An improved lower bound for packing seventeen unit squares in a square, via a deformation of Green's scaffold with certified defect charging

David R. MacIver

8 August 2026

Abstract

We prove that seventeen pairwise interior-disjoint unit squares cannot be packed in a closed square of side $(40\sqrt{2}+19)/17+1/200$, so that $s(17) > (40\sqrt{2}+19)/17+1/200 = 4.450208\ldots$; since s is nondecreasing, this also improves the lower bound for $s(18)$. The proof deforms Trevor Green's sixteen-point scaffold along its unit-circle constraint and charges the surviving empty-square families against an exact counting identity, discharging the case analysis through fourteen hash-pinned computational certificates and an exactly replayed 634,562-row ledger. The proof is computer-assisted in an essential way; its algebraic spine is formally verified in Lean 4. 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 containing n pairwise interior-disjoint unit squares (rotations allowed; all squares closed); trivially $\sqrt{n} \le s(n) \le \lceil\sqrt{n}\rceil$. The problem was posed by Göbel [4] and popularized by Erdős and Graham [1], who proved that the wasted area $W(s) = s^2 - \max\{n : s(n) \le s\}$ is $O(s^{7/11})$; the state of the art is Friedman's dynamic survey [2], which also records the known exact values (Kearney–Shiu [6], Stromquist [8], Nagamochi [7]) and the best constructions (Gensane–Ryckelynck [3], Hämäläinen). For $n = 17$ the best packing, found by Bidwell in 1998 and recorded in the survey, gives $s(17) \le 4.6756$; it is the smallest instance whose best known packing uses squares at three distinct angles. The best published lower bound, due to Green (2000, private communication to Friedman), is

$$s(17) \;\ge\; \Lambda := \frac{40\sqrt{2}+19}{17} = 4.445208382054341\ldots \tag{1}$$

The survey records the value and a sixteen-point figure only, with no proof on record, and lists the same value as its lower bound for $s(18)$.

Theorem 1.1 (Main Theorem)**.** *Seventeen pairwise interior-disjoint unit squares cannot be contained in a closed square of side* $\Lambda + \frac{1}{200}$*. Consequently*

$$s(17) > \frac{40\sqrt{2}+19}{17} + \frac{1}{200} = 4.450208382054341291\ldots$$

Corollary 1.2 (Eighteen squares)**.** $s(18) \ge s(17) > \dfrac{40\sqrt{2}+19}{17} + \dfrac{1}{200}$.

Proof. Deleting one square from a packing of eighteen leaves a packing of seventeen in the same container, so s is nondecreasing. □

The improvement $1/200$ is the first movement of either bound since 2000.

The proof runs as follows. Green's scaffold is a sixteen-point set P placed so that, at side Λ, almost every contained unit square covers a point of P in its interior: the failures form exactly six one-parameter families of empty squares (*defect tubes*) threading the six mesh edges longer than one. Rather than deny these failures we charge them: a packing with E empty squares, U unused scaffold points, and R excess incidences satisfies the exact identity $E = U + R + 1$ (Section 4), so certifying $U + R \geq E$ in every case yields a contradiction. To gain an increment δ we deform the scaffold along its unit-circle constraint, keeping the left, right, and bottom fences tight, so that its width is exactly $\Lambda + \delta$ and all new slack, an amount $\eta \in (2\delta, 4\delta)$, accumulates at the top fence (Section 3). A classification theorem (Section 5) shows that at the target $\delta = 1/200$ every empty contained square is a defect tube, a carrier (one of 23 endpoint-pair labels), or a square in one of exactly two narrow channels $T12$, $T13$ that the slack opens over the top row's unit gaps. Certified geometry (Sections 6 and 7) supplies loss and deletion clauses for the tubes, the carriers, and the right channel $T13$. The genuinely new difficulty is the left channel $T12$, whose unconditional endpoint loss is false at the target; Section 8 replaces it by a certified nested three-way alternative closed by a cascade of five further certificates. A finite ledger of 634,562 rows (Section 9) then verifies the charging inequality in every compatible case with zero failures, so no packing exists at side $\Lambda + 1/200$, and an attainment lemma converts this exclusion into the strict inequality of Theorem 1.1. Compactness enters exactly once, in that final step, so the argument is a repeatable mechanism for walking the bound upward by explicit rational increments (it was previously run at $481/250000$, $31/10000$, $7/2000$, $19/5000$, $1/250$, and $9/2000$ before $1/200$).

Every statement delegated to computation is stated as a Certified Theorem citing one of fourteen certificates C1–C14, with its branch structure described in prose and its census, minimum margin, and SHA-256 hash recorded in Appendix A. Replay follows a four-rule discipline (Section 6.1): exact arithmetic for all sign decisions, outward-rounded interval reconstruction of every certificate row, deterministic coverage of every bisection, and no optimizer in the replay path, with certificate files pinned by hash and their statement manifests checked against what the ledger consumes. The non-computational algebraic spine of the deformation is formalized in Lean 4 over mathlib (Appendix A.1).

During preparation of this paper I re-ran the exact audits and the complete ledger, recomputed the SHA-256 hashes of all fourteen certificate files, and ran the full aggregate replay to completion; all checks passed (Appendix A).

The proof is pinned to the target increment $\delta_* = 1/200$ and channel radius cutoff $R_\star = 27/100$ and asserts nothing at any other increment; several natural uniform strengthenings are false at this target (Appendix A, retired-hypotheses remark). The loss statement behind the right channel genuinely fails near $\delta \approx 0.005872$, which caps what this architecture can reach.

Section 2 fixes notation and proves the attainment lemma. Section 3 constructs the deformed scaffold and proves the deformation lemma; Section 4 proves the defect identity $E = U + R + 1$ and the matching refinement through which it is charged. Section 5 classifies the empty squares at the target. Section 6 fixes the conventions for reading certified theorems and proves the tube and carrier structure; Sections 7 and 8 prove the certified facts about the right and left channels respectively. Section 9 states the conditional ledger theorem (Theorem 9.1) and assembles the main theorem. Appendix A records the verification discipline, the certificate inventory with replay commands, and the Lean validation.

2 Preliminaries

This section fixes the vocabulary used throughout and proves the attainment lemma.

All squares are closed unless stated otherwise. A *unit square* S is determined by its centre $c_S \in \mathbb{R}^2$ and a rotation, represented by an orthonormal *frame* $e = (c, s)$, $n = (-s, c)$ with

$c^2 + s^2 = 1$:

$$S = \{c_S + \alpha e + \beta n : \ |\alpha| \le \tfrac{1}{2}, \ |\beta| \le \tfrac{1}{2}\}.$$

Quarter-turn invariance lets us always take the frame in the closed first quadrant $c, s \ge 0$. The *reduced orientation* of S is $r = c - s \in [-1, 1]$; we write $q = c + s$ and $\kappa = cs$, so that $q = \sqrt{2 - r^2}$ and $\kappa = (1 - r^2)/2$. The *projection halfwidth* of S on a unit vector v is $h(v) = \frac{1}{2}(|e \cdot v| + |n \cdot v|)$; on either coordinate axis of a frame at angle difference θ it equals $\frac{1}{2}(|\cos\theta| + |\sin\theta|)$.

Interior-disjointness is governed by the *separating-axis theorem*: two convex bodies have disjoint interiors if and only if some axis from a finite candidate list weakly separates them, the list for two unit squares being their four frame axes. Squares S_i, S_j with centres c_i, c_j are thus interior-disjoint if and only if some frame axis v satisfies $|(c_i - c_j) \cdot v| \ge h_i(v) + h_j(v)$. Each interior-disjoint pair therefore selects one of finitely many *signed* separating axes (4 axes $\times$ 2 signs = 8 choices), and if no candidate axis separates, the interiors meet. We use the theorem in both directions: to enumerate cases over the eight signed choices, and to derive contradictions.

A *packing at side* L is a family of pairwise interior-disjoint unit squares contained in $[0, L]^2$, and

$$s(n) = \inf\{L : \ n \text{ unit squares pack at side } L\}.$$

Lemma 2.1 (Attainment). *The infimum defining $s(n)$ is attained.*

Proof. Take packings at sides $L_k \downarrow s(n)$, label the squares $1, \ldots, n$ in each, and represent each orientation by an angle in the compact interval $[0, \pi/2]$. All centres lie in a common bounded square, so after passing to a subsequence every centre and every angle converges. Containment is a closed condition on centres and angles: the corners depend continuously on them and lie in $[0, L_k]^2$, so the limit squares lie in $[0, s(n)]^2$. Pairwise interior-disjointness is closed as well: by the separating-axis theorem it is the finite disjunction, over the four frame axes v of each pair, of the closed conditions $|(c_i - c_j) \cdot v| \ge h_i(v) + h_j(v)$. The limit configuration is therefore a packing at side $s(n)$. ☐

Attainment is the paper's sole use of compactness, applied exactly once, in Section 9.

3 The deformed scaffold

This section defines the sixteen-point configuration over which the whole proof quantifies and proves that its width is exactly $\Lambda + \delta$, with all new slack, an amount $\eta \in (2\delta, 4\delta)$, at the top fence.

Green's configuration is governed by the endpoint constants

$$t_0 = \frac{12 + 2\sqrt{2}}{17}, \qquad u_0 = \frac{8\sqrt{2} - 3}{17}, \qquad w_0 = u_0 - t_0 = \frac{6\sqrt{2} - 15}{17},$$

which satisfy $t_0^2 + u_0^2 = 1$, and

$$\Lambda = \frac{40\sqrt{2} + 19}{17} = 2\sqrt{2} + 2 + w_0.$$

To gain an increment we deform along the unit-circle constraint: for $0 \le \delta \le \frac{1}{200}$ set

$$w = w_0 + \delta, \qquad p = \sqrt{2 - w^2}, \qquad u = \frac{p + w}{2}, \qquad t = \frac{p - w}{2}, \tag{2}$$

so that

$$t^2 + u^2 = \frac{p^2 + w^2}{2} = 1, \qquad u - t = w.$$

The pair (t, u) moves on the unit circle so that $u - t$ grows by exactly δ; at $\delta = 0$ this is Green's configuration. Throughout, $\delta_* = 1/200$ is the target increment and the container is $[0, \Lambda + \delta_*]^2$, unless a uniform range $0 \le \delta \le \delta_*$ is stated.

Set the fence margins $m_x = \sqrt{2} - \frac{t}{2}$ and $m_y = \sqrt{2} - \frac{1}{2}$. The scaffold consists of sixteen points $p_0, \dots, p_{15}$, indexed row-major in four rows at ordinates $m_y + jt$ $(0 \le j \le 3)$, with abscissa offsets relative to m_x within each row

$$\text{even rows } (j = 0, 2)\colon\ (0,\ u,\ u+1,\ u+2), \qquad \text{odd rows } (j = 1, 3)\colon\ (0,\ 1,\ 2,\ u+2).$$

Thus row 0 is $p_0, \dots, p_3$, row 1 is $p_4, \dots, p_7$, row 2 is $p_8, \dots, p_{11}$, and the top row is $p_{12}, \dots, p_{15}$. We write $P = \{p_0, \dots, p_{15}\}$.

The natural triangulation of P (Figure 1) has 18 faces and 17 mesh edges of length exactly one; exactly six edges are longer than one, the *defect edges*

$$L0 = p_1p_5, \quad R0 = p_2p_6, \quad L1 = p_5p_9, \quad R1 = p_6p_{10}, \quad L2 = p_9p_{13}, \quad R2 = p_{10}p_{14}$$

(bottom to top, left before right). Their common squared length is

$$d^2 = (1-u)^2 + t^2 = 2 - 2u \in (1, 2), \tag{3}$$

and since u increases with δ, d decreases; $d < 51/50$ on the whole interval, and $d^2 = 1.015519\dots$ at $\delta = \delta_*$.

