hexagon:2611.00379v2

The Steiner ratio is at least 0.8564

TEJASVI SINGH TOMAR

Abstract. The Steiner ratio is the infimum, over finite sets of points in the Euclidean plane, of the length of a Steiner minimal tree divided by the length of a minimum spanning tree. Gilbert and Pollak conjectured in 1968 that it equals $\sqrt{3}/2 \approx 0.8660$, and the conjecture is open. We prove that it is at least $2141/2500 = 0.8564$, improving the bound 0.8559 of Ke, Huang, Shu, He, Gai and Wang. Their method reduces the induction step of Du and Hwang's 1983 proof of the bound 0.8 to finitely many inequalities on boxes in a space of local configurations of a Steiner minimal tree, which are checked by computer. These inequalities use estimates for the positions of terminals outside the local configuration. We prove all the geometric lemmas that our certificate uses, filling gaps in two proofs of Ke et al., and we add two such estimates that hold for more parameter values than the published ones. Within this method some new estimate is necessary: with the published estimates, no certificate of this kind reaches 0.856, whatever the boxes. The computer-assisted part is also verified in Lean 4, by a proof whose checks all run in Lean's kernel; the rest of the argument enters the formal proof as hypotheses.

1. Introduction

For a finite set V of points in the Euclidean plane, let $L_{\mathrm{S}}(V)$ be the length of a Steiner minimal tree for V, that is, of a shortest network that connects the points of V and may use further vertices, and let $L_{\mathrm{M}}(V)$ be the length of a minimum spanning tree of V. The *Steiner ratio* is

$$\rho_S = \inf\Big\{ \frac{L_{\mathrm{S}}(V)}{L_{\mathrm{M}}(V)} : V \subset \mathbb{R}^2,\ 2 \le |V| < \infty \Big\}.$$

Gilbert and Pollak conjectured that $\rho_S = \sqrt{3}/2 = 0.8660\ldots$, the value for the vertices of an equilateral triangle [11]. Lower bounds followed: $\frac{1}{2}$, attributed to Moore in [11]; $1/\sqrt{3} = 0.577\ldots$ by Graham and Hwang in 1976 [10]; 0.743 by Chung and Hwang in 1978 [2]; 0.8 by Du and Hwang in 1983 [3]; and 0.824 by Chung and Graham in 1985 [1]. The conjecture is known for sets of at most four points [17, 8], five points [7] and six points [18]. Du and Hwang announced a proof of the conjecture [4, 5]; Ivanov and Tuzhilin explained why the argument does not prove it [12]. In 2026, Ke, Huang, Shu, He, Gai and Wang gave a computer-assisted proof of $\rho_S \ge 0.8559$ [14].

Theorem 1.1. *For every finite set V of points in the plane, $L_{\mathrm{S}}(V) \ge \frac{2141}{2500} L_{\mathrm{M}}(V)$. Hence $\rho_S \ge 2141/2500 = 0.8564$.*

The proof follows the method of [14], which rests on the induction of Du and Hwang [3]. To prove $L_{\mathrm{S}}(V) \ge \rho L_{\mathrm{M}}(V)$ for all V and a constant ρ, one removes a few edges of a Steiner minimal tree near a leaf, connects the separated terminals to the other terminals by segments, and reconnects the rest of the tree by a new network; the segments, together with a minimum spanning tree of the other terminals, connect all the terminals. If the length removed is at least the length added plus ρ times the length of the connecting segments, that is, if the *cost* of this

2020 *Mathematics Subject Classification.* 05C05; 90C27, 65G40, 68V20.

Key words and phrases. Steiner ratio, Gilbert–Pollak conjecture, Steiner minimal tree, minimum spanning tree, computer-assisted proof, interval arithmetic, formal verification, Lean 4.

splitting is at most 0, then the inequality for sets with fewer points gives it for the whole set (Lemma 2.2). Near a suitable leaf, every full Steiner minimal tree with at least five terminals contains a local configuration of one of two types, which is determined up to similarity by six or eight edge lengths divided by the length of the edge at the leaf (Lemmas 2.4 and 2.5, Figure 1). For each type, [14] fixes a finite list of splittings and bounds the cost of each by an explicit function of these parameters. The parameter space is covered by boxes, and each box carries one splitting whose cost bound is at most 0 at the vertices of the box; this suffices because along every line parallel to a coordinate axis the cost bounds first do not increase and then do not decrease (Lemma 4.1). Some splittings connect to a terminal that lies outside the configuration, at the end of a chain of Steiner points. Each configuration has two such *hidden terminals*, E and F (Section 3.1), and geometric lemmas of [14] and [3] bound their positions under conditions on the parameters.

We use this method, the splitting tables of [14] and a modified version of the generator of boxes of [14]. Our proof differs from that of [14] in four ways. First, we prove the reduction to the two types of configuration and every bound for the hidden terminals that our certificate uses, relying only on standard facts about Steiner minimal trees and on the known case of the conjecture for at most four points. Our proofs of the published bounds rest on an exchange inequality and a turning-number argument (Lemmas 3.1 and 3.2). Two of the proofs in [14] have gaps (Remark 3.9); we give complete proofs of the statements that we use.

