The Center-Area Lemma

Three disjoint unit squares whose centers form a non-obtuse triangle span area at least $\frac{1}{2}$

David R. MacIver

david@drmaciver.com

Abstract

Let S_1, S_2, S_3 be unit squares, rotated arbitrarily, with pairwise disjoint interiors, and let P_1, P_2, P_3 be their centers. If the triangle $P_1P_2P_3$ is non-obtuse, its area is at least $\frac{1}{2}$. The constant is attained, and the angle hypothesis is necessary: obtuse center triangles can have arbitrarily small area. The proof freezes one owned separating inequality per pair of squares, minimizes the area of the center triangle subject to the three frozen constraints, converts first-order optimality at the minimum into a reciprocal diagram of contact forces in Maxwell's sense, and closes with an ownership tournament and two scalar inequalities. The proof is fully formalized in Lean 4, and the statement is independently confirmed by a second machine-checked proof. This paper was produced with a mix of ChatGPT and Claude Fable and has not yet been adequately human reviewed.

1 Introduction

Let $s(n)$ denote the side of the smallest square that contains n unit squares with pairwise disjoint interiors; for the state of the art on this packing problem see Friedman's dynamic survey [3], and for the asymptotic regime Erdős and Graham [1]. Lower bounds for $s(n)$ are counting arguments that track the centers of the packed squares, and their basic pairwise input is one-dimensional: two centers lie at distance at least 1. The first genuinely two-dimensional constraint concerns *triples* of centers, and its natural quantitative form is a lower bound on the area of the center triangle. This paper proves the sharp such bound.

Throughout, a *unit square* is a closed square of side 1, in any orientation, and a family of unit squares is *admissible* if their interiors are pairwise disjoint. A triangle is *non-obtuse* if all three of its angles are at most $\pi/2$; equivalently, for each vertex the squared length of the opposite side is at most the sum of the squared lengths of the two adjacent sides. (The formalization reported in Section 9 uses the squared-side-length form.)

Theorem 1.1 (Center-area lemma)**.** *Let S_1, S_2, S_3 be unit squares in the plane, rotated arbitrarily and independently, with pairwise disjoint interiors, and let P_1, P_2, P_3 be their centers. If $\triangle P_1P_2P_3$ is non-obtuse, then*

$$\operatorname{area}(P_1P_2P_3) \;\geq\; \tfrac{1}{2}.$$

We record at once the pairwise fact used throughout.

Lemma 1.2 (center separation)**.** *The centers of two unit squares with disjoint interiors are at distance at least* 1.

Proof. Each unit square contains the open disk of radius $\frac{1}{2}$ about its center in its interior. If the centers were at distance less than 1, their midpoint would lie in both disks, so the interiors would meet. $\square$

The constant and the hypotheses are tight, in senses made precise in Section 8. The constant is attained: three axis-parallel unit squares centered at three vertices of a unit square are admissible with center triangle of area exactly $\frac{1}{2}$. The angle hypothesis cannot be dropped: three axis-parallel squares in a row, the middle one slightly offset, give an obtuse center triangle of arbitrarily small positive area. Without any angle hypothesis one still has $\operatorname{area}(P_1P_2P_3) \geq \frac{1}{2}\sin\theta$ for every angle θ of the center triangle. And there is no four-center analogue: four centers in strict convex position can span convex hulls of arbitrarily small area. The theorem is also exactly the hard part of a natural three-center threshold governing when three centers of admissible unit squares fit in a $w \times h$ rectangle: granted the elementary narrow-rectangle regime of that threshold, its remaining wide branch is equivalent to Theorem 1.1, so a first-principles proof of that branch would not bypass the lemma but reprove it.

The proof runs by contradiction in four layers. Suppose some admissible triple has a non-obtuse center triangle of area less than $\frac{1}{2}$, and write $D_0 < 1$ for its doubled area. First, freeze one separating branch per pair: the separating-axis theorem gives, for each pair of squares, a separating direction that is, up to sign, a side normal of one of the two squares, called the *owner* of the branch; the resulting separating inequality is linear in the two centers, with fixed unit normal n_i and fixed threshold $H_i \geq 1$. Second, minimize the doubled area $D = 2\operatorname{area}(P_1P_2P_3)$ over center triples satisfying the three frozen linear inequalities together with non-obtuseness and $D \leq D_0$; the minimum exists, satisfies $D \leq D_0$, and has all three branch inequalities active. Third, first-order optimality at the minimum (a finite Farkas argument) produces multipliers $\lambda_i \geq 0$ whose contact forces $f_i = \lambda_i n_i$ close into a triangle congruent to the center triangle rotated a quarter turn, a reciprocal force diagram in Maxwell's sense, and Euler homogeneity converts multipliers into area at the minimum:

$$2D \;=\; \lambda_1 H_1 + \lambda_2 H_2 + \lambda_3 H_3,$$

where each threshold has the explicit form $H(t) = (1 + \cos t + \sin t)/2 \geq 1$, with $t \in [0, \pi/4]$ the angle between the orientations of the two squares reduced modulo quarter turns. Fourth, ownership directs each edge of the center triangle away from its owner; a tournament on three vertices is cyclic or transitive, and each case ends in a short scalar inequality (a symmetric rational inequality in half-angle cotangents in the cyclic case, an explicit one-variable dual certificate in the transitive case), forcing $2D \geq 2$ and contradicting $D \leq D_0 < 1$.

One distinction governs the language of the entire paper: *square contact* versus *active frozen branch.* The proof never shows that extremal squares touch, and it never moves or rotates a square. For each pair, one valid separating inequality is selected at the outset and kept fixed while the centers alone vary; throughout, *contact* (equivalently, an *active* branch) means equality in one of these selected inequalities, never geometric tangency of squares.

Sections 2–7 give the proof in the order of the plan: the branch formalism, the extremal configuration, the reciprocal force system with its sine equations, the cyclic case, the transitive case together with the origin of its certificate, and the assembly of the contradiction; Section 8 establishes sharpness, Section 9 reports the machine verification, and Appendix A proves the cyclic scalar inequality.

2 Frames, separating branches, and ownership

The main proof replaces the disjointness of each pair of squares—a nonconvex condition—by a single linear inequality between their centers. This section defines these inequalities (*branches*), proves that disjoint squares always admit one whose normal is an axis of one of the two squares (the *owner*, Lemma 2.2), and computes the threshold of such a branch: a function $H(t) \geq 1$ of the two frames alone (Lemma 2.3).

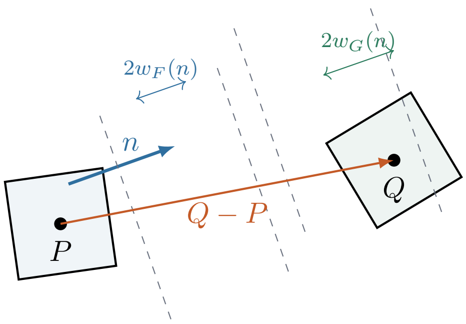

Figure 1: An oriented separating branch. The branch records only one-dimensional data: the squares project onto the line in direction n to intervals of lengths $2w_F(n)$ and $2w_G(n)$ (support lines dashed), and inequality (1) separates the interiors of these intervals, hence the interiors of the squares.

Definition 2.1. A *frame* is an ordered orthonormal pair $F = (e, f)$ of vectors in $\mathbb{R}^2$; the four unit vectors $\pm e, \pm f$ are its *axes*. The unit square with frame F and center P is

$$Q(F, P) = \left\{X \in \mathbb{R}^2 : \ |(X - P) \cdot e| \le \tfrac{1}{2}, \ |(X - P) \cdot f| \le \tfrac{1}{2}\right\}.$$

For a unit vector n, the *projected half-width* of F in direction n is

$$w_F(n) = \frac{|n \cdot e| + |n \cdot f|}{2},$$

so that the orthogonal projection of $Q(F, P)$ onto a line in direction n is an interval of length $2w_F(n)$ centered at the projection of P. An *oriented separating branch* for the squares $Q(F, P)$ and $Q(G, Q)$ is the inequality

$$(Q - P) \cdot n \ \ge \ H(F, G; n), \qquad H(F, G; n) := w_F(n) + w_G(n), \tag{1}$$

specified by its unit *normal* n; the number $H(F, G; n)$ is the *threshold* of the branch.

A branch is a sufficient condition for disjoint interiors. A linear functional attains its maximum over a square only on the boundary, so every $X \in \operatorname{int} Q(F, P)$ satisfies $(X - P) \cdot n < w_F(n)$, and every $X \in \operatorname{int} Q(G, Q)$ satisfies $(Q - X) \cdot n < w_G(n)$; a point in both interiors would give $(Q - P) \cdot n < w_F(n) + w_G(n)$, contradicting (1). The condition is only sufficient: a branch records nothing about the squares beyond their one-dimensional projections onto the line in direction n (Figure 1).

Disjointness does not supply an arbitrary branch: it supplies one of a special form.

Lemma 2.2 (Owned separating branch)**.** *Two unit squares with disjoint interiors admit an oriented separating branch* (1) *whose normal is, up to sign, an axis of one of the two frames.*

Proof. Write the squares as $Q(F, P)$ and $Q(G, Q)$ and form the Minkowski sum $K = Q(F, 0) + Q(G, 0)$, which by central symmetry of the summands is also the Minkowski difference $Q(F, 0) - Q(G, 0)$. A point X lies in both interiors exactly when $X = P + a = Q + b$ with $a \in \operatorname{int} Q(F, 0)$ and $b \in \operatorname{int} Q(G, 0)$, that is, exactly when $Q - P = a - b \in \operatorname{int} K$. Hence the interiors are disjoint if and only if $Q - P \notin \operatorname{int} K$.

Since (e, f) is an orthonormal basis, $Q(F, 0) = [-\frac{1}{2}, \frac{1}{2}]\, e + [-\frac{1}{2}, \frac{1}{2}]\, f$, and likewise for $Q(G, 0)$. Thus K is a Minkowski sum of four segments: a centrally symmetric convex polygon (a zonogon) whose edges are parallel to the four axis directions of F and G. The outward normal of an edge is a quarter turn of the edge direction, hence the other axis of the same frame; so the facet normals of K are frame axes. Now K is the intersection of its facet half-planes

$\{x : x \cdot n \le h_K(n)\}$, with n ranging over the outward facet normals and $h_K(n) = \max_{x \in K} x \cdot n$ the support function; so $Q - P \notin \operatorname{int} K$ forces $(Q - P) \cdot n \ge h_K(n)$ for some facet normal n. Support functions add under Minkowski sums, and $\max_{x \in Q(F,0)} x \cdot n = w_F(n)$; therefore $h_K(n) = w_F(n) + w_G(n)$, and (1) holds with this n. □

When the normal of a branch is, up to sign, an axis of one of the two frames, we call the square with that frame an *owner* of the branch. Lemma 2.2 says that every disjoint pair admits an owned branch. The threshold of an owned branch has an explicit form, best expressed through a notion of distance between frames. The set of axes of a frame is invariant under rotation by $\pi/2$, so a frame determines a point of the circle $\mathbb{R}/(\pi/2)\mathbb{Z}$, its axis direction modulo quarter turns. For frames F and G, let $d_4(F, G)$ be the geodesic distance between the corresponding points; it takes values in $[0, \pi/4]$ and equals the least angle between an axis of F and an axis of G.