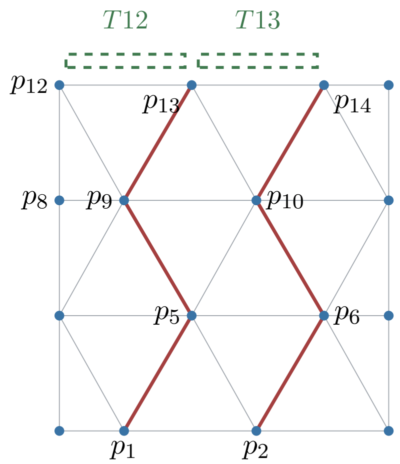

Figure 1: The scaffold and its exact eighteen-face triangulation, drawn schematically. The six red edges $L0, R0, L1, R1, L2, R2$ (from bottom to top, left before right) are the only mesh edges longer than one; each supports a one-parameter "tube" of empty squares. After deformation, the two dashed regions indicate the top channels $T12$ (over the gap $p_{12}p_{13}$) and $T13$ (over $p_{13}p_{14}$), the only new families of empty squares.

Lemma 3.1 (Deformation lemma: exact horizontal fit and top slack)**.** *For $0 \le \delta \le \delta_*$:*

1. *Width identity: $2m_x + u + 2 = \Lambda + \delta$; the deformed scaffold with its two horizontal fence margins has width exactly $\Lambda + \delta$.*

2. *With the bottom fence at ordinate m_y, the top-wall excess is*

$$\eta = (\Lambda + \delta) - (2\sqrt{2} - 1 + 3t) = \delta + 3(t_0 - t).$$

3. *For $0 < \delta \le \delta_*$:*

$$0 < t_0 - t < \delta, \qquad 0 < u - u_0 < \delta, \qquad 2\delta < \eta < 4\delta.$$

4. At $\delta = \delta_*$: $\eta = 0.0104035325656\ldots$

Proof. (1) Since $2m_x = 2\sqrt{2} - t$, the width is $2\sqrt{2} + 2 + (u - t) = 2\sqrt{2} + 2 + w_0 + \delta$, and

$$2\sqrt{2} + 2 + w_0 = \frac{34\sqrt{2} + 34 + 6\sqrt{2} - 15}{17} = \frac{40\sqrt{2} + 19}{17} = \Lambda.$$

(2) The natural fenced height is $2m_y + 3t = 2\sqrt{2} - 1 + 3t$. Compute

$$(\Lambda + \delta) - (2\sqrt{2} - 1 + 3t) - \big(\delta + 3(t_0 - t)\big) = \Lambda - (2\sqrt{2} - 1 + 3t_0) = 3 + w_0 - 3t_0 = 3 + u_0 - 4t_0 = 0,$$

using $\Lambda = 2\sqrt{2} + 2 + w_0$, $w_0 = u_0 - t_0$, and the collapse

$$u_0 - 4t_0 = \frac{(8\sqrt{2} - 3) - (48 + 8\sqrt{2})}{17} = -3.$$

(3) Write $\varepsilon = t_0 - t$ and $p_0' = \sqrt{2 - w_0^2}$. From (2), $2\varepsilon = (p_0' - p) + \delta$ and $2(u - u_0) = (p - p_0') + \delta$; also $p_0'^2 - p^2 = w^2 - w_0^2 = \delta(w + w_0)$, so $p_0' - p = \delta(w + w_0)/(p_0' + p)$. Hence the exact ratio identity

$$\frac{u - u_0}{\varepsilon} = \frac{(p_0' + p) - (w + w_0)}{(p_0' + p) + (w + w_0)} = \frac{t_0 + t}{u_0 + u}.$$

On the deformation interval $1 < (t_0 + t)/(u_0 + u) < 2$: numerator and denominator lie within δ of $2t_0 \approx 1.744$ and $2u_0 \approx 0.978$ respectively. Since $\delta = \varepsilon + (u - u_0)$, this gives $\delta/3 < \varepsilon < \delta/2$ and $\delta/2 < u - u_0 < 2\delta/3$, stronger than claimed, and then $\eta = \delta + 3\varepsilon$ gives $2\delta < \eta < \frac{5}{2}\delta < 4\delta$.

(4) The endpoint value of η is an exact evaluation replayed by the classification audit script (Section 5, Appendix A). $\square$

The deformation moves slack to exactly one place. The left, right, and bottom fences keep their endpoint equalities: the bottom-row gap lengths are still m_x, u, 1, 1, m_x, each at most one, and every vertical interior gap still has length t. Because $t^2 + u^2 = 1$ survives the deformation exactly, the triangulation keeps the same seventeen unit mesh edges and six defect edges. Only the top row moves away from its wall, by the slack η. Anchoring the scaffold at the lower-left corner instead would open five new top- and right-wall families of empty squares; the circle-constrained deformation opens exactly the two channels $T12$, $T13$ analysed in Section 5.

[The identities of this section are formalized in Lean as `span_identity`, `u0_collapse`, `height_collapse`, `eta_eq_top_excess`, `eta_bounds`, `t_sq_add_u_sq`, and `u_sub_t`; see Appendix A.1.]

4 The defect identity and the matching refinement

This section proves the exact counting identity $E = U + R + 1$ and its refinement through matchings in a co-hit graph, then validates that graph in the deformed scaffold. Every quantity the ledger of Section 9 consumes numerically is defined here.

Throughout, $\mathcal{A}$ is a family of seventeen pairwise interior-disjoint unit squares and $P = \{p_0, \ldots, p_{15}\}$ is the deformed scaffold of Section 3.

Definition 4.1. For a unit square S set $h(S) = \#(P \cap \operatorname{int} S)$; a packed square with $h(S) = 0$ is *empty*. A point of P is *used* if it lies in the interior of some packed square (interior-disjointness makes that square unique), and *unused* otherwise. A point is *lost* if it is unused or its unique using square contains at least two scaffold points in its interior; a point that is not lost has a unique *singleton owner*, a packed square containing it and no other scaffold point in its interior. Set $E = \#\{S \in \mathcal{A} : h(S) = 0\}$, let U be the number of unused points, let $R = \sum_{S:h(S)\geq 1}(h(S) - 1)$ be the excess incidences, and let $B \subseteq P$ be the set of lost points.

Lemma 4.2 (Defect identity)**.** *For every family $\mathcal{A}$ of seventeen pairwise interior-disjoint unit squares, $E = U + R + 1$.*

Proof. Put $H = \sum_{S \in \mathcal{A}} h(S)$. Disjoint interiors make the incidences disjoint, so $U = 16 - H$; the $17 - E$ nonempty squares give $R = H - (17 - E)$. Adding, $U + R = E - 1$. □

[The identity is formalized in Lean as `defect_identity`, an abstract counting lemma over a seventeen-element finite set; see Appendix A.1.] By Lemma 4.2, a certified bound $U + R \geq E$ in every case yields a contradiction; the rest of the proof supplies such bounds. Losses are counted through matchings.

Lemma 4.3 (Matching refinement)**.** *Let Γ be a graph on P such that for every packed square S with $h(S) \geq 2$ the set $P \cap \operatorname{int} S$ contains both endpoints of some edge of Γ. Then*

$$U + R \ \geq \ |B| - \nu(\Gamma[B]),$$

where ν is the matching number and $\Gamma[B]$ the induced subgraph. Moreover the bound is monotone: if $M \subseteq B$ then $|M| - \nu(\Gamma[M]) \leq |B| - \nu(\Gamma[B])$.

Proof. Let q be the number of packed squares with $h(S) \geq 2$. Each such square contributes all $h(S)$ of its interior scaffold points to B and $h(S) - 1$ to R; each unused point contributes itself to both U and B; no point is shared between these contributions. Hence $U + R = |B| - q$ exactly. Choosing one Γ-edge inside the hit-set of each square with $h(S) \geq 2$ produces q vertex-disjoint edges of $\Gamma[B]$, so $q \leq \nu(\Gamma[B])$, which gives the bound. For monotonicity it suffices to add one vertex at a time: adding a vertex to an induced subgraph raises $|\cdot|$ by one and the matching number by at most one, so $|M| - \nu(\Gamma[M])$ is nondecreasing in M. □

Monotonicity is what makes discarding uncertified losses safe in the ledger: any certified subset M of the true lost set B already yields the valid lower bound $U + R \geq |M| - \nu(\Gamma[M])$.

Lemma 4.3 is applied with a fixed concrete graph. The *conservative co-hit graph* Γ has vertex set P and edges exactly the pairs of points at Euclidean distance $\leq \sqrt{2}$. A unit square has diameter $\sqrt{2}$, so any two scaffold points interior to a common packed square span an edge of Γ, and Γ satisfies the hypothesis of Lemma 4.3. The edge set of Γ was computed for the endpoint scaffold ($\delta = 0$); the next proposition shows the deformation does not invalidate it.

Proposition 4.4 (Deformation validity of the co-hit graph)**.** *In the deformed scaffold at $\delta = \delta_*$, every pair of points forming one of the* 81 *non-edges of the endpoint co-hit graph has squared distance exceeding* 2*; the minimum excess over* 2 *among the* 81 *non-edges is* 0.226781706524. *Consequently the deformation creates no new co-hit pair: every pair of deformed scaffold points at distance $\leq \sqrt{2}$ is an edge of the endpoint graph, so the endpoint edge set serves as Γ at the target.*

Proof. This is a finite exact computation: each squared distance lies in $\mathbb{Q}(\sqrt{2}, p)$ and is compared with 2 by exact arithmetic. The audit is run directly by the theorem assembler; see Appendix A. □

The ledger sharpens the bound of Lemma 4.3 by two mechanisms. *Edge deletions*: a geometric theorem asserts that in the configuration at hand some co-hit pair cannot lie in the interior of a single packed square; the edge is removed from Γ before computing ν, and deleting edges can only decrease $\nu(\Gamma[M])$, so the bound improves. *Unused-point promotions*: a theorem certifies directly that a specific point is unused, enlarging the certified loss set. The theorems supplying both mechanisms occupy Sections 6–8; the ledger of Section 9 consumes them.

5 Classification of empty squares at the target

Throughout this section $\delta = \delta_* = 1/200$, the container is $[0, \Lambda + \delta_*]^2$, and $P = \{p_0, \ldots, p_{15}\}$ is the deformed scaffold of Section 3. We prove (Theorem 5.6) that every contained empty square is a defect tube, a carrier, or a channel square. Orientations of the square under study are reduced so that $0 \le s \le c$ (swap the frame vectors of Section 2 if necessary); then $r = c - s$, $q = c + s = \sqrt{2 - r^2}$ and $\kappa = cs = (1 - r^2)/2$.

Fences and chords

We record how container-parallel lines cut a unit square. A line parallel to a container side at perpendicular distance $\zeta \ge 0$ from the centre of a unit square cuts it in a chord of length $1/c \ge 1$ when the line crosses the central band of the square, and of length $(\frac{q}{2} - \zeta)/\kappa$ when it crosses a corner cap; in particular the chord has length at least ρ whenever $0 \le \zeta \le \frac{q}{2} - \rho\kappa$.

Lemma 5.1 (Fence inequalities)**.** *For every unit square,*

$$q - \kappa \;\ge\; m_y \qquad \text{and} \qquad q - t\kappa \;\ge\; m_x,$$

each with equality if and only if $q = \sqrt{2}$.

Proof. From $q = \sqrt{2 - r^2}$ we get $q - \sqrt{2} = -r^2/(q + \sqrt{2})$, and with $\kappa = (1 - r^2)/2$, $m_y = \sqrt{2} - \frac{1}{2}$, $m_x = \sqrt{2} - \frac{t}{2}$,

$$q - \kappa - m_y = r^2\Big(\frac{1}{2} - \frac{1}{q + \sqrt{2}}\Big), \qquad q - t\kappa - m_x = r^2\Big(\frac{t}{2} - \frac{1}{q + \sqrt{2}}\Big).$$

Since $q \ge 1$ we have $1/(q + \sqrt{2}) \le 1/(1 + \sqrt{2}) = \sqrt{2} - 1$, and both $\frac{1}{2} > \sqrt{2} - 1$ and $\frac{t}{2} > \sqrt{2} - 1$ (the latter is the exact comparison $t > 2\sqrt{2} - 2$, replayed by the classification audit). Both right-hand sides are therefore nonnegative and vanish exactly when $r = 0$, that is, when $q = \sqrt{2}$. $\square$

Lemma 5.1 converts wall containment into long chords. If the centre of a contained square has ordinate $y_c \le m_y$, containment above the bottom wall gives $y_c \ge q/2$, so the distance $\zeta = m_y - y_c$ to the bottom-row line satisfies $\zeta \le m_y - q/2 \le q/2 - \kappa$; the square cuts the bottom row in a chord of length at least 1. If the square is empty, the relative interior of that chord contains no scaffold point, so the chord lies in a closed bottom-row gap of length at least 1; the bottom-row gaps have lengths $m_x, u, 1, 1, m_x$, so only the unit gaps p_1p_2 and p_2p_3 can receive it. Likewise a centre at abscissa $x_c \le m_x$ (or $x_c \ge \Lambda + \delta_* - m_x$) forces a chord of length at least t across the side column. The degenerate equality cases are exactly the axis-parallel mesh carriers and the 45° wall diamonds defined before Theorem 5.6.

