Square Maximality for Spanning Trees of Induced Grid Subgraphs

August 2026

Abstract

For a finite $S \subset \mathbb{Z}^2$, let $\tau(S)$ be the number of spanning trees of its induced square-lattice graph. For positive integers r, s, set $\Lambda_{r,s} = \{0, \dots, r-1\} \times \{0, \dots, s-1\}$ and $\tau(r, s) = \tau(\Lambda_{r,s})$. We prove that, for every $n \geq 1$,

$$|S| = n^2 \quad \Longrightarrow \quad \tau(n, n) \geq \tau(S),$$

with equality exactly for $n \times n$ lattice squares. More precisely, suppose that $|S| = n^2$, that its induced graph is connected, and that $\Lambda_{r,s}$ is its minimal rectangular hull after translation; then $|\Lambda_{r,s} \setminus S| = rs - n^2$. Interchanging r and s if necessary,

$$\log \tau(r, s) - \log \tau(S) \;\geq\; \frac{4G_{\mathrm{Cat}}}{\pi}(rs - n^2) \;\geq\; \log \tau(r, s) - \log \tau(n, n),$$

where G_{Cat} is Catalan's constant. Filling in the hull is worth at least the square-lattice tree entropy per vertex added, and enlarging the $n \times n$ square to the hull at most that; the same number of vertices is added in each case, so cancelling $\log \tau(r, s)$ gives the theorem.

Background. In 1999, Chung obtained boundary-sensitive estimates for spanning-tree counts in lattice subgraphs [Chu99]. In 2010, Swanepoel proposed on MathOverflow that square shapes should maximize the count at fixed cardinality [Swa10]. Procaccia and Tucker-Foltz recorded the conjecture in the research literature in 2022 [PTF22], and Tapp established general boundary-sensitive grid-graph bounds in 2024 [Tap24]. In 2026, Zhang proved exact balancing for free rectangular grids [Zha26a], then extended the extremal result to pure Cartesian-product families with free and periodic boundary conditions and formulated the induced-subgraph conjecture in all dimensions [Zha26b]. The present note proves the square-maximality conjecture recorded by Procaccia and Tucker-Foltz in 2022, resolving for arbitrary finite induced subgraphs of the square lattice both the question proposed by Swanepoel in 2010 and the two-dimensional case of the induced-subgraph conjecture formulated by Zhang in 2026.

1 Introduction

Let $\mathcal{L}$ be the infinite square-lattice graph with vertex set $\mathbb{Z}^2$. Vertices $x = (x_1, x_2)$ and $y = (y_1, y_2)$ are adjacent when

$$|x_1 - y_1| + |x_2 - y_2| = 1.$$

We call every finite subgraph of $\mathcal{L}$ a square-lattice graph. For a finite graph G, write $V(G)$ and $E(G)$ for its vertex and edge sets, and let $\tau(G)$ denote its number of spanning trees. We set $\tau(G) = 0$ when G is disconnected and $\tau(G) = 1$ when G consists of a single vertex. For a finite set $S \subset \mathbb{Z}^2$, let $\mathcal{L}[S]$ be the subgraph induced by S, and write

$$\tau(S) := \tau(\mathcal{L}[S]).$$

For positive integers r, s, set

$$\Lambda_{r,s} := \{0, \dots, r-1\} \times \{0, \dots, s-1\}, \qquad R_{r,s} := \mathcal{L}[\Lambda_{r,s}], \qquad \tau(r,s) := \tau(R_{r,s}).$$

All logarithms are natural. We use

$$G_{\mathrm{Cat}} := \sum_{j \geq 0} \frac{(-1)^j}{(2j+1)^2}$$

for Catalan's constant.

Conjecture 1.1 ([PTF22])**.** *Let $n \geq 1$, and let H be a finite connected subgraph of $\mathcal{L}$ with $|V(H)| = n^2$. Then*

$$\tau(H) \leq \tau(n,n).$$

For a graph H in Conjecture 1.1, put $S := V(H)$. Adding all lattice edges with both endpoints in S produces $\mathcal{L}[S]$, and every spanning tree of H remains a spanning tree of $\mathcal{L}[S]$. Thus

$$H \subseteq \mathcal{L}[S], \qquad \tau(H) \leq \tau(S),$$

so it suffices to consider induced subgraphs. If $H \neq \mathcal{L}[S]$, adding a missing lattice edge strictly increases the spanning-tree count: adding that edge to any spanning tree of H and deleting one edge from the resulting cycle produces a new spanning tree containing the added edge. We prove the conjecture and characterize equality.

Theorem 1.2 (Square maximality)**.** *Let $n \geq 1$, and let $S \subset \mathbb{Z}^2$ be finite with $|S| = n^2$. Then*

$$\tau(S) \leq \tau(n,n).$$

Equality holds if and only if S is an $n \times n$ lattice square, up to translation and the symmetries of the square lattice.

1.1 Outline of the proof

Theorem 1.2 follows from two finite inequalities carrying the same constant: one compares rectangles of different areas, the other compares a set with its rectangular hull.

(1) *Rectangular area penalty* (**Lemma 3.1**). For positive integers n, r, s with $r \leq s$ and $rs \geq n^2$,

$$\log \tau(r,s) - \log \tau(n,n) \leq \frac{4G_{\mathrm{Cat}}}{\pi}(rs - n^2),$$

with equality only when $r = s = n$. The proof telescopes $\log \tau$ in one side length, drawing on the rectangular product and hyperbolic formulas, the trapezoidal estimate, and the residual bounds of [Zha26a]; the individual references are given in Section 3.

(2) *Normalized rectangular-hull inequality* (**Lemma 4.1**). Let $S \subset \mathbb{Z}^2$, suppose that $\mathcal{L}[S]$ is connected, and suppose that the smallest axis-parallel lattice vertex rectangle containing S is $\Lambda_{r,s}$ after translation. Then

$$\log \tau(S) - \log \tau(r,s) \leq -\frac{4G_{\mathrm{Cat}}}{\pi}|\Lambda_{r,s} \setminus S|.$$

Interior vacancy components are controlled by planar duality and Tapp's bound (**Theorem 2.1**). Boundary-attached vacancies are handled by Lemma 5.1 and the sliding-window and closed-walk arguments in Section 5.

(3) *Exact matching and square maximality* (**Theorem 1.2**). If $\mathcal{L}[S]$ is connected, $|S| = n^2$, and the smallest axis-parallel lattice vertex rectangle containing S is $\Lambda_{r,s}$, then

$$|\Lambda_{r,s} \setminus S| = rs - n^2,$$

and the two inequalities match exactly:

$$\log \tau(S) - \log \tau(n,n) \le -\frac{4G_{\mathrm{Cat}}}{\pi}(rs - n^2) + \frac{4G_{\mathrm{Cat}}}{\pi}(rs - n^2) = 0.$$

The strictness of the rectangular estimate forces $r = s = n$, after which equality of cardinalities forces $S = \Lambda_{n,n}$, up to lattice symmetry.

1.2 Related Work

Kirchhoff's Matrix–Tree Theorem is the classical starting point for exact spanning-tree enumeration [Kir47]. Square-lattice tree entropy and product formulas have been studied extensively; we use the standard closed-walk evaluation recorded, for example, in [SW00]. Tapp proved boundary-sensitive upper and lower bounds for arbitrary finite grid graphs [Tap24]. The rectangular balancing theorem gives the sharp finite comparison within the free rectangular family [Zha26a], while the product-grid extension studies other boundary conditions and higher-dimensional product families [Zha26b].

Organization. Section 2 fixes notation and records the external inputs. Section 3 proves the rectangular estimate. Sections 4 and 5 prove the normalized hull inequality. Section 6 completes the proof of Theorem 1.2, and Section 7 discusses the argument.