Lemma 2.3 (Threshold of an owned branch)**.** *Let $F = (e, f)$ and G be frames and let n be a unit vector.*

(i) If n is an axis of F, then $w_F(n) = \frac{1}{2}$.

(ii) If m is an axis of G, then $w_G(n) = \dfrac{|n \cdot m| + |\det(n, m)|}{2}$; in particular $w_G(n) \ge \frac{1}{2}$ for every unit n.

(iii) If n is an axis of F and $t = d_4(F, G)$, then

$$H(F, G; n) \;=\; H(t) \;:=\; \frac{1 + \cos t + \sin t}{2} \;\ge\; 1. \tag{2}$$

Consequently the threshold of an owned branch depends only on the two frames, not on which frame supplies the axis: whether n is an axis of F or of G, the threshold equals $H(d_4(F, G))$.

Proof. (i) If $n = \pm e$ then $|n \cdot e| = 1$ and $|n \cdot f| = 0$; symmetrically for $n = \pm f$.

(ii) The axes of G are $\pm m$ and $\pm m'$, where m' is a quarter turn of m; then $|n \cdot m'| = |\det(n, m)|$, so $w_G(n) = \big(|n \cdot m| + |n \cdot m'|\big)/2 = \big(|n \cdot m| + |\det(n, m)|\big)/2$. (The formula is independent of the choice of axis m: replacing m by m' swaps the two terms.) The numbers $a = n \cdot m$ and $b = n \cdot m'$ are the coordinates of n in an orthonormal basis, so $a^2 + b^2 = \|n\|^2 = 1$ and $|a| + |b| \ge \sqrt{a^2 + b^2} = 1$, giving $w_G(n) \ge \frac{1}{2}$.

(iii) Choose an axis m of G at angle $t = d_4(F, G)$ from n; then $|n \cdot m| = \cos t$ and $|\det(n, m)| = \sin t$, so by (i) and (ii)

$$H(F, G; n) = w_F(n) + w_G(n) = \frac{1}{2} + \frac{\cos t + \sin t}{2} = H(t).$$

On $[0, \pi/4]$ we have $\cos t + \sin t > 0$ and $(\cos t + \sin t)^2 = 1 + \sin 2t \ge 1$, so $\cos t + \sin t \ge 1$ and $H(t) \ge 1$.

The final claim follows from (2) and the symmetry of d_4. □

By (2), ownership carries no metric information: an owned branch has the same threshold whichever square owns it. Ownership matters for a discrete reason. Two branches owned by the same square use two axes of a single frame, so their normals are, up to sign, parallel or perpendicular. This fact closes the transitive case of the main argument.

3 Freezing the branches: the extremal configuration

The proof of Theorem 1.1 is by contradiction: we assume a configuration of admissible squares whose non-obtuse center triangle has doubled area less than 1. Disjointness is an awkward condition to minimize over, but each pair of squares certifies its disjointness by an owned branch (Lemma 2.2), which is a single linear inequality in the centers. We freeze one such branch per pair and minimize the area of the center triangle subject to the three frozen inequalities alone. This section constructs the resulting extremal configuration: the minimum exists, is acute, and has doubled area strictly between 0 and 1; and at the minimum all three branch inequalities are equalities. The remainder of the proof works exclusively with this minimizer.

3.1 The minimum exists

Assume, then, that Theorem 1.1 fails: there are admissible unit squares S_1, S_2, S_3 whose centers P_1, P_2, P_3 form a non-obtuse triangle with

$$D_0 \;=\; 2\,\mathrm{area}(P_1P_2P_3) \;<\; 1.$$

By Lemma 1.2 the centers are pairwise at distance at least 1, hence distinct; and three distinct collinear points give the middle one an angle of π, contradicting non-obtuseness. The centers are therefore noncollinear, so $D_0 > 0$. Write $D = 2\,\mathrm{area}(P_1P_2P_3)$ for the doubled area of a center triple and

$$\sigma \;=\; \operatorname{sign}\det(P_2 - P_1,\, P_3 - P_1) \;\in\; \{-1, 1\}, \qquad D \;=\; \sigma\det(P_2 - P_1,\, P_3 - P_1) \;>\; 0,$$

evaluated at the original configuration. The sign σ is retained from this configuration throughout: in the minimization below the objective is the signed determinant $\sigma\det(P_2 - P_1, P_3 - P_1)$, constrained to be nonnegative, and D denotes this quantity.

Compactness of the feasible set will come from an a priori bound on the sides of feasible triangles, which rests on one elementary fact about triangles that are non-obtuse, have long sides, and small area.

Lemma 3.1 (Subcritical triangles)**.** *Let XYZ be a triangle that is non-obtuse and has all sides of length at least 1. Then every altitude of XYZ has length at least $1/\sqrt{2}$. If moreover its doubled area is less than 1, then every side has length less than $\sqrt{2}$ and the triangle is acute.*

Proof. Consider the altitude from a vertex A to the opposite side. Since the two base angles are at most $\pi/2$, the foot lies on the side itself, not on its extension. Take coordinates with the foot at the origin and $A = (0, h)$, $h > 0$ the altitude, so the base endpoints are $B = (-x, 0)$ and $C = (y, 0)$ with $x, y \geq 0$. The adjacent sides have length at least 1: $x^2 + h^2 \geq 1$ and $y^2 + h^2 \geq 1$. The angle at A is at most $\pi/2$:

$$0 \;\leq\; (B - A)\cdot(C - A) \;=\; -xy + h^2,$$

so $h^2 \geq xy$. If $h^2 < \frac{1}{2}$, then $x^2 \geq 1 - h^2 > \frac{1}{2} > h^2$ forces $x > h$, and likewise $y > h$, whence $xy > h^2$, a contradiction. Hence $h \geq 1/\sqrt{2}$.

Now suppose the doubled area D satisfies $D < 1$. Each side times its corresponding altitude equals D, so each side has length $D/h < 1/(1/\sqrt{2}) = \sqrt{2}$. Finally, if some angle were exactly $\pi/2$, the two sides adjacent to it would be legs of length at least 1, giving $D = \mathrm{leg} \times \mathrm{leg} \geq 1$. So every angle is strictly less than $\pi/2$. $\square$

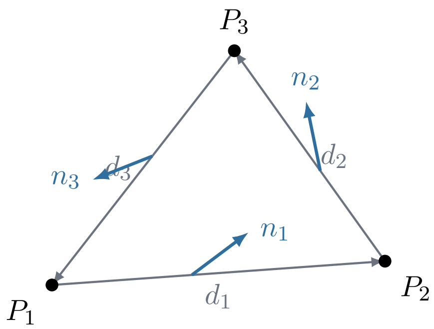

Figure 2: The frozen system: cyclic edges d_1, d_2, d_3 and branch normals n_1, n_2, n_3. Each branch is a linear inequality (3) in the centers and remains sufficient for separation as the centers move. The normals are axes of the frozen frames, not functions of the edges, and need not align with them.

The freeze. For the pairs (S_1, S_2), (S_2, S_3), (S_3, S_1), Lemma 2.2 supplies owned separating branches. Orient them cyclically: choose the sign of each normal so that the branch for (S_i, S_{i+1}) reads off the difference $P_{i+1} - P_i$ (the threshold is unchanged, since projected half-widths are even in the normal). This yields unit normals n_1, n_2, n_3 and thresholds H_1, H_2, H_3, and for each branch we fix once and for all one owner: a square whose frame has the normal as an axis. By Lemma 2.3 and (2), each $H_i \geq 1$. Now *freeze* all of this data: frames, normals, thresholds, and owners are constants from here on. Only the centers vary, subject to

$$(P_2 - P_1) \cdot n_1 \;\geq\; H_1, \qquad (P_3 - P_2) \cdot n_2 \;\geq\; H_2, \qquad (P_1 - P_3) \cdot n_3 \;\geq\; H_3. \tag{3}$$

With the cyclic edges $d_1 = P_2 - P_1$, $d_2 = P_3 - P_2$, $d_3 = P_1 - P_3$ (so $d_1 + d_2 + d_3 = 0$), the system (3) reads $d_i \cdot n_i \geq H_i$ for $i = 1, 2, 3$.

Two features of (3) carry the whole construction. First, the system remains *sufficient* for separation as the centers move: a branch certifies disjoint interiors from projection data of the two frames alone, and the frames are frozen, so any center triple satisfying (3) yields three admissible squares. The original configuration is one solution among many; we are free to minimize. Second, the normals n_i are axes of the frozen frames and have no relation to the edges d_i beyond the inequalities themselves: n_i need not be parallel or perpendicular to d_i, and Figure 2 draws them askew deliberately.

The minimization problem. Branches and area depend only on differences of centers, so we normalize $P_1 = 0$. Over $(P_2, P_3) \in \mathbb{R}^2 \times \mathbb{R}^2$, minimize

$$\sigma \det(P_2 - P_1,\; P_3 - P_1)$$

subject to the three branch inequalities (3), the three non-obtuseness inequalities (for each vertex, the square of the opposite side is at most the sum of the squares of the adjacent sides), and $0 \leq D \leq D_0$.

Lemma 3.2 (Existence and geometry of the minimum)**.** *The minimum is attained. Every minimizing configuration satisfies* $0 < D \leq D_0 < 1$*, and its center triangle is acute.*

Proof. We show the feasible set is nonempty and compact, and that Lemma 3.1 applies uniformly across it.

Nonempty. The original centers, translated so that $P_1 = 0$, are feasible: they satisfy (3) because the branches were constructed from them, they are non-obtuse by hypothesis, and their signed determinant is D_0.

Every feasible triple is a subcritical triangle. Let (P_2, P_3) be feasible. By Cauchy–Schwarz and (3), with indices cyclic,

$$\|P_{i+1} - P_i\| \;\geq\; (P_{i+1} - P_i) \cdot n_i \;\geq\; H_i \;\geq\; 1,$$

so all three sides are at least 1; in particular the centers are distinct. They are also noncollinear: for three distinct collinear points, the two edge vectors at the middle point are opposite in direction, so their inner product is negative and the non-obtuseness inequality at that point fails. This holds at *every* feasible point, not merely the original one. Noncollinearity makes the determinant nonzero, so the constraint $0 \leq D \leq D_0$ tightens to $0 < D \leq D_0 < 1$, and D is the doubled area of the triple. Lemma 3.1 now applies throughout the feasible set: every feasible triangle has all sides less than $\sqrt{2}$ and is acute.