The angular cost and the channel chart

The first identity in the proof of Lemma 5.1 is exact and governs the top fence. Define the *angular cost* $A(x) := \sqrt{2 - x} + \frac{x}{2} - \sqrt{2}$; then

$$q - \kappa - m_y \;=\; A(r^2) \;=\; r^2\Big(\frac{1}{2} - \frac{1}{\sqrt{2} + q}\Big), \qquad A'(x) = \frac{\sqrt{2 - x} - 1}{2\sqrt{2 - x}} > 0 \quad (0 \le x < 1), \qquad (4)$$

with $A(0) = 0$: A is strictly increasing and $A(r^2)$ is the extra top clearance consumed by tilting to reduced orientation r. [The monotonicity and the channel-radius margin below are Lean-formalized as `angularCost_strictMonoOn` and `angularCost_margin`; see Appendix A.1.] A square in the top cap region can be empty only if its angular cost plus its wall slack fits inside the top slack η of Lemma 3.1; this is the exact mechanism that limits the new empty-square families to two narrow channels.

We make the mechanism a closed chart. Translate and rotate so that a top gap has endpoints $O = (0, 0)$ and $H = (1, 0)$ and the top wall is the line $y = -(m_y + \eta)$; the channel opens downward in these coordinates. Consider a contained square with reduced orientation r whose centre lies between the wall and the gap line, with *wall slack* $\tau \geq 0$ (its centre is at distance $q/2 + \tau$ from the wall) and *endpoint clearances* $\alpha, \beta \geq 0$ (the distances from O and H to the two ends of its chord along the gap line). The distance from the centre to the gap line is $\zeta = (m_y + \eta) - q/2 - \tau$. If the gap line crosses the central band, the chord has length $1/c \geq 1$ and can only fit in the unit gap as the full gap itself, forcing $c = 1$: the axis-aligned mesh carrier. Otherwise the gap line crosses a corner cap, and the chord has length $(\frac{q}{2} - \zeta)/\kappa = 1 + (A(r^2) + \tau - \eta)/\kappa$ by (4) in the form $\kappa - q + m_y = -A(r^2)$. Subtracting from the gap length 1 gives the total clearance budget: the chart

$$0 \;\leq\; \tau + A(r^2) \;\leq\; \eta, \qquad \alpha + \beta \;=\; \frac{\eta - \tau - A(r^2)}{\kappa}, \tag{5}$$

with the centre an explicit affine function of $(r, \tau, \alpha - \beta)$. The *channel families* $T12$ and $T13$ are the closed families of contained squares satisfying (5) over the top gaps $p_{12}p_{13}$ and $p_{13}p_{14}$ respectively.

Lemma 5.2 (Channel radius, multiplicity, endpoint co-hits)**.** *At $\delta = \delta_*$:*

1. *Every channel square has reduced orientation $|r| < 27/100 = R_\star$.*
2. *An interior-disjoint family of unit squares contains at most one $T12$ square and at most one $T13$ square.*
3. *A $T12$ square shares an interior point with every closed unit square containing both p_{12} and p_{13}, and a $T13$ square with every closed unit square containing both p_{13} and p_{14}.*

Proof. (1) Since $\tau \geq 0$, the chart (5) forces $A(r^2) \leq \eta$, and A is strictly increasing by (4). Exact evaluation gives

$$A\big((27/100)^2\big) - \eta \;=\; 3.31787970482\ldots \times 10^{-5} \;>\; 0,$$

so $|r| < 27/100$. The exact feasible endpoint is $|r| \leq 0.2695602432185\ldots$, so the rational cutoff $R_\star = 27/100$ is tight: the positive tail of radii below it must be retained, not rounded away. $R_\star$ is shared by every certificate in the paper.

(2) By part (1), $\kappa = (1 - r^2)/2 \geq 9271/20000$, so $\alpha + \beta \leq \eta/\kappa < 0.02245$. Both clearances are below $\frac{1}{2}$, so the gap midpoint lies in the relative interior of the chord. The gap line has points of the square strictly on both sides (indeed $q/2 - \zeta \geq q - m_y - \eta = \kappa + A(r^2) - \eta > 0$ by (4)), so by convexity the relative interior of the chord lies in the interior of the square. Hence every channel square of a given label contains its gap midpoint in its interior, and two such squares cannot be interior-disjoint.

(3) Let D be a closed unit square containing both endpoints of the gap; by convexity D contains the gap segment, hence the gap midpoint g. Since D is a convex body, g lies in the closure of $\operatorname{int} D$, so every neighbourhood of g meets $\operatorname{int} D$. By the proof of part (2) a channel square of that label contains a neighbourhood of g in its interior; the two interiors therefore meet. This yields the ledger deletions $T12$: (p_{12}, p_{13}) and $T13$: (p_{13}, p_{14}). □

Two threading lemmas

A segment *weakly threads* a closed unit square Q if the line through it cuts Q in a chord whose endpoints lie on two opposite sides of Q and the chord is contained in the segment; the threading is *proper* if the relative interior of the chord lies in $\operatorname{int} Q$. The following two Euclidean lemmas drive the interior case of the classification.

Lemma 5.3 (Closed short-leg face threading)**.** *Let Q be a closed unit square and let A, B, C lie outside* $\operatorname{int} Q$ *with the centre of Q in the closed triangle ABC, and suppose $|A - C| \le 1$ and $|B - C| \le 1$. Then either the segment AB weakly threads Q, or one of AC, BC is a unit segment which weakly threads Q with endpoints in the relative interiors of opposite sides of Q.*

Proof. Normalize $Q = [-\frac{1}{2}, \frac{1}{2}]^2$ and give each vertex every closed exterior side label it carries: left ($x \le -\frac{1}{2}$), right ($x \ge \frac{1}{2}$), bottom ($y \le -\frac{1}{2}$), top ($y \ge \frac{1}{2}$). The proof splits on whether a short leg joins opposite labels.

Step 1: some opposite-labelled pair of vertices exists. Otherwise all labels lie in two adjacent classes, say (after symmetry) right and top. Every vertex then has $x > -\frac{1}{2}$, $y > -\frac{1}{2}$, and $x \ge \frac{1}{2}$ or $y \ge \frac{1}{2}$, hence $x + y > 0$; the convex hull of the vertices misses the origin, contradicting centre containment.

Step 2: a short pair is opposite-labelled, say A left and C right (horizontally; the other cases are symmetric). Then $|A - C| \ge C_x - A_x \ge 1$, so $|A - C| = 1$ and $A = (-\frac{1}{2}, y_0)$, $C = (\frac{1}{2}, y_0)$ for some y_0. If $|y_0| < \frac{1}{2}$, the horizontal line $y = y_0$ cuts Q in exactly the chord AC, whose endpoints lie in the relative interiors of the left and right sides: the required unit thread. If $|y_0| = \frac{1}{2}$ we are in the corner case, Step 4. Suppose $y_0 > \frac{1}{2}$. Only $|B - C| \le 1$ is assumed for B. It excludes $B_y \le -\frac{1}{2}$ (distance at least $y_0 + \frac{1}{2} > 1$); and if $B_x \le -\frac{1}{2}$ then $|B - C| \ge 1$ forces $B = (-\frac{1}{2}, y_0)$, whence $B_x + B_y = y_0 - \frac{1}{2} > 0$. Otherwise exteriority leaves $B_x \ge \frac{1}{2}$ or $B_y \ge \frac{1}{2}$. If $B_x \ge \frac{1}{2}$, write $B = C + (\delta_x, \delta_y)$ with $\delta_x \ge 0$ and $\delta_x^2 + \delta_y^2 \le 1$; the minimum of $\delta_x + \delta_y$ under these constraints is -1, so $B_x + B_y \ge \frac{1}{2} + y_0 - 1 = y_0 - \frac{1}{2} > 0$. If $B_y \ge \frac{1}{2}$ then $B_x + B_y > -\frac{1}{2} + \frac{1}{2} = 0$. Since also $A_x + A_y = y_0 - \frac{1}{2} > 0$ and $C_x + C_y = y_0 + \frac{1}{2} > 0$, all three vertices satisfy $x + y > 0$ and the hull misses the origin: contradiction. A reflection handles $y_0 < -\frac{1}{2}$.

Step 3: only the long pair A, B is opposite-labelled. After symmetries, $A_x \le -\frac{1}{2}$ and $B_x \ge \frac{1}{2}$. The short legs give $B_x - 1 \le C_x \le A_x + 1$, so $-\frac{1}{2} \le C_x \le \frac{1}{2}$; if $|C_x| = \frac{1}{2}$ then equality throughout produces a horizontal opposite-labelled short pair and Step 2 applies, so assume $|C_x| < \frac{1}{2}$ and, reflecting if necessary, $C_y \ge \frac{1}{2}$ by exteriority of C. The short legs then give $A_y, B_y > -\frac{1}{2}$: for instance $A_y \ge C_y - 1 \ge -\frac{1}{2}$, with equality forcing $A = (C_x, -\frac{1}{2})$ and $|C_x| < \frac{1}{2}$, against $A_x \le -\frac{1}{2}$. If the origin lies on AB, the chord of Q through the origin along that line is centrally symmetric, so its endpoints lie on opposite sides; the chord is contained in AB because $A_x \le -\frac{1}{2}$ and $B_x \ge \frac{1}{2}$ straddle the vertical supporting lines, so AB weakly threads. Otherwise the ray from C through the origin meets AB at a point $R = -\lambda C$ with $\lambda > 0$. Since $C_y \ge \frac{1}{2} > |C_x|$ we get $-\frac{1}{2} < R_y < -|R_x| \le 0$, the lower bound because R lies on AB and $A_y, B_y > -\frac{1}{2}$; in particular $R \in \operatorname{int} Q$, so the line AB cuts a genuine chord of Q whose relative interior contains R, and the chord is contained in AB as before. If that chord had endpoints on adjacent sides, then, passing through the sector $y < -|x|$ containing R, it would use the bottom side; continuing along the line from that endpoint to A or to B would force $A_y \le -\frac{1}{2}$ or $B_y \le -\frac{1}{2}$, both impossible. Hence the chord joins opposite sides and AB weakly threads.

