A proof of Nivat's conjecture

Student

Pranava Kumar

School location

North Carolina, USA

Instructor

Manolis Kellis

A proof of Nivat's conjecture

Pranava Kumar

14 September 2026

Abstract

We prove that a finite-alphabet configuration on $\mathbb{Z}^2$ with at most kn patterns on a $k \times n$ rectangle has a nonzero period. The main ingredient is a sharp lower bound: a nonperiodic finite-valued configuration that has a product of directional binomial annihilators and agrees with a fully periodic configuration on a rational wedge has at least $kn+1$ patterns on every $k \times n$ rectangle. We obtain this bound by supplementing the affine coordinate functions with local quadratic functions. Their independence follows from a support argument: directional differences produce nonzero pointwise products of half-plane-supported defects, confined to pointed cones on which no Laurent recurrence can hold. A selection argument based on Colle's construction supplies the wedge. Here we retain a finite boundary disagreement, so that the orbit-closure point to which the bound is applied remains nonperiodic.

Keywords: Nivat's conjecture; rectangular pattern complexity; periodicity; Laurent polynomial annihilators; symbolic dynamics.

Table of Contents

1 Introduction **1**

2 Annihilators and directional filters **3**

3 The pointed-cone obstruction and a local observable **6**

4 The exact rectangle count **8**

5 Generating sets and a periodic reference **9**

6 Boundary continuation in polygons **11**

7 From one regional period to two **13**

8 Half-strips and anchored compactness **14**

9 Proof of the selection proposition **18**

Acknowledgments and AI use **18**

A Relation to previous work **19**

B Boundary cases and examples **20**

References **21**

1 Introduction

Let $c : \mathbb{Z}^2 \to \mathcal{A}$ be a configuration over a finite alphabet. For a finite set $D \subset \mathbb{Z}^2$, write

$$\mathrm{Pat}_c(D) = \{(c(g+u))_{u\in D} : g \in \mathbb{Z}^2\}, \qquad P_c(D) = |\,\mathrm{Pat}_c(D)|.$$

Thus $P_c(D)$ counts the patterns seen by translating the observation window D. We use the rectangles

$$R_{k,n} = \{0,\dots,k-1\} \times \{0,\dots,n-1\}, \qquad P_c(k,n) = P_c(R_{k,n}).$$

A *period* of c is a nonzero vector $h \in \mathbb{Z}^2$ such that $c(g+h) = c(g)$ for all g. We call c *fully periodic* if it has two linearly independent periods, and *nonperiodic* if it has no period. The following is the classical Nivat statement, a two-dimensional analogue of the Morse–Hedlund theorem [9].

Theorem 1.1. *Let $\mathcal{A}$ be a finite alphabet and $c : \mathbb{Z}^2 \to \mathcal{A}$. If*

$$P_c(k,n) \le kn$$

for some positive integers k, n, then c has a period.

The threshold is sharp. A configuration equal to 1 at the origin and 0 elsewhere is nonperiodic and has exactly $kn+1$ patterns on $R_{k,n}$. The conclusion cannot be strengthened to full periodicity: if s is a finite-valued aperiodic sequence, then $c(x,y) = s(x)$ has a vertical period but no transverse period, while $P_c(1,n)$ is at most the alphabet size and hence at most n for sufficiently large n.

Cyr and Kra proved periodicity under the bound $P_c(k,n) \le kn/2$ [4, Theorem 1.2]. Their generating-set method was refined by Colle and Garibaldi to retain information from the alphabet size [2, Theorem 2.2]. Short rectangles admit the full threshold: Cyr and Kra treat short side at most three [5, Theorem 1.1], and Colle and Garibaldi treat short side at most four [3, Theorem 1.3].

A different approach, due to Kari and Szabados, associates a Laurent annihilator to a low-complexity configuration and obtains a product of directional differences [8, Lemma 5 and Corollary 13]. Szabados proves the conjecture when the resulting periodic decomposition has at most two components [10]. Kari and Moutot show that the orbit closure contains a periodic point, and obtain the conjecture for uniformly recurrent configurations [7, Theorem 5 and Corollary 8]. These conclusions do not remove the need to control the particular point to which a complexity estimate is applied. Our selection argument follows the periodic-region constructions of Colle [1, Sections 3–4]; their relation to the work of Colle and Garibaldi is described in Appendix A.

Even the single-defect example above has the zero configuration in its orbit closure: one can simply translate the defect to infinity. Thus a periodic limit, or agreement with a periodic background on a wedge, is not itself a contradiction. We need a point that retains nonperiodicity. On this point, the periodic region will be used to produce additional patterns, rather than to extend a period through the plane.

We use the shifts and differences

$$T^u c(g) = c(g+u), \qquad \Delta_u = T^u - 1,$$

and write $X_c = \overline{\{T^u c : u \in \mathbb{Z}^2\}}$ for the orbit closure in the product topology, with the finite alphabet discrete. A complex configuration is *finitary* if it has finite image. A *rational wedge* is a translate of a closed rational cone with nonempty interior and aperture less than π.

Theorem 1.2 (The periodic-wedge bound)**.** *Let $c : \mathbb{Z}^2 \to \mathbb{C}$ be finitary and nonperiodic. Suppose that*

$$Q_0 c = 0, \qquad Q_0 = \prod_{\ell=1}^{q} \Delta_{h_\ell},$$

where the nonzero integer vectors h_ℓ have pairwise distinct unoriented directions. If c agrees with a fully periodic configuration on a rational wedge, then

$$P_c(k,n) \geq kn + 1 \qquad (k,n \geq 1).$$

If the constant function and the kn coordinate functions were independent, they would give the desired bound. Annihilators can impose relations among them, so we look for nonlinear functions of the same pattern. Let ϕ be the common directional divisor of the annihilators. Applying all but one of its directional factors gives a singly periodic configuration whose agreement with the background spreads from the wedge to a half-plane. A transverse difference then vanishes on that half-plane, after a fixed inset.

With two distinct directions, we take the *pointwise* product of the resulting defects: for arrays $x, y : \mathbb{Z}^2 \to \mathbb{C}$, $(xy)(g) = x(g)y(g)$. Its support is the intersection of the two supports, hence is contained in a translated pointed cone. This rules out a nonzero Laurent annihilator provided the product itself is nonzero. There is a placement issue here: shifting arbitrary nonzero sites together could require shifts far larger than the rectangle. Short recurrences along the factor directions instead allow us to work inside the coordinate boxes of the factors of ϕ. The resulting product is the double difference of a quadratic local observable, whose fitting translates are independent modulo the affine coordinate functions.

The same factor boxes also control how many affine relations there can be. Write r, s for the coordinate widths of ϕ and D_R for the dimension of the affine relations on $R = R_{k,n}$. With $t_+ = \max\{t, 0\}$, we have

$$D_R \leq L_R := (k-r)_+(n-s)_+.$$

There are exactly L_R translates of the quadratic observable that fit in R. Adding them to the $kn + 1 - D_R$ dimensions supplied by the affine coordinate functions gives

$$P_c(k,n) \geq (kn + 1 - D_R) + L_R \geq kn + 1.$$

The details occupy Sections 2–4.

Proposition 1.3 (Selection of a nonperiodic point)**.** *Suppose that $c : \mathbb{Z}^2 \to \mathbb{Z}$ is finitary and nonperiodic and $P_c(k,n) \leq kn$. There exist a nonperiodic θ in the orbit closure X_c, a fully periodic configuration B, and a rational wedge K such that $\theta = B$ on $K \cap \mathbb{Z}^2$.*

We prove this in Sections 5–9, adapting the two cases in the proof of Colle's Lemma 4.6 [1]. The main difficulty in the compactness case is to retain nonperiodicity. We anchor maximal agreement polygons at a boundary point and keep a disagreement on a fixed finite set. A hypothetical period of the limit would then extend agreement to a reference half-plane, or give that half-plane a second period; either conclusion is ruled out by the construction. The agreement region spreads to that half-plane under the reference period. No recurrence hypothesis is needed.

Deduction of Theorem 1.1. Relabel the alphabet injectively by integers, preserving patterns and periods, and suppose for a contradiction that c is nonperiodic. If $n = 1$, every horizontal row has at most k words of length k. The Morse–Hedlund theorem [4, Theorem 2.2(iii)] gives each row a period at most k, so $\operatorname{lcm}(1, \ldots, k)$ is a common horizontal period. The case $k = 1$ follows in the same way from the columns.

We may therefore assume $k, n \geq 2$. Lemma 2.1 supplies a nonzero integral annihilator, and Theorem 2.2 supplies the product Q_0. This finite local identity holds on every point of X_c. Every

finite pattern of such a point also occurs in c, since its cylinder neighborhood meets the orbit. Apply Proposition 1.3 and then Theorem 1.2 to the selected nonperiodic point θ. We obtain

$$kn + 1 \le P_\theta(k, n) \le P_c(k, n) \le kn,$$

a contradiction. □

2 Annihilators and directional filters

We regard configurations as functions on $\mathbb{Z}^2$, rather than multiply them as formal Laurent series. Their values may be added and multiplied pointwise. In particular,

$$(xy)(g) = x(g)y(g), \qquad \operatorname{supp}(x) := \{g \in \mathbb{Z}^2 : x(g) \neq 0\}.$$

No summation is involved, even when both supports are infinite. Over $\mathbb{C}$,

$$\operatorname{supp}(xy) = \operatorname{supp}(x) \cap \operatorname{supp}(y), \qquad \operatorname{supp}(T^u x) = \operatorname{supp}(x) - u. \tag{2.1}$$

The second sign follows from our convention $T^u x(g) = x(g + u)$.