Compactness. With $P_1 = 0$, the side bounds confine P_2 and P_3 to the closed ball of radius $\sqrt{2}$. The constraints are finitely many non-strict polynomial inequalities, so the feasible set is closed, hence compact, and nonempty; the continuous objective attains its minimum.

At a minimizer, feasibility gives $0 < D \leq D_0 < 1$ and, by Lemma 3.1 again, acuteness. □

From now on *the minimizing configuration* means a fixed minimizer from Lemma 3.2; the symbols P_i, d_i, D refer to its data, while the frames, normals n_i, thresholds H_i, and owners remain the frozen data of the original counterexample.

3.2 Complete contact

At the minimizing configuration, call a frozen branch *active* if its inequality in (3) holds with equality. Activity is equality in a linear inequality on centers and nothing more: it does not assert that the squares touch, and no square is moved or rotated anywhere in the argument. This subsection proves that all three branches are active, so that the first-order optimality analysis to come has all three constraints to work with. The obstruction to descent at a vertex comes from its active branches only, and the dual description of a single half-plane is the key.

Lemma 3.3 (Half-plane duality)**.** *Let $n \in \mathbb{R}^2$ be a unit vector and let $g \in \mathbb{R}^2$ satisfy $g \cdot v \geq 0$ for every v with $n \cdot v \geq 0$. Then $g = \mu n$ with $\mu \geq 0$.*

Proof. If $v \perp n$ then both $\pm v$ satisfy $n \cdot (\pm v) = 0 \geq 0$, so $g \cdot v \geq 0$ and $g \cdot (-v) \geq 0$, forcing $g \cdot v = 0$. Thus g is orthogonal to $n^\perp$, so $g = \mu n$ for some $\mu \in \mathbb{R}$, and $v = n$ gives $\mu = g \cdot n \geq 0$. □

Two preliminaries make single-vertex variations available. First, the area gradient. Let J_σ denote rotation through $\sigma\pi/2$, so that $J_\sigma^2 = -I$ and $\langle J_\sigma u, v\rangle = \sigma \det(u, v)$. When one vertex P moves and the other two stay fixed, D is an affine function of P; its gradient g is the quarter turn of the opposite edge. Indeed, as a function of P_3,

$$D \;=\; \sigma \det(d_1,\, P_3 - P_1) \;=\; \langle J_\sigma d_1,\, P_3 - P_1\rangle,$$

with gradient $J_\sigma d_1$; as a function of P_2 the gradient is $-J_\sigma(P_3 - P_1) = J_\sigma d_3$; and since D is translation invariant the three gradients sum to zero, so the gradient in P_1 is $J_\sigma d_2$. In every case

$$\|g\| \;=\; \text{length of the edge opposite } P \;\geq\; 1,$$

the bound being that edge's *own* frozen branch together with Cauchy–Schwarz, as in the proof of Lemma 3.2. Second, the normalization $P_1 = 0$ does not forbid varying P_1: after moving P_1, translate all three centers back so that $P_1 = 0$ again. The system (3) and the area depend only on differences of centers, so feasibility and objective value are unchanged.

Lemma 3.4 (Complete contact)**.** *At the minimizing configuration, all three inequalities of* (3) *hold with equality:* $d_i \cdot n_i = H_i$ *for* $i = 1, 2, 3$.

Proof. Form the graph on $\{P_1, P_2, P_3\}$ in which the branch for a pair, if active, joins its two vertices; we rule out any vertex of active degree 0 or 1, and since a graph on three vertices with at most two edges has such a vertex, the graph must be complete. Suppose then that some vertex P has active degree at most 1, and let g be the gradient of D in P, so $\|g\| \geq 1$. Figure 3 shows the two cases.

Active degree 0. Both branches incident to P have positive slack. Move P to $P - tg$ for small $t > 0$: the incident branches are affine in P and keep positive slack, and the third branch does not involve P. Since D is affine in P,

$$D(P - tg) \;=\; D - t\|g\|^2,$$

which is strictly smaller. For small t the remaining constraints persist: acuteness and $D > 0$ are open conditions satisfied at the minimizer (Lemma 3.2), and $D \leq D_0$ persists because D decreased. This produces a feasible configuration of smaller objective, contradicting minimality.

Active degree 1. Let the unique active branch incident to P join P to Q, and orient its normal locally so that it reads $(P - Q) \cdot n \geq H$: here H is its frozen threshold, and n is the frozen normal or its negative, whichever makes the frozen inequality take this form. Activity says $(P - Q) \cdot n = H$. Consider any velocity v of P with $n \cdot v \geq 0$. Moving P to $P + tv$ preserves the active branch, since $(P + tv - Q) \cdot n = H + t\, n \cdot v \geq H$; the inactive branch at P tolerates small motions by positive slack; the third branch does not involve P. If $g \cdot v < 0$, then D strictly decreases along this motion, and exactly as in the previous case the open conditions and the upper bound $D \leq D_0$ persist for small t, contradicting minimality. Hence $g \cdot v \geq 0$ for every v with $n \cdot v \geq 0$, and Lemma 3.3 gives $g = \mu n$ with $\mu \geq 0$; since $\|g\| \geq 1$ and $\|n\| = 1$, in fact $\mu \geq 1$.

With the other two vertices fixed, the map sending a position X of P to the resulting value of D is affine with gradient g. It takes the value D at $X = P$, and it takes the value 0 when X is placed at *either* of the other two vertices, because the triangle then degenerates and the determinant vanishes. For an affine function the gradient pairs with a difference of arguments to give the difference of values, so

$$g \cdot (P - Q') \;=\; D - 0 \;=\; D \qquad \text{for either other vertex } Q'.$$

Both vertices qualify, so in particular $Q' = Q$, the vertex the active branch joins to P, whichever of the two it is. Setting $r = P - Q$, activity gives $n \cdot r = H$, and therefore

$$D \;=\; g \cdot r \;=\; \mu\, n \cdot r \;=\; \mu H \;\geq\; 1,$$

since $\mu \geq 1$ and $H \geq 1$ by (2). This contradicts $D \leq D_0 < 1$ (Lemma 3.2). Neither degree 0 nor degree 1 can occur, so every vertex has active degree 2 and all three branches are active. □

Remark (where first-order feasibility is used). Both cases of the proof rest on the same principle. Take a velocity v at a single vertex that weakly relaxes every active branch there, meaning $n \cdot v \geq 0$ for each locally oriented active normal n. Along the ray $P + tv$, $t > 0$ small, those affine inequalities continue to hold, and every inactive branch stays satisfied by positive slack. Because the minimizing triangle is acute with $0 < D < 1$, the strict conditions $D > 0$ and acuteness are open and persist, and whenever the area derivative $g \cdot v$ is negative the bound $D \leq D_0$ persists as well. So a feasible velocity with negative area derivative yields a genuinely smaller feasible configuration, which minimality forbids.

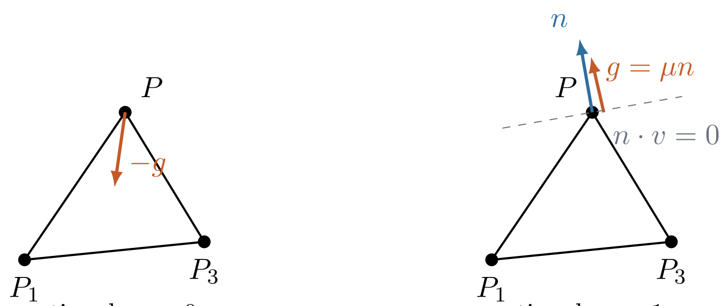

active degree 0:
move down the area gradient

active degree 1:
the gradient lies on the normal ray

Figure 3: The two ways contact could fail at a vertex P. With no active incident branch, P moves down the area gradient. With exactly one, first-order minimality over the half-plane $n \cdot v \geq 0$ forces the gradient onto the normal ray, $g = \mu n$ with $\mu \geq 1$, and then $D = \mu H \geq 1$.

4 The reciprocal triangle and the sine system

At the minimizing configuration of Lemma 3.2 all three frozen branches are active (Lemma 3.4). We now convert this first-order information into the central object of the proof: nonnegative multipliers $\lambda_1, \lambda_2, \lambda_3$ whose forces $f_i = \lambda_i n_i$ close into a quarter-turned copy of the center triangle, together with the identity $2D = \lambda_1 H_1 + \lambda_2 H_2 + \lambda_3 H_3$. Dotting the force equations with the normals then reduces the geometry to three scalar equations, whose case analysis occupies the rest of the paper.

Throughout we work at the minimizing configuration, with the normalization $P_1 = 0$ in force. Write $d_1 = P_2 - P_1$, $d_2 = P_3 - P_2$, $d_3 = P_1 - P_3$ for the cyclic edges, so that $d_1 + d_2 + d_3 = 0$ and $D = \sigma \det(d_1, d_2)$. Let J_σ denote rotation through $\sigma\pi/2$; we use exactly two facts about it:

$$J_\sigma^2 = -I \qquad \text{and} \qquad J_\sigma u \cdot v = \sigma \det(u, v).$$

Complete contact says that each frozen inequality of (3) holds with equality,

$$d_i \cdot n_i = H_i, \qquad H_i \geq 1 \quad (i = 1, 2, 3),$$

the lower bound because each frozen branch is owned (Lemma 2.3).

4.1 Farkas multipliers and the reciprocal triangle

First-order optimality under linear constraints is a finite Farkas statement [2]. We isolate the version we need; the strict dual vector in hypothesis (i) is what makes the generated cone closed, and our minimization problem supplies it for free.

Lemma 4.1 (Finite Farkas lemma)**.** *Let $a_1, \ldots, a_m$ and g be vectors in a Euclidean space. Suppose*

(i) there is a vector w with $a_i \cdot w > 0$ for every i;

(ii) every v with $a_i \cdot v \geq 0$ for all i satisfies $g \cdot v \geq 0$.

Then $g = \sum_{i=1}^m \lambda_i a_i$ for some $\lambda_1, \ldots, \lambda_m \geq 0$.

Proof. Let $K = \{\sum_i \lambda_i a_i : \lambda_i \geq 0\}$, the convex cone generated by the a_i. We first show K is closed. Set $c = \min_i a_i \cdot w > 0$ and suppose $x_k = \sum_i \lambda_{i,k} a_i \to x$ with all $\lambda_{i,k} \geq 0$. Dotting with w gives $x_k \cdot w \geq c \sum_i \lambda_{i,k}$, and the left side converges, so the coefficient sums are bounded.

