Finding Two Connected Factors of a Finite Dynamical System in Quasipolynomial Time

Nuutti Kuosa

22 November 2026

hexagon:2611.00849v1

Abstract

Doré, Perrot, Porreca, Riva and Rolland asked whether their tree methods could be used to find two nontrivial connected factors of a given finite dynamical system. We give a deterministic algorithm for this problem. Given the successor table of an N-state system B, the algorithm constructs connected systems X and Y, each with at least two states, such that $X \times Y \cong B$, or determines that no such pair exists. Its bit complexity is $\exp(O((\log(2N))^3 \log\log(4N)))$, even when B is disconnected. All ordered pairs of factor isomorphism types can be listed within the same total bound. The proof uses two predecessor paths to control branching across tree levels and bounds the arithmetic search at each level. Whether one factorization can be found in polynomial time remains open.

1 Problem and result

Doré et al. give a polynomial-time algorithm for division by a specified finite dynamical system when the quotient is connected [1, Algorithm 3.12 and Theorem 3.13]. They close their paper by asking whether their tree methods can find a nontrivial solution of $X \times Y \cong B$ when both X and Y are unknown and connected [1, Sec. 6, p. 14]. Trying all connected systems on at most $\sqrt{N}$ states as the smaller factor already gives a $2^{O(\sqrt{N})}\,\mathrm{poly}(N)$-time search on an N-state input, as shown below. To obtain a quasipolynomial bound, we search for both factors along two predecessor paths in B. Mahler measure and residues of differential forms control the number of arithmetic trials required at each extension.

Write $B = (V, f)$, where $f : V \to V$ is given explicitly and $N = |V| \geq 1$. The synchronous product advances both coordinates by one step. Connectedness refers to the underlying undirected graph.

Theorem 1. *Given B, one can deterministically construct successor tables of connected X, Y with $|X|, |Y| \geq 2$ and $X \times Y \cong B$, or decide that no such pair exists, in*

$$\exp\bigl(O((\log(2N))^3 \log\log(4N))\bigr)$$

bit operations. All ordered pairs of isomorphism types can be listed within the same total time.

Corollary (Reducibility of a connected target)**.** *For every explicit connected B, the same bound decides whether B is a product of two systems of at least two states, and constructs a pair whenever one exists. For B a dendron, both factors are dendrons.*

Proof. The coordinate projections of a nonempty product are surjective graph homomorphisms. Each sends a weakly connected graph onto a weakly connected graph, so connectedness of the target forces connectedness of both factors [3, pp. 39, 185]. Theorem 1 now applies to every nontrivial factorisation. If every product state eventually reaches the same fixed pair, projecting the orbit of any pair shows that each coordinate reaches the corresponding fixed point. □

Comparison with known division. To derive the elementary search bound, take a smaller factor on $m \le \sqrt{N}$ states. Delete its cycle edges, order the resulting rooted trees in cycle order, and choose an order on the children of every vertex. This plane forest has a balanced-parenthesis encoding of length $2m$, with one opening and one closing parenthesis per vertex. Conversely, the roots of each nonempty decoded forest can be linked in cyclic order. Try all words for $2 \le m \le \sqrt{N}$ with $m \mid N$, discard invalid words, and use the connected-quotient algorithm on each decoded factor. There are at most 4^m words at size m, and $N/m \ge 2$ for these trials. Theorem 1 improves this elementary consequence of known division to a quasipolynomial bound and also bounds the total time to list all pairs.

Related work. Dennunzio et al. formulated decomposition questions in the semiring of finite dynamical systems [4]. Unrolling was developed in work on transient trees and division by a specified system [5, 2, 1]. Naquin and Gadouleau give a polynomial division algorithm for dendrons and prove unique factorisation in certain submonoids; their Figure 7 exhibits nonunique factorisation for dendrons in general [2, Theorems 39 and 56]. For permutations, Bridoux et al. study division and give an $\exp(O(\sqrt{N}))$ test for unrestricted irreducibility of sums of cycles [6, author version, Sec. 6, p. 19 and Appendix C, Prop. 3, p. 22]. With connected factors a permutation target forces both factors to be single cycles, so their lengths can be tried directly. Crespelle, Phan and Richard study division and polynomial equations on the cycle part using the greatest-common-divisor and least-common-multiple rule for cycle products [7]. Porreca and Rolland obtain polynomial algorithms for injective and pseudo-injective univariate equations and identify the complexity of general reducibility as open [8, Sec. 7]. Ascone, Porreca and Rolland prove NP-hardness for linear equations with a constant side and three variables [9]; their reduction concerns a different family of equations from the single product studied here.

For undirected graphs, Imrich gives a polynomial factorisation algorithm for connected non-bipartite graphs [10], and Calderoni et al. reduce graph isomorphism to direct-product compositeness for general undirected graphs [11]. The directed polynomial algorithm of Imrich and Klöckl uses stronger connectivity conditions and thinness conditions on successive quotients [12, Theorem 3.6 and Corollary 3.7, p. 473]. For instance, distinct vertices of a directed cycle have neither a common successor nor a common predecessor, although the cycle is connected as a dynamical system. Among the works compared here we found no earlier quasipolynomial bound for two unknown connected factors on arbitrary explicit targets.

The input occupies $\Theta(N \log(2N))$ bits in the usual fixed-width encoding. For $N < 4$ the answer is negative. All logarithms below are natural. The search has three parts. A polynomial list describes how the periodic states could fit into a product grid. Two paths in the observed system then reduce every extension to factoring a multiset of whole sibling trees. Section 3 counts the resulting search tree, including every failed trial. Sections 4–6 prove the local arithmetic estimate used in that count.

2 Periodic states and predecessor paths

Let h be the largest distance of a state of B to its cycle. Write $T_j(B, v)$ for the rooted tree of predecessor walks of length at most j ending at v. Its width at depth j is $w_j(B, v) = |(f^j)^{-1}(v)|$. The root of the tree is the walk of length zero. Each state determines at most one walk of a given length, so the tree has at most $(j + 1)N$ vertices. Trees multiply by pairing vertices at equal depth. In particular $T_j(X \times Y, (u, v)) = T_j(X, u)T_j(Y, v)$, up to rooted isomorphism. A forest here is a multiset of rooted trees. Products distribute over multiplicities. Nonempty forests all of whose trees have height j admit cancellation and polynomial-time division by a specified forest of that height. Adjoin a single root over all the trees in each forest, apply finite tree division, remove the added root, and verify the product [2, Sec. 4]. We use this operation only with the specified height and verify each division by multiplication.

A polynomial list of possible phase grids

Two connected factors of cycle lengths p, q give $g = \gcd(p, q)$ components with the same cycle length $\lambda = \operatorname{lcm}(p, q)$. Count g from the target. Reject unequal observed cycle lengths. The cyclic states form a $p \times q$ grid whose transition increases both indices by one. Enumerate the ordered divisors $p, q \mid \lambda$ with $\operatorname{lcm}(p, q) = \lambda$ and $pq = g\lambda$. When $g = 1$, the grid is determined by forward iteration from any chosen cyclic state. The following construction handles unbounded g.