Laurent polynomials play a different role. Put $\mathscr{L} = \mathbb{C}[X^{\pm 1}, Y^{\pm 1}]$ and abbreviate $X^{h_1}Y^{h_2}$ by X^h. For $f = \sum_{u \in \operatorname{supp}(f)} f_u X^u \in \mathscr{L}$, define

$$f(T)x(g) = \sum_{u \in \operatorname{supp}(f)} f_u x(g + u), \qquad \operatorname{Ann}(x) = \{f \in \mathscr{L} : f(T)x = 0\}. \tag{2.2}$$

This sum is finite. Polynomial multiplication corresponds to composition, $(fg)(T)x = f(T)(g(T)x)$, so $\operatorname{Ann}(x)$ is an ideal. This is the annihilator formalism of Kari–Szabados [8, Section 2], with the shift convention reversed. In formal-series notation, our product of two configurations would be a coefficientwise (Hadamard) product, not a Cauchy product. No convolution of two infinitely supported arrays is used.

A shift preserves pointwise products: $T^u(xy) = (T^u x)(T^u y)$. A general operator $f(T)$ does not. We will expand the two differences of a product explicitly when this distinction matters. Here, *finitary* means finite image, not finite support. Finite sums and pointwise products of translates of a finitary configuration remain finitary.

For a polynomial f, the set $\operatorname{supp}(f)$ is the input stencil of $f(T)$; it need not contain the support of the output array $f(T)x$. Its Newton polygon is $\operatorname{Newt}(f) = \operatorname{conv}(\operatorname{supp}(f))$ for $f \neq 0$. The coordinate minima and maxima of nonzero Laurent polynomials add under multiplication: viewing a polynomial in X with coefficients in the integral domain $\mathbb{C}[Y^{\pm 1}]$ proves this for the two X-extrema, and interchanging the coordinates proves it for the other two. The *coordinate widths* are the differences between the corresponding maxima and minima. We call a nonconstant Laurent polynomial a *line polynomial* when its Newton polygon is a nondegenerate segment.

The initial linear-algebra argument is the one in Kari–Szabados [8, Lemma 5]. We include it to keep the equality case visible.

Lemma 2.1. *If c is finitary and integer-valued and $P_c(D) \le |D|$ for a nonempty finite set D, then c has a nonzero integral Laurent annihilator.*

Proof. Fix a D-pattern p_0. The at most $|D| - 1$ nonzero vectors $p - p_0$ have a nonzero orthogonal vector in $\mathbb{Q}^D$. Its coefficients define a nonzero polynomial f supported in D for which $f(T)c$ is constant. For any $w \in \mathbb{Z}^2 \setminus \{0\}$, the nonzero polynomial $(X^w - 1)f$ annihilates c. Clearing denominators gives an integral annihilator. □

Theorem 2.2 (Kari–Szabados, Corollary 13 of [8]). *A finitary integer configuration with a nonzero annihilator has an annihilator* $\prod_{\ell=1}^{q}(X^{h_\ell}-1)$ *with nonzero integer vectors in pairwise distinct directions.*

The source uses the opposite shift convention. Replacing every exponent by its negative gives the statement above, since reversing a period binomial changes it only by a Laurent unit. We use the product identity, not a bounded decomposition into periodic components.

The next proof uses the elimination argument of Kari–Moutot [6, Lemma 8]. We state the finite-image hypothesis explicitly: it is what turns the resulting recurrences into periods.

Lemma 2.3 (Coprime annihilators). *A finitary complex configuration with two coprime nonzero Laurent annihilators is fully periodic.*

Proof. Let f, g be the two annihilators. They are coprime in $\mathbb{C}(X)[Y^{\pm1}]$: a common nonunit factor there would give one in $\mathscr{L}$ after clearing denominators and taking primitive parts. A Bézout identity in this one-variable ring, with denominators cleared, gives a nonzero polynomial in X in (f, g). Interchanging the variables gives a nonzero polynomial in Y.

We remove the monomial factor from each eliminant. A nonconstant eliminant now gives a reversible recurrence of some order d along every parallel lattice line, since its constant and leading coefficients are nonzero. The length-d states occurring anywhere in the configuration form one finite set. Each state determines both its successor and its predecessor, so the successor map is a permutation, whose order is a period on all the parallel lines simultaneously. Applying this to the two eliminants gives independent coordinate periods. If either eliminant is a nonzero constant, the configuration is zero. □

For the remainder of Sections 2–4, assume the hypotheses of Theorem 1.2. An individual annihilator may have factors which another annihilator lacks. For the local count we need their common factors. We do not assume that their gcd annihilates c; the following lemma shows that applying it leaves only a fully periodic remainder (cf. [8, Corollaries 28–29]).

Lemma 2.4 (The common directional divisor). *The nonzero elements of* $\mathrm{Ann}(c)$ *have a gcd* ϕ*, defined up to a Laurent unit. It can be written*

$$\phi = \phi_1 \cdots \phi_m, \qquad m \geq 2,$$

where the ϕ_i *are nonconstant line polynomials in distinct directions, each dividing the corresponding binomial in* Q_0*. Choose their monomial representatives so that both coordinate minima of every* ϕ_i *are zero, and use their product as the representative of* ϕ*. Then* $\rho = \phi(T)c$ *is fully periodic. In addition,*

$$f(T)c \text{ constant} \quad \Longrightarrow \quad \phi \mid f.$$

Proof. Write $h = ap$ with $a \geq 1$ and p primitive. Then the factors of $X^h - 1$ over $\mathbb{C}$ are $X^p - \zeta$, with $\zeta^a = 1$, up to a unit. They are distinct. Binomials in different directions have no common nonunit factor. Thus Q_0 is squarefree.

Because $Q_0 \in \mathrm{Ann}(c)$, we may define ϕ as the product of the irreducible factors of Q_0 which divide every annihilator. This is a gcd even though the ideal contains infinitely many polynomials. For each factor omitted from ϕ, choose an annihilator not divisible by that factor. We can choose a finite linear combination g of these annihilators which avoids every omitted factor: each divisibility condition defines a proper linear subspace of the coefficient space, and a finite union of such subspaces does not cover a complex vector space. Thus $\gcd(Q_0, g) = \phi$. If no factor is omitted, take $g = Q_0$.

The coprime quotients Q_0/ϕ and g/ϕ annihilate $\rho = \phi(T)c$. This configuration is finitary, so Lemma 2.3 makes it fully periodic. Grouping the irreducible factors of ϕ by direction, we obtain the polynomials ϕ_i. Multiplying each by the prescribed monomial unit fixes its representative without changing the divisibility statements or periodicity.

There cannot be zero or one surviving direction. The first case makes c fully periodic. In the second, choose a nonzero multiple v of the vector in Q_0 which is also a period of ρ, and write $\Delta_v = \phi a$. The difference $\Delta_v c = a(T)\rho$ is v-periodic, so its increments along each v-line are constant: for $N \geq 0$,

$$c(g + Nv) = c(g) + N(a(T)\rho)(g).$$

Boundedness forces $a(T)\rho = 0$, making v a period of c. Both cases contradict nonperiodicity, so $m \geq 2$.

If $f(T)c$ is constant, choose $w \in \mathbb{Z}^2 \backslash \{0\}$ parallel to none of the directions of ϕ. Then $\Delta_w f \in \mathrm{Ann}(c)$ and $\gcd(\Delta_w, \phi) = 1$. Every irreducible factor of the squarefree polynomial ϕ therefore divides f, as required. □

Put

$$f_i = \phi/\phi_i, \qquad A_i = f_i(T)c.$$

The filter f_i applies every directional factor except ϕ_i. The phrase "direction" refers to a line, not a chosen period: a primitive vector spans all lattice points on that line, but need not itself be a period. Let B be the fully periodic background on the wedge. Choose a nonzero multiple v_i of the corresponding vector in Q_0 which is a period of both ρ and B. This is possible because both period lattices have finite index in $\mathbb{Z}^2$.

Lemma 2.5 (Directional filters)**.** *Each A_i has period v_i, is not fully periodic, and agrees with a fully periodic configuration on a rational lattice half-plane whose boundary is parallel to v_i.*

Proof. Write $\Delta_{v_i} = \phi_i a_i$. Then

$$\Delta_{v_i} A_i = a_i(T)\rho$$

is v_i-periodic. Applying the boundedness argument of Lemma 2.4 to A_i gives $\Delta_{v_i} A_i = 0$. A transverse period w would put $\Delta_w f_i$ in $\mathrm{Ann}(c)$, although it is coprime to ϕ_i. Thus A_i is not fully periodic.