Second, we add two bounds. The published bound for F requires $f \le d$, where d and f are the lengths of the edges YD and ZQ (Figure 1), and one of the published bounds for E requires a linear condition on the parameters. These restrictions stop the method of [14]: by Proposition 6.1, no certificate of this form built on the lemmas of [14] reaches 0.856, since at the configurations near a configuration w_* with f slightly larger than d, every splitting of the form used in [14], not only those in its tables, either requires a bound for F that these lemmas do not give or has a positive cost bound with them. Lemma 3.12 gives the published bound for F in that type also when $f > d$, under an explicit inequality between lengths, and Lemma 3.11 replaces the linear condition for E by a less restrictive condition with the same conclusion. Both apply at w_* with a margin, and hence near it. Their proofs compare the Steiner tree with networks that replace several of its edges at once, not only with networks obtained by exchanging one edge.

Third, when a splitting reconnects four points, we bound the length of the new network by that of an explicit network with two Steiner points, instead of the closed formula of [14] for the length of a full Steiner tree, which is an upper bound only where such a tree exists, so that [14] imposes conditions for its existence. Fourth, our certificate consists of 555,102,905 boxes. A checker that shares no code with the generator verifies, in ball arithmetic, the inequality of every box that meets the domain of its family, and a second program verifies that the boxes of each family tile the orthant $[0, \infty)^n$ of its coordinates.

The conditional bound 0.860 of [19], listed in Tao's repository of optimization constants [20], uses the generator and the lemmas of [14]; Remark 6.2 explains how Proposition 6.1 bears on it.

We have also verified the computer-assisted part in the Lean 4 proof assistant [15] with the Mathlib library [16]. A checking program written in Lean verifies the boxes of all four families, including the covering of the domains, and a Lean proof shows that whenever the program accepts, the inequalities of Section 4 hold on the whole domain of each family. From this, Lean derives Theorem 1.1, with L_{S} and L_{M} defined in Lean, under hypotheses that state the rest of the argument. In one formal proof the checks are run as compiled code; a second proof of the same statement checks the same boxes with a program in integer arithmetic whose runs are evaluated by Lean's kernel, so that it relies on the kernel instead of code generated by Lean's compiler.

Section 2 sets up the induction and the two types of configuration, and Section 3 proves the bounds for the hidden terminals. Sections 4 and 5 describe the functions checked on the boxes and the certificate, and complete the proof of Theorem 1.1. Section 6 proves Proposition 6.1, and Section 7 states the formal theorems and their hypotheses, and Appendix A describes the checking programs of the formal proofs.

2. Splittings and local configurations

2.1. **Steiner minimal trees.** We use the following facts [11]. Every finite set V with $|V| \geq 2$ has a Steiner minimal tree (SMT). It is a tree whose vertices are the points of V, called *terminals*, and possibly further points, called *Steiner points*; its edges are segments, and two edges meet only at a common endpoint. Every Steiner point has degree three, and its three edges make angles of 120°. There are at most $|V| - 2$ Steiner points. An SMT with exactly $|V| - 2$ Steiner points is *full*; then every terminal is a leaf. Conversely, if every terminal is a leaf, counting degrees shows that there are $|V| - 2$ Steiner points. We call a connected finite union of segments a *network*; no network that contains V is shorter than an SMT for V.

2.2. **Splittings.** For a tree S and a set S^- of its edges, $S \setminus S^-$ denotes the union of the edges of S that are not in S^-. For a finite set N of segments, $L(N)$ is the sum of their lengths; a network formed by these segments has length at most $L(N)$.

Definition 2.1. Let S be an SMT for V. A *splitting* of S is a tuple (V^*, S^-, S^+, t^*) such that

(a) V^* is a nonempty proper subset of V;
(b) S^- is a set of edges of S;
(c) S^+ is a finite set of segments such that $V \setminus V^*$ lies in one connected component of $(S \setminus S^-) \cup S^+$;
(d) t^* is a finite set of segments with endpoints in V such that every point of V^* is joined to a point of $V \setminus V^*$ by a path in t^*.

Its *cost* at ρ is $\rho L(t^*) + L(S^+) - L(S^-)$.

This is weaker than [14, Definition 13], which also asks that no point of V^* is incident to $S \setminus S^-$, that $(S \setminus S^-) \cup S^+$ is a Steiner tree, and that no path in t^* joins two points of $V \setminus V^*$; Lemma 2.2 holds without these conditions. The following lemma is stated in [14, Lemma 14], where it is attributed to [3, Lemma 3]. We include a proof for the weaker definition.