2 Preliminaries and Notation

Empty products and 0×0 determinants are equal to one. The free-boundary constant is

$$\alpha := \operatorname{arsinh}(1) = \log(1 + \sqrt{2}).$$

The square-lattice tree entropy is $4G_{\mathrm{Cat}}/\pi$.

For $v \in \mathbb{Z}^2$, let $\square_v$ be the unit-square graph having v as its bottom-right corner. Following [Tap24, Definition 6], the top-left boundary of a finite square-lattice graph H is

$$\widehat{\partial H} := \{v \in V(H) : \square_v \text{ is not a subgraph of } H\}.$$

Let $\mathrm{sq}(H)$ denote the number of unit lattice squares whose four boundary edges belong to H. Thus $\mathrm{sq}(H) = |V(H)| - |\widehat{\partial H}|$.

Theorem 2.1 ([Tap24, Theorem 9]). *If H is a finite connected square-lattice graph, then*

$$\log \tau(H) \ge \frac{4G_{\mathrm{Cat}}}{\pi}\,\mathrm{sq}(H).$$

We also use two standard facts about spanning trees. First, if G is a connected planar multigraph, let $G^\star$ denote its geometric dual, including the exterior dual vertex. Deletion of a primal edge corresponds to contraction of its dual edge, so the cycle matroids of G and $G^\star$ are dual. For an edge set $A \subseteq E(G)$, let G/A denote the multigraph obtained by simultaneously contracting the edges of A, discarding the resulting loops and retaining parallel-edge multiplicities.

Second, if T is a uniform spanning tree of a connected multigraph G and F is a forest, then

$$\Pr(F \subseteq T) = \frac{\tau(G/F)}{\tau(G)}.$$

The transfer-current theorem represents this probability as a principal minor of a positive semidefinite kernel [BP93]. Hadamard–Fischer therefore implies that, for an integer $d \geq 1$ and pairwise disjoint edge sets $F_1, \dots, F_d$ whose union is a forest,

$$\Pr\left(\bigcup_{i=1}^{d} F_i \subseteq T\right) \leq \prod_{i=1}^{d} \Pr(F_i \subseteq T). \tag{2.1}$$

3 Rectangular Area Penalty

We first compare a rectangle whose area is at least n^2 with the $n \times n$ square.

Lemma 3.1 (Rectangular area penalty)**.** *Let n, r, s be positive integers with $r \leq s$ and $rs \geq n^2$. Then*

$$\log \tau(r, s) - \log \tau(n, n) \leq \frac{4G_{\mathrm{Cat}}}{\pi}(rs - n^2). \tag{3.1}$$

Equality holds if and only if $r = s = n$.

Proof. Both cases below telescope $\log \tau$ in one side length. Each increment splits into a main term, depending only on the side length held fixed, and a residual. The main terms supply the entropy constant together with a deficit of at least α; each residual sum contributes less than α. The strict inequality comes from that margin.

For $0 \leq x \leq 1$ and every integer $a \geq 1$, define

$$c(x) = \operatorname{arcosh}(2 - \cos \pi x), \qquad C_a = \sum_{j=1}^{a-1} c(j/a).$$

Set

$$\varepsilon_a := a\frac{4G_{\mathrm{Cat}}}{\pi} - \alpha - C_a.$$

The rectangular product formula, its hyperbolic reduction, and the trapezoidal estimate in [Zha26a, Proposition 2.1, equations (3.1)–(3.2), Lemmas 3.1, 3.2, and 5.1, and equation (5.6)] give

$$C_a = a\frac{4G_{\mathrm{Cat}}}{\pi} - \alpha - \varepsilon_a, \qquad \varepsilon_a \geq 0, \qquad \frac{\varepsilon_a}{a} \text{ is nonincreasing.} \tag{3.2}$$

The identification of the integral in the main term with $4G_{\mathrm{Cat}}/\pi$ is [Zha26a, Appendix Lemma A.3].

For positive integers a, t, define

$$\mathcal{R}_a(t) := \sum_{j=1}^{a-1} \log \frac{1 - e^{-2(t+1)c(j/a)}}{1 - e^{-2tc(j/a)}}.$$

The same hyperbolic formula gives

$$\Delta_a(t) := \log\frac{\tau(a,t+1)}{\tau(a,t)} = C_a + \mathcal{R}_a(t), \tag{3.3}$$

and $\mathcal{R}_a(t) \geq 0$. These formulas include $a = 1$ by the empty-product convention.

Since c is concave, $c(0) = 0$, and $c(1) = 2\alpha$, for $1 \leq j \leq a-1$ we have

$$c(j/a) \geq \frac{2\alpha j}{a}.$$

For integers $u \geq t \geq a$, set

$$q := e^{-4\alpha} = (3-2\sqrt{2})^2, \qquad K := \sum_{j\geq 1} -\log(1-q^j).$$

The residuals telescope to give

$$\sum_{v=t}^{u-1} \mathcal{R}_a(v) = \sum_{j=1}^{a-1} \log\frac{1-e^{-2uc(j/a)}}{1-e^{-2tc(j/a)}}.$$

For $j \geq 1$, the estimate

$$-\log(1-q^j) \leq \frac{q^j}{1-q^j} \leq \frac{q^j}{1-q}$$

therefore gives the uniform bound

$$0 \leq \sum_{v=t}^{u-1} \mathcal{R}_a(v) < K \leq \frac{q}{(1-q)^2} = \frac{1}{32} < \alpha. \tag{3.4}$$

Suppose first that $r \leq n \leq s$. Put

$$k = n-r, \qquad \ell = s-n, \qquad \delta = rs - n^2 = \ell r - kn.$$

Telescoping (3.3) and using symmetry gives

$$\log\frac{\tau(r,s)}{\tau(n,n)} = \sum_{t=n}^{s-1} \Delta_r(t) - \sum_{t=r}^{n-1} \Delta_n(t). \tag{3.5}$$

Substituting (3.3) splits the right side of (3.5) into a C-part and a residual part:

$$\log\frac{\tau(r,s)}{\tau(n,n)} = \left(\ell C_r - kC_n\right) + \left(\sum_{t=n}^{s-1} \mathcal{R}_r(t) - \sum_{t=r}^{n-1} \mathcal{R}_n(t)\right).$$

We bound the two in turn.

The area condition $\delta \geq 0$ reads $\ell r \geq kn$. If the rectangles are distinct it forces $\ell - k \geq 1$: when $k = 0$, distinctness implies $\ell \geq 1$; when $k > 0$, the inequality $\ell r \geq kn$ together with $r < n$ gives $\ell > k$. The monotonicity in (3.2) gives $\varepsilon_r/r \geq \varepsilon_n/n$, whence

$$\ell\varepsilon_r \geq \frac{\ell r}{n}\varepsilon_n \geq k\varepsilon_n.$$

By (3.2) the C-part is therefore

$$\ell C_r - kC_n = \delta\frac{4G_{\mathrm{Cat}}}{\pi} - (\ell-k)\alpha - \ell\varepsilon_r + k\varepsilon_n \le \delta\frac{4G_{\mathrm{Cat}}}{\pi} - (\ell-k)\alpha.$$

In the residual part the subtracted sum is nonnegative, and the first sum is less than $K < \alpha$ by (3.4). For distinct rectangles $\ell - k \ge 1$, so the two parts together are less than $\delta\, 4G_{\mathrm{Cat}}/\pi$; thus (3.1) is strict unless $r = s = n$.