Lemma A (Anchored phase reconstruction)**.** *For fixed periods, choose a physical cyclic state b_* as cell $(0, 0)$. The two rooted types $P_{-1}Q_0, P_0Q_{-1}$ defined below determine at most one array of cell types. Any two placements of this array that respect the dynamics and anchor differ by an automorphism of the target fixing b_*.*

Proof. Choose a cyclic state b_*, put $J = h + \lambda + 2$, and form the multiset $\mathcal{U}$ containing $e_b = T_J(B, b)$ once for each physical cyclic state. Write $e_0 = e_{b_*}$. In a true product index $b_* = (x_0, y_0)$, and write $P_i = T_J(X, x_i), Q_l = T_J(Y, y_l)$. The two genuine axes are the sequences with terms $u_i = P_iQ_0$ and $v_l = P_0Q_l$. Every such term belongs to the support of $\mathcal{U}$ because the whole grid occurs in B.

For cyclic profiles e let $\kappa(e)$ be the cut at $J - 1$ and $\pi(e)$ its unique child reaching that height, the cycle predecessor. Both operations preserve products on these profiles. Furthermore κ separates distinct types in $\mathcal{U}$: tracing the cycle predecessor for $\lambda - 1$ steps displays the complete off-cycle decorations of all λ phases: after at most $\lambda - 1$ steps the remaining depth of the cut is at least $J - 1 - (\lambda - 1) = h + 2$. All transient branches at that phase have height at most $h - 1$ and the cyclic predecessor is recognizable. Explicitly,

$$\pi(P_iQ_l) = \kappa(P_{i-1})\,\kappa(Q_{l-1}).$$

Guess the single term $u_- = P_{-1}Q_0$ directly from $\mathcal{U}$. The predecessor u' of a term u must satisfy

$$\kappa(u')\pi(e_0) = \pi(u)\kappa(u_-).$$

Indeed for $u = P_iQ_0$ both sides are the product of the two cuts at the preceding X phase $i - 1$ and the preceding X phase -1, and the two Y phases $0, -1$. Cancellation and injectivity of κ give at most one u' in the entire dictionary $\mathcal{U}$. Trace p steps from e_0, checking the return. Recover the second axis by guessing P_0Q_{-1} from $\mathcal{U}$ and applying the analogous q-step test. Use $e_{i,l}e_0 = u_iv_l$ to read off all cell types. For each diagonal orbit match its word of types, up to rotation, to an unused physical component, respecting b_* on its orbit. Match the anchored orbit first. Each orbit uses one entire cycle of length λ. Matching within the remaining word classes succeeds exactly when their multiplicities agree. A placement is then a type-preserving bijection from the $pq = g\lambda$ cells to the cyclic states that commutes with the transition and sends $(0, 0)$ to b_*. For two placements σ, σ', the permutation $\sigma'\sigma^{-1}$ commutes with the transition and fixes b_*. At each cyclic state, the cyclic predecessor is the only child of T_J reaching height $J - 1 > h - 1$. Its other children give all the attached transient branches. Equal T_J-types therefore give an isomorphism between the attached rooted trees. These trees are disjoint at distinct cyclic states, so the isomorphisms extend the permutation to an automorphism of B. It commutes with the successor map and therefore carries possible guide endpoints and their predecessor levels to corresponding endpoints and levels of the same widths. Each feasible ordered guess thus needs only one greedy placement. There are at most $(g\lambda)^2$ guesses per period pair, and reconstruction takes polynomial time per guess. □

One extension on a fixed grid

Fix a grid and a cyclic cell b_*. Put $H = h + 1$. Each factor in a true product has transient depth at most h, by pairing each of its states with a cyclic state in the other factor. Enumerate

two ordered endpoints z_1, z_2 satisfying $f^H(z_i) = b_*$. They determine two paths $v_{i,j} = f^j(z_i)$, $0 \le j \le H$, in the physical input. Ties in any choice of a minimum below require no further enumeration, because every endpoint is already tried.

In a true product choose guide 1 by following the cycle backwards in X, while in Y choosing, at every step, a child with minimum positive width at the remaining bottom level. Guide 2 interchanges the two roles. For the step $j \to j+1$, start at the $(H-j-1)$-st cyclic predecessor of b_*, say (x_a, y_b). By induction the radius-j marks T_j at every factor phase are known. An active child is a predecessor whose tree reaches level j; we retain its whole tree cut at that level. Let $\mathcal{O}_0, \mathcal{O}_1, \mathcal{O}_2$ be the observed active child forests at the cyclic cell, at $v_{1,j+1}$, and at $v_{2,j+1}$, respectively. For the true guides,

$$\mathcal{O}_0 = AV, \qquad \mathcal{O}_1 = AC, \qquad \mathcal{O}_2 = VD.$$

Here A, V are the active child forests of x_a, y_b. C, D come from the moving Y, X coordinates of the respective guides. A designated member $c_0 \in C$ is obtained by dividing $T_j(B, v_{1,j})$ by the known mark at x_{a-1}. Likewise dividing $T_j(B, v_{2,j})$ by the mark at y_{b-1} gives $d_0 \in D$. These are tree types, so repeated copies require no choice of their identities. Compute

$$E = (\mathcal{O}_1\mathcal{O}_2)/\mathcal{O}_0.$$

In a true product $E = CD$. Reject if the anchored product tree $c_0 d_0$ does not occur in E. Its mass $m = \mathrm{mass}(E)$ counts trees with multiplicity. Each possible continuation is obtained by finding an ordered split $E = CD$ containing the specified members, then computing $A = \mathcal{O}_1/C$, $V = \mathcal{O}_2/D$. Failed divisions, missing members, and inconsistent products are rejected. For a fixed grid, endpoint pair and level, all three observed forests come from specified states of B. Thus E, when it exists, and its mass can be computed once for the entire level. The anchors supplied to the local splitter and the subsequent mark checks depend on the node.

To build $T_{j+1}(X, x_a)$, subtract the $(j-1)$-cuts of A from the children of the old j-mark. The remainder consists of complete shorter branches. Attach these and the trees of A to a new root. At $j = 0$ all children are active and there is no subtraction. Apply the same rule to V. The resulting marks of the distinguished cell determine every other mark by division along its row and column. Check the product at every cell, consistency with the preceding cuts, and the inclusion of the appropriate cyclic predecessor marks. All these operations are deterministic given one local split. Reject any prospective radius-k factor mark exceeding $(k+1)N$ vertices or maximum degree N; this also controls the cost of false partial paths.

Initially the 0-marks are single vertices. At radius H remove the cyclic predecessor child at each factor phase. It is the unique child reaching height h in a true factor mark. The remaining children are complete transient branches. Close the two cycles, check connectedness, sizes and the product against the original table. The last test uses the cyclic word of the canonically encoded attached trees in each component, up to rotation, and compares multisets of component types. These tests prove soundness. The grid, the two endpoints and the local split along every true factorization occur in the search, proving completeness once the local enumeration is established. Stopping at the first success gives one witness. Exhaustion certifies a negative answer.

3 From local choices to the whole search tree

The following calculation is for a fixed grid and one fixed pair of physical endpoints. It applies before any choices of factor marks.

Lemma 2 (Observable level budgets)**.** *Set* $W_{i,j} = w_j(B, v_{i,j})$ *and*

$$\beta_j = \prod_{i=1}^{2} \lfloor W_{i,j+1}/W_{i,j} \rfloor, \qquad 0 \le j < H.$$

For every pair of endpoints, $\beta_j \geq 1$ and $\prod_{j=0}^{H-1} \beta_j \leq w_H(B, b_*)^2 \leq N^2$. On the two minimum-width guides of a true factorization, $m \leq \beta_j$ at every level.