Lemma 2.2. *Let $\rho > 0$ and $n \geq 2$, and suppose that $L_{\mathrm{S}}(W) \geq \rho L_{\mathrm{M}}(W)$ for every set W with $2 \leq |W| < n$. Let S be an SMT for a set V of n points. If S has a splitting of cost at most 0 at ρ, then $L_{\mathrm{S}}(V) \geq \rho L_{\mathrm{M}}(V)$.*

Proof. Let (V^*, S^-, S^+, t^*) be the splitting and $V' = V \setminus V^*$, so $1 \leq |V'| < n$. By (c), a connected component of $(S \setminus S^-) \cup S^+$ is a network containing V'; hence $L_{\mathrm{S}}(V') \leq L(S) - L(S^-) + L(S^+)$. If $|V'| \geq 2$, the hypothesis gives $\rho L_{\mathrm{M}}(V') \leq L_{\mathrm{S}}(V')$; if $|V'| = 1$, both sides are 0. The edges of a minimum spanning tree of V' together with the segments of t^* form a connected graph on V by (d), whose edges are segments between points of V; so $L_{\mathrm{M}}(V) \leq L_{\mathrm{M}}(V') + L(t^*)$. Hence

$$\rho L_{\mathrm{M}}(V) \leq L_{\mathrm{S}}(V') + \rho L(t^*) \leq L(S) - \big(L(S^-) - L(S^+) - \rho L(t^*)\big) \leq L(S) = L_{\mathrm{S}}(V). \qquad \square$$

2.3. **Reduction to full trees.**

Lemma 2.3. *Let $\rho > 0$ and $n \geq 3$, and suppose that $L_{\mathrm{S}}(W) \geq \rho L_{\mathrm{M}}(W)$ for every set W with $2 \leq |W| < n$. If a set V of n points has an SMT S in which some terminal has degree at least two, then $L_{\mathrm{S}}(V) \geq \rho L_{\mathrm{M}}(V)$.*

Proof. Let x be such a terminal, let $S_1, \dots, S_k$ ($k \geq 2$) be the closures of the components of $S \setminus \{x\}$, and let $V_i = V \cap S_i$. Each S_i contains x and a leaf of S, which is a terminal, so $|V_i| \geq 2$; and V_i misses the leaves of S in the other branches; hence $|V_i| < n$. Each S_i is an SMT for V_i, since a shorter network for V_i could replace it. The minimum spanning trees of the V_i all contain x, so their union is a connected graph on V, and $L_{\mathrm{M}}(V) \leq \sum_i L_{\mathrm{M}}(V_i) \leq \sum_i L(S_i)/\rho = L_{\mathrm{S}}(V)/\rho$. □

2.4. **Local configurations.**

Lemma 2.4. *Let S be a full SMT for a set of $n \geq 5$ terminals. Then S has Steiner points X, Y, Z and terminals A, B such that the neighbours of X are A, B, Y; the neighbours of Y are X, Z and a vertex D that is either a terminal or a Steiner point whose two other neighbours U, V are terminals; and Z has two neighbours R, Q besides Y. If $n = 5$, then D, R and Q are terminals.*

Proof. The Steiner points span a subtree with $n - 2 \geq 3$ vertices, which contains a path with two edges. A longest path in this subtree ends at two of its leaves, each of which is adjacent to two terminals; so S has a path between two leaves with at least four edges. Let $v_0, v_1, \dots, v_k$ be a longest such path, so $k \geq 4$, and put $A = v_0$, $X = v_1$, $Y = v_2$, $Z = v_3$. The third neighbour B of X is a terminal: otherwise it would have a further neighbour, and a path from a leaf behind B through $B, X, Y, \dots, v_k$ would have at least $k + 1$ edges. Let D be the third neighbour of Y. If a neighbour of D other than Y were a Steiner point, a leaf behind it would give a path with at least $k + 1$ edges ending at v_k; so D is a terminal or a Steiner point whose other two neighbours are terminals. Since $k \geq 4$, $Z = v_3$ is not an end of the path and is therefore a Steiner point; let R, Q be its other neighbours. If $n = 5$, the only Steiner points are X, Y, Z, so D, R, Q are terminals. □

Following [14], we call the configuration *d_regular* if D is a terminal and *d_steiner* otherwise. We write T_R and T_Q for the closures of the components of $S \setminus [Z, R]$ and $S \setminus [Z, Q]$ that contain R and Q; every terminal other than A, B and D (or U, V) lies in $T_R \cup T_Q$.

We identify the plane with $\mathbb{C}$ and put $\omega = e^{i\pi/3}$, so that ω and $\bar{\omega}$ are the unit vectors at angles $60°$ and $-60°$.

Lemma 2.5. *In Lemma 2.4, after possibly exchanging the names A, B, the names R, Q and the names U, V, and after a similarity of the plane, possibly reversing orientation, we have*

$$\begin{gathered} A = 0, \quad X = 1, \quad B = 1 + b\omega, \quad Y = 1 + s\bar{\omega}, \quad D = Y + d, \quad Z = Y - c\omega, \\ R = Z - e, \quad Q = Z + f\bar{\omega}, \qquad \text{and in d_steiner} \quad U = D + u\omega, \quad V = D + v\bar{\omega}, \end{gathered} \tag{2.1}$$

with positive numbers b, s, c, d, e, f (and u, v), the lengths of the edges XB, XY, YZ, YD, ZR, ZQ (and DU, DV) divided by $|AX|$. In d_steiner one may moreover assume that $\max(u, v) \leq \max(b, 1)$.