It remains to consider $n \le r \le s$. For $a \ge 1$, symmetry and (3.3) give

$$\log\frac{\tau(a+1,a+1)}{\tau(a,a)} = \Delta_a(a) + \Delta_{a+1}(a).$$

The C-part is

$$(2a+1)\frac{4G_{\mathrm{Cat}}}{\pi} - 2\alpha - \varepsilon_a - \varepsilon_{a+1}.$$

The first residual, $\mathcal{R}_a(a)$, is less than $1/32$. For $1 \le j \le a$, concavity also gives

$$2a\,c(j/(a+1)) \ge \frac{4a}{a+1}\alpha j \ge 2\alpha j.$$

With $q_0 = e^{-2\alpha} = 3 - 2\sqrt{2}$, the second residual, $\mathcal{R}_{a+1}(a)$, is therefore less than

$$\sum_{j\ge 1} -\log(1-q_0^j) \le \frac{q_0}{(1-q_0)^2} = \frac{1}{4}.$$

Since $1/32 + 1/4 < 2\alpha$, we obtain

$$\log\tau(a+1,a+1) - \log\tau(a,a) < (2a+1)\frac{4G_{\mathrm{Cat}}}{\pi}.$$

If $r > n$, summing this strict inequality gives

$$\log\tau(r,r) - \log\tau(n,n) < \frac{4G_{\mathrm{Cat}}}{\pi}(r^2 - n^2); \tag{3.6}$$

when $r = n$, the left side and the right side are both zero.

If $s > r$, then (3.2) and (3.4) give

$$\begin{aligned}\log\frac{\tau(r,s)}{\tau(r,r)} &= (s-r)C_r + \sum_{t=r}^{s-1}\mathcal{R}_r(t)\\ &< \frac{4G_{\mathrm{Cat}}}{\pi}\,r(s-r). \end{aligned}\tag{3.7}$$

Indeed, before the last inequality the right side is at most

$$\frac{4G_{\mathrm{Cat}}}{\pi}\,r(s-r) - (s-r)\alpha + K,$$

which is strictly smaller because $s - r \ge 1$ and $K < \alpha$. When $s = r$, both sides of the corresponding non-strict inequality are zero. Adding (3.6) and (3.7),

$$\log\frac{\tau(r,s)}{\tau(n,n)} = \log\frac{\tau(r,r)}{\tau(n,n)} + \log\frac{\tau(r,s)}{\tau(r,r)} \le \frac{4G_{\mathrm{Cat}}}{\pi}(r^2-n^2) + \frac{4G_{\mathrm{Cat}}}{\pi}r(s-r) = \frac{4G_{\mathrm{Cat}}}{\pi}(rs-n^2),$$

which is (3.1). The first summand is strict when $r > n$ and the second when $s > r$, so the inequality is strict unless $r = s = n$. This also covers $n = 1$. □

4 Normalized Rectangular-Hull Inequality

For a finite nonempty $S \subset \mathbb{Z}^2$, define its coordinate extrema by

$$x_- := \min_{(x,y)\in S} x, \quad x_+ := \max_{(x,y)\in S} x, \quad y_- := \min_{(x,y)\in S} y, \quad y_+ := \max_{(x,y)\in S} y.$$

Its minimal rectangular vertex hull is

$$\{x_-, \dots, x_+\} \times \{y_-, \dots, y_+\}.$$

Thus, after translation, the hull is $\Lambda_{r,s}$, where $r = x_+ - x_- + 1$ and $s = y_+ - y_- + 1$. The central estimate is the following.

Lemma 4.1 (Normalized hull inequality)**.** *Let $S \subset \mathbb{Z}^2$ be finite, suppose that $\mathcal{L}[S]$ is connected, and suppose after translation that $\Lambda_{r,s}$ is its minimal rectangular vertex hull. Then*

$$\log \tau(S) - \log \tau(r,s) \le -\frac{4G_{\mathrm{Cat}}}{\pi} |\Lambda_{r,s} \setminus S|. \tag{4.1}$$

The proof runs to the end of Section 5. Passing to the planar dual converts the deletion of a set of vertices into the event that a prescribed forest is contained in a uniform spanning tree; the missing vertices are then split into king components, and negative correlation reduces the lemma to one bound per component. Components disjoint from the perimeter are handled below by counting the spanning trees of a union of unit squares in the dual, the rest by the determinant estimate of Section 5.

If $r = 1$ or $s = 1$, the connectedness of $\mathcal{L}[S]$ and the minimality of the hull imply $S = \Lambda_{r,s}$ and $\tau(S) = \tau(r,s)$, so (4.1) is immediate. We assume below that $r, s \ge 2$.

4.1 Dual Contractions

Let $D = R^\star_{r,s}$ be the geometric planar dual, including its exterior vertex. If $e \in E(R_{r,s})$, let $e^\star$ be its corresponding dual edge. For $X \subseteq \Lambda_{r,s}$, let

$$E^\star_X := \{e^\star : e \in E(R_{r,s}) \text{ has at least one endpoint in } X\},$$

and let H_X be the dual subgraph formed by $E^\star_X$, with isolated vertices suppressed.

Whenever $\mathcal{L}[\Lambda_{r,s} \setminus X]$ is connected, planar matroid duality gives

$$\tau(\Lambda_{r,s} \setminus X) = \tau(D/E^\star_X).$$

Planar tree duality also gives $\tau(D) = \tau(r,s)$. Indeed, deleting the primal edges incident with X leaves the connected graph $R_{r,s}[\Lambda_{r,s} \setminus X]$, together with the isolated vertices of X. Its cycle-matroid bases are exactly the spanning trees of the connected component; the isolated vertices contribute no edge choice. Deletion of the primal edges is contraction of their duals.

Call a component nontrivial if it contains at least one edge, and let F_X be a spanning forest of the nontrivial components of H_X. Contracting F_X makes every edge of $E^\star_X \setminus F_X$ a loop, so the preceding duality identity and the contraction identity give

$$\frac{\tau(\Lambda_{r,s} \setminus X)}{\tau(r,s)} = \Pr(F_X \subseteq T), \tag{4.2}$$

where T is a uniform spanning tree of D.

For $z = (z_1, z_2) \in \mathbb{Z}^2$, set $\|z\|_\infty := \max\{|z_1|, |z_2|\}$. Two lattice vertices u, v are king-adjacent when $\|u - v\|_\infty = 1$. The perimeter vertex set of the rectangular hull is

$$\partial\Lambda_{r,s} := \{(x, y) \in \Lambda_{r,s} : x \in \{0, r-1\} \text{ or } y \in \{0, s-1\}\}.$$

Partition $\Lambda_{r,s} \setminus S$ into components for king adjacency. Let $V_\partial^{\mathrm{vac}}$ be the union of those components that meet $\partial\Lambda_{r,s}$, and denote the remaining, interior components by

$$J_1, \ldots, J_m.$$

Here m may be zero. For a primal vertex v, call the subgraph formed by the duals of the primal edges incident with v its dual face-star. Distinct king components have disjoint dual edge sets. Two primal vertices can be incident with a common bounded unit face only if their ℓ^∞-distance is at most one; their dual face-stars are therefore vertex-disjoint, except that different boundary components can share the exterior dual vertex. Within one king component, the face-stars of successive vertices meet, so their union is connected; for a boundary component, one of these stars also meets the exterior vertex. Hence $H_{V_\partial^{\mathrm{vac}}}$ is connected when $V_\partial^{\mathrm{vac}} \neq \varnothing$, and each H_{J_i} is connected. Set $F_{V_\partial^{\mathrm{vac}}} = \varnothing$ when $V_\partial^{\mathrm{vac}} = \varnothing$; otherwise choose a spanning tree of $H_{V_\partial^{\mathrm{vac}}}$. For $1 \leq i \leq m$, choose a spanning tree F_{J_i} of H_{J_i}. These forest edge blocks are pairwise disjoint, and their union is a spanning forest of the nontrivial components of $H_{\Lambda_{r,s} \setminus S}$.