Proof. If $f^j(u) = v_{i,j}$, then $f^{j+1}(u) = v_{i,j+1}$. The sets being counted are nested on every physical path, including a wrong guess. Also $W_{i,0} = 1$ and $W_{i,H} = w_H(B, b_*)$. Drop the floors and telescope to obtain the first assertion. For guide 1 let $y_{\min}$ be its selected moving-coordinate child and u any active child of the same parent. Pair both with the cyclic predecessor x_{a-1}. Minimality gives

$$w_j(X, x_{a-1})\, w_j(Y, u) \;\geq\; w_j(X, x_{a-1})\, w_j(Y, y_{\min}) = W_{1,j}.$$

Distinct physical u's give disjoint sets of bottom vertices in the tree of $v_{1,j+1}$, even if their rooted types agree. Thus $\mathrm{mass}(C) \leq \lfloor W_{1,j+1}/W_{1,j} \rfloor$. Guide 2 gives the analogous bound for D. Multiply them and use $\mathrm{mass}(E) = \mathrm{mass}(C)\,\mathrm{mass}(D)$. □

The algorithm prunes the current node if E does not exist, is empty, or $m > \beta_j$, *before* calling the local splitting algorithm. This check preserves a minimum-width witness for every genuine factorization. An arbitrary endpoint need not itself satisfy the minimum-width rule. One can also use the slightly weaker integer cap $\widehat{\beta}_j = \lfloor \prod_i (W_{i,j+1}/W_{i,j}) \rfloor$. The same proof applies with $\widehat{\beta}_j$ throughout, since $\beta_j \leq \widehat{\beta}_j \leq \prod_i (W_{i,j+1}/W_{i,j})$.

Proposition 3 (Local trial count)**.** *In the pivot enumeration, a complete trial specifies a pivot tuple, a dictionary shift and a positive integer content. A pivot tuple or a shift rejected before the next choice counts as one trial ending there. If $m > 1$ and E has only one tree type, a trial tests one multiplicity. For $m = 1$ there is at most one trial, the product test for the two designated trees. For $m \geq 2$ the number of trials is at most*

$$m^{K_0} (1 + K_1 \ell)^{\, K_1 \ell^2 \log(2m)}, \qquad \ell = \log(2N).$$

Here K_0, K_1 are absolute constants. Setup and testing one trial each cost $N^{O(1)}$ bit operations, including a trial rejected after exact forest division. All genuine anchored splits are returned.

The proof is given below. Every forest in its input has polynomial size: each $\mathcal{O}_i$ is observed in B, and $\mathcal{O}_1\mathcal{O}_2$ is an explicit product. Its maximum number of children is at most N^2, since numbers of children multiply. Whenever local enumeration is reached, division has verified $\mathcal{O}_0 E = \mathcal{O}_1\mathcal{O}_2$, with nonempty $\mathcal{O}_0$ of the prescribed height j. Fix a length-j path in one tree of $\mathcal{O}_0$. Pair every vertex of E with the vertex at the same depth on this path. This embeds the forest vertices into the dividend. Every vertex of E with children is paired with a path vertex having at least one child. Its degree is therefore at most the degree of the paired product vertex. Consequently $\Delta(E) \leq N^2$ and the total size is bounded by that of the dividend, also on an incorrect grid or guide.

Proposition 4 (Whole-tree bound)**.** *The complete computation on one grid and one endpoint pair, including unsuccessful guesses, takes* $\exp(O(\ell^3 \log\log(4N)))$ *bit operations.*

Proof. Let $b_j = 1$ if $\beta_j = 1$, and otherwise set $b_j = \beta_j^{K_0} (1 + K_1 \ell)^{K_1 \ell^2 \log(2\beta_j)}$. By the rejection rule and Proposition 3, every node at level j makes at most b_j trials. A complete trial decodes at most one left forest. Exact division determines its partner, and the subsequent mark checks are deterministic. Each trial therefore creates at most one successor. Write Q_j for the number of nodes and A_j for the number of trials at depth j. Then $Q_0 = 1$ and

$$Q_{j+1} \leq A_j \leq Q_j b_j \leq \prod_{u=0}^{j} b_u.$$

If a trial yields an exact quotient and fails only at the right anchor, it contributes one to A_j and zero to Q_{j+1}. A node pruned before enumeration pays for setup. For a common polynomial P

bounding setup, one trial with its mark checks, and verification at a leaf, the total time is at most

$$P(N)\left(\sum_{j=0}^{H} Q_j + \sum_{j=0}^{H-1} A_j\right) \le P(N)(2H+1)\prod_{j=0}^{H-1} b_j.$$

In particular a run of unit budgets costs at most its length times the polynomial cost per node for each path entering the run.

By Lemma 2,

$$\sum_j \log \beta_j \le 2\log N, \qquad \sum_{j:\beta_j \ge 2} \log(2\beta_j) \le 2\sum_j \log \beta_j \le 4\log N.$$

Taking logarithms gives

$$\log \prod_j b_j \le 2K_0 \log N + 4K_1\ell^2 \log N \, \log(1+K_1\ell) = O(\ell^3 \log\log(4N)).$$

Since $H \le N$, the extra polynomial factor is absorbed. There are at most N^2 endpoint pairs per grid and polynomially many grids. This proves the time bound in Theorem 1 once Proposition 3 is proved. For enumeration, canonicalize the two output tables and remove duplicates by sorting their strings; the additional logarithm of the number of records preserves the time bound. □

Dependencies of the decision procedure. The final product test certifies every returned pair. Exhaustion certifies a negative answer because Lemma A preserves a placement for each true pair, Lemma 2 preserves suitable guides, and the anchored enumeration recovers each genuine local split by Lemma C, cancellation and exact division. In Section 6, for any fixed positive integer τ and separating grade γ, a genuine split supplies a solution within the actual specialized multiplicities; the dictionary shift and integer content recover the forest. In particular $\tau = 1$ suffices for exhaustive search with those ranges. The norm gap, the factor and dimension bounds, and the choice of a slice with short ranges give the stated running time. Section 6.1 also proves that the bounded search for such a slice always succeeds.

4 A finite additive encoding with a uniform gap

In the local forest E, let Δ denote the maximum number of children at a vertex. We construct labels for its observed tree types and prove the identities needed to encode a product of two unknown forests.

Lemma B (Uniform gap for rational-function coordinates)**.** *Put $\mathcal{H} = \mathbb{Q}(z)^\times/\{\pm z^k : k \in \mathbb{Z}\}$. Represent a class by A/B with nonzero integral Laurent polynomials coprime over $\mathbb{Q}$ and jointly primitive over $\mathbb{Z}$. With $\mu_{\mathbb{T}}$ the normalized Haar measure of the unit circle,*

$$\mathfrak{h}([A/B]) = \int_{\mathbb{T}} \log\max(|A(z)|, |B(z)|)\, d\mu_{\mathbb{T}}(z)$$

is a symmetric, subadditive length, homogeneous under integer powers. Every nonidentity class has length at least $\delta = \log\theta$, where $\theta > 1$ is the real root of $t^3 - t - 1$.