Proof. Traversing $X \to Y$, one of the edges YD, YZ turns left by $60°$ and the other turns right. Reflecting the plane if necessary, YD turns left. Exactly one of the two leaves at X is such that $A \to X \to Y$ turns right; we call it A and the other one B, which is allowed since both are terminals adjacent to X. A translation, a rotation and a scaling by $1/|AX|$ give $A = 0$ and $X = 1$. Since the edges at each Steiner point make angles of $120°$, the edge XY has direction $-60°$ and XB direction $60°$; YD has direction $0°$ and YZ direction $-120°$. At Z we call R the neighbour in direction $180°$ and Q the one in direction $-60°$, and in d_steiner we call U and V the neighbours of D in directions $60°$ and $-60°$. This gives (2.1).

In d_steiner, suppose $\max(u, v) > \max(b, 1)$. The vertex D with its leaves U, V and the vertex X with its leaves A, B are both attached to Y; thus exchanging their roles gives another labelling that satisfies Lemma 2.4 with D a Steiner point. After normalizing it as above, the

new unit length is $|DU|$ or $|DV|$, the new $\max(b,1)$ is the old $\max(u,v)$ divided by it, and the new $\max(u,v)$ is the old $\max(b,1)$ divided by it. Hence the new parameters satisfy $\max(u,v) < \max(b,1)$. □

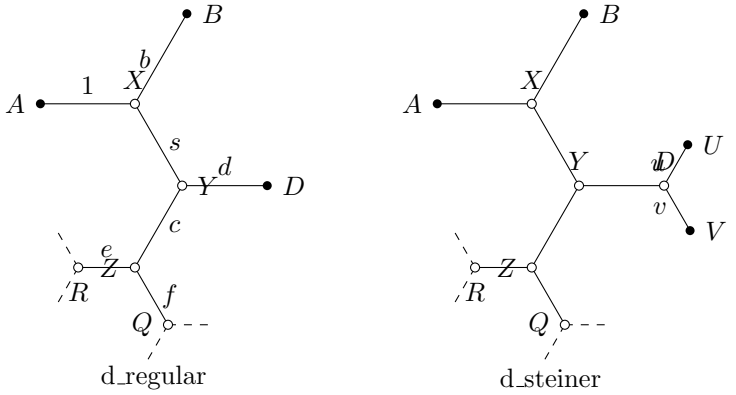

FIGURE 1. The two local configurations of Lemma 2.4 in the coordinates (2.1). Filled dots are terminals, open dots are Steiner points or, at R and Q, vertices that may be of either kind; the dashed edges continue into T_R and T_Q.

2.5. **The subcases.** Every configuration has $f \le d$ or $f \ge d$. The certificate treats four families: d_regular with $f \le d$ and with $f \ge d$, and d_steiner with $f \le d$ and with $f \ge d$. In the parameter space of a family, the coordinates are (b,c,d,s,e,f) for d_regular and (b,c,d,s,u,v,e,f) for d_steiner; for d_steiner with $f \ge d$ we use (b,c,d,s,u,v,e) and the value $f = d$, see Lemma 4.3.

2.6. **Small sets.** For $|V| = 2$ the ratio is 1. For three points with sides $a \le b' \le c'$, let γ be the angle opposite c', so $\gamma \ge 60^\circ$ and $L_{\mathrm{M}} = a + b'$. If $\gamma \ge 120^\circ$, the SMT is the minimum spanning tree. Otherwise

$$L_{\mathrm{S}}^2 = \tfrac{1}{2}\bigl(a^2 + b'^2 + c'^2\bigr) + 2\sqrt{3}\,\mathrm{area} = a^2 + b'^2 + 2ab'\sin(\gamma - 30^\circ) \ge a^2 + ab' + b'^2 \ge \tfrac{3}{4}(a+b')^2.$$

For four points the conjecture was proved by Pollak [17]; a shorter proof is in [8]. Thus $L_{\mathrm{S}}(V) \ge \frac{\sqrt{3}}{2} L_{\mathrm{M}}(V) > \frac{2141}{2500} L_{\mathrm{M}}(V)$ whenever $2 \le |V| \le 4$.

3. Hidden terminals

Throughout this section S is a full SMT in the position of Lemma 2.5, with at least five terminals. For a point p and a segment σ, $\mathrm{dist}(p,\sigma)$ is the distance between them, and $B(p,r)$ is the open disk of radius r about p.

3.1. **Chains.** Put $A_0 = A$, $A_1 = X$, $A_2 = Y$, $A_3 = Z$, $A_4 = R$. If A_k ($k \ge 4$) is a Steiner point, let A_{k+1} be its neighbour reached by turning right by 60°. The walk never goes back; hence it follows a path of the tree and ends at a leaf A_N, $N \ge 4$, which is a terminal; we call it E. We write $\sigma_k = [A_{k-1}, A_k]$ and $a_k = |\sigma_k|$, so $a_1 = 1$, $a_2 = s$, $a_3 = c$, $a_4 = e$, and σ_k has direction $\bar{\omega}^{\,k-1}$. Similarly the walk $D \to Y \to Z \to Q$ turns left by 60° at Y and at Z; continued by left turns at Steiner points, it ends at a terminal F. In d_steiner the walk $V \to D \to Y \to Z \to Q \to \cdots$ also turns left at every vertex and ends at the same F.

The terminal E lies in T_R and F lies in T_Q. These two trees are disjoint and contain none of A, B, D, U, V. In particular $E \ne F$, and both are terminals of $V \setminus \{A, B, D, U, V\}$.