Step 4: the corner case. Say $A = (-\frac{1}{2}, \frac{1}{2})$, $C = (\frac{1}{2}, \frac{1}{2})$. Containment of the origin in the hull forces $B_y \le 0$, and $B_y < -\frac{1}{2}$ is impossible, since then $|B - C| \ge \frac{1}{2} - B_y > 1$. If $B_y = -\frac{1}{2}$, then $|B - C| \le 1$ forces $B = (\frac{1}{2}, -\frac{1}{2})$ and AB is a diagonal joining opposite corners: AB weakly threads. If $-\frac{1}{2} < B_y \le 0$, then $B_x \le -\frac{1}{2}$ is excluded ($|B - C| \ge 1$ would force $B = A$), so exteriority forces $B_x \ge \frac{1}{2}$; then all three vertices satisfy $x + y \ge 0$ with equality only at A, so the affine function $x + y$ vanishes on the hull only at $A \ne 0$, excluding the origin: contradiction. Thus the corner case restores the AB alternative. □

Lemma 5.4 (A defect edge cannot coincide with a square side)**.** *If a defect edge weakly threads an empty unit square, it properly threads it.*

Proof. Each defect edge is the short diagonal, of length $d \in (1, \sqrt{2})$, of a rhombus with unit sides ($L0 = p_1p_5$ in the rhombus $p_1p_2p_5p_4$, and so on); this survives the deformation because $t^2 + u^2 = 1$. Suppose the threading of an empty square Q is not proper. The chord joins opposite sides of Q but its relative interior avoids $\operatorname{int} Q$, so the defect line supports Q and the chord is a full unit side of Q. That unit side lies inside the length-d defect segment, so its midpoint is within $(d-1)/2 < \frac{1}{2}$ of the segment midpoint, which is the rhombus centre. Consider the rhombus vertex on the long diagonal lying on the side of the defect line towards Q: its tangential offset from the side midpoint has absolute value less than $\frac{1}{2}$, and its inward normal offset is $\sqrt{4-d^2}/2 \in (0,1)$. It therefore lies strictly inside Q, contradicting emptiness. □

The classification theorem

Definition 5.5. An empty contained unit square is a *defect tube* square if it is properly threaded by one of the six defect edges $L0, \dots, R2$; a *carrier* if it realizes one of twenty-three endpoint pairs, namely one of the seventeen unit mesh edges cut as a chord normal to two opposite sides with both endpoints in the relative interiors of those sides (a *mesh carrier*), or a wall-tangent 45° *diamond* on one of the pairs $(p_0,p_4),(p_4,p_8),(p_8,p_{12}),(p_3,p_7),(p_7,p_{11}),(p_{11},p_{15}),(p_1,p_2),(p_2,p_3)$ (this gives $17+8=25$ tangent types but 23 labels: two labels are multiply typed); and a *channel* square if it belongs to the closed family $T12$ or $T13$ of (5).

Theorem 5.6 (Classification at the target)**.** *At $\delta=\delta_*$, every contained unit square whose interior contains none of $p_0,\dots,p_{15}$ is a defect tube, a carrier, or a channel square, boundary contacts included.*

Proof. Let the square be empty; the proof splits on whether its centre escapes a fence.

Fence cases. If the centre has ordinate at most m_y, or abscissa at most m_x or at least $\Lambda+\delta_*-m_x$, the chord argument after Lemma 5.1 applies, unchanged from the endpoint: the square cuts a bottom-row chord of length at least 1, which fits only in the unit gaps p_1p_2, p_2p_3, or a side-column chord of length at least t. The equality analysis shows the only realizations are the mesh carriers on the two unit bottom gaps and the wall diamonds: a chord coinciding with a square side is excluded by containment, and side-fence equality forces $q=\sqrt{2}$, the diamond.

For the top fence, (4) together with $\frac{1}{2}-\frac{1}{\sqrt{2+q}} \ge \frac{3}{2}-\sqrt{2} > \eta$ and $1-2\kappa = r^2$ gives

$$q-(1-2\eta)\kappa-(m_y+\eta) \;=\; A(r^2)-\eta r^2 \;=\; r^2\Big(\frac{1}{2}-\frac{1}{\sqrt{2+q}}-\eta\Big) \;\ge\; 0.$$

Hence, by the chord bookkeeping with $\rho=\ell_0 := 1-2\eta$, a square whose centre lies at or above the top row cuts a top-row chord of length at least $\ell_0 = 0.97919\ldots$. The top-row gaps have lengths $m_x, 1, 1, u, m_x$, and the exact reserves $\ell_0 - m_x > 0.0002089$ and $\ell_0 - u > 0.4869$ show such a chord fits only in the unit gaps $p_{12}p_{13}$ or $p_{13}p_{14}$. A central-band chord forces the axis-aligned mesh carrier; a cap chord satisfies exactly the closed channel chart (5), whose zero-clearance boundary is a channel limit, not a new carrier. In particular the two top-wall diamonds of the endpoint configuration are now impossible: a top-tangent diamond cuts the displaced top row in a chord of length $1-2\eta<1$, too short to span a unit gap.

Interior case. Otherwise the centre lies in the inner rectangle $[m_x, \Lambda+\delta_*-m_x]\times[m_y, m_y+3t]$, which is exactly triangulated by the 18 faces of the scaffold, with the same 17 unit edges and 6 defect edges as at the endpoint (Section 3; the radical comparisons are replayed by the audit). Apply Lemma 5.3 to the face containing the centre, taking the possible long edge as AB: every other edge of a face has length 1, t, or u, all at most 1, and the face vertices are scaffold points, hence outside the interior of the empty square. Either AB weakly threads the square, or a unit leg threads it with relative-interior side contacts. In the first alternative, if AB is the defect edge of the face, the threading is proper by Lemma 5.4 and the square is a defect tube; if AB is an edge of length at most 1, the chord between opposite sides has length at least 1 and is contained in AB, so AB is a unit mesh edge and the chord is AB itself, normal to the two sides: a mesh carrier. The second alternative is a mesh carrier directly. The corner branch of Lemma 5.3 restores the long-edge alternative, so no case is missed. □

All radical comparisons in this section and the full equality-case analysis (coincident sides, corner contacts, the zero-clearance channel boundary) are exact statements in $\mathbb{Q}(\sqrt{2}, p)$ replayed by `scripts/s17_green_deformed_classification.py` (PASS: $\eta = 0.0104035325656$, top_chord_lower = 0.979192934869, 18 faces, 17 unit edges, 6 defect edges). The repository's endpoint TEX documents record the corresponding endpoint analyses as ancillary material.

6 Certified theorems: conventions, tubes, and carriers

This section fixes how the paper's computer-certified statements are read and replayed, then proves the tube and carrier structure the ledger consumes: tube multiplicity and midpoint deletions, the adjacent-tube exclusion with the carrier combinatorics, and the four tube-loss and promotion packages.

6.1 Reading the certified theorems

Every certified theorem in this paper reduces, after explicit case choices, to the infeasibility of finitely many systems of linear inequalities whose coefficients are algebraic functions of at most a few orientation variables. The choices are always of three kinds:

1. a *signed separating axis* for each interior-disjoint pair of squares in play: by the separating-axis theorem (Section 2), four axes and two signs give eight choices per pair;
2. a closed *escape face* for each scaffold point required to avoid a square's interior: the point lies beyond one of the square's four sides, four choices per point;
3. where a channel or tube square participates, its *exact translation chart*: the channel chart (5) for channels, the threaded-rhombus rectangle for tubes. The translation variables are eliminated by exact support functions, never relaxed into boxes. For example, in the chart (5) write $y = \eta - \tau - A(r^2)$ and $\zeta = \alpha - \beta$: these are the chart's translation coordinates, and at fixed reduced orientation r they range over the triangle $0 \le y \le g$, $|\zeta| \le y/\kappa$ of height $g = \eta - A(r^2)$. C1 eliminates the channel translation over this triangle via the exact support identity

$$\max_{0\le y\le g,\ |\zeta|\le y/\kappa} \left(c_y y + c_\zeta \zeta\right) \;=\; g \max\Big\{0,\ c_y + \frac{|c_\zeta|}{\kappa}\Big\} \qquad \text{for all real } c_y, c_\zeta. \tag{6}$$

Each resulting branch is closed by an *interval–Farkas certificate*: a binary subdivision tree over the orientation box, at whose leaves nonnegative dyadic multipliers exhibit a strictly negative weighted combination of the necessary inequalities. Certificates are generated offline, where midpoint linear programming may *propose* multipliers. Replay is independent of generation: it rebuilds every inequality row in outward-rounded interval arithmetic with exact rational squaring at every square-root endpoint, checks the recorded multipliers, checks that the two children of every split cover their parent (so the leaves provably cover the stated case analysis), and imports no optimizer.

The *minimum margin* quoted for a certificate is the largest replay upper bound on any accepted leaf, negated: the smallest certified distance to failure. Under this discipline a margin of any positive size suffices, since every row is rounded outward with exact rational guards; even the 10^{-9}-scale minimum margin of C2 is therefore sound, but it is thin, and we flag it where it occurs.

Three artifacts are of simpler kinds under the same replay discipline: C8 is a finite outward-rounded separating-axis overlap cover, C9 serializes exact margins in $\mathbb{Q}(\sqrt{2}, p)$, and C11 is a family of exact interval Bernstein sign charts.

Every statement delegated to computation is stated in full as a *Certified Theorem*, with its branch structure described in prose and its certificate cited as (Cn); the full census (roots or branches, nodes, depth, minimum margin), SHA-256 hash, and replay command of each certificate are recorded in Appendix A, though proofs quote individual margins as diagnostics. The heading *Conditional Theorem* is reserved for the ledger theorem, whose hypotheses the certified package supplies.

6.2 Tubes and carriers

By Theorem 5.6, every empty square at the target is a defect tube, a carrier, or a channel square. This subsection bounds how tubes and carriers can combine.

Lemma 6.1 (One square per tube; generic midpoint deletions)**.** *For $0 \le \delta \le \delta_*$:*

1. *two interior-disjoint empty squares cannot be threaded by the same defect edge;*
2. *a square properly threaded by a defect edge contains the midpoint of that defect segment in its interior; consequently it is interior-incompatible with every closed unit square containing both endpoints of its defect edge, and in any packing containing a square threaded by a given tube the co-hit edge between that tube's endpoints is deleted from Γ in the matching bound of Lemma 4.3 (Table 1).*

tube	$L0$	$L1$	$L2$	$R0$	$R1$	$R2$
deleted edge	(p_1, p_5)	(p_5, p_9)	(p_9, p_{13})	(p_2, p_6)	(p_6, p_{10})	(p_{10}, p_{14})

Table 1: The generic midpoint deletions: the co-hit edge deleted from Γ when the given tube is occupied.

Proof. A properly threading square cuts an open chord of length at least one from the defect segment: the chord has length $1/|v \cdot n| \ge 1$, where v is the defect direction and n the normal of the crossed face. The deformed defect length satisfies $d < 51/50$ throughout the interval, by the exact monotonicity $d^2 = 2 - 2u$ with u increasing in δ.

(1) Two interior-disjoint threaded squares would cut two disjoint open chords of length at least one from the same segment, requiring length at least $2 > d$.

(2) In arclength coordinates the chord is (a, b) with $b - a \ge 1$ and $0 \le a < b \le d$, so

$$\frac{d}{2} - a \;\ge\; 1 - \frac{d}{2} \;>\; \frac{49}{100},$$

and likewise $b - \frac{d}{2} > \frac{49}{100}$: the midpoint lies in the open chord at distance more than 49/100 from its ends, hence strictly interior to the threaded square. A closed unit square containing both defect endpoints contains the midpoint by convexity, and its interior accumulates at the midpoint, so it meets the threaded square's interior. □

The exact reserves ($d < 51/50$; midpoint reserve $> 49/100$; the six deleted edges) are replayed by `scripts/s17_green_defect_midpoint_cohits.py` (PASS). (The analogous channel endpoint deletions $T12$: (p_{12}, p_{13}) and $T13$: (p_{13}, p_{14}) are Lemma 5.2(3).)

Certified Theorem 6.2 (Same and adjacent tubes; carrier graph (C8))**.** *For $0 \le \delta \le \delta_*$: no two interior-disjoint empty squares occupy the same tube or any of the adjacent pairs* $(L0, L1)$, $(L1, L2)$, $(R0, R1)$, $(R1, R2)$*; hence the possible tube-label sets of a packing are the twenty-five independent sets of the two three-edge chains $L0$–$L1$–$L2$ and $R0$–$R1$–$R2$* ($5 \times 5 = 25$)*. At $\delta = \delta_*$: every carrier square carries one of the same twenty-three endpoint-pair labels as at the endpoint;*

no label occurs twice; two labels can share an endpoint only along the conservative six-pair sharing graph; consequently every compatible family of $K \geq 1$ carriers has at least $K+1$ distinct endpoints, and each such endpoint is unused by every packed square. There are exactly 10,339 *conservative carrier sets, each a forest of paths.*

This statement merges one certificate with two exact enumerations, and we record which part each supplies. The adjacent-tube exclusion is C8: it places each threaded square in the exact unit-rhombus threading chart with $|\sin\beta| \leq 3/20$, where β is the chart angle of the threaded square against the defect direction: writing e, f for unit vectors along the short (defect) and long diagonals of the unit rhombus, the square's frame vector crossing the defect is $\cos\beta\, e + \sin\beta\, f$ (this β is unrelated to the clearance β of (5)), and checks, on an 8×8 outward-rounded subdivision per pair (256 boxes in all, after four sign and isometry reductions), that every candidate separating axis fails; the largest possible axis gap lies below -0.0159965429520. The carrier statements are an exact enumeration: the $17 + 8$ tangent types, with two multiply-typed labels, give the twenty-three labels, and of the 62 shared-type pairs only 5 necessary sharings survive, all inside the conservative six-pair graph; the enumeration, including all 10,339 carrier sets and the minimum endpoint surplus 1, is replayed by `scripts/s17_green_deformed_carrier_graph.py` (PASS). The unused-endpoint claim is the convexity-accumulation argument of Lemma 6.1 applied at a side contact: a packed square whose interior reached a shared carrier endpoint would meet the carrier's interior.