Proof. Changing representatives multiplies numerator or denominator by signed monomials of modulus one. Inversion interchanges them and reduced positive powers remain reduced. For multiplication, the pointwise maximum of the product pair is at most the product of the maxima. Cancelling a common integral factor subtracts its nonnegative logarithmic Mahler measure, including any common integer content. This proves subadditivity. Write $m_{\mathbb{T}}(P)$ for the mean of $\log|P|$ on the appropriate torus. Jensen's formula in y gives $\mathfrak{h}([A/B]) = m_{\mathbb{T}}(A(z) + yB(z))$.

Call a Laurent polynomial reciprocal if simultaneous inversion of its variables reproduces it up to a signed monomial. A nonreciprocal univariate integer Laurent polynomial has logarithmic measure at least δ: some primitive irreducible factor is nonreciprocal. After stripping off powers of z, a leading coefficient of absolute value at least two already gives this bound; otherwise Smyth's theorem applies to the monic associate [13]. All remaining integer factors have nonnegative logarithmic measure.

The multivariate nonreciprocal bound is due to Boyd [14]. We give the specialization argument for the bivariate case used here. For bivariate nonreciprocal P with support S_0, the maps $\lambda_n(i,j) = i + nj$ are injective on both S_0 and $S_0 + S_0$ for all sufficiently large n. Reciprocity of $P(t, t^n)$ would pair its support points by reflection in the midpoint of its extreme exponents, with a common coefficient sign. Let $p, q \in S_0$ lift those extremes. Every reflected pair u, v satisfies $\lambda_n(u+v) = \lambda_n(p+q)$, hence $u + v = p + q$. This would give reciprocity of P with the same sign. Consequently $m_{\mathbb{T}}(P(t,t^n)) \geq \delta$ for large n. Lawton's specialization theorem [15] gives convergence to $m_{\mathbb{T}}(P)$, since the shortest nonzero integral relation for $(1, n)$ has sup norm n.

For reciprocal $P = A + yB$, its y-support $\{0, 1\}$ forces $B(z) = \pm z^k A(z^{-1})$. Since the coefficients are real, $\mathfrak{h}([A/B]) = m_{\mathbb{T}}(A)$. If A is nonreciprocal, the univariate result applies. If A is reciprocal then A/B is a signed monomial. The identity has length zero. Finally, a polynomial with nonnegative integer coefficients of sum $d \geq 1$ represents a class of length at most $\log d$ by the pointwise bound on the circle. □

Lemma C (Encoding finite dictionaries of equal-height trees)**.** *For a finite family of explicitly given height-k tree types there is a deterministic polynomial-bit construction of integer coordinates for their labels with the following properties. The ideal label map is a monoid homomorphism on all trees of the given height. It is injective on the observed dictionary. Its norm on observed labels is at most* $\log \max(2, \Delta)$*, and every nonzero element of the subgroup generated by them has norm at least δ. The computed coordinates are global homomorphisms and are faithful on this subgroup, where they carry the transported norm.*

Proof. For each $j \leq k$, let $\mathcal{T}_j$ consist of the height-j cuts at all input vertices whose subtrees reach that height. A cut keeps every vertex within its radius, including leaves on shorter branches. Every iterated subtree and cut so obtained is specified by an original vertex and a remaining radius. The dictionaries have polynomial total size. Put $G_0 = \{0\}$, with the unique label at height zero. Inductively, at each previously constructed height j fix an integral homomorphism $g_j : G_j \to \mathbb{Z}$ distinguishing the labels of $\mathcal{T}_j$, as constructed below. Order G_j first by g_j and resolve ties by a fixed translation-invariant total order. For example, coordinates in $\mathcal{H}$ are the prime valuations of the absolute leading scalar and the multiplicities of the monic irreducibles over $\mathbb{Q}$ other than z. Lexicographic comparison in a fixed enumeration, iterated for finite products, gives such an order.

Write ϕ_j for the ideal height-j label. For a tree T of height $d > 0$, let $S_j(T)$ be the nonempty multiset of height-j cuts of its children that reach height j, for $j < d$. With $G_d = \prod_{j<d}(\mathcal{H} \times G_j)$, define $\phi_d(T)$ by putting in its j-th slot

$$\left(\left[\sum_{u \in S_j(T)} z^{g_j(\phi_j(u))} \right], \min_{u \in S_j(T)} \phi_j(u) \right).$$

The formulas apply to every height-d tree once the dictionary fixes the grades. The active cuts of a product are all pairwise products of the active cuts of the factors, with multiplicity. By induction the first entries multiply in $\mathcal{H}$, and translation invariance makes the minima add in G_j. The unit path has zero label in additive notation. This proves the global homomorphism property.

On an observed tree every active cut belongs to $\mathcal{T}_j$. Its grade identifies it in the preceding dictionary. The minimum label fixes the least exponent of the polynomial, recovering the graded

histogram from its class since the coefficients are positive. Thus all the multisets $S_j(T)$ are known. Descend from $j = d - 1$ to zero, subtracting from $S_j(T)$ the cuts of all taller children already recovered. The remainder consists exactly of the children of height j. Their union determines T. Give G_d the maximum of its coordinate lengths. Each histogram has coefficient sum at most Δ, and the minimum entry is an observed lower-height label. Induction gives the radius bound. A nonzero coordinate has length at least δ, proving the gap.

To construct the grades, at height d use the scalar and polynomial valuations of each histogram and the previously retained coordinates of each minimum. These are integer homomorphisms on all of G_d. Store only valuations present somewhere in $\mathcal{T}_d$. The rows are jointly injective on the subgroup generated by its labels: omitted valuations vanish there, and each minimum entry belongs to the subgroup generated at its earlier height. Take a maximal rationally independent subset of the rows evaluated on the observed columns. A vector vanishing on these rows also vanishes on all original rows on the generated subgroup. This proves faithfulness. The retained rows themselves remain defined on the entire ideal group, so using them as grades requires no extension from a finite-index subgroup.

For $n_d = |\mathcal{T}_d|$ distinct observed columns, take a linear combination of the retained rows with integer weights in $[0, \binom{n_d}{2}]$. Choose the weights sequentially: the last differing row of each pair forbids at most one weight there. This gives g_d. Each scalar valuation is bounded by Δ. A histogram with s' terms has multiplicity at most $s' - 1$ at any nonzero root, by the Vandermonde matrix of its exponents applied to Euler derivatives of orders $0, \dots, s' - 1$. All retained entries are consequently bounded by $\max(2, \Delta)$, including those of earlier minima, with at most n_d retained rows per height. Grades, shifted univariate degrees and coefficient bit lengths are polynomial in explicit input size uniformly over heights. Factor the leading integers by trial division and the univariate polynomials over $\mathbb{Q}$ by deterministic polynomial-bit factorization [21]. The algorithm stores the retained rows and histograms of the finite dictionaries. On every observed active cut, g_j already separates distinct dictionary types, so ties have identical labels. The ideal tie-breaking order defines the homomorphism on arbitrary trees. Computing labels of unobserved trees is unnecessary. □

In particular the identity $(cd)(c_0d_0) = (cd_0)(c_0d)$ passes to the computed coordinates. All four entries cd, c_0d_0, cd_0, c_0d are observed product types when an anchored split exists. The algorithm labels these observed types and recovers a proposed axis by tree division.