3.2. **Exchange and trapping.**

Lemma 3.1. *Let p, q be points of S and let τ be a segment contained in an edge of S and in the path of S from p to q. Then $|pq| \geq |\tau|$. In particular, for $0 \leq u \leq k-2$,*

$$\operatorname{dist}(A_u, \sigma_k) \geq \max(a_{u+1}, \dots, a_{k-1}), \tag{3.1}$$

and $|A_u A_v| \geq a_t$ for $u < t \leq v$.

Proof. Removing the relative interior of τ splits S into two parts, one containing p and the other q; no terminal is removed. Adding $[p, q]$ gives a network that contains all terminals and has length $L(S) - |\tau| + |pq|$, which is at least $L(S)$. For (3.1), the path from A_u to a point of σ_k contains $\sigma_{u+1}, \dots, \sigma_{k-1}$. □

A point p is to the *right* of an oriented line ℓ through p_0 with unit direction δ if $\operatorname{Im}((p-p_0)\bar{\delta}) < 0$. Below, $\lambda_\ell(p) = \operatorname{Im}((p-p_0)\bar{\delta})$ is the signed distance to the left of ℓ.

Lemma 3.2 (trap principle)**.** *Let $1 \leq m \leq N$, and let ℓ be the line through A_{m-1} and A_m, oriented by σ_m. Let P be a simple polygon in the closed half-plane to the right of ℓ whose boundary consists of pieces of the edges $\sigma_1, \dots, \sigma_m$, a segment $J = [A_m, z_0] \subset \ell$ with z_0 beyond A_m, and further segments called doors, such that A_m lies in the relative interior of the side of P that contains J. If no point of the tail $A_m A_{m+1} \cdots A_N$ other than A_m lies on a door outside J, then the tail lies in P; hence $E \in P$.*

Proof. If $m = N$, then $E = A_m \in P$. Let $m < N$. Near A_m the polygon coincides with the half-disk to the right of ℓ, and σ_{m+1} turns right by $60°$, so the tail enters the interior of P. Suppose it leaves P, and let z be the first point of the tail after A_m that lies on ∂P. The tail meets the edges $\sigma_1, \dots, \sigma_m$ only at A_m, since S is a tree whose edges meet only at common endpoints, and it does not meet a door outside J; so $z \in J \setminus \{A_m\}$. The tail from A_m to z followed by $[z, A_m]$ is a simple closed polygon in the closed right half-plane of ℓ. Its angle at A_m is $60°$; thus it is traversed counterclockwise and its turning angles add up to $360°$. They are $120°$ at A_m, $-60°$ at each vertex of the tail before z, and $180° - \varphi$ at z for some $\varphi \in (0°, 180°)$, since the edge that arrives at z does not lie on ℓ (its earlier points would lie on ∂P). Their sum is less than $300°$, a contradiction. □

Lemma 3.3. *Let $0 \leq u < t \leq v < N$, let ℓ be the line through A_v and A_{v+1}, oriented by σ_{v+1}, and let H be the foot of the perpendicular from A_u to ℓ. Suppose that*

- (T1) *H lies on the ray from A_v through A_{v+1};*
- (T2) *$|A_u H| < a_t$;*
- (T3) *$A_u, \dots, A_v$ lie in the closed half-plane to the right of ℓ, and $A_{u+1}, \dots, A_{v-1}$ in the open one;*
- (T4) *$[A_u, H]$ meets the edges $\sigma_{u+1}, \dots, \sigma_t$ only at A_u.*

Then H lies beyond A_{v+1}, the polygon $P = A_u A_{u+1} \cdots A_v H$ is simple, and $E \in P$.

Lemma 3.3 is a form of [14, Lemma 8] with the additional hypotheses (T3) and (T4); see Remark 3.9.

Proof. The door $[A_u, H]$ lies in $B(A_u, a_t)$ by (T2), and by (3.1) the edges σ_k with $k \geq t+1$ avoid this disk. In particular $H \notin \sigma_{v+1}$, so by (T1) H lies beyond A_{v+1}. By (T3) the edges $\sigma_{u+1}, \dots, \sigma_v$ meet ℓ only at A_v, and at A_u if $A_u \in \ell$, in which case $H = A_u$. The door is perpendicular to ℓ; hence it meets $[A_v, H]$ only at H; by (T4) and the disk it meets $\sigma_{u+1}, \dots, \sigma_v$ only at A_u. Therefore P is simple, and it lies in the closed right half-plane of ℓ. Lemma 3.2 with $m = v+1$, $J = [A_{v+1}, H]$ and the door $[H, A_u]$ applies, since the tail consists of edges σ_k with $k \geq v+2 \geq t+2$. □

3.3. **Bounds for E from the published lemmas.** The following lemmas are Lemmas 27, 29, 30, 33 and 36 of [14]; the first is due to Du and Hwang [3, Lemma 2]. Let ℓ_5 be the line through $A_4 = R$ in direction $\omega^2 = e^{2\pi i/3}$, which contains σ_5 if $N \ge 5$, and let $W = 1 - c - e$ be its intersection with the real axis; then $|RW| = s + c$. The pentagon $\Pi = WA_1A_2A_3A_4$ is convex, it is traversed clockwise, and it lies to the right of ℓ_5 oriented by ω^2.