6.3 Tube losses and promotions

The ledger charges each occupied tube through a certified loss or promotion clause; the four packages below cover the six tubes.

Certified Theorem 6.3 (Retained outer package (C9))**.** *For $0 \leq \delta \leq \delta_*$: an $L0$-tube square forces p_1 lost, and an $R2$-tube square forces p_{14} lost. Moreover an $L0$ square is interior-incompatible with every contained unit square whose closed set contains both p_1 and p_2, and an $R2$ square with every such square containing both p_{13} and p_{14}: the closed-doubleton deletions $L0$: (p_1, p_2) and $R2$: (p_{13}, p_{14}).*

The endpoint's analytic argument, a four-axis separating-axis contradiction against the threaded chart using the exact scaffold constants, is re-verified with exact intervals over the whole deformation range. The top statements transfer by the scaffold inversion $p_i \mapsto p_{15-i}$, which maps the top wall to a support at distance $m_y + \eta$; the lower proofs use only the wall-support inequality, uniform over $m_y \leq m_y + \eta \leq m_y + 4\delta$. The replay serializes twenty-one exact margins in $\mathbb{Q}(\sqrt{2}, p)$; the minimum uniform, $L0$, and doubleton margins are 0.00180117752187, 0.000393910788804, and 0.00533486413337.

The endpoint's remaining single-point outer clauses, $R0 \Rightarrow p_2$ lost and $L2 \Rightarrow p_{13}$ lost, are false at this target: their decisive boundary scalar changes sign near $\delta \approx 0.0014255$, and at $\delta = 3/2000$ an exact algebraic counterexample exhibits an $R0$ tube square coexisting with a singleton p_2-owner. See the remark on retired hypotheses in Appendix A; the certified replacement is disjunctive.

Certified Theorem 6.4 (Disjunctive outer losses (C12))**.** *At $\delta = \delta_*$: an $R0$-tube square forces p_2 lost or p_1 lost, and an $L2$-tube square forces p_{13} lost or p_{14} lost.*

If neither point of a disjunction is lost, both points have singleton owners, D of p_2 and E of p_1, with E excluding p_0 and p_4 from its interior. Eight signed axes for the pair (F, D), where F is the tube square, eight for (D, E), and four escape faces for each of p_0, p_4 give $8 \times 8 \times 4 \times 4 = 1{,}024$ branches of a thirteen-row system in the four owner coordinates; the tube translation is support-eliminated exactly. The top disjunction follows by the container half-turn H, using the exact pointwise identity $H(p_i) = p_{15-i} + (0, \eta)$ and the fact that the certificate retains only translation-invariant rows plus one containment row with slack $m_y + \eta + \sqrt{2} < \Lambda + \delta_*$. The tree has 222,128 nodes; the minimum reserve is $9.21935842511 \times 10^{-8}$.

Certified Theorem 6.5 (Simultaneous outer tubes (C13)). *At $\delta = \delta_*$: if both an $L0$ and an $R0$ tube square occur, then p_2 is unused; if both an $L2$ and an $R2$ tube square occur, then p_{13} is unused.*

A packed square whose interior contains the point is interior-disjoint from both tube squares; after exact support elimination of both tube translations this yields sixty-four branches of a ten-row system in the two owner coordinates, and the top pair again transfers by the half-turn identity. The tree has 91,630 nodes; the minimum reserve is $9.44878172180 \times 10^{-7}$.

Certified Theorem 6.6 (Upper-middle losses (C11)). *For $0 \le \delta \le \delta_*$: an $L1$-tube square forces p_9 lost or p_{12} lost, and an $R1$-tube square forces p_6 lost or p_3 lost.*

The endpoint's upper-domino argument, a chain of projection eliminations over the three squares in play, is made uniform in δ by partitioning $[0, 1/200]$ into twenty-eight exact $\mathbb{Q}(\sqrt{2})$-enclosed slabs and replaying 42 strict projection charts per slab (1,176 in all) by exact interval Bernstein conversion, together with 38 structural charts and 336 weak charts. The one deformation-sensitive discrete sign, $t^2 - (1-u)u > 0$, stays strict on the range, and the six endpoint-nonstrict charts are closed by an exact monotonicity in t. All strict charts lie below -0.0202065299533. (The endpoint's lower-domino clauses are not used at this target.)

7 The right channel and channel–tube interactions

This section proves the three certified facts about the right channel $T13$ and the channel–tube pairs that the ledger consumes: the endpoint loss $T13 \Rightarrow p_{14}$ (C1), the exclusions $T13 \perp R2$ and $T12 \perp L2$ (C10), and the conditional $L2{+}T13$ repair (C14). All three follow the reduction pattern of Section 6.1; census, margin, and hash data are recorded in Appendix A.

Certified Theorem 7.1 ($T13$ endpoint loss (C1)). *At $\delta = \delta_*$, a $T13$ channel square forces p_{14} lost: any contained square, interior-disjoint from the channel square, with p_{14} in its interior also contains p_{11} or p_{15} in its interior.*

A contained square with p_{14} but neither p_{11} nor p_{15} in its interior must let both points escape through a face and be separated from the channel square along a signed axis: four escape faces for each of p_{11}, p_{15} and eight signed owner–channel axes give $4 \times 4 \times 8 = 128$ branches. In each branch the channel square's translation variables, ranging over the channel triangle $0 \le y \le g$, $|\zeta| \le y/\kappa$ of the chart (5), are eliminated by the exact support identity (6), leaving an eleven-row system in the two owner coordinates, closed by an interval–Farkas certificate. The statement genuinely fails at $\delta \approx 0.005872$, by an explicit configuration rather than an interval artifact; see the remark in Appendix A.

Certified Theorem 7.2 (Channel–tube exclusions (C10)). *For $0 \le \delta \le \delta_*$, $T13$ is incompatible with $R2$, and $T12$ with $L2$: for either pair, every candidate separating-axis projection overlaps with margin greater than $1/50$.*

By the separating-axis theorem, no packing contains two squares of which one is a $T13$ channel square and the other an $R2$ tube square, or one a $T12$ channel square and the other an $L2$ tube square: exactly the pairwise form in which the exclusions enter the ledger.

After an isometry the tube square lies in the outer-defect chart, and the channel square's translation is eliminated by the exact affine-minimum identity over the channel triangle. Conservative lower bounds for the eight directed projection gaps reduce the claim to a three-variable sign proof over $\delta \in [0, \frac{1}{200}]$, $r \in [-\frac{7}{20}, \frac{7}{20}]$, $a \in [-1, 0]$, where r and a are the reduced orientations of the channel and tube squares respectively (after the isometry, the tube frame $e = (c, s)$ has $a = c - s \in [-1, 0]$); outward-rounded subdivision covers the box, and the smallest certified overlap exceeds 0.0200011262671.

Certified Theorem 7.3 (Conditional $L2$+$T13$ repair (C14))**.** *At $\delta = \delta_*$, with the channel radius $R_\star = 27/100$ of Lemma 5.2, if an empty square is properly threaded by L2 and a contained square lies in the closed $T13$ channel, then p_{13} is unused.*

A putative owner D of p_{13}, the channel square Q, and the tube square F give three signed separating-axis choices, hence $8^3 = 512$ branches of one six-variable system retaining the exact channel triangle and the exact $L2$ translation rectangle. The certificate is pinned to exactly $\delta = 1/200$ and $R_\star = 27/100$; it asserts no uniform range. This conditional statement replaces a retired unconditional $T13 \Rightarrow p_{13}$ claim, and an earlier "uniform" v2 of the certificate was retired when aggregate validation found its parameter reuse unjustified; see the remark in Appendix A.

8 The left channel: the nested $T12$ cascade

At the target the left channel forces no unconditional endpoint loss: the implication that $T12$ forces p_{13} lost fails above $\delta \approx 0.0019240$ (the sharp certified threshold satisfies $0.001924054 < \delta_{T12} < 0.001924055$), and the weaker disjunction "p_{12} or p_{13} lost" used at $\delta = 19/5000$ also fails at 1/200. The replacement is a nested handoff (C2) whose residual branches are closed one by one (C3–C7); Proposition 8.7 assembles the four exhaustive alternatives the ledger consumes. Throughout this section $\delta = \delta_*$ unless a range is stated, Q denotes a closed $T12$ channel square, r its reduced orientation ($|r| < 27/100$ by Lemma 5.2), and a_D the reduced orientation of a putative owner square D, written a_{12} when $D = D_{12}$. For a unit square with frame $e = (c, s)$, $n = (-s, c)$ we name its closed faces $e^\pm$ (outward normal $\pm e$) and $f^\pm$ (outward normal $\pm n$; the source charts write (e, f) for the frame, their f being our n); escape faces (Section 6.1) are named accordingly.

Certified Theorem 8.1 (Endpoint residual (C2))**.** *At $\delta = \delta_*$, in the presence of a closed $T12$ channel square Q, exactly one of the following nested alternatives holds.*

1. *p_{13} is lost; or*
2. *p_{13} has a canonical singleton owner D_{13}, and then necessarily*

$$\frac{57}{400} \le r \le \frac{27}{100},$$

p_{10} leaves D_{13} through its e^- face, p_{14} through its e^+ face, and the Q, D_{13} separating axis is one of two canonical choices; and, conditional on this handoff, either

(2a) p_{12} is lost, or

(2b) p_{12} also has a canonical singleton owner D_{12}, with

$$\frac{57}{400} \le r \le \frac{41}{200}, \qquad -\frac{41}{200} \le a_{12} \le -\frac{17}{200},$$

p_8 leaving D_{12} through its f^- face and the Q, D_{12} axis one of two canonical choices.

The channel square enters through the chart (5), its translations eliminated by exact supports. For each endpoint there are 128 branches: owner escape faces for its two neighbours times signed Q–owner axes. Of the p_{13} branches, 126 close over the full channel interval and two close for $r \le 57/400$, leaving exactly the displayed handoff; of the p_{12} branches, 120 close globally and the surviving eight group into two weaker systems which close for $r \ge 41/200$ and, on the middle slab, outside the displayed orientation window. The certificate contains 254 one-owner records and no coupled two-owner records: at this target the four coupled six-variable systems that would have boxed the two owners jointly are feasible, a certified dead end no subdivision can rescue, so the theorem exposes the double-singleton residual rather than suppressing it. Its minimum margin, $1.99418732628 \times 10^{-9}$, is the weakest in the entire proof.

Certified Theorem 8.2 (Singleton escape (C3))**.** *Assume the singleton-p_{13} handoff of alternative 2 of Theorem 8.1, and write $D = D_{13}$. Then $-\frac{27}{100} \le a_D \le -\frac{17}{100}$, and:*

1. *if $\frac{57}{400} \le r \le \frac{21}{100}$, then p_{10} is lost;*
2. *if $\frac{21}{100} \le r \le \frac{27}{100}$, then either p_{10} is lost or it has a canonical singleton owner S, in which case $-\frac{27}{100} \le a_D \le -\frac{867}{5000}$.*

Five owner-only guard systems establish the orientation window; then, for each of the two residual Q, D axes, eight signed D, S axes with support elimination close five, and the three survivors refine by escape faces of p_5, p_9 (and once p_6), for 114 roots. S is assumed only to exclude p_5, p_9, p_6: no wall constraint, no Q, S separation, no other point exclusion. The hardest branches eliminate the orientation of S by exact pointwise dual cancellations of the form $1 + A\cos d + B\sin d \ge 1 - \sqrt{A^2 + B^2} > 0$ rather than by relaxation.

