while failure means that the available certificate mechanisms did not certify the instance.

28 Conclusion

We have developed a certificate framework for effective zero-free tails in scalar and vector linear recurrence sequences. The main novelty is the vector toric viewpoint: for Simultaneous Skolem and

Subspace Orbit, it is enough to separate the dominant vector from zero. This can succeed even when scalar coordinatewise methods fail.

The framework includes:

1. global vector dominant-shell normal form;
2. torsion saturation and toric coset closures;
3. certified-tail soundness;
4. joint semialgebraic and SOS certificates;
5. Bezout–toric certificates;
6. scalar Baker character-polynomial atoms;
7. quotient-character reduction;
8. binomial-stratified lower bounds;
9. Baker separation from exact binomial hits;
10. automatic binomial extraction from dominant elimination ideals;
11. finite certificate trees;
12. a partial Factor–Reduce–Certify compiler with a recursive binomial-stratified mode;
13. applications to Simultaneous Skolem, Subspace Orbit, linear-loop reachability, algebraic target hitting, and ultimate positivity;
14. a high-rank vector family showing genuine advantage over scalar methods.

The resulting framework gives a rigorous, checkable, and partially automatic mechanism for proving effective Skolem tails in high-rank, high-order families. Its strongest automatic layer is the binomial-stratified class: semialgebraic Lojasiewicz inequalities control vanishing near the strata, Baker separation controls the actual toric orbit away from exact hits, and recursive restriction handles the exact-hit progressions.

References

[1] A. Baker. *Transcendental Number Theory*. Cambridge University Press, 1975.

[2] A. Baker and G. Wuestholz. *Logarithmic Forms and Diophantine Geometry*. Cambridge University Press, 2007.

[3] P. Bacik. Completing the picture for the Skolem Problem on order-4 linear recurrence sequences. arXiv:2409.01221.

[4] P. Bacik, J. Ouaknine, D. Purser, and J. Worrell. On the p-adic Skolem Problem. arXiv:2504.14413.

[5] P. Bacik and A. Varonka. On the Subspace Orbit Problem and the Simultaneous Skolem Problem. arXiv:2601.18349.

[6] S. Basu. Algorithms in real algebraic geometry: a survey. arXiv:1409.1534.

[7] Y. Bilu. Skolem Problem for linear recurrence sequences with 4 dominant roots, after Mignotte, Shorey, Tijdeman, Vereshchagin and Bacik. arXiv:2501.16290.

[8] Y. Bilu, F. Luca, J. Nieuwveld, J. Ouaknine, D. Purser, and J. Worrell. Skolem meets Schanuel. arXiv:2204.13417.

[9] V. Chonev, J. Ouaknine, and J. Worrell. The Orbit Problem in higher dimensions. In *Proceedings of LICS 2013*.

[10] J.-H. Evertse, H. P. Schlickewei, and W. M. Schmidt. Linear equations in variables which lie in a multiplicative group. *Annals of Mathematics* 155 (2002), 807–836.

[11] D. Eisenbud and B. Sturmfels. Binomial ideals. *Duke Mathematical Journal* 84 (1996), 1–45.

[12] G. Ge. *Algorithms Related to Multiplicative Representations of Algebraic Numbers*. PhD thesis, University of California, Berkeley, 1993.

[13] G. Jeronimo, D. Perrucci, and E. Tsigaridas. On the minimum of a polynomial function on a basic closed semialgebraic set and applications. arXiv:1112.0544.

[14] R. Kannan and R. Lipton. The Orbit Problem is decidable. In *Proceedings of STOC 1980*.

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

[16] H. P. Schlickewei and W. M. Schmidt. The number of solutions of polynomial-exponential equations. *Compositio Mathematica* 120 (2000), 193–225.

[17] T. N. Shorey and R. Tijdeman. *Exponential Diophantine Equations*. Cambridge University Press, 1986.

[18] N. K. Vereshchagin. The problem of the occurrence of zero in a linear recursive sequence. *Mathematical Notes*.