Lemma 3.4 ([3, Lemma 2], [14, Lemma 27])**.** *If $c \le 1$, then $E \in \Pi$ and $|AE| \le \max(|AY|, |AZ|)$.*

Proof. If $N = 4$, then $E = R$. Let $N \ge 5$. We first show $a_5 < s + c$, that is, A_5 lies strictly between R and W. If $c + e \le 1$, then $W \in \sigma_1$, and $a_5 \ge s + c$ would put W on σ_5, while $\operatorname{dist}(\sigma_1, \sigma_5) \ge a_2 > 0$ by Lemma 3.1. If $c + e > 1$ and $c < 1$, then $|AW| = c + e - 1 < e = a_4$, so $W \in B(A, a_4)$, which σ_5 avoids by (3.1). If $c = 1$, the foot of the perpendicular from A to ℓ_5 is at distance $1 + s - e/2$ from R, and $0 < e \le 1 + s$ because $|AR|^2 - e^2 = (1 + s)(1 + s - e)$ and $|AR| \ge e$ (Lemma 3.1); so the foot lies on $[R, W]$, at distance $\frac{\sqrt{3}}{2}e < e$ from A. By (3.1), σ_5 does not contain the foot; hence it ends before reaching it, and $a_5 < s + c$.

Now apply Lemma 3.2 with $m = 5$, $P = \Pi$ and $J = [A_5, W]$. If $c + e \le 1$, the boundary of Π consists of $[W, A_1] \subset \sigma_1$, σ_2, σ_3, σ_4 and $[A_4, W]$, and there is no door. If $c + e > 1$, the door is $[W, A]$. All its points other than W lie in $B(A, e)$, since $|AW| = c + e - 1 \le e$, and the tail consists of edges σ_k with $k \ge 6$, which avoid $B(A, a_4)$ by (3.1); the point W lies on J. Consequently $E \in \Pi$.

For the bound, $|AA_1| = 1 \le |AA_2|$. Let $x_3 = 1 + (s - c)/2 > 0$ be the real part of A_3. From $|AA_4|^2 = |AA_3|^2 - e(2x_3 - e)$ and $|AA_4| \ge e$ we get $e \le |AA_3|^2/(2x_3)$, and $4x_3^2 - |AA_3|^2 = 3(1 + s)(1 - c) \ge 0$ gives $e \le 2x_3$, hence $|AA_4| \le |AA_3|$. Finally $|AW| \le \max(1, e) \le |AA_4|$, because $|AA_4| \ge a_1$ and $|AA_4| \ge a_4$. □

Lemma 3.5 ([14, Lemmas 29 and 30])**.** *Suppose that $c + e \ge 1$ and that $e < 1 - c + \frac{2}{\sqrt{3}}s$ or $e < 1 + \big(\frac{2}{\sqrt{3}} - 1\big)c$. Let H be the foot of the perpendicular from A to ℓ_5. Then E lies in the polygon $AA_1A_2A_3A_4H$ and $|AE| \le \max\big(1, |AY|, |AZ|, |AR|, \frac{\sqrt{3}}{2}(c + e - 1)\big)$.*

Proof. If $N = 4$, then $E = R$. Otherwise we apply Lemma 3.3 with $(u, t, v) = (0, 2, 4)$ under the first condition and $(0, 3, 4)$ under the second; $\ell = \ell_5$. The signed distances to the left of ℓ_5 are

$$\lambda(A_3) = -\tfrac{\sqrt{3}}{2}e, \qquad \lambda(A_1) = \lambda(A_2) = -\tfrac{\sqrt{3}}{2}(c + e), \qquad \lambda(A) = -\tfrac{\sqrt{3}}{2}(c + e - 1),$$

so (T3) holds because $c + e \ge 1$, and $|AH| = \frac{\sqrt{3}}{2}(c + e - 1)$; (T2) is the respective condition. The point H is at signed distance $\frac{1}{2}(1 + 2s + c - e)$ from R along ℓ_5, and $1 + 2s + c - e$ exceeds $2c + (2 - \frac{2}{\sqrt{3}})s$ under the first condition and $2s + (2 - \frac{2}{\sqrt{3}})c$ under the second; this gives (T1). For (T4), $H = -(c + e - 1)\big(\frac{3}{4} + \frac{\sqrt{3}}{4}i\big)$, so the door $[A, H]$ lies in the half-plane $\operatorname{Re} \le 0$, which meets $\sigma_1 \cup \sigma_2$ only at A; and the lines of $[A, H]$ and σ_3 meet at the point $Y + (1 + 2s)\omega$, which is not on $\sigma_3 = \{Y - t\omega : 0 \le t \le c\}$. The bound is the maximum of $|A \cdot |$ over the vertices of the polygon. □

Lemma 3.6 ([14, Lemma 33])**.** *Suppose that $e < \kappa c$, where $\kappa = \frac{2}{\sqrt{3}} - 1$. Let $H = R - \frac{c-e}{2}\bar{\omega}$ be the foot of the perpendicular from Y to ℓ_5. Then E lies in the quadrilateral $YZRH$. Consequently $|pE| \le \max(|pY|, |pZ|, |pR|, |pH|)$ for every point p, in particular for $p = D$ in d_regular and $p \in \{U, V\}$ in d_steiner, and $|AE| \le \max(|AY|, |AZ|, |AR|)$.*