Sharpness on observed labels. Take two height-two trees each with three star children, of leaf counts $(1, 16, 32)$ and $(1, 2, 32)$, respectively. The height-one dictionary also contains the three-leaf star obtained by cutting the roots. The grade $v_2 + 2v_3$ on star degrees separates the whole dictionary and agrees with v_2 on the four relevant powers of two. In the height-two labels the root counts and minimum entries agree. The remaining histograms are $1 + z^4 + z^5$ and $1 + z + z^5$. Cancelling their common factor $1 + z + z^2$ leaves P/P^*, where

$$P = 1 - z + z^3, \qquad P^* = 1 - z^2 + z^3 = z^3 P(1/z).$$

They are coprime, since $P - P^* = z(z - 1)$ and $P(0) = P(1) = 1$. On the unit circle their moduli agree. The zeros of P are $-\theta$ and a conjugate pair of modulus $\theta^{-1/2}$. Hence the norm of this difference of observed labels is exactly $\log \theta = \delta$. The same holds for any positive integer multiple of this grade by Haar invariance under $z \mapsto z^e$.

5 Bounding the number of factors by a Newton polytope

Let Λ be the lattice generated by the observed label differences. Translate the support by one of its points. The difference vectors have norm at most $R = 2 \log \max(2, \Delta)$. The group norm extends to a genuine norm on $\Lambda \otimes_{\mathbb{Z}} \mathbb{R}$: extend first homogeneously and continuously from rational

vectors. If there were a nonzero kernel direction, the unit ball of the seminorm would contain cylinders of arbitrarily large volume, contradicting the lattice gap by Minkowski's convex body theorem applied to a ball of radius below δ. Only the finite-dimensional space spanned by the input is used.

Lemma 5. *If a Laurent polynomial on Λ has s support points in a norm ball of radius R and nonzero lattice vectors have length at least δ, the number k of nonmonomial irreducible factors over $\mathbb{C}$, counted with multiplicity, is at most*

$$8(R/\delta)^2 \log(2s).$$

Proof. From each of the k nonpoint Newton polytopes choose a segment with distinct lattice endpoints and difference v_i. Their Minkowski sum, up to translation, lies inside the Newton polytope P of the product. Indeed Newton polytopes add under multiplication over a field. Take real dual norm-one functionals β_i satisfying $\beta_i(v_i) = \|v_i\| \geq \delta$. For independent random signs define $\beta = \sum_i \epsilon_i \beta_i$. Conditioning on all signs except ϵ_j gives $\mathbb{E}|\beta(v_j)| \geq |\beta_j(v_j)|$. The expected width of the sum of segments is therefore at least $k\delta$. At each of the s input points the random sum has exponential-moment bound $\mathbb{E}e^{t\beta(x)} \leq e^{t^2 k R^2/2}$. The maximum of the $2s$ signed point evaluations has expectation at most $R\sqrt{2k\log(2s)}$. Twice this bounds the expected width of P. Squaring yields the stated bound. Random signs are used only in this proof. □

6 A deterministic linear system for local multiplicities

Ruppert and Gao relate polynomial factorisation to linear differential equations in two variables [16, 17]; Shaker uses closed forms for reduced polynomials in several variables [18]. Here the unknowns used for enumeration are multiplicities after specialization. Proposition 6 bounds the dimension of this projected space in terms of the factors over the original lattice.

Represent the auxiliary forest E by its Laurent polynomial $\mathcal{P} = \sum_{v \in S} n_v X^v \in \mathbb{Q}[\Lambda]$, after translation taking the observed type $c_0 d_0$ to the origin. For a genuine split $E = CD$, apply the additive labels to $(cd)(c_0 d_0) = (cd_0)(c_0 d)$. The two axes are then supported at differences of observed points from the origin, so their positive Laurent polynomials multiply to $\mathcal{P}$ in this ring. Dividing each observed type by d_0 gives the dictionary $\mathcal{D} \subseteq S$ of possible left-axis points and their physical trees. An axis tree c can itself be unobserved. Its relative label $\phi(c) - \phi(c_0)$ equals the observed difference $\phi(cd_0) - \phi(c_0 d_0)$ in Λ, where ϕ denotes the ideal label. If two candidate trees give the same computed point, faithfulness on Λ and injectivity on observed types give $cd_0 = c'd_0$. Cancelling the fixed height-j tree d_0 yields $c = c'$. The other axis uses c_0 in the same way. Thus the lookup is unique even when the decoded tree lies outside the dictionary of observations.

The fixed-anchor argument separates the active sibling cuts compared within each minimum in Lemma C. When constructing the observed dictionary, cuts of the same radius may come from different depths, and the separating grade distinguishes their observed types directly. Fix a path $w_0, \ldots, w_j$ of length j in d_0, a depth s, and a cut radius $k \leq j - s$. For any active k-cut u at a node x of depth s in c, let v be the k-cut below w_s. Then v has height k and uv is the observed cut at (x, w_s) in cd_0. The same v works for every such u. If $g_k(u) = g_k(u')$, additivity gives $g_k(uv) = g_k(u'v)$. Separation on $\mathcal{T}_k$ and cancellation of v give $u = u'$ up to isomorphism. Thus a minimum among active siblings can tie only on copies of the same type. Cuts at different depths may have equal grades.

The intrinsic lattice. Let $r = \operatorname{rank} \Lambda$. The case $r = 0$ is the singleton case described below. Retain r integer coordinate rows $\alpha_i : \Lambda \to \mathbb{Z}$ independent over $\mathbb{Q}$, and denote their combined map by α. The table of point sums in $\alpha(\Lambda) \subseteq \mathbb{Z}^r$ is an isomorphic representation of Λ. Its index may exceed one. For example, with $\Lambda = 2\mathbb{Z}$, the element $1 + X^2 = 1 + u$ has one absolute factor

in the intrinsic coordinate $u = X^2$, and two absolute factors after adjoining X. All factor counts below use $\mathbb{C}[\Lambda]$. Computations use integer row values and equality of point sums.

On this group algebra let $\mathcal{E}_i X^v = \alpha_i(v) X^v$. These commuting Euler derivations span the invariant derivations. Their dual invariant forms η_i satisfy $dX^v/X^v = \sum_i \alpha_i(v)\eta_i$. Choose a small integer grade $\gamma(v) = \sum_i \gamma_i \alpha_i(v)$ separating S. Write σ for a second integer grade and τ for a positive integer. The curve used in the algorithm has pullback

$$\psi(X^v) = \tau^{\sigma(v)} t^{\gamma(v)}.$$

The single pair (σ, τ) used for enumeration is chosen in Section 6.1. Positivity gives $\psi(\mathcal{P}) \neq 0$. After removing its monomial and scalar, write its rational factorization as $\prod_\nu \pi_\nu^{m_\nu}$. No π_ν has zero as a root.