We record why all deletion graphs used here are connected. Every king component C has a nearest-neighbour edge to S: on a nearest-neighbour path from C to S, the first vertex outside C cannot lie in another king component. Moreover, a diagonal king step between $u, v \in C$ can be replaced by a two-edge nearest-neighbour path through an orthogonal corner of their unit square. If the corner is in $\Lambda_{r,s} \setminus S$, it is king-adjacent to u and hence belongs to C; otherwise it belongs to S. Thus $\mathcal{L}[S \cup C]$ is connected. It follows that every undeleted king component is joined to the common connected subgraph $\mathcal{L}[S]$. Hence $\mathcal{L}[\Lambda_{r,s} \setminus V_\partial^{\mathrm{vac}}]$ and every $\mathcal{L}[\Lambda_{r,s} \setminus J_i]$ are connected.

Applying (4.2) and (2.1) to the disjoint forest edge blocks gives

$$\frac{\tau(S)}{\tau(r,s)} \leq \frac{\tau(\Lambda_{r,s} \setminus V_\partial^{\mathrm{vac}})}{\tau(r,s)} \prod_{i=1}^{m} \frac{\tau(\Lambda_{r,s} \setminus J_i)}{\tau(r,s)}. \tag{4.3}$$

The factor indexed by $V_\partial^{\mathrm{vac}}$ is one when $V_\partial^{\mathrm{vac}} = \varnothing$.

4.2 Interior Vacancy Components

For each $i \in \{1, \ldots, m\}$, put $J := J_i$. The following argument is vacuous when $m = 0$. The dual graph H_J is a finite connected square-lattice graph. Every $v \in J$ contributes a complete unit square to H_J, formed by the four duals of the primal edges incident with v. Different vertices give different squares, so

$$\mathrm{sq}(H_J) \geq |J|.$$

By Theorem 2.1,

$$\tau(H_J) \geq \exp\left(\frac{4G_{\mathrm{Cat}}}{\pi}|J|\right).$$

For every spanning tree F of H_J, (4.2) gives

$$\Pr(F \subseteq T) = \frac{\tau(\Lambda_{r,s} \setminus J)}{\tau(r,s)}.$$

These events are pairwise incompatible: the union of two distinct spanning trees on the same vertex set contains a cycle, while T is acyclic. Summing over the spanning trees of H_J therefore gives

$$\tau(H_J)\,\frac{\tau(\Lambda_{r,s}\setminus J)}{\tau(r,s)} = \sum_F \Pr(F\subseteq T)\le 1,$$

which with the preceding consequence of Tapp's theorem yields

$$\frac{\tau(\Lambda_{r,s}\setminus J)}{\tau(r,s)} \le \frac{1}{\tau(H_J)} \le \exp\left(-\frac{4G_{\mathrm{Cat}}}{\pi}|J|\right). \tag{4.4}$$

The argument used that J misses the perimeter: only then does each $v\in J$ have four incident primal edges, whose duals close into a unit square. A vacancy on $\partial\Lambda_{r,s}$ has fewer, so the count above does not apply to it, and the same exponential factor for the union $V_\partial^{\mathrm{vac}}$ of the boundary-attached components is obtained differently, in Section 5.

5 Boundary-Attached Vacancies

What is proved here is a comparison of two determinants. Both $\tau(r,s)$ and $\tau(\Lambda_{r,s}\setminus V_\partial^{\mathrm{vac}})$ have the form $\det(4I-\mathsf{A})$ for the adjacency matrix of a set of unit faces, and expanding $\log\det$ in closed walks turns their ratio into a sum, over closed walks, of the number of placements meeting a face with a corner in $V_\partial^{\mathrm{vac}}$. The first three subsections bound that number by $|V_\partial^{\mathrm{vac}}|$ plus a multiple of the number of vacancy runs interior to a row or column; the last assembles the estimate.

Retain the boundary-vacancy set $V_\partial^{\mathrm{vac}}\subseteq\Lambda_{r,s}\setminus S$ from the preceding section. For $z\in\mathbb{Z}^2$ and $A\subseteq\mathbb{Z}^2$, define the translate

$$z+A := \{z+a : a\in A\}.$$

Let

$$\Phi = \{0,\dots,r-2\}\times\{0,\dots,s-2\}$$

index the unit faces of $R_{r,s}$. The face indexed by $z\in\Phi$ has corner set $z+\{0,1\}^2$. Define the marked and unmarked face sets by

$$M=\{z\in\Phi : (z+\{0,1\}^2)\cap V_\partial^{\mathrm{vac}}\neq\varnothing\},\qquad U=\Phi\setminus M.$$

An interval component is a maximal nonempty set of consecutive integers. Let ρ_x be the total number of interval components of $\Lambda_{r,s}\setminus V_\partial^{\mathrm{vac}}$ in all horizontal rows, and let ρ_y be the corresponding total in all vertical columns. Put

$$g_x=\rho_x-s,\qquad g_y=\rho_y-r.$$

Every row and column contains a vertex of S. Indeed, minimality gives a vertex of S on each of the two extreme columns, and a nearest-neighbour path in S between them visits every intermediate column. The row statement is identical. Consequently $g_x, g_y\ge 0$. Here a run of $V_\partial^{\mathrm{vac}}$ is an interval component of $V_\partial^{\mathrm{vac}}$ in a row or column, and it is internal when it touches neither endpoint of that row or column. Then g_x is exactly the total number of horizontally internal runs of $V_\partial^{\mathrm{vac}}$, and g_y is the corresponding number of vertically internal runs.

5.1 One-Dimensional Windows

We begin with an elementary interval estimate.

Lemma 5.1. *Let $L \geq 0$ be an integer, let $B \subseteq \{0, \dots, L\}$, and let $i(B)$ be the number of interval components of B that touch neither endpoint. For an integer p with $0 \leq p \leq L$, the number of integer intervals*

$$[t, t+p] := \{t, t+1, \dots, t+p\} \subseteq \{0, \dots, L\}, \qquad t \in \mathbb{Z},$$

that meet B is at most

$$|B| + p\, i(B).$$

Proof. A boundary run of length a is met by at most a such intervals, whereas an internal run of length a is met by at most $a + p$. Summing over the runs proves the claim; any double counting only strengthens the upper bound. □

For an integer p with $0 \leq p \leq r-1$, define

$$Y_p = \left\{(x, y) \in \{0, \dots, r-1-p\} \times \{0, \dots, s-1\} : ([x, x+p] \times \{y\}) \cap V_\partial^{\mathrm{vac}} \neq \varnothing\right\}.$$

Applying Lemma 5.1 in each row gives

$$|Y_p| \leq |V_\partial^{\mathrm{vac}}| + p\, g_x. \tag{5.1}$$

5.2 A Sliding-Window Injection

For an integer x with $0 \leq x \leq r-1-p$, let

$$Y_{p,x} = \{y : (x, y) \in Y_p\}.$$

The second ingredient is

$$\sum_{x=0}^{r-1-p} i(Y_{p,x}) \leq g_y. \tag{5.2}$$

It is proved in two steps: each internal interval component of $Y_{p,x}$ contains the projection of a connected piece of vacancy runs lying in the window and carrying no run that touches row 0 or row $s-1$; and those pieces, taken over all windows, inject into the vertically internal runs of $V_\partial^{\mathrm{vac}}$.

To prove it, form the full run graph, whose vertices are the maximal vertical runs of $V_\partial^{\mathrm{vac}}$. Join two runs in adjacent columns when their integer intervals overlap or are adjacent, and mark a run if it touches row 0 or row $s-1$. The connected components of the full run graph are precisely the king components of $V_\partial^{\mathrm{vac}}$. Since every such component meets $\partial\Lambda_{r,s}$, it contains either a marked run or a run in column 0 or $r-1$.