A finite stencil only costs a translation of the wedge into its interior. After making this translation, $A_i = B_i$ on a smaller wedge K', where $B_i = f_i(T)B$. Their common period v_i spreads agreement to $(K' \cap \mathbb{Z}^2) + \mathbb{Z}v_i$. For a nonprimitive period this set can miss lattice points near its boundary; what we need, and will prove, is that it contains an inset lattice half-plane.

Write $K' = \{g : \lambda_1(g) \geq a_1,\ \lambda_2(g) \geq a_2\}$ with independent rational linear forms. A lattice point g belongs to this saturation whenever some integer t satisfies

$$\lambda_1(g) + t\lambda_1(v_i) \geq a_1, \qquad \lambda_2(g) + t\lambda_2(v_i) \geq a_2.$$

If $\lambda_1(v_i)$ and $\lambda_2(v_i)$ have the same strict sign, an integer t exists for every g. This would give $A_i = B_i$ globally, contrary to non-full-periodicity. If one coefficient vanishes, the other is nonzero, and the inequalities can be solved throughout the half-plane imposed by the inequality with zero coefficient of t.

In the remaining case, after reversing v_i if necessary, write $a = \lambda_1(v_i) > 0$ and $b = -\lambda_2(v_i) > 0$. The admissible interval is

$$\frac{a_1 - \lambda_1(g)}{a} \leq t \leq \frac{\lambda_2(g) - a_2}{b}.$$

Put $\nu = \lambda_1/a + \lambda_2/b$. This is a nonzero rational form with $\nu(v_i) = 0$. On the half-plane

$$\nu(g) \geq \frac{a_1}{a} + \frac{a_2}{b} + 1,$$

the interval has length at least 1 and contains an integer. This proves the assertion for nonprimitive periods as well; the inset avoids any missing lattice points near the boundary of the real saturation. □

3 The pointed-cone obstruction and a local observable

An array supported in a half-plane may still be periodic along its boundary. A pointwise product of two such arrays has a different support restriction: by (2.1), it vanishes outside the intersection of their half-planes. When the boundaries are nonparallel this intersection is a translated pointed cone. The product need not be nonzero; Lemma 3.2 will deal with that issue. We first record what the support restriction gives once nonvanishing is known.

Lemma 3.1 (The pointed-cone obstruction)**.** *A nonzero complex array supported in a translate of a closed pointed cone of aperture less than π has no nonzero Laurent annihilator.*

Proof. Choose a linear functional λ in the interior of the dual cone whose kernel contains no nonzero lattice vector; this requires avoiding only countably many lines. Its sublevel sets in the translated cone are bounded and contain finitely many lattice points. Hence the nonempty support of the array has a unique λ-minimum g_*. For a nonzero operator $f = \sum_u a_u T^u$, choose the unique λ-maximum u_* of its exponent support. In

$$f(T)e(g_* - u_*) = \sum_u a_u e(g_* - u_* + u),$$

all terms with $u \neq u_*$ sample points below the minimum g_* and hence vanish. The remaining term is $a_{u_*}e(g_*) \neq 0$. We have used a maximum for the operator and a minimum for the array because T^u samples at $g + u$. □

Let r_i, s_i be the coordinate widths of the normalized ϕ_i, and put $r = \sum_i r_i$, $s = \sum_i s_i$. Let $[E_i^-, E_i^+]$ be its Newton segment, of primitive lattice length $t_i \geq 1$. The support of f_i lies in

$$[0, r - r_i] \times [0, s - s_i].$$

Fix distinct indices i, j. Increase v_i, v_j by common multiples so that they fix ρ and all the fully periodic B_i in Lemma 2.5. Define the transverse defects

$$b_i = \Delta_{v_j} A_i, \qquad b_j = \Delta_{v_i} A_j.$$

Neither is zero, since otherwise the corresponding filtered configuration would be fully periodic. To locate the support, write an agreement half-plane for A_i, B_i as $\nu_i(g) \geq \beta_i$, with $\nu_i(v_i) = 0$. Since v_j is a period of B_i,

$$b_i(g) = 0 \qquad \text{if} \qquad \nu_i(g) \geq \beta_i + \max\{0, -\nu_i(v_j)\}.$$

Thus b_i is supported in the opposite closed half-plane. The same argument applies to b_j. Their support boundaries are parallel to v_i, v_j, respectively. Also

$$\phi_i(T)b_i = \Delta_{v_j}\rho = 0, \qquad \phi_j(T)b_j = \Delta_{v_i}\rho = 0. \tag{3.1}$$

The useful length scale is that of these recurrences, not the possibly large periods v_i, v_j. On a line where a defect is nonzero, a short recurrence prevents a long run of zeros. We can therefore move nonzero values together using only the factor boxes.

Lemma 3.2 (Placement in the factor boxes)**.** *There are integer vectors $u_i \in [0, r_i] \times [0, s_i]$ and $u_j \in [0, r_j] \times [0, s_j]$ for which*

$$e = (T^{u_i} b_i)(T^{u_j} b_j)$$

is nonzero.

Proof. Choose a real affine line L_i in the direction of ϕ_i containing a lattice point where b_i is nonzero, and choose L_j in the same way. The translated lines intersect at

$$z = (L_i - E_i^-) \cap (L_j - E_j^-).$$

Although z need not be integral, the segment $z + [E_i^-, E_i^+]$ lies on L_i and contains at least t_i consecutive points of $L_i \cap \mathbb{Z}^2$: an interval of length t_i in primitive lattice units contains at least t_i lattice points, regardless of its endpoints.

Translate the operator by $-E_i^-$. Equation (3.1) then gives a recurrence of order t_i on $L_i \cap \mathbb{Z}^2$ with nonzero first and last coefficients. A run of t_i zeros would propagate in both directions and make the entire line zero. We may therefore choose p_i on the segment with $b_i(p_i) \neq 0$, and similarly choose p_j for b_j.

Set $g = \lfloor z \rfloor$ coordinatewise and $u_i = p_i - g$, $u_j = p_j - g$. Since $p_i - z$ belongs to the integral coordinate box $[0, r_i] \times [0, s_i]$,

$$u_i = p_i - \lfloor z \rfloor = \lceil p_i - z \rceil$$

also belongs to that box. The same holds for j, and $e(g) = b_i(p_i) b_j(p_j) \neq 0$. □

If H_i, H_j are the half-planes containing the supports of b_i, b_j, then

$$\operatorname{supp}(e) = (\operatorname{supp}(b_i) - u_i) \cap (\operatorname{supp}(b_j) - u_j) \subset (H_i - u_i) \cap (H_j - u_j).$$

The containing half-planes, not the supports themselves, intersect in a translated pointed cone. Since $e \neq 0$, Lemma 3.1 applies. We do not use e directly as the local observable, since its definition involves the large shifts v_i, v_j. Instead, form the pointwise product before taking those differences:

$$d = (T^{u_i} A_i)(T^{u_j} A_j).$$

Its input sites, relative to the evaluation site, belong to

$$(u_i + \operatorname{supp}(f_i)) \cup (u_j + \operatorname{supp}(f_j)) \subset W := [0, r] \times [0, s] \cap \mathbb{Z}^2. \tag{3.2}$$

Thus $d(g)$ is a quadratic function of the c-values on $g + W$. This is a statement about its input stencil, not a claim that the output array d has finite support.

To recover e, write $U(g) = A_i(g + u_i)$ and $V(g) = A_j(g + u_j)$, so that $d(g) = U(g)V(g)$. Since $T^{v_i} U = U$ and $T^{v_j} V = V$,

$$\begin{aligned} \Delta_{v_i} \Delta_{v_j} (UV) &= (T^{v_j} U)(T^{v_i} V) - U(T^{v_i} V) - (T^{v_j} U)V + UV \\ &= (\Delta_{v_j} U)(\Delta_{v_i} V) = e. \end{aligned}$$

Set $Q = \Delta_{v_i} \Delta_{v_j} Q_0$. This operator annihilates both c and constants, while

$$Qd = Q_0 e \tag{3.3}$$

has no nonzero annihilator. Indeed, an annihilator H of $Q_0 e$ would make the nonzero polynomial HQ_0 annihilate e.

4 The exact rectangle count

Kari–Szabados relate linear independence of patterns to the absence of annihilators fitting in a window [8, Lemma 36]. We keep the constant coordinate as well, rather than first normalizing the configuration as in [8, Lemma 33]. The extra step is to supplement this affine space with translates of the quadratic function just constructed.

Proof of Theorem 1.2. Fix $R = R_{k,n}$ and let

$$F_R = \{f \in \mathscr{L} : \operatorname{supp}(f) \subset R,\ f(T)c \text{ is constant}\}, \qquad D_R = \dim F_R.$$

Lemma 2.4 makes every member of F_R divisible by ϕ. Since the latter has coordinate minima zero and widths r, s, additivity of coordinate extrema confines every nonzero quotient f/ϕ to

$$[0, k-r-1] \times [0, n-s-1].$$

Multiplication by ϕ is injective, so

$$D_R \le L_R := (k-r)_+(n-s)_+. \tag{4.1}$$

If $k \le r$ or $n \le s$, the quotient box has no lattice sites and $D_R = 0$.