Proof. If $N = 4$, then $E = R$. Otherwise we apply Lemma 3.3 with $(u, t, v) = (2, 3, 4)$. The point H is at signed distance $\frac{c-e}{2} > 0$ from R along ℓ_5, which gives (T1), and $|YH| = \frac{\sqrt{3}}{2}(c + e)$; thus (T2) is the condition $e < \kappa c$. The signed distances of Y and Z to the left of ℓ_5 are $-\frac{\sqrt{3}}{2}(c + e)$

and $-\frac{\sqrt{3}}{2}e$, which gives (T3). The door $[Y,H]$ leaves Y in direction 210° and σ_3 leaves Y in direction 240°, which gives (T4). For the bound on $|AE|$, the vertex H can be dropped, because $|AH|^2-|AZ|^2=\left(\frac{3}{4}(c+e)^2-c^2\right)-\frac{c}{2}-\frac{3e}{2}-cs<0$. □

Lemma 3.7 ([14, Lemma 36])**.** *If $e<s+1$, then $E\in\Pi$ and $|AE|\le\max(|AY|,|AZ|,|AR|)$.*

Proof. If $N=4$, then $E=R$; let $N\ge5$. If $c+e\le1$, the proof of Lemma 3.4 for this case applies; it did not use $c\le1$. Let $c+e>1$, so $W=1-c-e<0$. First, $a_5<s+c$: otherwise $W\in\sigma_5$, and Lemma 3.1 with $\tau=[R,W]$ gives $c+e-1=|AW|\ge s+c$, contrary to $e<s+1$. If $N=5$ we are done. Let $N\ge6$ and let ℓ_6 be the line through A_5 and A_6, oriented by σ_6, which has direction ω. The ray from A_5 in direction ω leaves Π through $[W,A_1]$ or through σ_2, because σ_3 is parallel to it and σ_4 lies below it; let M' be the exit point and $M=1+s-e-a_5$ the intersection of ℓ_6 with the real axis; hence $|A_5M|=s+c-a_5$. The chord $[A_5,M']$ cuts from Π the convex polygon P' that contains A_2, A_3, A_4; the signed distances to the left of ℓ_6 are

$$\lambda(A_4)=-\tfrac{\sqrt{3}}{2}a_5,\quad \lambda(A_3)=\lambda(A_2)=-\tfrac{\sqrt{3}}{2}(a_5+e),\quad \lambda(A_1)=\tfrac{\sqrt{3}}{2}(M-1),\quad \lambda(A)=\tfrac{\sqrt{3}}{2}M,$$

and P' lies to the right of ℓ_6. Its boundary consists of pieces of $\sigma_1,\dots,\sigma_5$, the segment $[A_5,M']\subset\ell_6$ and, if $M<0$, the door $[M,A]$. We have $a_6<|A_5M'|$, since σ_6 does not meet σ_1 or σ_2 (Lemma 3.1) and, if $M<0$, $|AM|=a_5-(s+1-e)<a_5$, so the door lies in $B(A,a_5)$, which $\sigma_6,\sigma_7,\dots$ avoid by (3.1). Lemma 3.2 with $m=6$ and $J=[A_6,M']$ gives $E\in P'\subset\Pi$.

For the bound, $|AA_1|=1\le|AA_2|$, and $|AW|$ is dropped: if $c+e\le1$ then $|AW|\le1$, and otherwise $|AA_4|^2-(c+e-1)^2=(s+c)(s+1-e)>0$. □

3.4. **Bounds for F when $f\le d$.** The proofs of Lemmas 3.2–3.7 use only that $A_0,A_1,\dots$ are consecutive vertices of S, that A_1, A_2, A_3 are Steiner points, that the chain turns right by 60° at each Steiner point, and Lemma 3.1; they do not use that A_0 is a terminal. Hence they apply, after a similarity that maps its first edge to $[0,1]$, to every chain of S with these properties. The map $z\mapsto\overline{(D-z)}/d$ sends D,Y,Z,Q to 0, 1, $1+\frac{c}{d}\bar\omega$, $1+\frac{c}{d}\bar\omega-\frac{f}{d}\omega$, and the chain from D to F to a chain that turns right at every vertex, with edge lengths $1,\frac{c}{d},\frac{f}{d},\dots$. In d_steiner the map $z\mapsto\bar\omega\,\overline{(V-z)}/v$ sends the chain $V,D,Y,Z,Q,\dots$ to a right-turning chain with edge lengths $1,\frac{d}{v},\frac{c}{v},\frac{f}{v},\dots$. Both maps are similarities; hence they map S to an SMT.