Fix once and for all a rooted spanning tree in every component of the full run graph. Use a marked run as the root whenever possible, and otherwise use a run in column 0 or $r-1$. For the window of columns $[x, x+p]$, consider the subgraph of the full run graph induced by the run vertices whose columns lie in this window. In each of its components containing no marked run, choose a run of minimum depth, where depth is measured in the previously fixed rooted tree.

If the chosen run is not the root, its parent lies outside the window; otherwise the parent would belong to the same induced component at a smaller depth. Parent and child lie in adjacent columns. If the parent is to the left, then x is the child's column; if the parent is to the right, then x is the child's column minus p. Thus the oriented parent–child pair determines the window start uniquely.

If the chosen run is an unmarked root, column 0 forces $x = 0$, while column $r-1$ forces $x = r-1-p$. Therefore the pairs

$$\text{(window, unmarked induced component)}$$

inject into the unmarked vertices of the full run graph. These vertices are exactly the vertically internal runs of $V_\partial^{\mathrm{vac}}$, of which there are g_y.

Every internal interval component of $Y_{p,x}$ contains the vertical projection of an induced run-graph component containing no marked run. The projection of a connected induced run-graph component is itself an integer interval, because adjacent runs overlap or are one unit apart. Distinct interval components therefore contain distinct such induced components. This proves (5.2).

For integers a, b with $1 \le a \le r$ and $1 \le b \le s$, define

$$\mathrm{win}_{a,b}(V_\partial^{\mathrm{vac}}) := \left| \left\{ (x,y) \in \mathbb{Z}^2 : \begin{array}{l} 0 \le x \le r-a, \quad 0 \le y \le s-b, \\ ([x, x+a-1] \times [y, y+b-1]) \cap V_\partial^{\mathrm{vac}} \neq \varnothing \end{array} \right\} \right|.$$

Applying Lemma 5.1 vertically and then using (5.1) and (5.2) gives, for every integer q with $0 \le q \le s-1$,

$$\mathrm{win}_{p+1,q+1}(V_\partial^{\mathrm{vac}}) \le |V_\partial^{\mathrm{vac}}| + p g_x + q g_y. \tag{5.3}$$

5.3 Closed-Walk Charging

For an integer $\ell \ge 1$, put $e_1 = (1,0)$ and $e_2 = (0,1)$. A closed step word of length ℓ is a sequence

$$\omega = (d_1, \dots, d_\ell), \qquad d_i \in \{\pm e_1, \pm e_2\}, \qquad \sum_{i=1}^{\ell} d_i = 0.$$

Define its partial-sum set by

$$P_\omega = \{0, d_1, d_1 + d_2, \dots, d_1 + \cdots + d_\ell\},$$

and set

$$r_x(\omega) = \max_{z \in P_\omega} z_1 - \min_{z \in P_\omega} z_1, \qquad r_y(\omega) = \max_{z \in P_\omega} z_2 - \min_{z \in P_\omega} z_2.$$

Let C_ℓ be the set of closed words of length ℓ, and put $w_\ell = |C_\ell|$.

We use two exact closed-walk sums. The first is the standard square-lattice tree-entropy identity [SW00]:

$$\sum_{\ell \ge 1} \frac{w_\ell}{\ell 4^\ell} = \log 4 - \frac{4G_{\mathrm{Cat}}}{\pi}. \tag{5.4}$$

This series is absolutely convergent: $w_{2a} = \binom{2a}{a}^2$ for $a \ge 1$, so its nonzero terms are $O(a^{-2})$. The second is

$$\sum_{\ell \ge 1} \frac{1}{\ell 4^\ell} \sum_{\omega \in C_\ell} (r_x(\omega) + 1) = \log \frac{4}{1 + \sqrt{2}}. \tag{5.5}$$

We include the proof of (5.5). Let $a \ge 1$ and $h = 2a$. A one-dimensional bridge of length h is a sequence $\gamma = (\gamma_0, \dots, \gamma_h)$ such that $\gamma_0 = \gamma_h = 0$ and $\gamma_t - \gamma_{t-1} \in \{-1, 1\}$ for $1 \le t \le h$. Define

$$\mathrm{span}(\gamma) := \max_{0 \le t \le h} \gamma_t - \min_{0 \le t \le h} \gamma_t.$$

Thus $\mathrm{span}(\gamma)+1$ is the number of integer levels visited by γ. For an integer $k>0$ and a bridge that visits k, let σ_k be its first visit to k, and set

$$\gamma'_t = \begin{cases} \gamma_t, & t \le \sigma_k, \\ 2k-\gamma_t, & t \ge \sigma_k. \end{cases}$$

This reflection is a bijection with unrestricted walks of length h from 0 to $2k$. The negative levels are symmetric, while level zero is visited by every bridge. Summing over levels rather than over bridges,

$$\begin{aligned} \sum_{\substack{\gamma_0=\gamma_h=0 \\ \gamma_t-\gamma_{t-1}\in\{-1,1\},\ 1\le t\le h}} (\mathrm{span}(\gamma)+1) &= \binom{h}{a} + \sum_{k=1}^{a}\left[\binom{h}{a+k}+\binom{h}{a-k}\right] \\ &= 2^h. \end{aligned}$$

For even ℓ, condition a two-dimensional closed word on the even number h of its horizontal steps and on their positions. Write $[\zeta^0]$ for extraction of the constant term in the formal variable ζ. The bridge identity is also immediate when $h=0$. Using it gives

$$\begin{aligned} \sum_{\omega\in C_\ell}(r_x(\omega)+1) &= \sum_{\substack{0\le h\le \ell \\ h \text{ even}}} \binom{\ell}{h} 2^h \binom{\ell-h}{(\ell-h)/2} \\ &= [\zeta^0](2+\zeta+\zeta^{-1})^\ell = \binom{2\ell}{\ell}. \end{aligned}$$

There are no closed words for odd ℓ. For $|t|<1/4$, define

$$\mathcal{B}(t) := \sum_{\ell\ge 1}\binom{2\ell}{\ell}\frac{t^\ell}{\ell} = 2\log\frac{2}{1+\sqrt{1-4t}}.$$

At $t=\pm 1/4$, its terms are $O(\ell^{-3/2})$, so both boundary series converge. Abel's theorem extends the displayed closed form to both endpoints, applied to $\mathcal{B}(-t)$ for the negative endpoint. Taking the even part gives

$$\frac{\mathcal{B}(1/4)+\mathcal{B}(-1/4)}{2} = \log\frac{4}{1+\sqrt{2}},$$

which proves (5.5).

Define

$$N(M;\omega) := \left|\{z\in\mathbb{Z}^2 : z+P_\omega\subseteq\Phi,\ (z+P_\omega)\cap M\neq\varnothing\}\right|.$$

Its bounding face rectangle has $r_x(\omega)+1$ face columns and $r_y(\omega)+1$ face rows. The corresponding primal vertex window has $r_x(\omega)+2$ columns and $r_y(\omega)+2$ rows and meets $V_\partial^{\mathrm{vac}}$. For fixed ω, translation by its coordinatewise minima makes the map from a start z to this window's lower-left corner injective. If either span is too large for Φ, there is no admissible start and the desired estimate is immediate. Otherwise, taking $p=r_x(\omega)+1$ and $q=r_y(\omega)+1$ in (5.3) yields

$$N(M;\omega) \le |V_\partial^{\mathrm{vac}}| + g_x(r_x(\omega)+1) + g_y(r_y(\omega)+1). \tag{5.6}$$

5.4 The Determinant Estimate

Recall that M and $U = \Phi \setminus M$ are respectively the marked and unmarked face sets. Every face in M belongs to a face-star indexed by $V_\partial^{\mathrm{vac}}$. Since every king component of $V_\partial^{\mathrm{vac}}$ meets the boundary, contraction of $E^\star_{V_\partial^{\mathrm{vac}}}$ joins all these face vertices to the exterior dual vertex. On the other hand, a face indexed by U has no corner in $V_\partial^{\mathrm{vac}}$, so all four primal boundary edges remain. Thus the bounded faces of $R_{r,s}[\Lambda_{r,s} \setminus V_\partial^{\mathrm{vac}}]$ are exactly those indexed by U.