Let V_R be the subspace of all complex functions on $\mathrm{Pat}_c(R)$ spanned by 1 and the coordinate evaluations $p \mapsto p(u)$, $u \in R$. The map from their $kn+1$ coefficients to this function space has kernel

$$\{(-\kappa(f), f) : f \in F_R\}, \qquad f(T)c = \kappa(f).$$

Since $\kappa(f)$ is uniquely determined by f, including for $f = 0$, we obtain

$$\dim V_R = kn + 1 - D_R. \tag{4.2}$$

If $L_R = 0$, this already proves the theorem. For integral configurations, this thin-rectangle conclusion also appears in Kari–Szabados [8, Lemma 39].

Suppose $L_R > 0$. Write $f_i = \sum_a \alpha_{i,a} X^a$ and $f_j = \sum_b \alpha_{j,b} X^b$, with both sums restricted to the finite polynomial supports. The function on W-patterns which gives d is

$$\mathcal{D}(p) = \left(\sum_a \alpha_{i,a} p(u_i + a)\right)\left(\sum_b \alpha_{j,b} p(u_j + b)\right).$$

In other words, $d(g) = \mathcal{D}((c(g+z))_{z\in W})$. This also makes the word "local" precise. For each t with $t + W \subset R$, the function $p \mapsto \mathcal{D}((p(t+z))_{z\in W})$ is defined on $\mathrm{Pat}_c(R)$; evaluating it at the pattern based at g gives $T^t d(g)$. There are exactly L_R such translations.

We claim that these functions are independent modulo V_R. A relation on the finite set $\mathrm{Pat}_c(R)$, evaluated at every translate of c, gives the global array identity

$$\sum_t \alpha_t T^t d = \alpha + f(T)c, \qquad \operatorname{supp}(f) \subset R.$$

The identity holds at every $g \in \mathbb{Z}^2$, so it remains valid at all sites sampled by Q. Applying Q removes the affine terms and gives

$$\Big(\sum_t \alpha_t T^t\Big)(Qd) = 0.$$

By (3.3), all α_t vanish. These L_R functions are thus independent modulo V_R. The full function space on the finite set $\mathrm{Pat}_c(R)$ has dimension $P_c(R)$, and consequently

$$P_c(k,n) \geq \dim V_R + L_R = kn + 1 - D_R + L_R \geq kn + 1.$$

The distinction between defining the functions and testing a relation among them is worth keeping in view. Each function counted above uses only the R-pattern; Q acts afterwards on the global identity, so its large shifts do not enlarge the observation window. □

5 Generating sets and a periodic reference

For Sections 5–9, fix a nonperiodic finitary integer configuration c with $P_c(k,n) \leq kn$ and a product annihilator Q_0 as in Theorem 2.2. Since c is nonperiodic, Q_0 has at least two factor directions. Every point of X_c satisfies $Q_0 z = 0$ and uses only finite patterns occurring in c. The reference we construct below will have one period, but its restriction to a chosen half-plane must not be fully periodic. This distinction will later exclude a transverse period of the selected point.

A finite set is *lattice-convex* if it consists of all lattice points of its convex hull. A vertex v of a lattice-convex set S is *generated* if any two occurring S-patterns agreeing on $S \setminus \{v\}$ also agree at v. We call S a *generating set* when every vertex is generated, and an array *S-legal* when each of its translated S-patterns occurs in c. An edge e in a pattern count denotes its lattice sites; $\mathrm{conv}(e)$ denotes the corresponding geometric segment. We use the generating-set construction of Cyr–Kra [4, Lemmas 2.5 and 2.8], in the extremal form stated by Colle [1, Section 2]. The short discrepancy proof below records the edge bound needed later.

Lemma 5.1 (A generating polygon)**.** *There is a lattice-convex $S \subset R_{k,n}$ of positive area such that every vertex is generated and, for every edge e of S,*

$$P_c(S) - P_c(S \setminus e) \leq |e| - 1. \tag{5.1}$$

Proof. Minimize the discrepancy $P_c(S) - |S|$ over nonempty lattice-convex subsets of $R_{k,n}$, and choose an inclusion-minimal minimizer. Its discrepancy is nonpositive, since $R_{k,n}$ is a competitor.

If S were a segment of t consecutive lattice sites, every parallel bi-infinite line would have at most t words of length t. Morse–Hedlund would give each line a period at most t, and their common multiple would be a global period of c. This also excludes a singleton, so S has positive area.

Every proper nonempty lattice-convex subset of S has larger discrepancy. Deleting a vertex v therefore gives $P_c(S) - P_c(S \setminus \{v\}) < 1$. This difference is a nonnegative integer, since restriction of patterns is surjective, so it is zero and v is generated. Deleting an edge also leaves a nonempty lattice-convex set, and the same comparison gives (5.1). □

For an edge e, an occurring $S \setminus e$-pattern is *ambiguous* if it has at least two occurring S-extensions. If $N(\gamma)$ is the number of such extensions, then

$$\sum_{\gamma} (N(\gamma) - 1) = P_c(S) - P_c(S \setminus e) \leq |e| - 1. \tag{5.2}$$

In particular, there are at most $|e| - 1$ ambiguous base patterns.

Following Colle [1, Section 1], an oriented line through the origin is *one-sided nonexpansive* for X_c if two distinct points of X_c agree on the closed lattice half-plane on its left. The two orientations specify the two sides of the line. The next result is the contrapositive of Colle's theorem.

Theorem 5.2 (Colle, Theorem 1.9 of [1]). *If a finitary integer configuration with a nonzero annihilator is not fully periodic, then some line has both orientations one-sided nonexpansive on its orbit closure.*

Both supporting faces of the generating polygon in this direction must be edges. Indeed, suppose the face were a single generated vertex v, and let λ increase into the half-plane. The positive minimum of $\lambda(u - v)$ over $u \in S \setminus \{v\}$ gives a slab thickness through which agreement propagates. Iterating would determine the entire plane, contradicting one-sided nonexpansiveness. In particular the line is rational. Applying the same argument to the nonzero extreme coefficient of Q_0 shows that the line is parallel to a factor of Q_0: otherwise its Newton polygon has a unique minimizing vertex.

A nonzero vector h is a *regional period* of x on U if $x(g + h) = x(g)$ whenever $g, g + h \in U$. The restriction is *fully periodic* if it has two linearly independent regional periods (compare [1, Definition 2.12]). If $U + h \subset U$, a regional period h is equivalently expressed by $\Delta_h x = 0$ on U. We use this equivalence only when the indicated translation stays in the region.

For a tangent half-plane period, we can work directly with its difference. Taking a suitable multiple of the period cancels any parallel factor of the annihilator; the remaining factors then determine successive normal levels, extending the period to the plane (cf. [1, Proposition 2.14]).

Lemma 5.3 (Extending a tangent period). *Suppose $Q_0 z = 0$ and z has a period parallel to the boundary of a rational half-plane on that half-plane. Then z has a global period in the same direction.*

Proof. If a factor Δ_{h_j} is parallel to the boundary, choose a nonzero common multiple h of h_j and the regional period, and write $\Delta_h = a\Delta_{h_j}$. Then

$$E = \Delta_h z, \qquad \left(Q_0/\Delta_{h_j}\right)E = aQ_0 z = 0.$$

If there is no parallel factor, take the given period difference as E and retain Q_0. In either case E vanishes on the half-plane and has a product annihilator whose directions are all transverse to the boundary.

Choose a primitive integer normal ν pointing into the half-plane, and reverse the remaining factor vectors so that their ν-values are positive. This changes the product only by a Laurent unit. Its constant coefficient is nonzero, while every other exponent has positive integer ν-value. The recurrence therefore propagates zero one outward lattice level at a time, giving $E = 0$ on the plane. If the remaining product is empty, the same conclusion is immediate. □

The following construction adapts the ambiguity argument of Cyr–Kra [4, Lemma 2.24], as developed in Colle [1, Proposition 2.13].

Lemma 5.4 (A periodic reference). *There are a primitive integer vector a, an integer half-plane H with boundary parallel to a, and a globally a-direction-periodic $w \in X_c$ such that $w|_H$ is not fully periodic. The generating polygon S has two opposite edges parallel to a, and this direction occurs in Q_0.*

Proof. Choose the line supplied by Theorem 5.2, and orient its primitive vector a so that the facing edge e of S is no longer than the opposite edge. Set $r = |e| - 1 \geq 1$. Subtracting r primitive units from these two opposite edges leaves nonnegative cyclic edge vectors with sum zero. Hence $[0, ra]$ is a Minkowski summand of $\operatorname{conv} S$, and every parallel lattice row through S contains at least r consecutive sites.

Choose $x, y \in X_c$ agreeing on the associated half-plane and differing on the first outer lattice line where agreement fails; translate the pair so that this is the next line. The first such line exists

because the normal levels are integral. Place e on it. Every a-translate of the common $S \setminus e$-pattern is ambiguous: a unique extension at one position would give agreement on an edge block, and generation of the two endpoints would extend this agreement along the entire line.

Let d be the width of S in primitive normal levels. On each of the d rows of $(S \setminus e) + \mathbb{Z}a$, choose the initial site of a run of r sites in $S \setminus e$. Read these initial sites as one tuple and translate the tuple along a. Every length-r word of the resulting bi-infinite sequence is a restriction of an ambiguous base pattern. There are at most r such patterns by (5.2), so Morse–Hedlund gives the whole strip a common period pa, with $p \le r$.