The coefficient vectors $(\lambda_{1,k}, \ldots, \lambda_{m,k})$ therefore have a convergent subsequence, whose limit is a nonnegative coefficient vector representing x. Hence $x \in K$.

Now suppose $g \notin K$. Separating the point g from the closed convex set K yields v and β with $x \cdot v \geq \beta$ for all $x \in K$ and $g \cdot v < \beta$. Since $0 \in K$, $\beta \leq 0$. If some $x_0 \in K$ had $x_0 \cdot v < 0$, then $tx_0 \in K$ for all $t > 0$ and $tx_0 \cdot v \to -\infty$, contradicting the lower bound; so $x \cdot v \geq 0$ on all of K. In particular $a_i \cdot v \geq 0$ for every i, yet $g \cdot v < \beta \leq 0$, contradicting (ii). □

We apply the lemma in explicit coordinates. With $P_1 = 0$, let $x = (P_2, P_3) \in \mathbb{R}^4$ and define the branch functions

$$b_1(x) = P_2 \cdot n_1, \qquad b_2(x) = (P_3 - P_2) \cdot n_2, \qquad b_3(x) = (-P_3) \cdot n_3,$$

so that the frozen inequalities (3) read $b_i(x) \geq H_i$. Each b_i is *linear* in x — not merely affine — and this is the point of the normalization $P_1 = 0$: linearity is used twice below, once for the strict dual vector and once for Euler homogeneity. The gradients are the constant vectors

$$a_1 = (n_1, 0), \qquad a_2 = (-n_2, n_2), \qquad a_3 = (0, -n_3),$$

and we write g for the gradient at the minimizer of the objective $x \mapsto \sigma \det(P_2, P_3)$, the doubled signed area, whose value there is D.

We check the two hypotheses of Lemma 4.1. For (ii), let $v \in \mathbb{R}^4$ satisfy $a_i \cdot v \geq 0$ for $i = 1, 2, 3$. By linearity, $b_i(x + tv) = b_i(x) + t\, a_i \cdot v \geq H_i$ for all $t \geq 0$, so the three branch constraints hold along the entire ray. Acuteness (Lemma 3.2) and $D > 0$ hold strictly at the minimizer, so both persist for small $t > 0$. If $g \cdot v < 0$, then for small $t > 0$ the point $x + tv$ is still acute, still has positive doubled area, and has strictly smaller objective value — in particular still at most D_0 — so it is feasible and beats the minimum. This contradiction shows $g \cdot v \geq 0$. For (i), take $w = x$ itself, the uniform dilation direction: linearity gives $a_i \cdot x = b_i(x)$, and complete contact gives $b_i(x) = H_i \geq 1 > 0$. Lemma 4.1 therefore produces

$$g = \lambda_1 a_1 + \lambda_2 a_2 + \lambda_3 a_3, \qquad \lambda_1, \lambda_2, \lambda_3 \geq 0.$$

Define the *forces* $f_i = \lambda_i n_i$.

The multiplier equation lives in $\mathbb{R}^4$; we resolve it into plane components. From $\nabla_a \det(a, b) = -Jb$ and $\nabla_b \det(a, b) = Ja$, the gradient of $\sigma \det(P_2, P_3)$ has components $\nabla_{P_2} = -J_\sigma P_3$ and $\nabla_{P_3} = J_\sigma P_2$. Reading $g = \sum_i \lambda_i a_i$ in the two components:

$$-J_\sigma P_3 = \lambda_1 n_1 - \lambda_2 n_2 = f_1 - f_2, \qquad J_\sigma P_2 = \lambda_2 n_2 - \lambda_3 n_3 = f_2 - f_3.$$

Since $P_1 = 0$ we have $d_1 = P_2$ and $d_3 = -P_3$, so these are the first and third of the equations

$$J_\sigma d_1 = f_2 - f_3, \qquad J_\sigma d_2 = f_3 - f_1, \qquad J_\sigma d_3 = f_1 - f_2, \tag{4}$$

and the middle one follows from $d_1 + d_2 + d_3 = 0$, since the three right-hand sides also sum to zero; the vanishing of both sums is the consistency check on the resolution.

Equations (4) say that f_1, f_2, f_3 are the vertices of a triangle whose sides are the sides of the center triangle rotated through a quarter turn, while each vertex f_i lies on the ray from the origin through the frozen normal n_i, at distance λ_i (Figure 4). This is a reciprocal diagram in Maxwell's classical sense [6]: the multipliers are contact forces along the frozen normals, and the force polygon closes. This quarter-turned copy of the center triangle, pinned to the three normal rays, is the central object of the proof.

The multipliers convert area into thresholds. The objective is homogeneous of degree 2 in x and each b_i is homogeneous of degree 1 — the second use of linearity. Euler's identity gives $g \cdot x = 2D$, and $a_i \cdot x = b_i(x) = H_i$ at contact, so

$$2D = g \cdot x = \sum_{i=1}^{3} \lambda_i\, a_i \cdot x = \lambda_1 H_1 + \lambda_2 H_2 + \lambda_3 H_3. \tag{5}$$

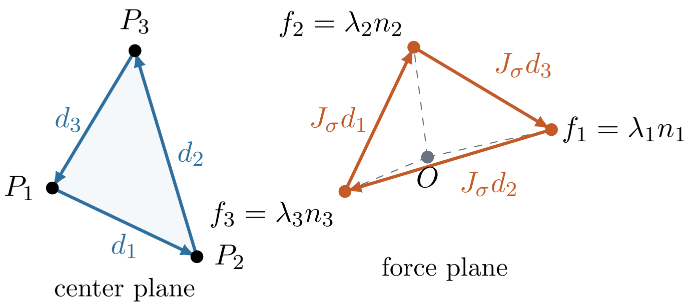

Figure 4: The reciprocal diagram. Left: the center triangle with cyclic edges d_1, d_2, d_3. Right: the force plane. The points $f_i = \lambda_i n_i$ lie on the rays spanned by the frozen normals, and by (4) the sides of the triangle $f_1 f_2 f_3$ are the quarter-turned edges $J_\sigma d_i$.

Since $0 < D \leq D_0 < 1$ at the minimizer, the remaining task of the entire paper is to prove $\lambda_1 H_1 + \lambda_2 H_2 + \lambda_3 H_3 \geq 2$: by (5) this forces $D \geq 1$, the desired contradiction.

4.2 The sine system and the owner tournament

The reciprocal equations are vector equations; the case analysis needs scalars. Dotting with the normals produces them. Define the *oriented sines*

$$s_1 = \sigma \det(n_1, n_2), \qquad s_2 = \sigma \det(n_2, n_3), \qquad s_3 = \sigma \det(n_3, n_1).$$

Since $J_\sigma^2 = -I$, the first equation of (4) is equivalent to $d_1 = J_\sigma(f_3 - f_2)$. Dot with n_1, use $J_\sigma u \cdot v = \sigma \det(u, v)$, and use the contact equality $d_1 \cdot n_1 = H_1$:

$$H_1 = d_1 \cdot n_1 = J_\sigma(f_3 - f_2) \cdot n_1 = \lambda_3 \sigma \det(n_3, n_1) - \lambda_2 \sigma \det(n_2, n_1) = \lambda_2 s_1 + \lambda_3 s_3.$$

The second and third equations yield, in identical fashion, the other two lines of the *sine system*

$$H_1 = \lambda_2 s_1 + \lambda_3 s_3, \qquad H_2 = \lambda_3 s_2 + \lambda_1 s_1, \qquad H_3 = \lambda_1 s_3 + \lambda_2 s_2. \tag{6}$$

One case dies immediately.

Proposition 4.2 (Nonpositive sine). *If $s_i \leq 0$ for some i, then $\lambda_1 + \lambda_2 + \lambda_3 \geq 2$, and consequently $\lambda_1 H_1 + \lambda_2 H_2 + \lambda_3 H_3 \geq 2$.*

Proof. The oriented sine of two unit vectors lies in $[-1, 1]$; in particular every $s_i \leq 1$. Suppose first $s_3 \leq 0$. In the first equation of (6) drop the nonpositive term $\lambda_3 s_3$ and bound s_1 by 1:

$$H_1 = \lambda_2 s_1 + \lambda_3 s_3 \leq \lambda_2 s_1 \leq \lambda_2.$$

In the second, bound both sines by 1: $H_2 = \lambda_3 s_2 + \lambda_1 s_1 \leq \lambda_1 + \lambda_3$. Since $H_1, H_2 \geq 1$ this gives $\lambda_2 \geq 1$ and $\lambda_1 + \lambda_3 \geq 1$, hence $\lambda_1 + \lambda_2 + \lambda_3 \geq 2$. As every $H_i \geq 1$ and every $\lambda_i \geq 0$, $\sum_i \lambda_i H_i \geq \sum_i \lambda_i \geq 2$. The cases $s_1 \leq 0$ and $s_2 \leq 0$ reduce to this one by cyclic relabeling: shifting the vertex labels cyclically shifts every index by one, leaves σ unchanged, and preserves the displayed form of (6) (see the relabeling conventions at the end of this section). $\square$

Standing assumption. *From here through Section 6 we assume $s_1, s_2, s_3 > 0$.* The three normals then occur in positive cyclic order: each consecutive pair (n_1, n_2), (n_2, n_3), (n_3, n_1) is positively oriented for the orientation σ.

The surviving case is organized by ownership. The freeze of Section 3 fixed, for each pair of squares, a branch and a choice of its owner. Direct each edge of the center triangle *away from* *the owner* of its branch. The result is a tournament on three vertices, and a tournament on three vertices has exactly two forms up to isomorphism: a directed 3-cycle, or a transitive order with a source S of out-degree 2, a middle vertex M, and a sink T of in-degree 2 (Figure 5). There are eight possible owner assignments — the fixed owner of each branch is one or the other square of its pair — and cyclic relabeling together with reflection reduces them to these two geometries. Section 5 closes the cyclic case, in which the normals wind once around the circle. Section 6 closes the transitive case, in which the source owns two branches and one frame supplies two perpendicular normals. Section 7 assembles the three cases into the contradiction.

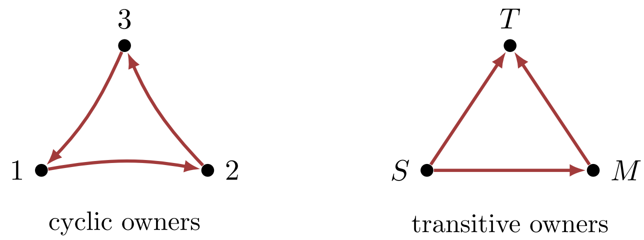

Figure 5: The two owner tournaments on three vertices, up to cyclic relabeling and reflection: a directed cycle (left) and a transitive order with source S, middle M, sink T (right). Each arrow points from the owner of a branch to the other vertex of that center-triangle edge; it is *not* the direction of the branch normal.