Certified Theorem 8.3 (High-radius $L1$-exclusion tail (C4))**.** *In the high-radius handoff of Theorem 8.2 ($\frac{21}{100} \le r \le \frac{27}{100}$, $-\frac{27}{100} \le a_D \le -\frac{867}{5000}$, S the canonical singleton owner of p_{10}), no $L1$-threaded empty square is interior-disjoint from S. Consequently this branch excludes the tube $L1$.*

There are $2 \times 8 \times 4 \times 4 \times 8 = 2{,}048$ cells: the residual Q, D axis, a signed D, S axis, escape faces of p_5 and p_9, and a signed F, S axis, where F is the $L1$-threaded square. Twenty-one sensitive cells are refined by the four p_6 escape faces, for 2,111 roots. The $L1$ translation rectangle is support-eliminated only in its own row; all orientation dependencies are retained. This is the largest certificate in the proof.

Certified Theorem 8.4 (Special midpoint deletion (C5))**.** *In the singleton escape of Theorems 8.1 and 8.2, the midpoint of $p_{10}p_{14}$ lies in* $\operatorname{int} D_{13}$ *with face depth exceeding* $13/100$. *Hence no square interior-disjoint from D_{13} contains both p_{10}, p_{14} in its closed set: the co-hit edge (p_{10}, p_{14}) is deleted whenever the p_{13}-owner survives.*

Eight branches: owner-face bounds, the two residual axes, the orientation guard, and $\frac{57}{400} \le r \le \frac{27}{100}$. The convexity and accumulation step is that of Lemma 6.1.

Certified Theorem 8.5 (Conditional $L1$+$T12$ co-hit deletion (C6))**.** *For $0 \le \delta \le \delta_*$ and channel radius $R_\star$, if a $T12$ channel square and an $L1$-threaded empty square are both present, no unit square whose closed set contains both p_9, p_{13} is interior-disjoint from both. The edge (p_9, p_{13}) is deleted whenever $L1$ and $T12$ are both selected.*

There are 64 branch pairs; the channel and $L1$ translations are eliminated by exact supports as in (6), and each branch closes a ten-row system in the comparison square's coordinates. The unconditional deletion $T12 \Rightarrow \neg(p_9, p_{13})$ is false, by an exact tangential counterexample (Appendix A), so the $L1$ hypothesis is essential.

Certified Theorem 8.6 (Double-singleton p_8 loss (C7))**.** *In the double-singleton branch (2b) of Theorem 8.1, if an $L1$-threaded empty square is selected then p_8 is lost.*

A putative singleton owner S of p_8 excludes p_4; the packed pairs D_{12}, S and F, S contribute signed axes, with F the $L1$-threaded square, and p_4 contributes four escape faces: $2\times8\times8\times4 = 512$ roots. The certificate deliberately drops the Q, S, D_{12}, F, and Q, F disjointness constraints, all further containment, and every S-point exclusion except p_4, and still closes.

branch	forced lost	deleted edge	conditional effect
standard (p_{13} lost)	p_{13}	—	—
singleton escape (2a, low r)	p_{12}, p_{10}	(p_{10}, p_{14})	—
double singleton (2b)	p_{10}	(p_{10}, p_{14})	$L1 \Rightarrow p_8$ lost
high-r tail (2a, high r)	p_{12}	(p_{10}, p_{14})	$L1$ excluded

(7)

Proposition 8.7. *At $\delta = \delta_*$, in the presence of a closed $T12$ channel square, at least one of the four alternatives of* (7) *holds, with the listed forced losses, deleted edge, and conditional effect.*

Proof. Theorem 8.1 yields alternative 1, (2a), or (2b). Alternative 1 is the standard branch. In (2b) the radius bookkeeping is decisive: $r \leq \frac{41}{200} = 0.205 < 0.21 = \frac{21}{100}$, so Theorem 8.2(1) applies and p_{10} is lost; the p_{13}-owner survives, so Theorem 8.4 deletes (p_{10}, p_{14}), and Theorem 8.6 supplies the conditional p_8 loss under $L1$. In (2a), p_{12} is lost and the p_{13}-owner survives, so Theorem 8.4 deletes (p_{10}, p_{14}); either $r \leq \frac{21}{100}$ and Theorem 8.2(1) loses p_{10} (the singleton-escape row), or $r \geq \frac{21}{100}$ and Theorem 8.2(2) either loses p_{10} (the same row) or hands off a canonical singleton p_{10}-owner S, in which case Theorem 8.3 excludes $L1$ (the high-radius tail row). □

9 The ledger and the proof of the main theorem

This section states the Conditional Theorem, whose hypotheses collect in one interface every geometric statement proved so far, proves it by exact finite enumeration, and derives Theorem 1.1 and Corollary 1.2.

Fix $\delta = \delta_*$, the container $[0, \Lambda + \delta_*]^2$, and the deformed scaffold P. The *labels* are the possible kinds of empty square supplied by Theorem 5.6: the six defect tubes $L0, \ldots, R2$, the twenty-three carrier labels, and the two channels $T12$, $T13$; a label is *realized* when some packed empty square is of that kind (for tubes we also say the tube is *threaded*). The hypotheses use four kinds of clause, with Γ the conservative co-hit graph of Section 4. A *loss clause* $X \Rightarrow p$ (or $X \Rightarrow p \vee p'$) asserts that whenever X is realized the point p is lost (respectively, p or p' is lost). A *deletion* $X\colon (p, q)$ asserts that whenever X is realized, no packed square other than the clause's witness squares contains both p and q in its *closed* set; the witnesses are the square realizing X for the tube and channel deletions of (H3)–(H5), the surviving p_{13}-owner for the (p_{10}, p_{14}) deletion of (H8), and the realizing pair for (H9). This closed form is what the cited theorems prove, and both of its consequences are used. First, no packed square contains both p and q in its interior: such a square would contain at least two scaffold points, while every witness is an empty square or a singleton owner; the edge pq is therefore removed from Γ before matching numbers are computed, and the hypothesis of Lemma 4.3 survives the removal (every pair inside a hit-set is at distance $\leq \sqrt{2}$, hence a Γ-edge, and a deleted pair lies in no hit-set). Second, no packed square other than a witness realizes the carrier label (p, q), since a carrier square contains both its endpoints in its closed set; the enumeration uses this to discard rows. An *exclusion* $X \perp Y$ asserts that no packing contains two squares realizing X and Y respectively. A *promotion* $X + Y \Rightarrow p$ *unused* asserts that whenever X and Y are both realized the point p is unused.

Conditional Theorem 9.1 (The $T12$-escape ledger)**.** *At $\delta = \delta_*$, assume:*

(H1) *(classification and collision rules) every empty contained square realizes a label; no two packed squares thread the same defect tube or an adjacent pair, so the threaded tubes form one of the twenty-five independent sets of the two chains $L0$–$L1$–$L2$ and $R0$–$R1$–$R2$; the carrier labels realized form one of the* 10,339 *conservative carrier sets, no label twice; and every family of $K \geq 1$ realized carriers has at least $K+1$ distinct endpoints, each unused by every packed square;*

(H2) *the tube-loss clauses*

$$L0 \Rightarrow p_1, \quad R0 \Rightarrow p_2 \vee p_1, \quad L1 \Rightarrow p_9 \vee p_{12}, \quad R1 \Rightarrow p_6 \vee p_3, \quad L2 \Rightarrow p_{13} \vee p_{14}, \quad R2 \Rightarrow p_{14};$$

(H3) *the closed-doubleton deletions* $L0\colon (p_1, p_2)$ *and* $R2\colon (p_{13}, p_{14})$*;*

(H4) *the generic midpoint deletions: for each threaded tube, the deletion of the edge joining its two defect endpoints;*

(H5) *the channel facts: each channel is realized by at most one packed square, and the deletions* $T12\colon (p_{12}, p_{13})$*,* $T13\colon (p_{13}, p_{14})$*;*

(H6) *the exclusions* $T12 \perp L2$ *and* $T13 \perp R2$*;*

(H7) *the loss* $T13 \Rightarrow p_{14}$*;*

(H8) *the four exhaustive* $T12$ *alternatives: if* $T12$ *is realized, at least one of the following holds:* (a) standard*:* p_{13} *is lost;* (b) singleton escape*:* p_{12} *and* p_{10} *are lost, with the deletion* (p_{10}, p_{14})*;* (c) double singleton*:* p_{10} *is lost, with the deletion* (p_{10}, p_{14})*, and if* $L1$ *is threaded then* p_8 *is lost;* (d) high-radius tail*:* p_{12} *is lost, with the deletion* (p_{10}, p_{14})*, and* $L1$ *is not threaded;*

(H9) *the conditional deletion: if* $L1$ *and* $T12$ *are both realized, the deletion* (p_9, p_{13})*;*

(H10) *the unused-point promotions* $L0{+}R0 \Rightarrow p_2$ *unused,* $L2{+}R2 \Rightarrow p_{13}$ *unused, and* $L2{+}T13 \Rightarrow p_{13}$ *unused.*

Then seventeen pairwise interior-disjoint unit squares do not fit in the closed square of side $\Lambda + \delta_*$*.*

Proof. The proof has two halves: an exact enumeration of the finitely many label combinations a packing could realize, each checked against a charging inequality, and a counting step converting the per-row inequality into $U + R \geq E$.

Suppose seventeen pairwise interior-disjoint unit squares occupy the closed square of side $\Lambda + \delta_*$, and let E, U, R, and the lost set B be as in Lemma 4.2 and Lemma 4.3. Lemma 4.2 gives $U + R = E - 1$, so it suffices to prove $U + R \geq E$.

A *row* is a tube set among the twenty-five independent sets, a carrier set among the 10,339 conservative carrier sets, a subset of $\{T12, T13\}$, and, when $T12$ is present, a choice of one alternative (H8)(a)–(d), subject to three compatibility conditions: the exclusions (H6); in branch (d), the absence of $L1$; and no carrier label of the row coinciding with an edge deleted under the row's clauses (only the closed-doubleton and channel deletions of (H3) and (H5) can so coincide; the other deleted edges are defect edges, which are not carrier labels). Its *label count* λ is the number of labels in the combination.

The packing realizes a compatible row with $\lambda = E$. Assign to each empty square one label it realizes (Theorem 5.6); since no label is realized by two packed squares — tubes and channels by (H1) and (H5), carriers by (H1) — distinct squares receive distinct labels, so the assigned labels, together with the (H8) alternative that holds, form a row with $\lambda = E$ all of whose labels are realized, and any two of its labels are realized by *distinct* packed squares. The compatibility conditions follow: a pair violating (H6) would put its two labels on two distinct packed squares, which (H6) forbids; in branch (d) no square threads $L1$ at all, so $L1$ is not among the assigned labels; and a carrier label coinciding with a deleted edge would put both endpoints of that edge in the closed set of a packed square distinct from the clause's witnesses, which the deletion forbids.

For each row form:

- A: the carrier endpoints forced unused by (H1), at least $K + 1$ of them for K carriers, augmented by the points promoted by (H10) when the promoting labels are present;
- W: a *witness tuple*, selecting one point from each applicable loss clause (namely (H2) for each tube of the row, (H7) when $T13$ is present, the forced losses of the row's $T12$ branch

(H8), and, when the row has both $L1$ and branch (c), the p_8 loss), with a free choice of disjunct wherever a clause is disjunctive; W is the *set* of selected points, so coincident selections count once; and $M := W \setminus A$, the selected witnesses outside A;

- $\Gamma[M]$: the subgraph of Γ induced on M after removing the edges deleted under (H3), (H4), (H5), (H8), and (H9), as licensed by the row's labels.

The exact finite computation enumerates every row and verifies

$$|A| + |M| - \nu(\Gamma[M]) \;\geq\; \lambda$$

for every witness tuple W, with matching numbers evaluated by exact integer combinatorics; it is replayed by `scripts/s17_green_t12_escape_conditional_ledger.py` (Appendix A). The left side equals $|A \cup W| - \nu(\Gamma[W \setminus A])$, so a selected witness lying in A is counted once, through $|A|$. There are 263,483 rows without $T12$; 102,601 in each of the standard, singleton-escape, and double-singleton branches; and 63,276 in the high-radius tail, where $L1$ is excluded:

$$263{,}483 + 3 \cdot 102{,}601 + 63{,}276 \;=\; 634{,}562 \tag{8}$$

rows in all. The computation reports zero failures, and the minimum reserve $|A|+|M|-\nu(\Gamma[M])-\lambda$ over all rows and witness tuples is zero.