Take numerator polynomials H_i with unknown rational coefficients on S, an unknown scalar λ_0, and unknown z_ν. Use the rational linear system

$$\begin{aligned} \mathcal{P}(\mathcal{E}_i H_j - \mathcal{E}_j H_i) &= H_j \mathcal{E}_i \mathcal{P} - H_i \mathcal{E}_j \mathcal{P} \quad (i < j), \\ \sum_i \gamma_i \, \psi(H_i)/\psi(\mathcal{P}) &= \lambda_0 + \sum_\nu z_\nu \, t\pi_\nu'/\pi_\nu. \end{aligned}$$

Compare the first equations on $S + S$, and the second after multiplication by $\psi(\mathcal{P})$. A Laurent monomial clears negative univariate exponents. For a positive Laurent factor V from a genuine split, $H_i = (\mathcal{P}/V)\mathcal{E}_i V$ is supported in S, by positivity of the undifferentiated convolution. The form $\omega = \sum_i (H_i/\mathcal{P})\eta_i$ is then dV/V; the second equations hold with z_ν the multiplicities in $\psi(V)$.

Lemma D (Closed forms on the intrinsic torus)**.** *Choose an integral basis $e_1, \ldots, e_r$ of Λ and put $x_i = X^{e_i}$. Let $f_1, \ldots, f_k$ represent the distinct nonmonomial irreducible factors of $\mathcal{P}$ in $\mathbb{C}[\Lambda]$. On $A = \mathbb{C}[\Lambda][1/\mathcal{P}]$, each closed regular one-form has an expression*

$$dg + \sum_{i=1}^{r} a_i \frac{dx_i}{x_i} + \sum_{j=1}^{k} b_j \frac{df_j}{f_j}, \qquad g \in A, \quad a_i, b_j \in \mathbb{C}.$$

Proof. Multiply each f_j by a Laurent monomial so that it is a polynomial p_j with minimum exponent zero in every variable. Each p_j is irreducible in $\mathbb{C}[x_1, \ldots, x_r]$: it has no coordinate factor, and any other polynomial factorization would survive localization at the coordinates. Distinct p_j's are nonassociate. For $g_0 = x_1 \cdots x_r \prod_j p_j$ we have

$$A = \mathbb{C}[x_1, \ldots, x_r, 1/g_0].$$

Indeed inverting $\mathcal{P}$ in the Laurent ring inverts every p_j; inverting g_0 in the polynomial ring also inverts every coordinate. This identifies the open torus complement with the affine principal open set $g_0 \neq 0$, including its complex of regular differentials. Bürgisser–Scheiblechner [19, Sec. 5.1.1, Theorem 5.3], over $\mathbb{C}$, give a basis of its first algebraic de Rham cohomology by the logarithmic differentials of the distinct irreducible factors of g_0. The thesis states this basis theorem on p. 107 and proves it in Section 8.3, pp. 111–112 [20, Theorem 8.3]. Their statement allows positive powers of these factors; the localization already depends only on the reduced zero set. Subtracting a constant linear combination of the basis from a closed form leaves dg with g in the localized ring, by the definition of that cohomology. Replacing p_j by f_j only changes the coordinate logarithms. □

Proposition 6. *The rational dimension of the projection of the linear system to the z coordinates is at most the number of distinct nonmonomial absolute irreducibles of $\mathcal{P}$ in $\mathbb{C}[\Lambda]$.*

Proof. Extend scalars to $\mathbb{C}$. The first equations say $d\omega = 0$, as the $\mathcal{E}_i$ commute and the η_i are closed. Apply Lemma D. The curve maps A to $\mathbb{C}(t)$: the coordinates map to nonzero Laurent monomials, and $\psi(\mathcal{P}) \neq 0$. In particular every $\psi(f_j)$ is a nonzero rational function. In the intrinsic basis, $d\log\psi(x_l) = \gamma(e_l)\,dt/t$. The matrix $(\alpha_i(e_l))_{i,l}$ is invertible over $\mathbb{Q}$; its defining equations for the η_i therefore give $\psi^*\eta_i = \gamma_i\,dt/t$. This identity includes the weighting by $\tau^{\sigma(v)}$.

For each π_ν, take one complex root $\xi_\nu \neq 0$. Rational irreducibles in characteristic zero are separable and distinct ones have disjoint root sets. The second equation, multiplied by dt/t, therefore gives

$$z_\nu = \operatorname{res}_{t=\xi_\nu} \psi^*\omega = \sum_{j=1}^{k} b_j \operatorname{ord}_{t=\xi_\nu} \psi(f_j).$$

An exact rational differential has zero residue even when its primitive has a pole there; each coordinate logarithm is a multiple of dt/t. The displayed formula bounds the dimension for *every* solution by k. It allows distinct f_j's to share specialized roots and one f_j to split or acquire repeated roots. Projection of a rational kernel commutes with extension of scalars, giving the same dimension over $\mathbb{Q}$. The integral basis and the absolute factors were used only in this proof. □

A rational basis of the projected space determines a set of coordinate pivots. A pivot tuple fixes one full vector z. For a genuine split z_ν is an integer in $[0, \min(m_\nu, m-1)]$, by the sparse Vandermonde bound. There are $O(r^2s^2)$ multivariate coefficient equations and polynomially many univariate coefficients. Their entries have polynomial bit length, and fraction-free elimination has polynomial bit cost. If $s = 1$, dividing the sole observed type by d_0 gives the only possible C-type. Try its positive multiplicity up to m and verify the opposite factor. If $m = 1$, test c_0d_0 directly and charge one trial.

6.1 A second grade shortens the multiplicity ranges

Choose a second small linear integer grade σ so that (γ, σ) preserves noncollinearity of all noncollinear triples in S. For differences u, v of such a triple impose $\sigma(\gamma(u)v - \gamma(v)u) \neq 0$; polynomially many hyperplanes can be avoided by the same separating-weight procedure. Clear a Laurent monomial from the planar polynomial

$$\mathcal{B}(t, y) = \sum_{v \in S} n_v\, t^{\gamma(v)} y^{\sigma(v)}.$$

Every edge of its Newton polygon has at most $1 + 2R/\delta$ support points, since their lifts lie on one intrinsic lattice line in the ball of radius R. If the polygon is a segment, take the entire segment as its face. Let μ be the maximum multiplicity of a nonzero root of any edge polynomial in the primitive edge variable. The Vandermonde bound gives $\mu \leq 2R/\delta$. Compute μ by univariate squarefree decomposition on the faces. In rank one take $\sigma = 0$. Every nonmonomial factor has an edge (or is supported on a segment). Choose a nonzero normal exposing that edge or segment. Initial forms multiply, and the exposed face of the planar product contains a nontrivial segment parallel to the chosen edge. Its edge polynomial contains the corresponding initial factor, with its multiplicity. This bounds all nonmonomial factor multiplicities of the planar polynomial by μ.