For $Y \subseteq \Phi$, let A_Y be the adjacency matrix of the face-grid graph induced by Y, and let I denote the identity matrix of the appropriate order. Rooted planar duality and the Matrix–Tree Theorem [Kir47] give

$$\tau(r,s) = \det(4I - \mathsf{A}_\Phi), \qquad \tau(\Lambda_{r,s} \setminus V_\partial^{\mathrm{vac}}) = \det(4I - \mathsf{A}_U). \tag{5.7}$$

Primal bridges become dual loops and do not affect the determinant. When $U = \varnothing$, the second determinant is the 0×0 determinant and equals one.

The horizontal edges of $\Lambda_{r,s} \setminus V_\partial^{\mathrm{vac}}$ number $|\Lambda_{r,s} \setminus V_\partial^{\mathrm{vac}}| - \rho_x$, and the vertical edges number $|\Lambda_{r,s} \setminus V_\partial^{\mathrm{vac}}| - \rho_y$. Therefore

$$|E(R_{r,s}[\Lambda_{r,s} \setminus V_\partial^{\mathrm{vac}}])| = 2|\Lambda_{r,s} \setminus V_\partial^{\mathrm{vac}}| - \rho_x - \rho_y.$$

Since the graph is connected, Euler's formula, together with $g_x = \rho_x - s$, $g_y = \rho_y - r$, and $|\Phi| = (r-1)(s-1)$, gives

$$|U| = |E(R_{r,s}[\Lambda_{r,s} \setminus V_\partial^{\mathrm{vac}}])| - |\Lambda_{r,s} \setminus V_\partial^{\mathrm{vac}}| + 1 = |\Lambda_{r,s} \setminus V_\partial^{\mathrm{vac}}| - \rho_x - \rho_y + 1,$$

and therefore

$$|M| = |V_\partial^{\mathrm{vac}}| + g_x + g_y. \tag{5.8}$$

The spectral radius of the rectangular face adjacency is

$$2\cos\frac{\pi}{r} + 2\cos\frac{\pi}{s} < 4.$$

Since A_U is a principal submatrix, its spectral radius is also below 4, with the empty matrix assigned spectral radius zero. Hence, for $Y \in \{\Phi, U\}$, the matrix trace tr satisfies

$$\log\det(4I - \mathsf{A}_Y) = |Y|\log 4 - \sum_{\ell \geq 1} \frac{\operatorname{tr}(\mathsf{A}_Y^\ell)}{\ell 4^\ell}. \tag{5.9}$$

Both log-determinant series are absolutely convergent by the preceding spectral-radius bound.

The trace difference counts precisely the closed-walk starts in Φ whose walk meets M:

$$\operatorname{tr}(\mathsf{A}_\Phi^\ell) - \operatorname{tr}(\mathsf{A}_U^\ell) = \sum_{\omega \in C_\ell} N(M;\omega). \tag{5.10}$$

All terms used below are nonnegative, and (5.4)–(5.5) show that the dominating series are finite. Thus Tonelli's theorem justifies the ensuing interchanges of sums. Subtract the two instances of (5.9), substitute (5.10), and apply the charging bound (5.6). Evaluating the three resulting series by

(5.4) and (5.5), the g_y term by symmetry between the coordinates, and then inserting (5.8) and the identifications (5.7), gives

$$\begin{aligned}\log \frac{\tau(r,s)}{\tau(\Lambda_{r,s} \setminus V_\partial^{\mathrm{vac}})} &\geq |M| \log 4 - |V_\partial^{\mathrm{vac}}| \left(\log 4 - \frac{4G_{\mathrm{Cat}}}{\pi} \right) - (g_x + g_y) \log \frac{4}{1+\sqrt{2}} \\ &= \frac{4G_{\mathrm{Cat}}}{\pi} |V_\partial^{\mathrm{vac}}| + \alpha(g_x + g_y) \\ &\geq \frac{4G_{\mathrm{Cat}}}{\pi} |V_\partial^{\mathrm{vac}}|.\end{aligned}$$

We have proved the boundary factor

$$\frac{\tau(\Lambda_{r,s} \setminus V_\partial^{\mathrm{vac}})}{\tau(r,s)} \leq \exp\left(-\frac{4G_{\mathrm{Cat}}}{\pi} |V_\partial^{\mathrm{vac}}| \right). \tag{5.11}$$

Completion of the proof of Lemma 4.1. Substitute (4.4) and (5.11) into (4.3). Since $\Lambda_{r,s} \setminus S$ is the disjoint union of $V_\partial^{\mathrm{vac}}, J_1, \ldots, J_m$, we obtain

$$\frac{\tau(S)}{\tau(r,s)} \leq \exp\left(-\frac{4G_{\mathrm{Cat}}}{\pi} |\Lambda_{r,s} \setminus S| \right).$$

Both tree counts are positive because $\mathcal{L}[S]$ and $R_{r,s}$ are connected. Taking logarithms gives (4.1). □

6 Proof of Square Maximality

Proof of Theorem 1.2. Let $|S| = n^2$. If $\mathcal{L}[S]$ is disconnected, then $\tau(S) = 0 < \tau(n,n)$.

Assume that $\mathcal{L}[S]$ is connected. After translation and an interchange of coordinates, let its minimal rectangular vertex hull be $\Lambda_{r,s}$, where $r \leq s$. Then

$$rs \geq n^2, \qquad |\Lambda_{r,s} \setminus S| = rs - n^2.$$

Applying Lemmas 3.1 and 4.1 and adding their log-difference estimates gives

$$\log \tau(S) - \log \tau(n,n) \leq -\frac{4G_{\mathrm{Cat}}}{\pi}(rs - n^2) + \frac{4G_{\mathrm{Cat}}}{\pi}(rs - n^2) = 0.$$

This proves the inequality.

If $(r,s) \neq (n,n)$, Lemma 3.1 is strict, so equality is impossible. If $r = s = n$, then the hull and S both have n^2 vertices, hence $S = \Lambda_{n,n}$. Undoing the translation and lattice symmetry gives the stated equality characterization. □

7 Discussion

The rectangular area penalty bounds the gain available from enlarging the hull, and the normalized hull inequality charges every omitted vertex by exactly the same entropy constant. In log-difference form the two estimates add, their constants cancel, and strictness comes entirely from the rectangular estimate.

Use of AI tools

A substantial part of the reasoning in this paper was developed with the assistance of a large language model. Both runs used Work mode with Speed mode enabled. GPT-5.6 Sol, at the Extra High reasoning setting, worked on the initial prompt for 87 minutes 7 seconds and returned incomplete results. GPT-5.6 Sol Ultra continued from those results under the shorter follow-up prompt and produced its result after 199 minutes 9 seconds. The raw output became the basis for a substantial part of the present argument, in particular Lemmas 3.1 and 4.1. I manually checked the argument, prepared the bibliography, checked each citation against its source, and wrote the present note. Both prompts are published in Appendix A.

I am deeply grateful to Sloan Nietert for many helpful suggestions and insightful discussions. I also thank him for collaborating with me on the prompting strategy for this problem and, in particular, for his "breakthrough" prompt.

Final verification (August 22, 2026). The square-maximality theorem as stated here for arbitrary finite square-lattice subgraphs, together with its equality characterization and complete proof chain, has been formalized in Lean 4 and checked by the Lean kernel. The formalization and its reproducible trust audit are available in the public repository. I also completed a paper-to-Lean specification audit matching the definitions, hypotheses, intermediate results, and conclusions in the formalization to the mathematical claims in this manuscript.