The reductions just invoked, and those of the next two sections, permute vertex labels; we fix the bookkeeping once. Under a relabeling of the vertices, σ is recomputed from the new vertex order. Any branch whose ordered pair of centers is reversed by the relabeling is reoriented: its normal changes to $-n$, while its threshold and its owner are unchanged. With these conventions the contact equalities $d_i \cdot n_i = H_i$ and the sine system (6) keep their displayed form under every relabeling used in the case analysis. For a cyclic shift this is immediate: every index shifts by one, σ is unchanged, and no branch is reversed. For the transposition of P_2 and P_3, σ flips and all three branches are reversed; the sign flips cancel in pairs, the sines s_1 and s_2 are exchanged while s_3 is fixed, and both systems again keep their form. Every relabeling we use is a composition of these two.

5 Cyclic owners

We now settle the cyclic tournament: each center owns the branch to the next center in cyclic order. The outcome is Proposition 5.2, $\lambda_1 + \lambda_2 + \lambda_3 \geq 2$, and hence $\lambda_1 H_1 + \lambda_2 H_2 + \lambda_3 H_3 \geq 2$, the bound the main proof requires. The route is geometric: the three normals wind once around the circle, and the directed gaps between them, in half-angle cotangent coordinates, obey a single algebraic constraint. The coordinates express the thresholds through one rational function, and the constraint turns the sine system (6) into a weighting identity for $\lambda_1 + \lambda_2 + \lambda_3$; a scalar inequality, proved in Appendix A, finishes the case. The geometry lives here; the algebra lives in the appendix.

By the relabeling conventions of Section 4 we may assume that P_i owns branch i, the branch joining P_i to P_{i+1} (indices mod 3, here and throughout). Each square then supplies exactly one normal: n_i is an axis of the frame at P_i, while the other square of that pair, centered at P_{i+1}, supplies the axis n_{i+1} of the next branch. The standing assumption of Section 4 remains in force: $s_1, s_2, s_3 > 0$.

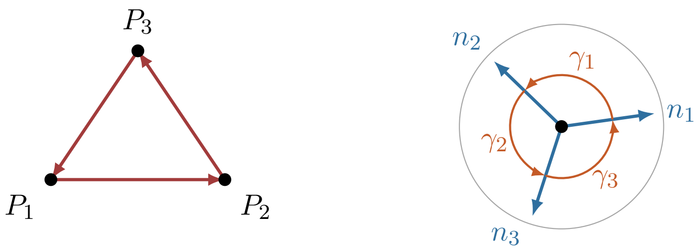

Figure 6: The cyclic case. Left: each center owns the branch to the next center; the arrows point from owner to other vertex and are not the directions of the normals. Right: the normals wind once around the circle, with directed gaps γ_i positive, less than π, and summing to 2π.

Gaps. Let γ_i be the directed angle from n_i to n_{i+1}, measured in the orientation selected by σ, so that $s_i = \sigma \det(n_i, n_{i+1}) = \sin\gamma_i$ and $n_i \cdot n_{i+1} = \cos\gamma_i$. Strict positivity of the sines places each γ_i strictly in $(0, \pi)$, and this strictness is what pins the winding: following the three gaps returns the direction to n_1, so $\gamma_1 + \gamma_2 + \gamma_3$ is a multiple of 2π, and it lies strictly between 0 and 3π, hence

$$\gamma_1 + \gamma_2 + \gamma_3 = 2\pi.$$

The normals wind exactly once around the circle (Figure 6).

Half-angle cotangents. Since each half-gap $\gamma_i/2$ lies in $(0, \pi/2)$, the half-angle cotangents

$$u = \cot\frac{\gamma_1}{2}, \qquad v = \cot\frac{\gamma_2}{2}, \qquad w = \cot\frac{\gamma_3}{2}$$

are positive coordinates for the gaps. The half-gaps sum to π, so the cotangent addition formula gives

$$\frac{uv-1}{u+v} = \cot\left(\frac{\gamma_1}{2} + \frac{\gamma_2}{2}\right) = \cot\left(\pi - \frac{\gamma_3}{2}\right) = -w,$$

and clearing the positive denominator $u + v$ gives

$$uv + vw + wu = 1. \tag{7}$$

Expanding $(u + v)(u + w) = u^2 + (uv + vw + wu)$, the constraint (7) yields the identity

$$(u + v)(u + w) = 1 + u^2 \qquad \text{(and its cyclic variants).}$$

Lemma 5.1. *With the cyclic factor* $F(q) = \dfrac{(1+q)\max\{1, q\}}{1+q^2}$,

$$H_i = \frac{1 + |\cos\gamma_i| + \sin\gamma_i}{2}, \qquad H_1 = F(u), \quad H_2 = F(v), \quad H_3 = F(w).$$

Proof. Branch i is owned by the frame at P_i, and n_i is one of its axes, so that frame contributes half-width $\frac{1}{2}$ (Lemma 2.3). The other frame of the pair, at P_{i+1}, has n_{i+1} as an axis, so the other-frame width formula of Lemma 2.3 gives

$$H_i = \frac{1}{2} + \frac{|n_i \cdot n_{i+1}| + |\det(n_i, n_{i+1})|}{2} = \frac{1 + |\cos\gamma_i| + \sin\gamma_i}{2},$$

using $|\det(n_i, n_{i+1})| = s_i = \sin\gamma_i > 0$.

For the second claim, substitute the half-angle formulas: with $q = \cot(\gamma/2)$,

$$\sin\gamma = \frac{2q}{1+q^2}, \qquad \cos\gamma = \frac{q^2-1}{1+q^2}.$$

Two regimes. If $q \le 1$, then $\gamma \ge \pi/2$, so $|\cos\gamma| = (1-q^2)/(1+q^2)$ and

$$\frac{1+|\cos\gamma|+\sin\gamma}{2} = \frac{1}{2}\cdot\frac{(1+q^2)+(1-q^2)+2q}{1+q^2} = \frac{1+q}{1+q^2} = F(q),$$

since $\max\{1,q\} = 1$. If $q \ge 1$, then $|\cos\gamma| = (q^2-1)/(1+q^2)$ and

$$\frac{1+|\cos\gamma|+\sin\gamma}{2} = \frac{1}{2}\cdot\frac{(1+q^2)+(q^2-1)+2q}{1+q^2} = \frac{q(1+q)}{1+q^2} = F(q),$$

since $\max\{1,q\} = q$. Applying this with $q = u, v, w$ at $\gamma = \gamma_1, \gamma_2, \gamma_3$ gives the three threshold formulas. □

The weighting identity. The sine system (6) expresses the thresholds in the multipliers; to extract $\lambda_1+\lambda_2+\lambda_3$ we seek weights on H_1, H_2, H_3 under which each λ_i collects a constant coefficient. The cotangent coordinates supply them:

$$2(\lambda_1+\lambda_2+\lambda_3) = H_1(u+w) + H_2(u+v) + H_3(v+w). \tag{8}$$

The pairing here is rigid: H_1 carries $u+w$, H_2 carries $u+v$, H_3 carries $v+w$, and the assignment is forced by which equations of (6) contain each λ_i; permuting the weights destroys the identity.

To prove (8), first express the sines in the coordinates. In positive cyclic order the oriented sines are the sines of the gaps, $s_i = \sin\gamma_i$, so the half-angle formulas and the identity $(u+v)(u+w) = 1+u^2$ with its cyclic variants give

$$s_1 = \frac{2u}{(u+v)(u+w)}, \qquad s_2 = \frac{2v}{(v+w)(v+u)}, \qquad s_3 = \frac{2w}{(w+u)(w+v)}.$$

Now substitute (6) into the right-hand side of (8) and collect the coefficient of each multiplier. The multiplier λ_1 appears in the equation for H_2 (term $\lambda_1 s_1$, weight $u+v$) and in the equation for H_3 (term $\lambda_1 s_3$, weight $v+w$), so its coefficient is

$$s_1(u+v) + s_3(v+w) = \frac{2u}{u+w} + \frac{2w}{w+u} = 2.$$

The coefficients of λ_2 and λ_3 are the cyclic images of this computation and equal 2 as well. This proves (8).

Proposition 5.2. *In the cyclic case $\lambda_1+\lambda_2+\lambda_3 \ge 2$, and hence*

$$\lambda_1 H_1 + \lambda_2 H_2 + \lambda_3 H_3 \;\ge\; 2.$$

Proof. By (8) and Lemma 5.1,

$$2(\lambda_1+\lambda_2+\lambda_3) = F(u)(u+w) + F(v)(u+v) + F(w)(v+w).$$

The numbers u, v, w are positive and satisfy (7), so Lemma A.1 (Appendix A) bounds the right-hand side below by 4; hence $\lambda_1+\lambda_2+\lambda_3 \ge 2$. Since $\lambda_i \ge 0$ and $H_i \ge 1$ for each i,

$$\lambda_1 H_1 + \lambda_2 H_2 + \lambda_3 H_3 \;\ge\; \lambda_1+\lambda_2+\lambda_3 \;\ge\; 2.$$ □

6 Transitive owners

We close the transitive case: when the owner tournament has a source, that source owns two branches, so both of their normals are axes of a single frame. This one discrete fact makes the case rigid. Perpendicularity of the two source normals collapses the sine system to a one-parameter family indexed by an angle $t \in (0, \pi/2)$, and a triangle inequality on the frame circle together with an explicit dual certificate delivers the bound $\lambda_1 H_1 + \lambda_2 H_2 + \lambda_3 H_3 \geq 2$ required by Section 4.

By the relabeling conventions of Section 4 we may label the vertices $P_1 = S$, $P_2 = M$, $P_3 = T$: the source, middle, and sink of the tournament. The source owns the branches $S \to M$ and $T \to S$, with normals n_1 and n_3; the middle owns the branch $M \to T$, with normal n_2. The standing assumption $s_1, s_2, s_3 > 0$ remains in force, and the relabeling keeps the contact equalities and the sine system (6) in their displayed form.