There is a deterministically findable $\tau > 0$ for which $\mathcal{B}(t, \tau)$ has all nonmonomial multiplicities at most μ. Write d_t, d_y for the two degrees. The extreme t-coefficients are monomials in y, so the t-degree is preserved at each positive specialization and every common vertical factor is a Laurent monomial. Let $\mathcal{B}_{\mathrm{red}}$ be the primitive squarefree part in t over $\mathbb{C}(y)$, chosen as the product of the distinct positive-t-degree polynomial factors over $\mathbb{C}[t, y]$. Put $a = \deg_t \mathcal{B}_{\mathrm{red}} \geq 1$. Its extreme t-coefficients divide the corresponding monomials of $\mathcal{B}$, so its degree stays a at every positive τ. Its resultant with $\partial_t\mathcal{B}_{\mathrm{red}}$ is nonzero by squarefreeness over $\mathbb{C}(y)$. The Sylvester matrix has $2a - 1$

rows with entries of y-degree at most d_y, giving a degree bound of $(2a-1)d_y \le (2d_t-1)d_y$. Try the integers $\tau = 1, \dots, 2d_t d_y + 1$, testing the actual multiplicities in $\mathcal{B}(t,\tau)$ by univariate gcds. Some trial avoids the resultant's roots. There the reduced slice is squarefree of full degree, so the specialized factors stay pairwise coprime and each multiplicity is at most μ. When $d_y = 0$, after the monomial shift the support is horizontal and the entire segment is the face. Its edge polynomial gives the slice itself in the primitive horizontal variable. If the exponents are compressed by their difference gcd g, the substitution $u = t^g$ preserves nonzero root multiplicities because $gt^{g-1} \neq 0$ there. Thus $\tau = 1$ is accepted even when $\mu > 1$. All degrees and bit lengths are polynomial. The resultant itself is not needed.

Factor the chosen slice, compute its projected space and enumerate only its pivot tuples, solving linearly for all other z_ν. After clearing denominators, Hadamard's bound gives polynomial bit length for the minors in elimination and hence for the rational basis in pivot coordinates. Its product with any trial tuple therefore has polynomial bit length, including a vector that fails the following test. Before taking any polynomial power, reject unless every coordinate is integral and $0 \le z_\nu \le \min(m_\nu, m-1)$. Take the candidate univariate product up to scalar and monomial. Try its least support point $v_0 \in \mathcal{D}$, translate its exponents by the shift determined by $\gamma(v_0)$, and look up their unique dictionary preimages v; reject if any is absent. Divide the coefficient at v by $\tau^{\sigma(v)}$, clear denominators, check the common sign and normalize to a positive primitive integer vector. Now try the original forest's integer content $1, \dots, m$. Reject a vector of total mass exceeding m before expanding copies as trees. For the resulting left forest C, check $c_0 \in C$, compute $D = E/C$ by exact forest division, and check $d_0 \in D$ and $CD = E$.

The weights $\sigma(v)$, degrees and $\log \tau$ are polynomially bounded, so undoing the weighting and taking gcds has polynomial bit cost, even for a rejected vector. Let P_* be the primitive integral specialization with its monomial removed. Every bounded subproduct has a primitive integral representative $Q \mid P_*$. Gauss's lemma and multiplicativity of Mahler measure give $\|Q\|_\infty \le 2^{\deg P_*} M(Q) \le 2^{\deg P_*} \|P_*\|_1$. This controls partial products as well. Monic representatives introduce a denominator no larger than the leading coefficient of their primitive representative.

There are $O(\ell^2 \log(2m))$ coordinate pivots by Lemma 5 and Proposition 6, since $s \le m$ and $\Delta = N^{O(1)}$. Each pivot has at most $1 + \mu = 1 + O(\ell)$ choices. For each pivot tuple there are at most $|\mathcal{D}|m \le m^2$ terminal trials. Indeed, a rejected prefix stands for all its completions among the dictionary shifts and the contents $1, \dots, m$. Distinct terminal trials have disjoint sets of formal completions, and $\mathcal{D}$ contains the point for c_0 because $c_0 d_0$ was checked in E. The work on the at most three choice steps of one trial has polynomial bit cost. Write $|F|$ for the number of vertices in a forest F. Pairing a dictionary tree with a full-height path in d_0 embeds it in its observed product tree, so it has at most $|E|$ vertices. Thus $|C| \le m|E|$, and any exact quotient has $|D| \le |E|$ by the same path argument in C. Exact division, the right-anchor test and remultiplication all have polynomial bit cost even when the trial fails. Every true split supplies a closed form dV/V, so its pivot tuple is tried. Its least dictionary point and original content are among the choices, and exact division returns its partner. Each output has passed multiplication of the actual forests. This proves Proposition 3.

A sharper count for continuations that pass the product test

Let k be the number of nonmonomial absolute factors of the local Laurent polynomial, counted with multiplicity. In $\mathbb{C}[\Lambda]$, unique factorization gives at most $\prod_i (e_i + 1) \le 2^{\sum_i e_i} = 2^k$ divisors modulo scalar and monomial units. For an admissible divisor, the least support point under the separating grade belongs to a dictionary of at most m points. This fixes the monomial shift once one support point of the divisor is aligned. Primitive integral normalization followed by an integer content between 1 and m fixes its scale. Conversely, a verified C fixes its specialized multiplicity vector, least support point and content, so it is reached by a unique complete trial. Verified division determines D. Thus the number of genuine anchored children is at most $m^2 2^k$ when

$m \geq 2$, and at most one when $m = 1$. Lemma 5 and the same level-budget sum give $\exp(O(\ell^3))$ nodes in the tree of verified local continuations per grid and endpoint pair. This gives a second cost calculation. Multiply this node bound by the largest one-node cost in Proposition 3, using $m \leq N^2$. The product again gives Theorem 1, including failed trials. The sharper node count thus gives no faster algorithm without a corresponding improvement in local enumeration.

7 Scope and reproducibility

The bounds concern explicitly written successor tables. In particular the argument allows both factor depths, the number of sibling types, and the number of target components to grow. For an arbitrary disconnected input, a negative answer rules out precisely the pairs of connected factors. The complexity of unrestricted reducibility is posed as open in [8, Sec. 7] and [3, pp. 185–186]. A uniform polynomial bound for constructing one connected pair would require a further argument for choosing a continuing local split or for finding all candidate splits efficiently. For example, a local procedure with polynomial-bit setup and testing, a single trial at $m = 1$, and at most m^c trials otherwise for one absolute constant c, would give a polynomial bound even for the full search: $\prod_j \beta_j^c \leq N^{2c}$. This is a sufficient condition. The number of local splits alone gives no lower bound for finding one complete pair.

The representation matters even for a single cycle C_n, $n \geq 2$. If $C_n = XY$, both factor maps are surjective and both factors are connected, hence both are cycles. The product $C_a C_b$ consists of $\gcd(a, b)$ cycles of length $\operatorname{lcm}(a, b)$. Thus a nontrivial split exists exactly when n has at least two distinct prime divisors [6, author version, Appendix A, p. 19]. For n given in binary this decision takes polynomial time in $\log n$: test whether $n = p^k$ using exact integer roots for $1 \leq k \leq \lfloor \log_2 n \rfloor$ and deterministic primality testing [22]. Constructing the two *lengths* already recovers p, q when $n = pq$ for distinct primes. Gaze-Maillot and Porreca establish a related search reduction for binary-encoded profiles, which record only the state counts at each distance from a cycle [23, author version, Theorem 20, p. 8 and encoding on p. 10].

Common anchors and short branches. We tested the product identity used in Lemma C and in the fixed-anchor argument of Section 6 on all 9,633 ordered pairs of equal positive height with at most eight vertices per rooted tree. There were 108,748 nonempty groups of active siblings at a given radius, including groups with one child. At every group the common anchor cut had the required height, all the paired cuts occurred in the observed product, and different sibling types had different products with this anchor. At radius two, 136 groups contained distinct types. For a concrete example, take a root whose two child trees are a two-edge path and the same path with an additional leaf attached to its root. The three-edge path is an anchor. The ordinary radius-two cut distinguishes the siblings by that leaf, which belongs to a branch ending before the cut radius.