The period can change as we add an outer row. It is enough to keep its length at most $2r$ on each strip of d rows, since we can take a common multiple afterwards. A unique next-line extension among S-legal configurations preserves the current period. Otherwise every translated base is ambiguous by the endpoint argument, and the tuple argument gives a period $p_0 a$ with $p_0 \le r$. The number of full S-patterns extending the observed bases is then at most

$$\sum_{\gamma \text{ observed}} N(\gamma) \le r + \#\{\gamma \text{ observed}\} \le 2r.$$

Among $2r+1$ successive full S-patterns of x, two therefore agree at a gap $q \in \{1, \dots, 2r\}$. The agreement contains r consecutive sites on every old row. Since these rows are $p_0 a$-periodic and $p_0 \le r$, qa is a period on the old strip. Generation of the endpoints of e propagates the equality in both directions along the new line, so the extended strip has period qa.

Iterating gives each row in the outward half-plane a period of length at most $2r$. Thus $\operatorname{lcm}(1, \dots, 2r)a$ is a common period there. The argument applies to both x and y, and Lemma 5.3 makes each globally periodic in direction a.

Let H be the original agreement half-plane enlarged to include the line where x and y differ. At least one of $x|_H, y|_H$ is not fully periodic. To see this, a fully periodic restriction on a half-plane extends to a fully periodic array: choose an inward regional period and transport each point far enough into the half-plane. The values stabilize, and both regional periods become global. If both restrictions had such extensions, the extensions would agree on the smaller half-plane. An inward vector common to their two finite-index period lattices would then give agreement on the plane, contradicting the differing boundary line. Choose w to be a member whose restriction to H is not fully periodic, and retain this choice throughout the selection argument. □

6 Boundary continuation in polygons

Let $\mathcal{Z} = \operatorname{Newt}(Q_0)$. Then $\mathcal{Z}$ is the Minkowski sum of the segments $[0, h_\ell]$, hence a lattice zonotope. Interior subset sums may cancel in Q_0; the Newton polygon identity is unaffected. Reversing a factor translates $\mathcal{Z}$.

A finite lattice polygon is *$\mathcal{Z}$-enveloped* if it has the same oriented edge directions as $\mathcal{Z}$, with each primitive edge length at least the corresponding length of $\mathcal{Z}$. Subtracting the cyclic edge-vector lists writes it as $\mathcal{Z} + C$ for a compact convex polygon C, possibly a segment or a point. A *two-ray region* is a rational convex polyhedron with pointed two-dimensional recession cone; its boundary consists of two rays joined by a finite chain of edges. In the continuation arguments below, these regions have the form $\mathcal{Z} + C$ with C a rational polyhedron. Here $A + B = \{a + b : a \in A,\ b \in B\}$ denotes a Minkowski sum, and $\operatorname{rec}(C) = \{v : C + tv \subset C \text{ for all } t \ge 0\}$ is the recession cone. A facet is an exposed one-dimensional face, which may be a finite edge or an infinite ray. Restrictions of arrays to real regions always mean restrictions to their lattice points. The enveloped regions follow Colle's geometric construction [1, Definition 3.3]; we retain the Minkowski summand because it guarantees that the recurrence stencil fits at the boundary.

Once the interior values vanish, the recurrence on a new layer involves only an exposed face of the Newton polygon. It becomes a one-dimensional recurrence along that layer. We will use this to control a growing boundary from a fixed finite seed.

Lemma 6.1 (Continuation across a facet)**.** *Let f be a nonzero Laurent polynomial with Newton polygon $\mathcal{Z}$, and let $f(T)\delta = 0$. Suppose $P^+ = \mathcal{Z} + C$ has a facet on $\nu = \beta$, with primitive integer inward normal ν and integer β, and $\delta = 0$ on*

$$P = P^+ \cap \{\nu \geq \beta + 1\}.$$

Orient the facet by a primitive integer vector b, pointing towards infinity when the facet is a ray. Write the corresponding exposed face of $\mathcal{Z}$ as $[q_0, q_1]$, with $q_1 - q_0 = tb$ and $t \geq 0$. The case $t = 0$ means that the exposed face is a single vertex.

If the facet of P^+ is finite and δ vanishes on its first t consecutive lattice points in direction b, it vanishes on the entire facet. If the facet is a ray in direction b, vanishing on a tail of the ray implies vanishing on the entire ray. When $t = 0$, no boundary seed is needed in either assertion.

Proof. Let F_C be the corresponding exposed face of C. If $t = 0$, additivity of faces under Minkowski summation gives $x - q_0 \in F_C$ for every facet point x. In the recurrence at $x - q_0$, only the nonzero extreme coefficient times $\delta(x)$ lies on the face; all other inputs are in P. Thus $\delta(x) = 0$ without a seed. Assume henceforth that $t \geq 1$.

On a finite facet, a point x at least t primitive units from its initial endpoint satisfies $x - q_1 \in F_C$. Thus

$$x - q_1 + \mathcal{Z} \subset P^+.$$

In the recurrence at $x - q_1$, every nonface input lies in P, its normal level having increased by a positive integer. The other face inputs precede x along the facet, so the nonzero coefficient at q_1 determines $\delta(x)$. Starting with the first t zeros proves the assertion. This also covers rational endpoints: each subsequent lattice point is at least t primitive units from the initial endpoint.

For a ray, the exposed face of P^+ is $q_0 + F_C$. The recurrence at $x - q_0$ is therefore valid for every lattice point x on the ray, with all nonface inputs in P. Its first and last face coefficients are nonzero. An all-zero tail extends backwards along the ray by solving for the term at q_0. □

To expand a region outwards, we relax a signed supporting inequality. For an infinite facet on the integer level $\nu = \beta$ of a primitive inward normal, let $\mathcal{R}^{(j)}$ be obtained by replacing that inequality with $\nu \geq \beta - j$ and keeping all others. Thus $\mathcal{R}^{(0)} = \mathcal{R} \subset \mathcal{R}^{(1)} \subset \cdots$. Moving the infinite facet lengthens the neighboring finite edge and retains the $\mathcal{Z}$ summand. For a finite polygon, any fixed number of such layers also retains the summand once the moving edge is sufficiently long: its neighbors lengthen, while its own length changes by a bounded amount.

Lemma 6.2 (A failed extension is ambiguous)**.** *Let $x, y \in X_c$ agree on a two-ray region $\mathcal{R}$ with Minkowski summand $\mathcal{Z}$, but not on its next outward lattice layer $\mathcal{R}^{(1)}$ along the infinite edge in primitive direction b. Assume the moving facet is on an integer level of its primitive inward normal. Then S has a supporting edge e_b in that direction. After translating S so that e_b lies on the new layer, every sufficiently far positive b-translate of the common $S \setminus e_b$-pattern is ambiguous.*

Proof. Place the minimizing face of S on the new lattice line. Every other irredundant inward normal of $\mathcal{R}$ is strictly positive on b, so after a sufficiently large positive b-translation all remaining sites of S lie in $\mathcal{R}$. They remain there under further positive translations.

A singleton supporting face would give agreement on a tail of the new line by generation of that vertex. For an edge, a unique extension of one translated base pattern would give agreement on an

edge block; generation of its final endpoint would then propagate agreement along the remaining tail. In either case, the ray part of Lemma 6.1, applied to $x - y$, would extend agreement to the entire new layer. This contradiction proves that the face is an edge and every sufficiently far base pattern is ambiguous. □

Two periods pointing along the recession rays of a two-ray region $\mathcal{R}$ determine one fully periodic background. If these periods are h_1, h_2, then $t = h_1 + h_2$ lies in the interior of the recession cone, and we may define

$$B(g) = x(g + Nt) \quad \text{for all sufficiently large } N. \tag{6.1}$$

The value stabilizes: an h_1 step followed by an h_2 step stays in the region and preserves the value. For either i, move both g and $g + h_i$ into $\mathcal{R}$ by the same sufficiently large multiple of t. The regional period gives $B(g + h_i) = B(g)$, so B is fully periodic. For $g \in \mathcal{R}$, the same steps show $B(g) = x(g)$. Since $\mathcal{R}$ contains a translate of its recession cone, agreement holds on a rational wedge.

7 From one regional period to two

Following Colle [1, Lemma 4.2], we use the existing period and generated vertices to extend equalities from an eventually periodic transverse half-strip into the widening region. A downward sweep requires the same transverse period on arbitrarily high starting bands; repetitions of one finite pattern provide this choice.

Lemma 7.1 (A second regional period)**.** *Let a, b be primitive integer vectors spanning a pointed cone of positive aperture, and let $\mathcal{R}$ be a rational two-ray polyhedron with recession cone $\mathbb{R}_{\geq 0}a + \mathbb{R}_{\geq 0}b$. Suppose $x \in X_c$ has a period pa, $p \geq 1$, on $\mathcal{R}$. Let S be a lattice-convex generating set for the language of c, with opposite edges parallel to a and a supporting edge e_b parallel to b. Assume*

$$P_c(S) - P_c(S \setminus e_b) \leq |e_b| - 1,$$

and that the inward normal of e_b is strictly positive on a. Suppose S can be placed so that every sufficiently far b-translate of $S \setminus e_b$ lies in $\mathcal{R}$ and is ambiguous. Then x has two independent periods, parallel to a, b, on a smaller two-ray region with the same recession cone.