The row inequality bounds $U + R$ as in Lemma 4.3. Each applicable clause holds for the packing, so some disjunct of it is lost; selecting one per clause yields a realized witness tuple W. Its points are lost, and every point of A is unused, hence lost; so $M = W \setminus A$ satisfies $A \cup M \subseteq B$ with A unused. Then $U \geq |A| + \mu$, where μ is the number of unused points of M. Group the used points of M by their using squares; each such square contains at least two scaffold points in its interior, because its M-points are lost, and distinct squares have disjoint interior hit-sets. A square whose hit-set is not contained in M contributes at least its number of M-points to R. A square whose hit-set lies inside M contributes one fewer, and its hit-set spans an edge of $\Gamma[M]$ that survives the deletions (by the first consequence of the deletion clauses, no packed square contains a deleted pair in its interior), so these squares carry vertex-disjoint edges of $\Gamma[M]$ and number at most $\nu(\Gamma[M])$. Summing, $R \geq (|M| - \mu) - \nu(\Gamma[M])$, whence

$$U + R \;\geq\; |A| + |M| - \nu(\Gamma[M]) \;\geq\; \lambda \;=\; E,$$

contradicting $U + R = E - 1$. Only certified losses enter M: as in the monotonicity clause of Lemma 4.3, counting a subset of the true lost set yields a valid bound, so discarding uncertified losses is safe. □

The enumeration is deliberately over-permissive (it forgets realization flags, allows both channels simultaneously, and minimizes over witness choices) and is therefore a sound upper approximation of what a packing could do.

Proof of Theorem 1.1. Every hypothesis of Theorem 9.1 is a proved statement at $\delta = \delta_*$. (H1) is Theorem 5.6 together with Theorem 6.2, which supplies the same- and adjacent-tube exclusions (hence the twenty-five tube cases), the 10,339 conservative carrier sets, and the unused-endpoint surplus; the conservative co-hit graph itself is valid at the target by Proposition 4.4. (H2): the $L0$ and $R2$ clauses are Theorem 6.3, the $R0$ and $L2$ disjunctions are Theorem 6.4, and the $L1$ and $R1$ disjunctions are Theorem 6.6. (H3) is Theorem 6.3. (H4) is Lemma 6.1. (H5) is Lemma 5.2. (H6) is Theorem 7.2. (H7) is Theorem 7.1. (H8) is Proposition 8.7, assembled from Theorems 8.1, 8.2, 8.3, 8.4, and 8.6. (H9) is Theorem 8.5. (H10): the two-tube promotions are Theorem 6.5, and the $L2$+$T13$ promotion is Theorem 7.3.

Theorem 9.1 therefore applies unconditionally: seventeen pairwise interior-disjoint unit squares cannot be contained in the closed square of side $\Lambda + \delta_*$, which is the first assertion of the theorem.

For strictness, suppose $s(17) \le \Lambda + \delta_*$. By Lemma 2.1 the infimum defining $s(17)$ is attained: some seventeen pairwise interior-disjoint unit squares fit in $[0, s(17)]^2 \subseteq [0, \Lambda + \delta_*]^2$, contradicting the exclusion just proved. Hence

$$s(17) \; > \; \Lambda + \delta_* \; = \; \frac{40\sqrt{2}+19}{17} + \frac{1}{200}. \qquad \square$$

Corollary 1.2 follows as in the introduction: deleting one square from a packing of eighteen leaves a packing of seventeen in the same container, so s is nondecreasing and $s(18) \ge s(17) > \Lambda + \frac{1}{200}$.

A What a referee must trust; Lean validation

This appendix is the complete interface between the paper and its computations: the verification discipline, the inventory of the fourteen certificates, the exact audits and the theorem assembler, the replay commands, and the Lean validation of the algebraic spine. The intended trust model is one sentence: *read* the mathematical sections as ordinary mathematics; *replay* the artifacts; trust nothing else.

Verification discipline

All computations follow four rules.

1. **Exact arithmetic for signs.** Every sign decision on numbers of the form $a + b\sqrt{2}$ (and, for the deformed parameters, in $\mathbb{Q}(\sqrt{2}, p)$ with $p = \sqrt{2 - w^2}$) is made by exact rational squaring, never floating point. Decimal values printed in this paper are diagnostics.

2. **Outward-rounded interval replay.** Certificate leaves are checked by rebuilding every inequality row in outward-rounded binary64 interval arithmetic with exact rational guards at square-root endpoints; a leaf is accepted only when its certified combination is strictly negative at the outward upper endpoint.

3. **Deterministic coverage.** Every internal node of a certificate tree records a bisection; replay verifies that the two closed children cover the parent and change only the recorded coordinate, so the leaves provably cover the root boxes, which cover the case analysis stated in the theorem.

4. **No optimizer in the replay path; hash pinning.** Linear programming may propose dyadic Farkas multipliers during offline generation, but replay only *checks* the recorded multipliers. Certificates are committed files; the theorem assembler loads each one and verifies its statement manifest (exact target $\delta = 1/200$, channel radius 27/100, domains, hypotheses) against what the ledger consumes; matching filenames are not accepted as evidence. The files are pinned by the SHA-256 hashes below.

The fourteen certificates

All certificate files live in `doc/certificates/` of the source repository. "Margin" is the minimum strict Farkas margin (or the stated overlap/gap bound) certified at replay. Three artifacts are of simpler kinds under the same replay discipline: C8 is a finite outward-rounded separating-axis overlap cover, C9 serializes exact margins in $\mathbb{Q}(\sqrt{2}, p)$, and C11 is a family of exact interval Bernstein sign charts.

C1 `s17-green-t13-loss-correlated.cert.json`
Role: $T13 \Rightarrow p_{14}$ lost (Theorem 7.1).

Census: 128 branches; 39,348 nodes; depth 22.
Margin: 3.58×10^{-7} (precise: $3.58063113964 \times 10^{-7}$).
SHA-256: `1c2a236ee0e92de2de9808a134b91849e14ad962f969402a8371e7096f43f05f`
Replay: `make verify-s17-green-t13-loss`

C2 `s17-green-t12-endpoint-disjunction-correlated.cert.json`
Role: the nested three-way $T12$ residual; 254 one-owner records, no coupled records (Theorem 8.1).
Census: 299,078 nodes; 148,878 leaves; depth 30.
Margin: 1.99×10^{-9}, weakest in the proof (precise: $1.99418732628 \times 10^{-9}$).
SHA-256: `05bbf48c493ba22cb3724716dab974ea8c5db2c841d086fba718cd73578774f7`
Replay: `make verify-s17-green-t12-endpoint-disjunction`

C3 `s17-green-t12-p10-loss-correlated.cert.json`
Role: singleton escape B: orientation guard, low-radius p_{10} loss, high-radius handoff (Theorem 8.2).
Census: 114 roots; 854,336 nodes; depth 32.
Margin: 8.65×10^{-9} (precise: $8.64956128943 \times 10^{-9}$).
SHA-256: `ad72af169d817de2ae34a45b0da457ad80074172b9a3d75f5a6a9f27db2cf0d5`
Replay: `make verify-s17-green-t12-p10-loss`

C4 `s17-green-t12-p10-l1-tail-correlated.cert.json.gz`
Role: high-radius $L1$-exclusion tail; deterministic-gzip certificate (Theorem 8.3).
Census: 2,111 roots (2,027 generic + 84 refined); 4,519,791 nodes; depth 24; no prunes; the largest certificate.
Margin: 2.26×10^{-8} (precise: $2.26107571347 \times 10^{-8}$).
SHA-256: `4a02034353d3a83408c65f97532b29cf35d2c4067f68f4b5f7fa2629f2c1c48f`
Replay: `make verify-s17-green-t12-p10-l1-tail`

C5 `s17-green-t12-p13-midpoint-cohit-correlated.cert.json`
Role: special (p_{10}, p_{14}) midpoint deletion (Theorem 8.4).
Census: 8 roots; 2,474 nodes; depth 22.
Margin: 1.13×10^{-6} (precise: $1.12739721085 \times 10^{-6}$).
SHA-256: `27ea36114914c582180383cfbd359a886f97c8814514964c76f763b51eb22bff`
Replay: `make verify-s17-green-t12-p13-midpoint-cohit`

C6 `s17-green-t12-l1-cohit-correlated-range.cert.json`
Role: conditional $L1{+}T12 \Rightarrow \neg(p_9, p_{13})$ over the full range $\delta \in [0, 1/200]$ (Theorem 8.5).
Census: 64 branch pairs; 163,386 nodes; depth 18.
Margin: 1.17×10^{-7} (precise: $1.17396176462 \times 10^{-7}$).
SHA-256: `80bde90134def1286ed8a9f33c72cc3e2535a59c810546431ad6bfa7309bb830`
Replay: `make verify-s17-green-t12-l1-cohit`

C7 `s17-green-t12-l1-p8-loss-correlated.cert.json`
Role: double-singleton $L1 \Rightarrow p_8$ lost (Theorem 8.6).
Census: 512 roots; 21,274 nodes; depth 12; no prunes.
Margin: 2.53×10^{-7} (precise: $2.52578073345 \times 10^{-7}$).
SHA-256: `5baa8feb7fdf389323a13dbda378c708b83581ab9cab81869ae8ab176532f03c`
Replay: `make verify-s17-green-t12-l1-p8-loss`

C8 `s17-green-adjacent-threading-perturbation.cert.json`
Role: same/adjacent-tube exclusion over $[0, 1/200]$ (Theorem 6.2).
Census: 4 pairs $\times$ 64 boxes.
Margin: largest axis gap < -0.01599654 (precise bound: below -0.0159965429520).
SHA-256: `2e4b574ecb1913e8fee819bca40435b00ffe8f28523e8535dcbc1d83518e5ca0`
Replay: `make verify-s17-green-adjacent-threading-perturbation`

C9 `s17-green-deformed-outer-retained.cert.json`
Role: retained $L0, R2$ losses and both closed doubletons (Theorem 6.3).
Census: 21 serialized exact margins.
Margin: 1.80×10^{-3} / 3.94×10^{-4} / 5.33×10^{-3} (minimum uniform / $L0$ / doubleton; precise: 0.00180117752187, 0.000393910788804, 0.00533486413337).
SHA-256: `eade49da442bf701f2a268fadb62243dc241ebd0c0ed7743c1b2b049e3683075`
Replay: `make verify-s17-green-deformed-outer-packages`

C10 `s17-green-top-channel-tube-direct.cert.json`
Role: $T13 \perp R2$ and $T12 \perp L2$ over $[0, 1/200]$ (Theorem 7.2).
Census: 11,437 leaves; 4,745 prunes; 16,181 splits.
Margin: overlap > 0.0200011262671.
SHA-256: `b0774718c004772ddd9c4d4b7a6f34d04d4536e6dd792c457f56c7bc2e6857e5`
Replay: `make verify-s17-green-top-channel-tube`

C11 `s17-green-upper-middle-correlated-range.cert.json`
Role: upper-middle disjunctions over $[0, 1/200]$, 28 exact slabs (Theorem 6.6).
Census: 1,176 strict + 38 structural + 336 weak charts.
Margin: 2.02×10^{-2} (all strict charts lie below -0.0202065299533).
SHA-256: `a1f850427ace6a6ca1b2afdd10643003a969865028bff45da8206ef31c83e330`
Replay: `make verify-s17-green-upper-middle-perturbation`

C12 `s17-green-r0-disjunctive-loss-correlated.cert.json`
Role: $R0$ and (by half-turn) $L2$ disjunctive losses (Theorem 6.4).
Census: 1,024 branches; 222,128 nodes; depth 20.
Margin: 9.22×10^{-8} (precise: $9.21935842511 \times 10^{-8}$).
SHA-256: `6989e0f45d4cd54376493386922a58a89e4e532299e1d32bdd97322423ba83b3`
Replay: `make verify-s17-green-r0-disjunctive-loss`

C13 `s17-green-l0-r0-p2-unused-correlated.cert.json`
Role: $L0{+}R0 \Rightarrow p_2$ unused; $L2{+}R2 \Rightarrow p_{13}$ unused (Theorem 6.5).
Census: 64 branches; 91,630 nodes; depth 18.
Margin: 9.45×10^{-7} (precise: $9.44878172180 \times 10^{-7}$).
SHA-256: `3e8920b428af96e6e38204c50bcca7825adc54b115b2c535b8961c7ff1c9b02e`
Replay: `make verify-s17-green-l0-r0-p2-unused`

C14 `s17-green-l2-t13-p13-unused-correlated.cert.json`
Role: $L2{+}T13 \Rightarrow p_{13}$ unused, pinned to $\delta = 1/200$, $R_\star = 27/100$; asserts no uniform range (Theorem 7.3).
Census: 512 branches; 98,394 nodes; depth 24.
Margin: 7.74×10^{-6} (precise: $7.74068054793 \times 10^{-6}$).
SHA-256: `1dc7a1adf834684368a788e6f63cff72d4a9ad507318e7c7dd281541e99affc9`
Replay: `make verify-s17-green-l2-t13-p13-unused`

Exact (certificate-free) audits and the ledger

Five scripts recompute their claims from scratch in exact arithmetic, without certificates:

- `scripts/s17_green_deformed_target.py` — the target constants (δ_*, $R_\star$, the deformed t, u, η).
- `scripts/s17_green_deformed_classification.py` — all radical comparisons behind Theorem 5.6: fences, gaps, η, ℓ_0, the 18 faces, 17 unit edges, 6 defect edges. Replay: `make verify-s17-green-deformed-classification`.