Palomar registration (August 23, 2026). The Lean verification has been registered in the Palomar Registry at a fixed commit of the public submission repository. The registration metadata records the role of the model and the subsequent human and kernel-level checks.

References

[BP93] Robert Burton and Robin Pemantle. Local characteristics, entropy and limit theorems for spanning trees and domino tilings via transfer-impedances. *The Annals of Probability*, 21(3):1329–1371, 1993.

[Chu99] Fan R. K. Chung. Spanning trees in subgraphs of lattices. In Michael D. Fried, editor, *Applications of Curves over Finite Fields*, volume 245 of *Contemporary Mathematics*, pages 201–219. American Mathematical Society, Providence, RI, 1999.

[Kir47] Gustav Kirchhoff. Ueber die auflösung der gleichungen, auf welche man bei der untersuchung der linearen vertheilung galvanischer ströme geführt wird. *Annalen der Physik*, 148(12):497–508, 1847.

[PTF22] Ariel D. Procaccia and Jamie Tucker-Foltz. Compact redistricting plans have many spanning trees. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 3754–3771. SIAM, 2022.

[SW00] Robert Shrock and F. Y. Wu. Spanning trees on graphs and lattices in d dimensions. *Journal of Physics A: Mathematical and General*, 33(21):3881–3902, 2000.

[Swa10] Konrad J. Swanepoel. Comment on "number of spanning trees in a grid". MathOverflow, 2010. Comment posted January 7, 2010.

[Tap24] Kristopher Tapp. Spanning tree bounds for grid graphs. *The Electronic Journal of Combinatorics*, 31(1):P1.26, 2024.

[Zha26a] Jiechen Zhang. The balancing theorem for spanning trees of rectangular grid graphs. *arXiv preprint arXiv:2605.23773*, 2026.

[Zha26b] Jiechen Zhang. Extremal spanning trees in product grid graphs. *arXiv preprint arXiv:2606.24016*, 2026.

A Prompts

Initial research prompt.

```
Current task statement

Let \(\mathcal L\) be the standard nearest-neighbor square-lattice graph with vertex set \(\mathbb
Z^2\): distinct vertices \(x,y\in\mathbb Z^2\) are adjacent exactly when \(\|x-y\|_1=1\). For a
finite set \(S\subset\mathbb Z^2\), let \(\mathcal L[S]\) be the induced subgraph. For a finite
connected graph \(G\), let \(\tau(G)\) denote its number of spanning trees. Set \(\tau(G)=0\) when
\(G\) is disconnected. Let

  \[
  Q_n=P_n\square P_n=\mathcal L[\{0,1,\ldots,n-1\}^2]
  \]

be the \(n\times n\) square grid graph, which has \(n^2\) vertices.

Resolve the following conjecture completely in dimension two:

SQUARE-MAXIMALITY THEOREM. For every integer \(n\ge 1\) and every finite set \(S\subset\mathbb
Z^2\) with \(|S|=n^2\),

  \[
  \tau(\mathcal L[S])\le \tau(Q_n).
  \]

Moreover, equality holds if and only if \(S\) is an \(n\times n\) lattice square, up to an
automorphism of the square lattice; equivalently, up to translation, reflections, and interchange
of the two coordinates.

The disconnected case is immediate under the convention \(\tau=0\), so the substantive case is that
\(\mathcal L[S]\) is connected. The theorem must cover arbitrary connected induced subgraphs.
Shapes may be nonconvex, have holes, have cutvertices or bridges, have disconnected intersections
with individual rows or columns, and have long tentacles.

REQUIRED BACKGROUND

Read the following two papers in full before committing to a proof route:
  1. https://arxiv.org/abs/2605.23773
  2. https://arxiv.org/abs/2606.24016

Treat their proved theorems as available background, with precise citation and hypothesis checking.
In particular:
  - For rectangular grids, if \(AB=ab\) and \(A\le a\le b\le B\), the first paper proves

  \[
  \log\tau(P_a\square P_b)-\log\tau(P_A\square P_B) \ge
  \operatorname{arsinh}(1)\bigl((A+B)-(a+b)\bigr),
  \]

with strict inequality unless the two rectangles agree. Consequently, \(Q_n\) is already the unique
maximizer among rectangular grids with \(n^2\) vertices.
  - The second paper proves the analogous perfect-power result within pure Cartesian-product
  families and states the induced-subgraph conjecture above as its \(d=2\) case.
  - Every connected induced \(n^2\)-vertex lattice subgraph satisfies

  \[
  |E|\le 2n(n-1), \qquad |E|-|V|+1\le (n-1)^2,
  \]
```

with equality only for Q_n. These edge-count and cycle-rank facts are useful but are explicitly too coarse by themselves.

- The product-spectrum and hyperbolic-residual arguments in the papers rely essentially on Cartesian-product structure. Do not silently extend them to arbitrary induced subgraphs without a new comparison theorem and complete error control.

You may use standard results such as Kirchhoff's Matrix-Tree Theorem, the weighted matrix-tree theorem, Rayleigh monotonicity, the transfer-current theorem, Wilson's and Aldous-Broder algorithms, planar duality for spanning trees, Schur complements, determinant identities, Cauchy-Binet, and standard majorization or rearrangement theorems. Every invoked result must nevertheless be stated precisely enough that its hypotheses and the direction of its implication can be audited.

For this benchmark, work under the premise that an exact affirmative finite proof is available to be found. A response that only labels the conjecture open, difficult, plausible, asymptotically supported, or computationally verified is not responsive. If exact integer calculations produce a genuine counterexample, that supersedes the premise and gives a complete negative resolution; approximate numerical evidence does not.

The finished argument must establish the stated theorem for every $n\ge1$, including all equality cases. It may not impose extra conditions such as simple connectivity, orthogonal convexity, row- or column-convexity, bridgelessness, bounded aspect ratio, or product structure.

WHAT DOES NOT COUNT

Intermediate advances are useful internally, but the delivered solution must settle the whole theorem. None of the following, standing alone, is an acceptable endpoint:

- proving only the rectangular, Ferrers, Young-diagram, convex, simply connected, or hole-free cases;
- proving an asymptotic inequality, an inequality only up to $e^{o(n^2)}$, or a statement only for sufficiently large n;
- observing that the square minimizes boundary or maximizes edges, average degree, face count, or cycle rank;
- applying a general upper bound on $\tau(G)$ that remains above $\tau(Q_n)$;
- asserting without proof that a compression, symmetrization, hole-filling operation, or boundary-vertex transfer increases the number of spanning trees;
- reducing the theorem to a new lemma that is equivalent or nearly equivalent in strength, without proving that lemma;
- proving a comparison for one eigenvalue or a few spectral moments while the full determinant comparison remains missing;
- computational verification through any fixed n, or floating-point determinant comparisons;
- omitting strictness or the equality case;
- invoking a theorem whose hypotheses fail for induced lattice subgraphs with holes, bridges, cutvertices, or nonsimple planar duals.

MULTIAGENT RESEARCH MANAGEMENT

Exploit multiagent v2 at full scale, coordinating as many as 64 simultaneous agents when useful. Allocate agents adaptively rather than reserving a predetermined number for each strategy. Launch the search with a broad collection of mathematically independent directions. In the early stages, keep the leading hypothesis private from most agents so that the group does not collapse onto one seductive but incomplete route. The initial collection should span genuinely different mechanisms, including:

- Minimal-counterexample and structural induction. Try to prove that an extremizer has no holes, no bridges, interval row or column sections, orthogonal convexity, or some other structure that eventually forces a rectangle.