Proof. Let $r \geq 1$ be the smaller primitive length of the two a-parallel edges. As in Lemma 5.4, the segment $[0, ra]$ is a Minkowski summand of $\operatorname{conv} S$. Since the inward normal of e_b is positive on a, the parallelogram

$$J = \operatorname{conv}(e_b) + [0, ra]$$

lies in $\operatorname{conv} S$. Put $t = |e_b| - 1$. The existence of an ambiguous base pattern and the stated edge bound give $t \geq 1$.

Each b-parallel lattice line through J, other than the line of e_b, contains at least t consecutive sites of $S \setminus e_b$. Along its sufficiently far positive ray, the length-t words are restrictions of the ambiguous base patterns. There are at most t such patterns by (5.2), so the one-sided Morse–Hedlund theorem gives an eventually periodic tail with period at most t [4, Theorem 2.2(ii)]. Taking a common multiple over the finitely many lines through J, and discarding their initial segments, gives a half-strip $H_0 \subset \mathcal{R}$ with period $h_b = lb$, $l \geq 1$.

To start the row extensions, we need consecutive sites on the a-rows, not just on the b-lines. Choose a primitive integer normal ν with $\nu(a) = 0$ and $\nu(b) > 0$, and let g_0 be the initial endpoint of e_b. Removing the line of e_b from J and extending in the positive b-direction gives

$$J^+ = \{g_0 + \alpha a + \beta b : 0 < \alpha \leq r,\ \beta \geq 0\}.$$

Every sufficiently high ν-section of J^+ is a semi-open interval of length r in primitive a-units, so it contains r consecutive lattice points. Each such point lies on a b-line whose intersection with J has primitive length $t \geq 1$ and hence contains a lattice point. The difference of the two points is an integer multiple of the primitive vector b. At sufficiently high levels this multiple is positive and exceeds the corresponding eventual-periodicity threshold. There are only finitely many b-lines through J, so a single threshold j_0 ensures that $H_0 \cap \{\nu = j\}$ contains r consecutive a-sites whenever $j \geq j_0$. No assumption that $|\det(a, b)| = 1$ is needed.

The period pa transports the same h_b-equalities to $H_0 + Mpa$ for every integer $M \geq 0$. These translates supply r consecutive equality sites arbitrarily far along each a-row above j_0. The same sites give equality for every positive multiple of h_b.

Let d be the ν-width of S. Choose a finite lattice parallelogram $F \subset \mathcal{R}$ containing p consecutive a-sites on each of $d + 1$ successive levels above j_0. Complete a to a unimodular basis to construct this finite set, then place it sufficiently far inside $\mathcal{R}$. This auxiliary choice of coordinates does not change the observation rectangle. Every $F + qh_b$, $q \geq 0$, remains in $\mathcal{R}$.

A repeated F-pattern at indices $q < q'$ gives agreement between x and $T^{(q'-q)h_b}x$ on $F + qh_b$. On each of its rows, both arrays are pa-periodic on the row half-line in $\mathcal{R}$. Agreement at p consecutive sites therefore gives agreement on the whole row, producing a band of $d + 1$ rows.

Suppose first that the lower a-parallel edge has primitive length r. We extend such a band downwards to level j_0. On the next row, choose an r-site equality seed far enough along a translate $H_0 + Mpa$. Place the lower edge of S there, with all remaining sites in the known band. Generation of the final endpoint gives the next equality on the row. Iterating along a gives a tail of equalities, and the common period pa extends them to the whole row half-line. The same argument applies to each successive lower row.

A different gap for each band would not suffice. Let N be the number of possible F-patterns. Every disjoint index block $[m(N + 1), m(N + 1) + N]$ contains a repeat with gap in $\{1, \ldots, N\}$, so one gap D occurs in infinitely many blocks. For this fixed D, the corresponding bands have heights tending to infinity, since $\nu(h_b) > 0$. Their downward sweeps cover every row above j_0 and all concern the same difference $\Delta_{Dh_b}x$. Hence Dh_b is a period on $\mathcal{R} \cap \{\nu \geq j_0\}$.

If the upper a-parallel edge has the smaller length r, one repeated band suffices. Using that edge, the same generated-endpoint argument extends agreement upwards through every row above the band. Its fixed gap again gives a period Dh_b on a smaller two-ray region. In both cases pa remains a period, and the recession cone is unchanged. □

8 Half-strips and anchored compactness

A *half-strip* in a primitive direction a, based on a finite polygon F, is $(F + \mathbb{R}_{\geq 0}a) \cap \mathbb{Z}^2$. If F is $\mathcal{Z}$-enveloped and a is a factor direction, this is also the union of the nonnegative integer a-translates of $F \cap \mathbb{Z}^2$. Indeed, every parallel lattice fiber of F has length at least one primitive unit and contains a lattice site.

Lemma 8.1 (A periodic half-strip supplies a wedge)**.** *Assume that Q_0 has at least two factor directions, and let a be primitive in one of them. Suppose $Q_0 z = 0$ and z agrees with a globally a-direction-periodic array on a forward a-half-strip based on a $\mathcal{Z}$-enveloped polygon. Then z has a period h in direction a on a translated region $\mathcal{Z} + \mathbb{R}_{\geq 0}a + \mathbb{R}_{\geq 0}b$, where b is the next factor direction on one side of a.*

Proof. Choose a common multiple h of the reference period and the parallel factor vector of Q_0, with its sign pointing along the forward half-strip. The difference $E = \Delta_h z$ vanishes on that half-strip.

As in Lemma 5.3, cancelling the parallel factor gives $Q_1 E = 0$, where Q_1 is the nonempty product of the remaining factors. We use this recurrence to extend the zero set into an adjacent wedge.

Complete a to a unimodular basis and use integral coordinates (τ, ν) with $\tau(a) = 1$ and $\nu(a) = 0$. Reverse each remaining factor vector p_j so that $\nu(p_j) > 0$. Let

$$d = \sum_j \nu(p_j), \qquad C = \max_j \frac{-\tau(p_j)}{\nu(p_j)}.$$

After multiplication by a Laurent unit, Q_1 has nonzero constant coefficient. Every other exponent u satisfies

$$1 \le \nu(u) \le d, \qquad \tau(u) + C\nu(u) \ge 0.$$

These inequalities hold for every subset sum of the factor vectors, and thus for the actual support even when coefficients cancel.

For some integers M, s, the half-strip contains $\{\tau \ge M,\ s \le \nu \le s + d - 1\}$. To see this, take a translate of $\mathcal{Z}$ in its base polygon. This zonotope has ν-width d, and its factor segment parallel to a gives a lattice site on each integer level in the indicated range. Forward a-translation then fills a tail on each of these finitely many levels, allowing a common threshold M. Put

$$A = M + \max_{s \le j \le s+d-1} Cj.$$

The recurrence of Q_1 propagates $E = 0$ downwards through all integer ν-levels in

$$W = \{\nu \le s + d - 1,\ \tau + C\nu \ge A\}.$$

The initial d levels of W lie in $\{\tau \ge M\}$ by the choice of A. At any lower level, every nonconstant recurrence input has a larger ν-value, at most $s + d - 1$, and no smaller value of $\tau + C\nu$. All these inputs are already zero, so the nonzero constant coefficient determines the next level. Induction gives $E = 0$ throughout W.

The recession rays of W are a and a primitive vector b antiparallel to a factor vector attaining the maximum in the definition of C. They are nonparallel. By this extremal choice of slope, no factor direction lies strictly between them. An integer translate of $\mathcal{Z}$ sufficiently far inside W remains in W after either recession ray is added. In particular, $\Delta_h z = 0$ on the resulting region, as required. □

Maximality gives a disagreement somewhere outside each agreement polygon. That site may move with the polygon, and then compactness need not preserve the disagreement. We anchor a corner before taking a limit and use the facet recurrence to locate a witness in a fixed finite set near that corner.

Lemma 8.2 (Anchored agreement with a retained boundary failure)**.** *Let a be primitive in a factor direction of Q_0, let H be a lattice half-plane whose boundary is oriented by a with H on its left, and let $w \in X_c$ have a period in direction a. Suppose that no shift of c agrees with w on any forward a-half-strip based on a $\mathcal{Z}$-enveloped polygon. Then there are a fixed phase $\widehat{w}$ of w, a point $\theta \in X_c$, and a two-ray region $\mathcal{R} \subset H$ with Minkowski summand $\mathcal{Z}$ such that:*

(i) its recession rays are $-a$ and a nonparallel primitive factor direction b;

(ii) $\theta = \widehat{w}$ on $\mathcal{R}$, but not on its next outward lattice layer $\mathcal{R}^{(1)}$ along the b-edge;

(iii) $\mathcal{R}^{(1)} \subset H$, and every finite subset of H can be moved into $\mathcal{R}$ by a sufficiently large negative multiple of a.