Ownership now does its work. The normals n_1 and n_3 are both, up to sign, axes of the source's frame, so they are parallel or perpendicular. If they were parallel then $s_3 = \sigma \det(n_3, n_1)$ would vanish, contrary to $s_3 > 0$; so $n_1 \perp n_3$, hence $\det(n_3, n_1) = \pm 1$, and positivity forces

$$s_3 = 1.$$

Thus (n_1, n_3) is an orthonormal basis with $\det(n_3, n_1) = \sigma$. Writing $n_2 = a\, n_1 + b\, n_3$ and expanding the determinants in this basis gives $s_1 = \sigma \det(n_1, n_2) = -b$ and $s_2 = \sigma \det(n_2, n_3) = -a$, so

$$s_1^2 + s_2^2 = a^2 + b^2 = \|n_2\|^2 = 1.$$

Since $s_1, s_2 > 0$, there is a unique $t \in (0, \pi/2)$ with

$$s_3 = 1, \qquad s_1 = c = \cos t, \qquad s_2 = s = \sin t.$$

Geometrically, the identity $J_\sigma u \cdot v = \sigma \det(u, v)$ of Section 4 gives $c = s_1 = J_\sigma n_1 \cdot n_2$, so t is the angle between $J_\sigma n_1$ and n_2.

Write

$$A = H(d_4(S, M)), \qquad B = H(d_4(M, T)), \qquad C = H(d_4(T, S)),$$

where H is the owned-branch threshold (2), and $d_4(S, M)$ abbreviates the frame distance between the frames at S and M. All three branches are owned, and by Lemma 2.3 the threshold of an owned branch depends only on the two frames of its pair, not on which frame supplies the axis. This independence is what the identification needs: the branch joining T to S is owned by S, not by T, yet its threshold is still $H(d_4(T, S))$. Hence

$$H_1 = A, \qquad H_2 = B, \qquad H_3 = C,$$

and $A, B, C \geq 1$ by (2).

The source geometry determines A explicitly. Since $J_\sigma n_1$ is a unit vector perpendicular to n_1, we have $J_\sigma n_1 = \pm n_3$, so the axes of the source's frame point along n_1 modulo quarter turns; and n_2 is an axis of the frame at M making angle t with $J_\sigma n_1$. Therefore $d_4(S, M) = \min\{t, \pi/2 - t\}$, which need not equal t. But $\cos x + \sin x$ is symmetric under $x \mapsto \pi/2 - x$, so the threshold formula is unambiguous:

$$A = H\big(\min\{t, \tfrac{\pi}{2} - t\}\big) = \frac{1 + c + s}{2}.$$

Substituting $s_3 = 1$, $s_1 = c$, $s_2 = s$ and $H_1 = A$, $H_2 = B$, $H_3 = C$ into the sine system (6) yields the *transitive system*

$$A = c\lambda_2 + \lambda_3, \qquad B = c\lambda_1 + s\lambda_3, \qquad C = \lambda_1 + s\lambda_2. \tag{9}$$

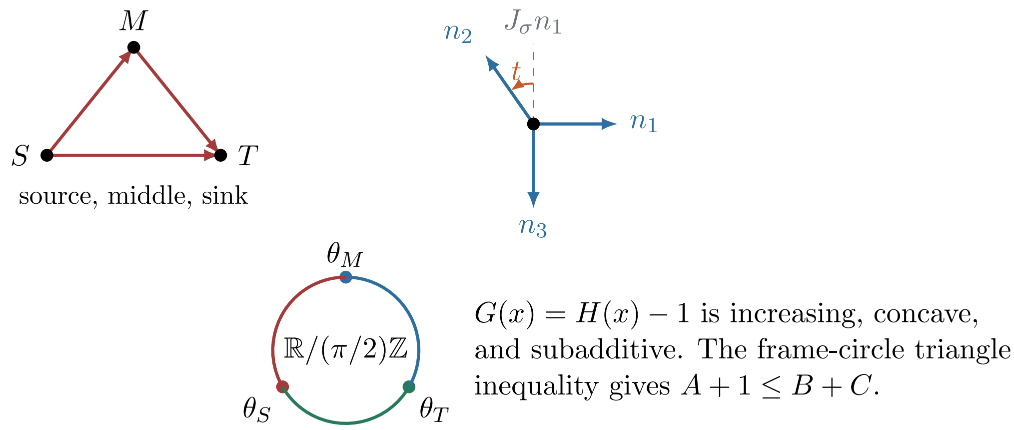

Figure 7: The transitive case. Top left: the owner tournament with source S, middle M, sink T; each arrow points from the owner of a branch to the other vertex of its edge, not along the branch normal. Top right: the source-owned normals $n_1 \perp n_3$, with n_2 at angle t from $J_\sigma n_1$. Bottom: the three frames on the frame circle $\mathbb{R}/(\pi/2)\mathbb{Z}$; since $G = H - 1$ is increasing, concave, and subadditive, the triangle inequality for d_4 gives $A + 1 \le B + C$ (Lemma 6.1).

The remaining task is the bound $A\lambda_1 + B\lambda_2 + C\lambda_3 \ge 2$.

The first ingredient is a triangle inequality among the three thresholds, measured in excess over the universal lower bound 1.

Lemma 6.1 (Frame excess)**.** *$A + 1 \le B + C$.*

Proof. We extend $G = H - 1$ to an increasing concave function vanishing at 0, deduce subadditivity, and apply the triangle inequality for d_4.

On $[0, \pi/4]$ set

$$G(x) = H(x) - 1 = \frac{\cos x + \sin x - 1}{2}.$$

Then $G(0) = 0$, $G'(x) = (\cos x - \sin x)/2 \ge 0$ with equality exactly at $x = \pi/4$, and $G''(x) = -(\cos x + \sin x)/2 < 0$; so G is increasing and concave on $[0, \pi/4]$. Extend G to $[0, \infty)$ by the constant value $G(\pi/4)$ on $[\pi/4, \infty)$. Because $G'(\pi/4) = 0$, the extension is C^1: the graph flattens tangentially rather than with a corner, and the derivative of the extension is nonnegative and nonincreasing on all of $[0, \infty)$. The extension is therefore still increasing and concave. The extension is needed because the sum of two frame distances below can be as large as $\pi/2$.

A nonnegative increasing concave function G on $[0, \infty)$ with $G(0) = 0$ is subadditive. Indeed, for $x, y \ge 0$ with $x + y > 0$, concavity along the segment from 0 to $x + y$ gives

$$G(x) = G\Big(\frac{x}{x+y}(x+y) + \frac{y}{x+y} \cdot 0\Big) \;\ge\; \frac{x}{x+y} G(x+y) + \frac{y}{x+y} G(0) = \frac{x}{x+y} G(x+y),$$

and symmetrically $G(y) \ge \frac{y}{x+y} G(x+y)$; adding the two gives $G(x) + G(y) \ge G(x+y)$. The case $x = y = 0$ is trivial.

Since d_4 is a metric on the frame circle, $d_4(S, M) \le d_4(S, T) + d_4(T, M)$. Monotonicity and subadditivity of the extension now give

$$\begin{aligned} A - 1 = G(d_4(S, M)) &\le G\big(d_4(S, T) + d_4(T, M)\big) \\ &\le G(d_4(S, T)) + G(d_4(T, M)) = (C - 1) + (B - 1), \end{aligned}$$

which rearranges to $A + 1 \le B + C$. $\square$

The second ingredient is an identity. It is stated in the half-angle variable

$$r = \tan(t/2) \in (0,1), \qquad c = \frac{1-r^2}{1+r^2}, \qquad s = \frac{2r}{1+r^2},$$

in which every quantity in sight becomes rational. The certificate data are

$$\alpha = \frac{r}{(1+r)(1+r^2)}, \qquad \beta = \frac{1+r+r^2}{2(1+r)}, \qquad K = \frac{r(1-r)(3-2r+2r^2-r^3)}{2(1+r^2)^2}.$$

Lemma 6.2 (Transitive dual certificate)**.** *Under the transitive system* (9)*,*

$$A\lambda_1 + \lambda_2 + \lambda_3 - 2 = \alpha(C-1) + \beta(B+C-A-1) + K. \tag{10}$$

Consequently $A\lambda_1 + \lambda_2 + \lambda_3 \geq 2$.

The multipliers enter the right side of (10) only through the two slacks $C-1$ and $B+C-A-1$; the constant K depends on t alone. That is the point of the certificate: it writes the target, minus 2, as a nonnegative combination of quantities already known to be nonnegative, plus a nonnegative constant.

Proof. After substituting (9) and the r-parametrization, (10) becomes an identity of polynomials; the mechanical checking recipe is given in the subsection below. Every term on the right is nonnegative: $\alpha, \beta \geq 0$ by inspection, since $r \in (0,1)$; $C - 1 \geq 0$ by (2); $B + C - A - 1 \geq 0$ by Lemma 6.1; and $K \geq 0$ because $r(1-r) > 0$ on $(0,1)$, while $2r + r^3 \leq 3$ on $[0,1]$ gives $3 - 2r + 2r^2 - r^3 \geq 2r^2 \geq 0$. Hence the left side of (10) is nonnegative. □

Proposition 6.3. *In the transitive case,* $\lambda_1 H_1 + \lambda_2 H_2 + \lambda_3 H_3 \geq 2$.

Proof. Using $H_1 = A$, $H_2 = B$, $H_3 = C$, then $B, C \geq 1$ with $\lambda_2, \lambda_3 \geq 0$, and finally Lemma 6.2,

$$\lambda_1 H_1 + \lambda_2 H_2 + \lambda_3 H_3 = A\lambda_1 + B\lambda_2 + C\lambda_3 \;\geq\; A\lambda_1 + \lambda_2 + \lambda_3 \;\geq\; 2.$$ □

6.1 Where the certificate comes from

Identity (10) can be checked in a few lines of computer algebra or by a patient hand computation. This subsection records the exact recipe, then derives α, β, and K, so that the certificate reads as the dual solution of a small linear program rather than as a rational accident.

Checking recipe. Form the difference of the two sides of (10). Substitute $B = c\lambda_1 + s\lambda_3$ and $C = \lambda_1 + s\lambda_2$ into the slacks, and eliminate λ_3 via the first equation of (9), $\lambda_3 = A - c\lambda_2$. Only the first equation is used for elimination; the second and third enter solely through B and C in the slacks. Then substitute

$$c = \frac{1-r^2}{1+r^2}, \qquad s = \frac{2r}{1+r^2}, \qquad A = \frac{1+c+s}{2} = \frac{1+r}{1+r^2},$$

together with the displayed formulas for α, β, K. Every denominator divides $2(1+r)(1+r^2)^2$, which is positive; clearing it leaves a polynomial in r and the remaining multipliers λ_1, λ_2, which expands to zero.