- Lattice-specific compressions and local vertex transfers. Derive exact determinant or effective-resistance formulas for moving an exposed vertex into a concavity while preserving cardinality, connectedness, and inducedness.
- Grounded-Laplacian determinant comparison. Explore Schur complements, rank-one or low-rank updates, log-determinant inequalities, Gaussian-free-field integral representations, and weighted interpolations between shapes.
- Heat-trace or spectral-zeta comparison. Seek a valid rearrangement inequality strong enough to integrate to the determinant, while explicitly testing whether pointwise heat-trace domination is even true.
- Planar methods. Explore dual spanning trees, deletion-contraction in the primal and dual, the Temperley bijection, Kasteleyn determinants, dimers, electrical networks, and face-by-face transformations, with explicit treatment of holes and bridges.
- Uniform-spanning-tree and random-walk methods. Explore Wilson and Aldous-Broder, transfer currents, loop-erased random walk, and a sharp extension of Tapp's multiplier and escape-probability framework.
- Graphic-matroid and tree-polynomial methods. Explore basis-generating polynomials, strong log-concavity, negative dependence, injections between spanning-tree sets, and lattice-aware edge-exchange maps.
- Extremal graph theory constrained by the lattice. Explore degree sequences, planar embeddings, block decompositions, and bounds that exploit lattice geometry rather than merely maximum degree.
- Exact computation for discovery and falsification. Enumerate small lattice animals up to translation and lattice symmetry, compute τ by exact integer determinants, test proposed local moves, and locate the first failure of every candidate sublemma.

Regard this as a menu of independent research programs, not as an indication that one route is intended. Sustain several mutually incompatible proof architectures over multiple rounds.

Two especially clean possible endgames are:

1. Construct a finite sequence of shape operations that preserves the admissible class, never lowers τ, is strictly improving away from a square, and terminates at Q_n.
2. Prove a direct determinant or tree-count inequality valid for every admissible S, with equality precisely at Q_n.

Any other fully rigorous architecture is equally valid.

Keep a live research ledger indexed by genuine mathematical mechanism. Each entry should record:

- the key statement in fully quantified form;
- the exact role that statement would play in a complete proof;
- every hypothesis on which it depends;
- exact computational experiments already run;
- counterexamples to natural stronger versions;
- the narrowest unresolved step;
- a status tag such as active, blocked, disproved, or merged.

Classify routes by mathematical content rather than cosmetic differences in presentation. When too many agents cluster around one idea, move some of them to neglected directions. Elegance of a reduction is not evidence of completion: a route that merely renames the main difficulty as a “compression theorem,” “rearrangement principle,” or “extremal-domain lemma” has made little progress until that statement is proved. If a program reaches a missing lemma essentially as hard as the original theorem, freeze that program. Resume it only after a new invariant, construction, exact identity, or induction mechanism appears.

Delay cross-fertilization until independently developed routes are mature enough that their actual advantages and defects are visible. Demand mathematical artifacts from every agent: complete lemma proofs, exact formulas, determinant quotients, resistance estimates, explicit tree maps, certified counterexamples to subsidiary claims, or finite certificates that another agent can check. Discard vague progress summaries, unsupported confidence, continuum analogies with no finite discretization, and assertions that a global compatibility step is automatic.

All computational evidence used in reasoning must be exact. Suitable tools include fraction-free Bareiss elimination, symbolic determinants, and independently reproduced integer Matrix-Tree calculations. Use computation to reveal patterns and destroy false claims, not as a substitute for the uniform proof. Archive the coordinates and exact tree counts of every shape that refutes a proposed local rule.

ADVERSARIAL AUDIT REQUIREMENTS

Maintain dedicated skeptical review throughout the search. Every prospective proof must survive checks for the following errors:

- confusing a vertex set in $\mathbb Z^2$ with a polyomino made of unit cells, thereby shifting the graph or miscounting vertices, faces, or boundary;
- inferring a tree-count comparison from boundary size or edge count when the graphs are not nested;
- invoking Rayleigh monotonicity without a valid common network and an actual conductance increase;
- using a vertex relocation that changes cardinality, breaks connectedness, fails to produce an induced subgraph, leaves an isolated component, or has no well-founded termination measure;
- comparing cofactors based on inconsistent roots or vertex identifications;
- dropping the Laplacian zero mode, the Matrix-Tree normalization, or a sign in a determinant or heat-kernel integral;
- asserting coordinatewise eigenvalue order or spectral majorization rather than proving it;
- using planar duality while ignoring bridges, dual loops, parallel dual edges, holes, the unbounded face, or disconnected collections of faces;
- proposing a combinatorial tree map without a complete proof of well-definedness and of injectivity or surjectivity as required;
- establishing weak monotonicity but leaving unexplained nonsquare equality cases;
- replacing the target by an equivalent “domain monotonicity,” “compression,” or “entropy” assertion and treating that assertion as known;
- using an asymptotic expansion to decide finite changes whose logarithmic gain can stay bounded;
- relying on computed positivity without a symbolic estimate uniform in n.

Every proposed geometric move must come with a full certificate containing:

1. its precise definition on lattice vertex sets;
2. preservation of $|S|$ and of the induced-subgraph interpretation;
3. preservation of connectedness, or a complete analysis of any exceptional case;
4. an exact formula for the resulting tree-count ratio or difference;
5. a proof of its sign in every allowed configuration;
6. the full characterization of equality;
7. a well-founded quantity that guarantees termination;
8. proof that terminal configurations are exactly the squares, or a justified final passage to the established rectangular theorem.

Every determinant, spectral, or heat-kernel route must supply:

1. the precise matrices or operators;
2. a common ambient space, or a rigorous embedding or intertwining construction;
3. correct handling of the zero eigenvalue and all normalization factors;
4. a proved Loewner, majorization, trace, or integral comparison with the needed orientation;
5. convergence and endpoint estimates for all integral formulas;
6. equality conditions propagated through every step.

The coordinating agent should repeatedly combine results, attack them, redirect effort, and initiate further rounds. A failed opening round is not a stopping condition. Once a serious lemma emerges, send independent agents to seek separate proofs, construct exact counterexamples, analyze equality, and test whether the lemma merely restates the target. Retain multiple distinct routes until one closes the entire argument.

Before endorsing a proof, commission a clean-room review by agents who did not build it. They must reconstruct the dependency graph, verify citations and algebra, test local assertions on exact small examples, and actively try to break the proof. Correct all defects and repeat this review until no gap remains.

FINAL-ANSWER REQUIREMENTS

Return a self-contained, publication-quality mathematical proof, not a research diary. The final response must:

- state the theorem precisely;
- define all notation;
- state and prove every new lemma in dependency order;
- identify precisely which results are cited from the two supplied papers or standard literature;
- prove the inequality for every \(n\ge1\);
- separately prove the equality characterization;
- include exact finite verification only where genuinely needed for finitely many base cases;
- contain no unsupported “clearly,” “routine,” or “by symmetry” at a substantive step;
- contain no appeal to the multiagent process, numerical evidence, or the premise that a proof exists.

Failure of the current collection of ideas is not a reason to terminate. Start additional rounds and revisit frozen routes only when a substantively new device appears. Deliver the main answer only after a complete argument has passed the independent audit.

If a nonnegotiable execution limit interrupts the investigation, never disguise a conjectural chain as a proof. Report only rigorously established lemmas, exact failures of attempted routes, and a single precisely formulated outstanding obstacle, all clearly marked incomplete. This is an emergency fallback, not a chosen endpoint.

Allocate no less than eight hours of sustained search before considering termination. Internet lookup is permitted for the two supplied papers, routine background, and standard named results. Do not use web search to hunt for a preexisting solution to this exact benchmark, and do not substitute an “open problem” report for the requested investigation.

Follow-up research prompt.

Please build upon these initial results and resolve the conjecture. Aim for a breakthrough and consider a diverse set of approaches, not limited by previous work. Please try for at least 8 hours before returning an incomplete result.