- `scripts/s17_green_deformed_carrier_graph.py` — the $17 + 8$ tangent types, 23 labels, sharing graph, and the enumeration of all 10,339 conservative carrier sets with the $K + 1$ endpoint surplus. Replay: `make verify-s17-green-deformed-carrier-graph`.

- `scripts/s17_green_defect_midpoint_cohits.py` — the exact reserves $d < 51/50$, midpoint reserve $> 49/100$. Replay: `make verify-s17-green-defect-midpoint-cohits`.

- `scripts/s17_green_t12_escape_conditional_ledger.py` — the finite bridge: 25 tube cases $\times$ 10,339 carrier sets $\times$ 4 top-channel cases, 634,562 compatible rows ($263{,}483 + 3 \cdot 102{,}601 + 63{,}276$), each satisfying $|A| + |M| - \nu(\Gamma[M]) \geq E$, zero failures, minimum reserve zero. The matching computation is exact integer combinatorics. Replay: `make verify-s17-green-t12-escape-ledger`.

The theorem assembler `scripts/s17_green_deformed_bound.py` does three jobs: it checks every certificate's statement manifest against the exact hypothesis consumed by the ledger at the exact target; it audits the incidence identity and directly re-checks the 81 co-hit non-edges in the deformed scaffold (minimum squared-distance excess over 2: 0.226781706524); and it enumerates the complete ledger. Replay: `make verify-s17-green-deformed-bound`.

The ledger branch row counts are fixed:

`none`	263,483
`p13_lost`	102,601
`p13_singleton_escape`	102,601
`p12_p13_double_singleton_escape`	102,601
`p13_p10_double_singleton_l1_excluded_tail`	63,276

Total 634,562.

Replay commands

Table 2 lists every replay command. The target `verify-s17-green-deformed-bound` is the authoritative theorem-level replay; passing a local certificate or the ledger alone is not a theorem-level check, because only the assembler verifies the statement manifests against what the ledger consumes.

Status of the replays

The aggregate replay is long-running, dominated by C4's 4,519,791-node tree and C3's 854,336 nodes. The source repository's research record contains a complete passing aggregate run pinning all fourteen manifests and the theorem wiring to $\delta = 1/200$. While preparing this paper I re-ran the exact classification, carrier-graph, and defect-midpoint audits and the complete finite ledger; all passed, the ledger in under ten minutes on commodity hardware, confirming the row split with zero failures and minimum reserve zero. I independently recomputed the SHA-256 hashes of all fourteen certificate files; all match the values printed above. I then ran the full aggregate replay `scripts/s17_green_deformed_bound.py` to completion on my own machine, and it passed: 634,562 rows, `final_failures=0`, `minimum_reserve=0`, `minimum_farkas_margin=1.99418732628e-09`. A referee who replays the displayed commands is independent of every claim in this subsection.

Remark A.1 (Retired and corrected hypotheses). Several natural unconditional claims are false at the target and were replaced by certified statements. $T12 \Rightarrow p_{13}$ fails above $\delta \approx 0.0019240$; Theorem 8.1 certifies the nested residual instead. $R0 \Rightarrow p_2$ and $L2 \Rightarrow p_{13}$ fail above $\delta \approx 0.0014255$; Theorem 6.4 certifies the disjunctions. $T12 \Rightarrow \neg(p_9, p_{13})$ has an exact tangential counterexample; the deletion holds only under $L1$ (Theorem 8.5). The coupled two-owner relaxations behind C2 are feasible, hence uncertifiable; the certificate exposes the double-singleton residual rather than

From the repository root:

```
make verify-s17-green-t12-endpoint-disjunction   # C2
make verify-s17-green-t12-p10-loss               # C3
make verify-s17-green-t12-p10-l1-tail            # C4
make verify-s17-green-t12-p13-midpoint-cohit     # C5
make verify-s17-green-t12-l1-p8-loss             # C7
make verify-s17-green-defect-midpoint-cohits     # exact midpoint audit
make verify-s17-green-t12-escape-ledger          # the 634,562-row ledger
make verify-s17-green-deformed-bound             # authoritative aggregate
```

together with the remaining focused targets:

`verify-s17-green-t13-loss`	(C1)
`verify-s17-green-t12-l1-cohit`	(C6)
`verify-s17-green-adjacent-threading-perturbation`	(C8)
`verify-s17-green-deformed-outer-packages`	(C9)
`verify-s17-green-top-channel-tube`	(C10)
`verify-s17-green-upper-middle-perturbation`	(C11)
`verify-s17-green-r0-disjunctive-loss`	(C12)
`verify-s17-green-l0-r0-p2-unused`	(C13)
`verify-s17-green-l2-t13-p13-unused`	(C14)
`verify-s17-green-deformed-classification`	(exact audit)
`verify-s17-green-deformed-carrier-graph`	(exact audit)

Table 2: Replay commands. `verify-s17-green-deformed-bound` is the authoritative theorem-level replay.

suppressing it. An earlier "uniform" v2 of C14 was retired; C14 is target-pinned, an episode that motivates statement-manifest checking. The proof is tied to $\delta_* = 1/200$ and $R_\star = 27/100$, with no continuation radius, not even at the 10^{-9} margin; and C1 genuinely fails at $\delta \approx 0.005872$, capping this architecture.

Remark A.2 (Effectivity contrast). The general semialgebraic separation bound of Jeronimo–Perrucci–Tsigaridas [5] applies here: encoding the problem with $n = 65$ variables, $m = 236$ degree-2 integer constraints, height $H = 150$, and $G(R) = (17L - 19)^2 - 3200$ yields $s(17) \geq (19 + \sqrt{3200 + 2^{-E_*}})/17 > \Lambda + 2^{-(E_*+11)}$ with $E_* = 206953526228295625811891392653598 3$, a positive but geometrically meaningless increment $\sim 10^{-6.2\times 10^{32}}$. The deformation route is larger by about 6.2×10^{32} orders of magnitude.

A.1 Lean validation

The non-computational algebraic spine, everything about the deformation constants over $\mathbb{R}$ as opposed to the certified computational packages, is formalized in Lean 4 over mathlib: library `PaperProofs`, file `PaperProofs/S17Paper.lean`, namespace `PaperProofs.S17`. The development was rebuilt end to end on 2026-08-09 with Lean v4.29.1 and mathlib v4.29.1; it contains no `sorry`, and `#print axioms` on the listed theorems reports only `propext`, `Classical.choice`, and `Quot.sound`. Proved from first principles over the reals:

- the deformation constants (Λ, w_0, $\delta = 1/200$, $w = w_0 + \delta$, $p = \sqrt{2 - w^2}$, $t = (p - w)/2$, $u = (p + w)/2$, t_0, u_0) and survival of the circle constraint: $t^2 + u^2 = 1$ and $u - t = w$ (`t_sq_add_u_sq`, `u_sub_t`), together with $t_0^2 + u_0^2 = 1$;
- the exact horizontal fit $2m_x + u + 2 = \Lambda + \delta$ (`span_identity`), through the endpoint collapses $u_0 - 4t_0 = -3$ (`u0_collapse`) and $2\sqrt{2} - 1 + 3t_0 = \Lambda$ (`height_collapse`), and the top-excess identity $\eta = (\Lambda + \delta) - (2\sqrt{2} - 1 + 3t) = \delta + 3(t_0 - t)$ (`eta_eq_top_excess`);
- the deformation bounds $2\delta < \eta < 4\delta$ at $\delta = 1/200$ (`eta_bounds`), established through explicit rational bounds on $\sqrt{2}$ and $\sqrt{2 - w^2}$;
- the exact monotonicity on $[0, 1)$ of the angular cost $A(x) = \sqrt{2 - x} + x/2 - \sqrt{2}$ (`angularCost_strictMonoOn`) and the channel-radius margin $A((27/100)^2) > \eta$ (`angularCost_margin`);
- the defect identity $E = U + R + 1$, as an abstract counting lemma over a seventeen-element finite set (`defect_identity`).

What is not formalized in Lean is exactly the certified computational content: the fourteen interval–Farkas certificate trees and the 634,562-row finite ledger. Those are validated solely by the exact-arithmetic replay pipeline of the four rules above, which remains the sole authority for them. (In the build, `eta_bounds` used the rational sandwiches $1.41421356237 < \sqrt{2} < 1.41421356238$ and $1.3626997404 < p < 1.3626997405$, with binding margin $t_0 - t - \delta/3 \approx 1.3 \times 10^{-4}$; the `angularCost` margin $\approx 3.3 \times 10^{-5}$ matches the paper's $3.31787970482 \times 10^{-5}$.)

Acknowledgments

This proof is computer-assisted in an essential way, and the assistance went well beyond arithmetic: the certificate generators, the replay infrastructure, and much of the case analysis were developed with substantial help from automated tools, including large-language-model assistance in drafting and exploring the geometry. Every mathematical claim ultimately rests on the deterministic, hash-pinned replays of Appendix A, not on any tool's say-so. I thank Erich Friedman for maintaining the survey, and Trevor Green, whose quarter-century-old bound concealed exactly enough structure to be moved.

References

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

[2] E. Friedman, *Packing unit squares in squares: a survey and new results*, Electron. J. Combin., Dynamic Survey **DS7** (2009), doi:10.37236/28. The lower-bound table attributes the value $(40\sqrt{2}+19)/17$ for $n = 17$ and $n = 18$ to T. Green (2000, private communication), and the upper bound 4.6756 for $n = 17$ to J. Bidwell (1998, private communication).

[3] T. Gensane and P. Ryckelynck, *Improved dense packings of congruent squares in a square*, Discrete Comput. Geom. **34** (2005), 97–109.

[4] F. Göbel, *Geometrical packing and covering problems*, in: *Packing and Covering in Combinatorics* (A. Schrijver, ed.), Math. Centre Tracts **106**, Amsterdam, 1979, 179–199.

[5] G. Jeronimo, D. Perrucci, and E. Tsigaridas, *On the minimum of a polynomial function on a basic closed semialgebraic set and applications*, SIAM J. Optim. **23** (2013), 241–255, doi:10.1137/110857751.

[6] M. J. Kearney and P. Shiu, *Efficient packing of unit squares in a square*, Electron. J. Combin. **9** (2002), #R14.

[7] H. Nagamochi, *Packing unit squares in a rectangle*, Electron. J. Combin. **12** (2005), #R37.

[8] W. Stromquist, *Packing 10 or 11 unit squares in a square*, Electron. J. Combin. **10** (2003), #R8.