Proof. Maximal agreement polygons. After completing a to an integral basis, call the forward a-direction rightwards; we use these auxiliary coordinates only to describe the polygons. Fix a lattice point g_1 on the boundary of H. Inductively choose a $\mathcal{Z}$-enveloped lattice polygon F_i, supported on this boundary, which contains M_{i-1} and $[-i, i]^2 \cap H$. For the first step, omit M_{i-1}. Such polygons are obtained by sufficiently large integral dilations of $\mathcal{Z}$, translated along their nondegenerate a-edge. Since $w \in X_c$, choose u_i so that $T^{u_i} c = w$ on F_i. Among the $\mathcal{Z}$-enveloped polygons between F_i and its forward half-strip on which these arrays agree, choose an inclusion-maximal polygon M_i.

To justify this choice, the half-strip fixes the transverse width. Along the finitely many nonhorizontal edge directions, the longitudinal change is bounded by that width times a constant depending only on the slopes. Thus if any right endpoint were unbounded, all right endpoints would tend to infinity. The left boundary is fixed by F_i, while some lattice site of the half-strip is a disagreement by hypothesis. This bounds all right endpoints, so there are only finitely many admissible lattice polygons. In particular M_i exists, and

$$M_{i-1} \subset F_i \subset M_i.$$

Anchoring and convergence. Translate the final endpoint of the oriented a-edge of M_i to g_1, writing the translation as $-k_i a$. Set

$$P_i = M_i - k_i a, \qquad \theta_i = T^{u_i + k_i a} c.$$

Then $\theta_i = T^{k_i a} w$ on P_i. The reference has a period parallel to a, so after passing to a subsequence $T^{k_i a} w$ is one fixed array $\widehat{w}$.

The different anchor translations need not preserve the nesting of the M_i. Pass to a further subsequence on which every primitive edge length is nondecreasing. This is possible for a sequence in $\mathbb{N}^{2q}$: in each coordinate, stabilize a bounded infinite subsequence if there is one, and otherwise take a strictly increasing subsequence. Now the anchored P_i are nested. The differences of successive cyclic edge-vector lists have nonnegative lengths and sum to zero, so they bound a convex polygon, possibly a segment or point, anchored at the origin. Minkowski addition adds both these lists and the anchored vertices. Each later P_i is therefore the earlier polygon plus this residual polygon, which contains the origin.

The a-edge lengths and the transverse widths tend to infinity. In the cyclic list following the a-edge, let b be the first direction with unbounded length before the opposite a-edge. Since the widths diverge, such a direction exists; all intervening edge lengths stabilize. The increasing union of the P_i is a rational two-ray polyhedron R_∞ with recession rays $-a, b$ and a fixed finite boundary chain in H.

More explicitly, the supporting lines of that finite chain stabilize. Every other supporting normal is negative on at least one of $-a, b$, so its support level tends to minus infinity. Precisely the

finite-chain inequalities remain, and every point satisfying them eventually lies in P_i. Thus the union is the resulting closed polyhedron. Subtracting $\mathcal{Z}$ from the edge lists, or taking the increasing union of the corresponding summands, also shows that R_∞ retains a Minkowski summand $\mathcal{Z}$.

Passing to a subsequence, we may assume that $\theta_i \to \theta$. Every fixed finite part of R_∞ eventually lies in P_i, so $\theta = \widehat{w}$ there. Moreover $\theta \in X_c$, and hence $Q_0\theta = 0$ and every finite pattern of θ occurs in c.

A fixed boundary witness. Write $\nu_b = \beta$ for the infinite b-facet of R_∞, with primitive integer inward normal. For large i, the b-facet of P_i lies on this same line. Let $P_i^{(j)}$ be obtained by replacing its inequality with $\nu_b \geq \beta - j$ and retaining all other inequalities. For each fixed j, these polygons retain a $\mathcal{Z}$ summand for large i: the moving edge has length tending to infinity, its neighbors lengthen, and its own length changes by a bounded amount.

An expansion by one lattice level need not leave lattice vertices. Let μ_-, μ_+ be the primitive inward normals of the neighboring facets, and choose $m \geq 1$ divisible by the fixed nonzero integers $|\det(\mu_-, \nu_b)|$ and $|\det(\mu_+, \nu_b)|$. To see why this choice works, an endpoint v incident to the neighbor with normal μ moves according to

$$\begin{pmatrix} \mu \\ \nu_b \end{pmatrix} (v^{(j)} - v) = \begin{pmatrix} 0 \\ -j \end{pmatrix}.$$

By Cramer's rule, the displacement is integral at $j = m$. All other vertices are unchanged, so $P_i^{(m)}$ is a lattice polygon. Intermediate expansions may have rational vertices, as allowed in Lemma 6.1. For large i, the final polygon is $\mathcal{Z}$-enveloped and strictly larger than P_i. It still lies in the allowed translated half-strip: the relaxed normal is negative on a, so its inequality does not define the half-strip, while every defining inequality of the half-strip was retained.

Let t be the primitive length of the b-face of $\mathcal{Z}$. For each $1 \leq j \leq m$, take the first t lattice points, in direction b, on the new facet of $P_i^{(j)}$. Their union E is independent of all sufficiently large i. Indeed, the initial endpoint is determined by the stabilized preceding supporting line and the fixed equation $\nu_b = \beta - j$; only the far endpoint escapes.

For every sufficiently large i, the arrays θ_i and $\widehat{w}$ must disagree somewhere on E. Otherwise apply Lemma 6.1 successively to $\delta_i = \theta_i - \widehat{w}$, which satisfies $Q_0\delta_i = 0$. Agreement on P_i, together with the first t zeros on each successive facet, would give agreement on $P_i^{(m)}$, contradicting the maximality of M_i after undoing the anchor translation.

This is now a condition on one fixed finite pattern. Convergence on E means eventual equality of entire E-patterns, so

$$\theta|_E \neq \widehat{w}|_E.$$

Since $E \subset R_\infty^{(m)}$, one of the first m outward layers contains a disagreement. Take the first such layer and let $\mathcal{R}$ be the preceding region. Then $\theta = \widehat{w}$ on $\mathcal{R}$ but not on $\mathcal{R}^{(1)}$. The defining inequality of H was never moved, so $\mathcal{R}^{(1)} \subset H$. Every other surviving inward normal is strictly positive on $-a$; hence a sufficiently large negative a-translation moves any fixed finite subset of H into $\mathcal{R}$. □

Lemma 8.3 (The selected point cannot acquire a period)**.** *In Lemma 8.2, suppose in addition that $w|_H$ is not fully periodic. Then θ is nonperiodic.*

Proof. We test a hypothetical period against the fixed reference, rather than the approximating arrays. The phase shift is parallel to the boundary of H, so $\widehat{w}|_H$ is not fully periodic. Let $a_* = la$, $l \geq 1$, be a period of $\widehat{w}$, and suppose that $p \neq 0$ is a period of θ. Given $g, g + p \in H$, subtracting Na_* puts both points in $\mathcal{R}$ for large N, whence

$$\widehat{w}(g) = \theta(g - Na_*) = \theta(g + p - Na_*) = \widehat{w}(g + p).$$

This makes p a regional period of $\widehat{w}|_H$. Together with a_* it would give full regional periodicity unless p were parallel to a.

In the remaining case, primitivity gives $p = ra$ for some nonzero integer r. Both arrays then have period $\mathrm{lcm}(l, |r|)a$. Transporting any point of H into $\mathcal{R}$ by a sufficiently large negative multiple of this common period extends their agreement to all of H, contradicting the disagreement on $\mathcal{R}^{(1)} \subset H$. □

9 Proof of the selection proposition

Proof of Proposition 1.3. Fix S, Q_0, and the periodic reference w from Lemmas 5.1 and 5.4. The product Q_0 has at least two factors, since a single factor would give a period of c. Its directions include a, and S has two opposite a-parallel edges and satisfies (5.1). We keep this reference fixed throughout. We divide into two cases, depending on whether one of its forward half-strips already occurs in a shift of c.

A matching half-strip occurs. Suppose that some shift $z = T^u c$ agrees with w on a forward a-half-strip based on a $\mathcal{Z}$-enveloped polygon. By Lemma 8.1, z has a period h parallel to a on a translate $\mathcal{R}_0$ of $\mathcal{Z} + \mathbb{R}_{\geq 0}a + \mathbb{R}_{\geq 0}b$, where a, b are neighboring factor directions.

We can enlarge the region towards a half-plane by adding successive factor rays on the side of b until reaching $-a$. Each new ray is added to the entire preceding region, with the earlier rays retained. Adjacency in the cyclic list makes each step remove precisely the inequality of the current infinite facet. The period h cannot persist to the final half-plane, whose boundary is parallel to a: Lemma 5.3 would give a global period of z and hence of c.

At the first enlargement on which the period fails, we relax the moving inequality through successive integer levels. These finite relaxations exhaust the enlargement, so there is a last region on which $\Delta_h z = 0$ and a next layer on which it is nonzero. The points $z, T^h z \in X_c$ agree on the former region and not on the latter layer. Lemma 6.2 gives ambiguity in that boundary direction, say b'. Its inward normal is positive on a, so Lemma 7.1 supplies periods parallel to a, b' on a smaller two-ray region. We take $\theta = z$, which is nonperiodic as a shift of c.