Derivation. The certificate draws on two nonnegative slacks: $C - 1$ (owned thresholds are ≥ 1) and $B + C - A - 1$ (Lemma 6.1). We seek coefficients α, β making

$$A\lambda_1 + \lambda_2 + \lambda_3 - 2 - \alpha(C-1) - \beta(B+C-A-1)$$

independent of the multipliers modulo the first equation of the system. Eliminating λ_3 as in the recipe, the coefficient of λ_1 in this expression is $A - \alpha - (1+c)\beta$ and the coefficient of λ_2 is $(1-c) - s\alpha - s(1-c)\beta$. Setting both to zero gives the linear system

$$\alpha + (1+c)\beta = A, \qquad s\alpha + s(1-c)\beta = 1 - c,$$

whose solution, after the half-angle substitution, is the displayed α and β. What remains is a constant in the multipliers; factored, it is exactly K. The one thing that could still have failed is the sign of that constant, and $K \geq 0$ on $(0, 1)$ as checked in the proof of Lemma 6.2. The certificate is thus the dual solution of a small linear program: (α, β) is the combination of the two slacks whose multiplier dependence matches the target exactly, and K is the unmatched constant.

7 Proof of the main theorem

The pieces are in place. It remains to check that the three cases treated by Propositions 4.2, 5.2 and 6.3 exhaust all possibilities, and to convert their common conclusion $\lambda_1 H_1 + \lambda_2 H_2 + \lambda_3 H_3 \geq 2$ into the contradiction.

Proof of Theorem 1.1. Suppose the theorem fails: there are unit squares S_1, S_2, S_3 with pairwise disjoint interiors whose centers P_1, P_2, P_3 form a non-obtuse triangle of area less than $\frac{1}{2}$. Write D_0 for its doubled area, so $D_0 < 1$. Freeze one owned branch per pair and minimize doubled area as in Lemma 3.2: a minimizing configuration exists, is acute, and satisfies $0 < D \leq D_0 < 1$. By Lemma 3.4 all three frozen branches are active there, so the Farkas multipliers $\lambda_1, \lambda_2, \lambda_3 \geq 0$ exist and satisfy the Euler identity (5),

$$2D \;=\; \lambda_1 H_1 + \lambda_2 H_2 + \lambda_3 H_3.$$

We claim this sum is at least 2. The oriented sines s_1, s_2, s_3 admit exactly three possibilities: some $s_i \leq 0$; all $s_i > 0$ with cyclic owner tournament; all $s_i > 0$ with transitive owner tournament. The list is exhaustive: when every sine is positive, directing each edge of the center triangle away from its owner yields a tournament on three vertices, and such a tournament is either a directed cycle or transitive; the relabeling conventions of Section 4 bring either geometry into the displayed form treated there, leaving thresholds and owners unchanged. In the first case Proposition 4.2 gives $\lambda_1 H_1 + \lambda_2 H_2 + \lambda_3 H_3 \geq 2$; in the cyclic case Proposition 5.2 gives the same bound; in the transitive case Proposition 6.3 does. The claim holds in every case.

By (5), then, $2D \geq 2$, so $D \geq 1$. This inequality concerns the doubled area of the *minimizing* configuration, not of the assumed counterexample; but Lemma 3.2 bounds precisely that quantity by $D \leq D_0 < 1$. The two bounds are incompatible, so no counterexample exists: every admissible configuration with a non-obtuse center triangle satisfies $\operatorname{area}(P_1P_2P_3) \geq \frac{1}{2}$. □

8 Sharpness and the limits of the statement

Theorem 1.1 is best possible in every direction: the constant $\frac{1}{2}$ is attained, the angle hypothesis cannot be dropped, and the statement does not extend to four centers. This section records these three boundaries, together with the correct bound when the angle hypothesis is dropped entirely.

Proposition 8.1 (Equality). *The constant $\frac{1}{2}$ in Theorem 1.1 is attained.*

Proof. Take axis-parallel unit squares centered at $(0,0)$, $(1,0)$, and $(0,1)$. The first two share only the edge on the line $x = \frac{1}{2}$, the first and third share only the edge on $y = \frac{1}{2}$, and the last two share only the corner $(\frac{1}{2}, \frac{1}{2})$; the interiors are pairwise disjoint. The center triangle is right isosceles with legs of length 1: non-obtuse, of area exactly $\frac{1}{2}$. $\square$

The angle hypothesis is not an artifact of the proof. Without it the center-triangle area can be made arbitrarily small.

Example 8.2 (Obtuse degeneration)**.** Fix $0 < \varepsilon < \frac{1}{2}$ and take axis-parallel unit squares centered at $(0,0)$, $(1,\varepsilon)$, and $(2,0)$. Their projections to the horizontal axis are the intervals $[-\frac{1}{2}, \frac{1}{2}]$, $[\frac{1}{2}, \frac{3}{2}]$, $[\frac{3}{2}, \frac{5}{2}]$, whose interiors are pairwise disjoint, so the squares have pairwise disjoint interiors. The center triangle has area $\frac{1}{2}|\det\big((1,\varepsilon),(2,0)\big)| = \varepsilon$. There is no contradiction with Theorem 1.1: at the middle vertex the angle between $(-1,-\varepsilon)$ and $(1,-\varepsilon)$ has cosine proportional to $\varepsilon^2 - 1 < 0$, so the triangle is obtuse there.

For arbitrary center triangles the correct general statement replaces the constant by a factor recording how far the triangle is from degenerate.

Proposition 8.3 (All-angle bound)**.** *Let S_1, S_2, S_3 be unit squares with pairwise disjoint interiors and centers P_1, P_2, P_3. Then for every angle θ of the center triangle,*

$$\mathrm{area}(P_1P_2P_3) \;\geq\; \tfrac{1}{2}\sin\theta,$$

and in particular this holds for the largest angle.

Proof. Suppose first that the centers are not collinear, and let a, b be the lengths of the two sides adjacent to θ. By Lemma 1.2 both are at least 1, so

$$\mathrm{area}(P_1P_2P_3) = \tfrac{1}{2}\,ab\sin\theta \;\geq\; \tfrac{1}{2}\sin\theta.$$

If the centers are collinear, every angle is 0 or π, the right side is 0, and the inequality holds trivially. $\square$

The obtuse degeneration makes $\sin\theta$ small only by making an angle obtuse; under the hypothesis of Theorem 1.1 the theorem beats Proposition 8.3 by removing the angle factor entirely.

Finally, no four-center analogue exists. The natural candidate quantity is the area of the convex hull of four centers in strict convex position (all four turn determinants of the cyclic boundary order strictly positive, so that the quadrilateral is genuinely two-dimensional), and no positive lower bound holds for it, even for axis-parallel unit squares that are pairwise disjoint as closed sets.

Proposition 8.4 (Four-center hull infimum)**.** *Over all quadruples of unit squares with pairwise disjoint interiors whose centers are in strict convex position, the infimum of the area of the convex hull of the centers is* 0*. This holds even restricted to axis-parallel unit squares that are pairwise disjoint as closed sets. The infimum is not attained.*

Proof. Fix $L > 1$ and $\varepsilon > 0$, and take axis-parallel unit squares centered at

$$p_0 = (0,0), \qquad p_1 = (L,\varepsilon), \qquad p_2 = (2L,-\varepsilon), \qquad p_3 = (3L,0).$$

The horizontal projections of the squares are intervals of length 1 centered at 0, L, $2L$, $3L$; since $L > 1$, consecutive centers are more than 1 apart, so these intervals are pairwise disjoint and the closed squares are pairwise disjoint.

In the cyclic order p_0, p_2, p_3, p_1 the four turn determinants are

$$\begin{aligned}
\det(p_2 - p_0,\ p_3 - p_2) &= \det\big((2L, -\varepsilon), (L, \varepsilon)\big) = 3L\varepsilon,\\
\det(p_3 - p_2,\ p_1 - p_3) &= \det\big((L, \varepsilon), (-2L, \varepsilon)\big) = 3L\varepsilon,\\
\det(p_1 - p_3,\ p_0 - p_1) &= \det\big((-2L, \varepsilon), (-L, -\varepsilon)\big) = 3L\varepsilon,\\
\det(p_0 - p_1,\ p_2 - p_0) &= \det\big((-L, -\varepsilon), (2L, -\varepsilon)\big) = 3L\varepsilon,
\end{aligned}$$

all strictly positive, so the four centers are in strict convex position with counterclockwise boundary order p_0, p_2, p_3, p_1. (The cyclic order matters: in the order p_0, p_1, p_2, p_3 the turn determinants are $-3L\varepsilon$, $+3L\varepsilon$, $+3L\varepsilon$, $-3L\varepsilon$, not all of one sign, so that quadrilateral is not convex.) The shoelace formula in the boundary order sums the vertex cross products over consecutive pairs; the two terms involving $p_0 = (0, 0)$ vanish, while $\det(p_2, p_3) = \det\big((2L, -\varepsilon), (3L, 0)\big) = 3L\varepsilon$ and $\det(p_3, p_1) = \det\big((3L, 0), (L, \varepsilon)\big) = 3L\varepsilon$. Hence

$$\operatorname{area}\operatorname{conv}\{p_0, p_1, p_2, p_3\} = \tfrac{1}{2}\,|\,0 + 3L\varepsilon + 3L\varepsilon + 0\,| = 3L\varepsilon,$$

which tends to 0 as $\varepsilon \to 0$ with L fixed. The infimum is not attained: a quadrilateral in strict convex position has positive area. □

9 Lean validation

The proof of Theorem 1.1 is complete as written: every mathematical step appears in the preceding sections, and no step defers to a machine computation. The formalization reported here is independent verification, not an ingredient. The theorem has been machine-checked twice, by two unrelated proofs.

The reciprocal development. The proof of this paper is formalized in Lean 4 [4] over mathlib [5] in roughly four thousand lines of dedicated modules, organized to mirror the paper section by section:

Module	Paper counterpart
`FixedBranches`	owned separating branches (Section 2)
`Extremal`	the compact minimum (Section 3)
`CompleteContact`, `FirstVariation`	complete branch contact (Section 3)
`Farkas`, `Stationarity`	the Farkas lemma and reciprocal equations (Section 4)
`Forces` and its conclusions	the sine system and nonpositive-sine case (Section 4)
`Cyclic`, `CyclicGeometry`	the cyclic case (Section 5 and Appendix A)
`Transitive`, `TransitiveGeometry`	the transitive case (Section 6)
`OwnerConclusion`	the eight owner assignments reduced to the two tournaments
`Proof`	assembly (Section 7)

The public statement is `CenterAreaLemma.Reciprocal.center_area_lemma_reciprocal`, formulated for closed unit squares given by orthonormal side normals, with pairwise disjoint interiors and non-obtuseness expressed as the three squared-side-length inequalities. The development contains no `sorry` and uses no axioms beyond the standard classical and quotient principles already present in mathlib.

An independent second proof. The statement is also verified by an entirely different machine-checked proof: a corner-trap localization argument, which normalizes one square to be axis-parallel, localizes the other two centers near a corner of its rounded exclusion neighbourhood, and closes with a large family of local algebraic estimates. That development runs to 44,389 lines of Lean (the import closure of its localization argument) with no `sorry`; its top-level theorem, `CenterAreaLemma.center_area_lemma_from_axis_aligned_normalized`, proves the statement of `CenterAreaLemma.center_area_lemma` without invoking the reciprocal development. The two proofs share no mathematical content beyond elementary facts, so the theorem has independent verification along two unrelated routes.