Guide budgets and complete pairs. For all 3,413 labelled functions on one through five states, direct calculation checked the budgets of Lemma 2 on 45,233 ordered endpoint pairs with a common H-step image, giving 139,777 level budgets. Separately, direct product generation gave a catalogue of all 203 targets on at most 12 states admitting two nontrivial connected factors. The global search with exhaustive local splitting recovered all 428 ordered pairs. At its 17,343 local calls of mass at least two (206 distinct target-and-anchor triples), the weighted splitter of Section 6.1 agreed with the exhaustive local splitter. At each extension the global search checks the cyclic predecessor marks and the size and degree bounds. It verifies complete pairs by multiplying their successor tables. The catalogue search enumerates its local splits exhaustively and compares them with the weighted splitter. The bit-complexity estimate in Proposition 4 uses the weighted enumeration of Proposition 3.

Local splitting and specialized multiplicities. An independent oracle directly multiplies rooted trees and enumerates admissible ordered forest pairs. Across 407 calls (347 distinct triples), it agreed with the weighted splitter on every set of ordered splits. Of these calls, 153 had no split. The examples include unequal branch lengths at heights two through five and a projected space of rank three.

On 106 further anchored star-forest instances of mass at most ten (103 distinct targets), fixing $\tau = 1$ and using the actual specialized multiplicities again recovered every split found by direct enumeration. Seven instances had maximum multiplicity two at $\tau = 1$ and edge bound $\mu = 1$. The bounded search chose $\tau = 2$ in all seven. In six instances two coprime factors acquired a common root. In the other instance the factor $1 + 2x^2 + x^2y$ specialized to $(1 + t^2)^2$ under $(x, y) = (t, t^2)$, coprime to the other specialized factor $1 + t^3$. The four ordered forest splits in this example were recovered at both $\tau = 1$ and $\tau = 2$. The residue formula in Proposition 6 counts both sources of multiplicity.

A trial rejected after exact division. On a 420-state target we followed one node with anchors S_1, S_2, where S_d is the rooted star with d children. Its local mass was 16 and its level budget was 595. At projected rank four and $\tau = 2$, the pivot tuple $(1, 0, 0, 0)$ decoded to $C = S_1 + S_2$ at one shift and content. Exact forest division returned

$$D = S_1 + S_3 + S_4 + S_6 + S_{12} + S_{18} + S_{24} + S_{72}.$$

Independent remultiplication gave $CD = E$. The missing right anchor S_2 rejected this trial. Another trial at the same node extended to factors of 21 and 20 states. The supplemental scripts also checked selected nodes on targets of 60, 360, 2,880 and 34,560 states, finding respectively 2, 4, 8 and 16 distinct candidate left forests with exact positive quotients missing the right anchor. These are targeted tests of the trials counted in A_j in Proposition 4.

Reproducibility. The companion repository contains input generators, independent enumerators and expected outputs. From its `finite-dynamics-factorisation` directory, run `python -B reproduce.py` to repeat the checks. The experiments compare outputs on finite inputs. The bit-complexity bound in Theorem 1 uses deterministic polynomial-bit factorization over $\mathbb{Q}$.

References

[1] F. Doré, K. Perrot, A. E. Porreca, S. Riva and M. Rolland. Infinite trees for division and roots over finite discrete-time dynamical systems. *Natural Computing* 25 (2026), article 16. DOI; author version.

[2] É. Naquin and M. Gadouleau. Factorisation in the semiring of finite dynamical systems. *Theoretical Computer Science* 998 (2024), 114509. DOI.

[3] M. Rolland. *Des unrolls aux systèmes dynamiques totaux et partiels: algorithmes de résolution et énumérations des (in)équations polynomiales*. Doctoral thesis, Aix-Marseille Université, 2026, NNT 2026AIXM0179. Thesis.

[4] A. Dennunzio, V. Dorigatti, E. Formenti, L. Manzoni and A. E. Porreca. Polynomial equations over finite, discrete-time dynamical systems. *ACRI 2018*, LNCS 11115, 298–306. DOI.

[5] F. Doré, E. Formenti, A. E. Porreca and S. Riva. Decomposition and factorisation of transients in functional graphs. *Theoretical Computer Science* 999 (2024), 114514. Author version.

[6] F. Bridoux, C. Crespelle, T. H. D. Phan and A. Richard. Dividing sum of cycles in the semiring of functional digraphs. *Natural Computing* 25 (2026), article 11. DOI; author version.

[7] C. Crespelle, T. H. D. Phan and A. Richard. Tensor product of functional graphs: algorithms for factorization and polynomial root-finding restricted to the cycle part. Preprint (2025). HAL.

[8] A. E. Porreca and M. Rolland. Injective and pseudo-injective polynomial equations: from permutations to dynamical systems. Preprint (2026). arXiv.

[9] R. Ascone, A. E. Porreca and M. Rolland. Polynomial equations over bijective dynamical systems and their semiring extensions. Extended abstract, Workshop on Cellular Automata and Boolean Automata Networks, UCNC 2026. DOI.

[10] W. Imrich. Factoring cardinal product graphs in polynomial time. *Discrete Math.* 192 (1998), 119–144. DOI.

[11] L. Calderoni, L. Margara and M. Marzolla. Direct product primality testing of graphs is GI-hard. *Theoretical Computer Science* 860 (2021), 72–83. DOI.

[12] W. Imrich and W. Klöckl. Factoring directed graphs with respect to the cardinal product in polynomial time II. *Discuss. Math. Graph Theory* 30 (2010), 461–474. DOI.

[13] C. J. Smyth. On the product of the conjugates outside the unit circle of an algebraic integer. *Bull. London Math. Soc.* 3 (1971), 169–175. DOI.

[14] D. W. Boyd. Kronecker's theorem and Lehmer's problem for polynomials in several variables. *J. Number Theory* 13 (1981), 116–121. DOI.

[15] W. M. Lawton. A problem of Boyd concerning geometric means of polynomials. *J. Number Theory* 16 (1983), 356–362. DOI.

[16] W. Ruppert. Reduzibilität ebener Kurven. *J. Reine Angew. Math.* 369 (1986), 167–191. DOI.

[17] S. Gao. Factoring multivariate polynomials via partial differential equations. *Math. Comp.* 72 (2003), 801–822. DOI.

[18] H. Shaker. Topology and factorization of polynomials. *Math. Scand.* 104 (2009), 51–59. DOI.

[19] P. Bürgisser and P. Scheiblechner. Differential forms in computational algebraic geometry. *ISSAC 2007*, 61–68, Theorem 5.3. Author PDF.

[20] P. Scheiblechner. *On the Complexity of Counting Irreducible Components and Computing Betti Numbers of Algebraic Varieties*. Doctoral thesis, Universität Paderborn, 2007. Thesis.

[21] A. K. Lenstra, H. W. Lenstra, Jr., and L. Lovász. Factoring polynomials with rational coefficients. *Math. Ann.* 261 (1982), 515–534. DOI.

[22] M. Agrawal, N. Kayal and N. Saxena. PRIMES is in P. *Ann. of Math.* 160 (2004), 781–793. DOI.

[23] C. Gaze-Maillot and A. E. Porreca. Profiles of dynamical systems and their algebra. Preprint (2020). arXiv.