No matching half-strip occurs. We apply Lemma 8.2 to w and H. It gives a point $\theta \in X_c$ agreeing with a fixed phase $\widehat{w}$ on a two-ray region with recession rays $-a, b$, and disagreeing on the next outward layer. By Lemma 8.3, this point is nonperiodic. Lemma 6.2 gives ambiguity along b. Agreement with the reference supplies a regional period pointing along $-a$, and the inward normal of the b-edge is positive on $-a$. Lemma 7.1 therefore gives two independent periods on a smaller region.

In either case, choose the signs of the two periods along the recession rays. Formula (6.1) extends the restriction to a single fully periodic background B. The region contains a translate K of its rational recession cone, of positive aperture less than π, and $\theta = B$ on $K \cap \mathbb{Z}^2$. Finally, $\theta \in X_c$ gives $P_\theta(k, n) \leq P_c(k, n)$ on the original rectangle. □

Acknowledgments and AI use

I used Qwen 3.8 27B, GLM 5.2, GLM 5.3 Flash, and GLM 5.3 through Draupnir, a custom agentic framework I developed, for proofreading, preliminary literature searches, and feedback on selected mathematical arguments. The paper was developed separately from these models. I thank the MIT Office of Research Computing and Data (ORCD) Engaging Cluster and the NCShare computing cluster for providing the computational resources to run them. I have checked and take full responsibility for the final wording, mathematical assertions, arguments, and references.

A Relation to previous work

We use two structural results from the literature: Kari–Szabados, Corollary 13 [8], and Colle, Theorem 1.9 [1], both quoted without proof. The one-dimensional arguments use the bi-infinite and one-sided forms of the Morse–Hedlund theorem recorded by Cyr–Kra, Theorem 2.2(ii)–(iii) [4]; see also the original paper [9].

The proof of Lemma 2.1 follows the affine argument in [8, Lemma 5], while Lemma 2.3 uses the elimination argument of [6, Lemma 8], with the finite-image hypothesis stated explicitly. For integral configurations, the directional structure underlying Lemma 2.4 is contained in [8, Corollaries 28–29]. Our proof starts from the product annihilator and also covers the finite-valued complex configurations in Theorem 1.2. The affine pattern count is related to [8, Lemmas 33 and 36]. Products of directional layers also occur in the four-dot argument of Kari–Moutot [6, proof of Theorem 11]. The contribution of Sections 3–4 is the construction of a quadratic observable with a small stencil and the proof that its fitting translates are independent modulo the affine functions on the same rectangle.

We use Colle's half-plane formulation of one-sided nonexpansiveness [1, Section 1]; Cyr–Kra [4, Definition 3.1 and Lemma 3.2] give the finite-extension formulation and its relation to half-plane ambiguity. Colle [1, Proposition 2.14], following Kari–Szabados, obtains the tangent-period conclusion from a periodic decomposition; Lemma 5.3 proves it directly from the product annihilator. The period-difference construction in Lemma 8.1 also appears in Colle [1, proof of Lemma 4.6, Claim 4.7]; our proof writes out the propagation into the wedge in integral coordinates.

Lemma 5.4 is based on the ambiguity argument of Cyr–Kra [4, Lemma 2.24] and the outward-extension argument from the proof of their Proposition 4.8, as presented by Colle [1, Proposition 2.13]. Since the full Cyr–Kra proposition assumes $P_c(k, n) \leq kn/2$, we prove the extension step directly from the edge ambiguity bound at the kn threshold. Related bounded-period propagation appears in Colle–Garibaldi [2, Lemmas 4.11–4.13].

The argument in Lemma 7.1 follows Colle's Lemma 4.2 [1], with the hypotheses arranged in the form needed here. In Proposition 1.3, we adapt the two-case region-selection argument in his proof of Lemma 4.6 to obtain a periodic region in a configuration that remains nonperiodic. For this purpose, the later conclusions about the angular positions of the nonexpansive directions, as well as Theorem 1.10 and its constant decomposition-order hypothesis, are not needed.

Related periodic-region arguments appear in Colle–Garibaldi [2, Sections 5–6] under an alphabet-sensitive complexity hypothesis. The more recent modular structure theorem [3, Theorem 1.2] concerns reductions of periodic components on pairs of half-planes, and the proof of Theorem 3.5 in that paper again uses the two-case construction. Here we keep a single orbit point over the original alphabet and prove its nonperiodicity directly, without imposing the earlier alphabet-sensitive hypothesis. The short-side-four result provides further context, though it is not actually used in the proof.

The anchored construction in Lemma 8.2 is based on Colle [1, proof of Lemma 3.5], with an extra boundary argument to retain a disagreement in the limit. We shift the approximants by $u_i + k_i a$, reducing only the periodic reference phase to a fixed residue. After passing to a componentwise monotone subsequence of edge lengths, the anchored polygons are nested by the residual Minkowski-sum argument. Outward layers are formed by relaxing signed supporting inequalities, so that each new region contains the preceding one; an unsigned distance restriction on the entire region would clip it. Lemma 6.1 propagates agreement from the fixed finite set E, supplying the boundary control that keeps disagreement from escaping under compactness. Finally, Lemma 8.3 avoids the two ray-limit arguments by transporting a hypothetical period of θ to the reference half-plane and using the retained disagreement, without transferring a period backwards from a limit.

B Boundary cases and examples

Finite image is needed in Lemma 2.3, although boundedness suffices for the later arithmetic-progression arguments applied to c and A_i. For example, if $1, \alpha/(2\pi), \beta/(2\pi)$ are rationally independent, then $e^{i(\alpha x+\beta y)}$ is bounded and has two coprime univariate annihilators, but no period. Its image is infinite.

The single defect has exactly $kn+1$ patterns but no Laurent annihilator by Lemma 3.1: it gives sharpness without satisfying the algebraic hypothesis of Theorem 1.2. A lifted aperiodic sequence has one period, though not two independent periods, and is excluded by the nonperiodicity hypothesis.

An example meeting the algebraic and wedge hypotheses is

$$c(x,y) = \mathbf{1}_{x\geq 0} + \mathbf{1}_{y\geq 0} - \mathbf{1}_{x+y\geq 0} \in \{0,1\}.$$

It is annihilated by $\Delta_{(1,0)}\Delta_{(0,1)}\Delta_{(1,-1)}$ and is 1 on the first quadrant. It is nonperiodic: for $y>0$, the zero interval in its row is $[-y,-1]$. Comparing these intervals at sufficiently large heights under a hypothetical period (a,b) forces $a=0$ from the right endpoints and then $b=0$ from the left endpoints. Thus the wedge hypotheses allow nonperiodic binary configurations with three factor directions, including a negatively sloped direction.

Section 4 covers every positive rectangle size, including $L_R = 0$, in the original coordinates; the auxiliary bases used in the selection do not alter the window. The cone argument does not extend to dimension three, since two half-spaces can intersect in a set containing a line.

References

[1] C. F. Colle, *On periodic decompositions, one-sided nonexpansive directions and Nivat's conjecture*, Discrete and Continuous Dynamical Systems **43** (2023), no. 12, 4299–4327. doi:10.3934/dcds.2023088.

[2] C. F. Colle and E. Garibaldi, *An alphabetical approach to Nivat's conjecture*, Nonlinearity **33** (2020), no. 7, 3620–3652. doi:10.1088/1361-6544/ab7638.

[3] C. F. Colle and E. Garibaldi, *A modular structure theorem for minimal periodic decompositions and periodicity of configurations with* $P_\eta(4, n) \leq 4n$, preprint, arXiv:2606.10193 (2026).

[4] V. Cyr and B. Kra, *Nonexpansive* $\mathbb{Z}^2$*-subdynamics and Nivat's conjecture*, Transactions of the American Mathematical Society **367** (2015), no. 9, 6487–6537. doi:10.1090/S0002-9947-2015-06391-0.

[5] V. Cyr and B. Kra, *Complexity of short rectangles and periodicity*, European Journal of Combinatorics **52** (2016), 146–173. doi:10.1016/j.ejc.2015.10.003.

[6] J. Kari and E. Moutot, *Nivat's conjecture and pattern complexity in algebraic subshifts*, Theoretical Computer Science **777** (2019), 379–386. doi:10.1016/j.tcs.2018.12.029.

[7] J. Kari and E. Moutot, *Decidability and periodicity of low complexity tilings*, Theory of Computing Systems **67** (2023), 125–148. doi:10.1007/s00224-021-10063-8.

[8] J. Kari and M. Szabados, *An algebraic geometric approach to Nivat's conjecture*, Information and Computation **271** (2020), article 104481. doi:10.1016/j.ic.2019.104481.

[9] M. Morse and G. A. Hedlund, *Symbolic dynamics*, American Journal of Mathematics **60** (1938), no. 4, 815–866. doi:10.2307/2371264.

[10] M. Szabados, *Nivat's conjecture holds for sums of two periodic configurations*, in *SOFSEM 2018: Theory and Practice of Computer Science*, Lecture Notes in Computer Science, vol. 10706, Springer, 2018, pp. 539–551. doi:10.1007/978-3-319-73117-9_38.