A cautionary counterexample. The corner-trap route began with a natural simplification that turned out to be wrong, and the formal development preserves the refutation. Replace the two moving squares by the center-only relaxation: each center must merely avoid the rounded exclusion neighbourhood of the fixed square. A first attempt classified all bad relaxed pairs (mutual distance ≥ 1, non-obtuse with the origin, area $< \frac{1}{2}$) into six coordinate-local side and corner blocks up to the symmetries of the square. That classification is false: the pair

$$B_0 = \left(\tfrac{1}{20}, 1\right), \qquad C_0 = \left(-\tfrac{19}{20}, \tfrac{4}{5}\right)$$

is a relaxed bad pair lying in none of the six blocks, and the file `Counterexample.lean` certifies this. The relaxation loses exactly the information the localization needs: which side of the actual tilted square does the excluding. We record the counterexample because the failed claim was plausible enough to survive informal scrutiny for some time; it is a concrete instance of why machine verification earns its keep here.

Paper-statement wrappers. A library `PaperProofs`, kept in the same repository as this paper, restates Theorem 1.1 in the paper's exact logical form (unit squares by orthonormal side normals, disjointness of interiors, non-obtuseness as the three squared-side-length inequalities, area as the halved absolute shoelace determinant) and derives it from each entry point:

`PaperProofs.paper_center_area_lemma` from `CenterAreaLemma.center_area_lemma`,
`PaperProofs.paper_center_area_lemma_reciprocal` from
`CenterAreaLemma.Reciprocal.center_area_lemma_reciprocal`.

The wrappers are thin by design; their value is the machine-checked confirmation that the theorem proved formally is the theorem asserted here, with every project-level definition unfolded. `PaperProofs` also formalizes the all-angle bound of Proposition 8.3 as `PaperProofs.paper_all_angle_bound`.

Build status. The complete development (both proofs, the wrapper library, and all supporting modules) was rebuilt end to end on 2026-08-09 with Lean v4.29.1 and mathlib v4.29.1, in 8,367 build jobs. No `sorry` occurs anywhere in the build, and `#print axioms` on the wrapper theorems and on the corner-trap entry points reports only `propext`, `Classical.choice`, and `Quot.sound`.

What is not formalized. The equality construction, the obtuse degeneration example, and the four-center hull infimum are not formalized. Each is an explicit construction requiring interior computations of concrete squares. Every step of the proof of Theorem 1.1 itself is covered.

Acknowledgments

This work relied on substantial computer assistance. The proof was developed alongside its Lean formalization, and AI-based tools were used both in the formal development and in drafting this paper. The mathematics was checked independently of that assistance, by hand and by the Lean proof checker.

A The cyclic scalar inequality

This appendix proves the scalar inequality to which the cyclic case of the owner tournament was reduced. It is pure algebra: the proof uses nothing from the rest of the paper, and the rest of the paper uses nothing from this appendix beyond the statement below.

Lemma A.1. *Let $u, v, w > 0$ satisfy $uv + vw + wu = 1$, and set*

$$F(q) = \frac{(1+q)\max\{1,q\}}{1+q^2}.$$

Then

$$F(u)\,(u+w) + F(v)\,(u+v) + F(w)\,(v+w) \;\geq\; 4.$$

The pairing matters: $F(u)$ carries the factor $u+w$, $F(v)$ carries $u+v$, and $F(w)$ carries $v+w$. This is the pairing produced by the weighting identity of the cyclic case. The sum is invariant under the cyclic relabeling $(u, v, w) \mapsto (v, w, u)$, which permutes its three terms.

Proof. We first use the constraint to replace each denominator $1+q^2$ by a product of pairwise sums, then split into two cases: when $u, v, w \leq 1$ a single weighted Cauchy–Schwarz estimate suffices, and when one variable exceeds 1 three termwise estimates with visibly nonnegative errors reduce the claim to $u + 1/u \geq 2$.

Structural reductions. The constraint gives

$$(u+v)(u+w) = u^2 + uv + uw + vw = 1 + u^2$$

together with its cyclic variants $(v+w)(v+u) = 1+v^2$ and $(w+u)(w+v) = 1+w^2$. Dividing, the three terms become

$$F(u)(u+w) = \frac{(1+u)\max\{1,u\}}{u+v}, \qquad F(v)(u+v) = \frac{(1+v)\max\{1,v\}}{v+w},$$
$$F(w)(v+w) = \frac{(1+w)\max\{1,w\}}{w+u},$$

so, writing R for their sum, we must prove $R \geq 4$. At most one of u, v, w exceeds 1: if two of them did, their product alone would exceed $uv + vw + wu = 1$. Since the sum R and the constraint are invariant under cyclic relabeling, two cases remain: either $u, v, w \leq 1$, or $u > 1$ and $v, w < 1$.

Case 1: $u, v, w \leq 1$. Here $\max\{1,q\} = 1$ in every term, so

$$R = \frac{1+u}{u+v} + \frac{1+v}{v+w} + \frac{1+w}{w+u}.$$

Substitute $U = 1-u$, $V = 1-v$, $W = 1-w$ and $T = U+V+W$, so $U, V, W \in [0, 1)$. Each term exceeds 1 by a controlled amount,

$$\frac{1+u}{u+v} - 1 = \frac{1-v}{u+v} = \frac{V}{u+v}$$

and cyclically, so

$$R - 3 = \frac{V}{u+v} + \frac{W}{v+w} + \frac{U}{w+u}.$$

In the new variables the constraint reads

$$1 = (1-U)(1-V) + (1-V)(1-W) + (1-W)(1-U) = 3 - 2T + (UV + VW + WU),$$

that is,

$$UV + VW + WU = 2(T-1).$$

Moreover $T > 0$: if $U = V = W = 0$ then $u = v = w = 1$ and $uv + vw + wu = 3 \neq 1$. This strict positivity is what makes the Cauchy–Schwarz denominator below positive.

Assume first $U, V, W > 0$. The Engel form of the Cauchy–Schwarz inequality (write each term as $V^2/\big(V(u+v)\big)$, and cyclically) gives

$$R - 3 \;\geq\; \frac{(U+V+W)^2}{V(u+v) + W(v+w) + U(w+u)} = \frac{T^2}{4T - T^2 - 2},$$

where the denominator simplifies by substituting $u = 1-U$, $v = 1-V$, $w = 1-W$ and using $UV + VW + WU = 2(T-1)$:

$$\begin{aligned}
V(u+v) + W(v+w) + U(w+u) &= V(2-U-V) + W(2-V-W) + U(2-W-U) \\
&= 2T - (U^2+V^2+W^2) - (UV+VW+WU) \\
&= 2T - \big(T^2 - 2\cdot 2(T-1)\big) - 2(T-1) \\
&= 4T - T^2 - 2.
\end{aligned}$$

The denominator is a sum of three nonnegative terms, not all zero because $T > 0$, hence positive. Finally

$$T^2 - (4T - T^2 - 2) = 2(T-1)^2 \;\geq\; 0,$$

so $T^2/(4T - T^2 - 2) \geq 1$ and $R - 3 \geq 1$, as required. If one of U, V, W vanishes, its term in $R - 3$ vanishes as well; applying the same estimate to the remaining two terms leaves both the numerator T^2 and the denominator $4T - T^2 - 2$ unchanged, since the missing summand of each is zero, and the conclusion is the same.

Case 2: $u > 1$, $v, w < 1$. Now $\max\{1, u\} = u$ while $\max\{1, v\} = \max\{1, w\} = 1$, so

$$R = \frac{u(1+u)}{u+v} + \frac{1+v}{v+w} + \frac{1+w}{w+u}.$$

We bound the first and third terms by estimates whose errors are visibly nonnegative:

$$\frac{u(1+u)}{u+v} - u - \frac{1-v}{1+v} = \frac{v(1-v)(u-1)}{(u+v)(1+v)} \;\geq\; 0, \qquad \frac{1+w}{w+u} - \frac{1}{u} = \frac{w(u-1)}{u(w+u)} \;\geq\; 0.$$

For the first: $\frac{u(1+u)}{u+v} - u = \frac{u(1-v)}{u+v}$, and subtracting $\frac{1-v}{1+v}$ leaves $(1-v)\,\frac{u(1+v)-(u+v)}{(u+v)(1+v)} = (1-v)\,\frac{v(u-1)}{(u+v)(1+v)}$. The second is a direct computation.

For the middle term, write the constraint as $u(v+w) + vw = 1$. Then

$$\begin{aligned}
1 + v^2 - (v+w)(1+v) &= u(v+w) + vw + v^2 - (v+w)(1+v) \\
&= (v+w)(u+v-1-v) = (u-1)(v+w) > 0,
\end{aligned}$$

so $1 + v^2 > (v+w)(1+v)$; since all four quantities are positive, cross-multiplying gives

$$\frac{1+v}{v+w} \;\geq\; \frac{(1+v)^2}{1+v^2}.$$

Furthermore

$$\frac{(1+v)^2}{1+v^2} + \frac{1-v}{1+v} \;\geq\; 2,$$

because clearing the positive denominators leaves the excess

$$\frac{(1+v)^2}{1+v^2} + \frac{1-v}{1+v} - 2 = \frac{2v^2(1-v)}{(1+v^2)(1+v)} \;\geq\; 0.$$

Combining the four estimates,

$$R \ \geq \ u + \frac{1-v}{1+v} + \frac{(1+v)^2}{1+v^2} + \frac{1}{u} \ \geq \ u + \frac{1}{u} + 2 \ \geq \ 4,$$

the last step by $u + 1/u \geq 2$. $\square$

References

[1] P. Erdős and R. L. Graham. On packing squares with equal squares. *J. Combin. Theory Ser. A*, 19:119–123, 1975.

[2] J. Farkas. Theorie der einfachen Ungleichungen. *J. Reine Angew. Math.*, 124:1–27, 1902.

[3] E. Friedman. Packing unit squares in squares: a survey and new results. *Electron. J. Combin.*, Dynamic Survey DS7 (version of 2009). doi:10.37236/28.

[4] L. de Moura and S. Ullrich. The Lean 4 theorem prover and programming language. In *Automated Deduction – CADE 28*, volume 12699 of *Lecture Notes in Computer Science*, pages 625–635. Springer, 2021.

[5] The mathlib community. The Lean mathematical library. In *Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs (CPP 2020)*, pages 367–381. ACM, 2020.

[6] J. C. Maxwell. On reciprocal figures and diagrams of forces. *Philos. Mag. (4)*, 27:250–261, 1864.