Lemma 3.8 ([14, Lemma 28 and Corollary 26])**.** *Suppose $f\le d$. In d_regular, $|DF|\le\max(|DZ|,|DQ|)$. In d_steiner, $|VF|\le\max(|VY|,|VZ|,|VQ|)$.*

Proof. If Q is a terminal, then $F=Q$. Otherwise the first three vertices after the start of the image chain are Steiner points: the images of Y, Z, Q in d_regular and of D, Y, Z in d_steiner. In d_regular, the hypothesis $c\le1$ of Lemma 3.4 for the image chain reads $f/d\le1$, and its bound, mapped back, is $\max(|DZ|,|DQ|)$. In d_steiner, the hypothesis $e<s+1$ of Lemma 3.7 for the image chain reads $f/v<d/v+1$, which holds since $f\le d$, and its bound, mapped back, is $\max(|VY|,|VZ|,|VQ|)$. □

Remark 3.9. (a) [14, Lemma 8] states, for $0\le u<t\le v<\min(6,N)$, that if the foot H of the perpendicular from A_u to the line A_vA_{v+1} lies on the ray from A_v through A_{v+1} and $|A_uH|<a_t$, then E lies in the polygon $A_u\cdots A_vH$. Its proof does not show that $A_u\cdots A_vH$ is simple, and it excludes exits through the sides $[A_{v+1},H]$ and $[H,A_u]$, which are not part of the chain, by appeal to the fact that the chain does not cross itself. For chains that satisfy every exchange inequality of Lemma 3.1 between their points, the polygon need not be simple: for $(u,t,v)=(0,1,4)$ and edge lengths $(a_1,\dots,a_6)=(1,2,\frac{1}{5},\frac{1}{10},\frac{1}{20},\frac{1}{40})$, the hypotheses hold and the side $[A_4,H]$ crosses σ_1. We do not know whether this happens in an SMT. We use Lemma 3.3 instead and check its hypotheses in each instance.

(b) In the proof of [14, Lemma 36], in the notation there, $N = A_5$ and M is the point where the ray from A_5 in the direction of σ_6 meets $[A, W]$, and the inequality $|AM| < |NM|$ is derived from $e < s + 1$. This implication is not valid: for $(s, c, e, a_5) = (1, \frac{3}{2}, \frac{3}{2}, \frac{17}{10})$, a chain that satisfies every exchange inequality of Lemma 3.1 up to A_5, we get $|AM| = \frac{6}{5} > \frac{4}{5} = |NM|$. The case in question does not occur, because $|AM| < a_5$ (proof of Lemma 3.7).

3.5. **Two new bounds.**

Lemma 3.10 (removal exchange)**.** *Let z_0 be a Steiner point of S with neighbours x, y, y', and let S_x, S_y, $S_{y'}$ be the closures of the components of S minus the three edges at z_0 that contain x, y, y'. Then $|x'K| \geq |z_0x| + |z_0y| + |z_0y'| - |yy'|$ for every point $x' \in S_x$ and every point $K \in S_y \cup S_{y'}$.*

Proof. Removing the three edges at z_0 and adding the segments $[y, y']$ and $[x', K]$ gives a network that contains all terminals, since z_0 is not a terminal. Its length is $L(S) - |z_0x| - |z_0y| - |z_0y'| + |yy'| + |x'K| \geq L(S)$. □

Lemma 3.11. *Let $\Phi = c + e + f - |RQ| - \frac{\sqrt{3}}{2}(c + e)$, where $|RQ| = \sqrt{e^2 + ef + f^2}$. If $e < c$ and $\Phi > 0$, the conclusions of Lemma 3.6 hold: E lies in the quadrilateral $YZRH$, with H as there.*

Proof. If $N = 4$, then $E = R$; let $N \geq 5$. Lemma 3.10 with $z_0 = Z$, $x = x' = Y$, $\{y, y'\} = \{R, Q\}$ gives $|YK| \geq c+e+f-|RQ|$ for every point K of T_R, and T_R contains R, σ_5 and the whole tail. As in the proof of Lemma 3.6, H is at signed distance $\frac{c-e}{2} > 0$ from R along ℓ_5, and $|YH| = \frac{\sqrt{3}}{2}(c+e)$. Since $\Phi > 0$, every point of the door $[Y, H]$ is at distance less than $c + e + f - |RQ|$ from Y, so no point of T_R lies on the door; in particular A_5 lies strictly between R and H. The door leaves Y in direction 210° and σ_3 in direction 240°, and every point of the door lies above the line ZR, at height at least $\frac{\sqrt{3}}{4}(c - e) > 0$; so $YZRH$ is a convex quadrilateral, and it lies to the right of ℓ_5. Lemma 3.2 with $m = 5$, $J = [A_5, H]$ and the door $[H, Y]$ gives $E \in YZRH$. The bounds follow as in Lemma 3.6; $|AH| < |AR|$ because $|AR|^2 - |AH|^2 = \frac{1}{4}(c - e)(2 + 4s + c - e)$. □

If $e < \kappa c$, then $e < c$ and $\Phi \geq c - \frac{\sqrt{3}}{2}(c + e) > 0$, since $|RQ| \leq e + f$. Thus Lemma 3.11 contains Lemma 3.6, and its condition defines a convex set, because Φ is concave. Moreover $\partial\Phi/\partial f = 1 - (\frac{e}{2} + f)/|RQ| \geq 0$.

Lemma 3.12. *Let the configuration be d_regular. Put $Q^* = Q + (2d + c - f)$ and*

$$\Psi = 1 + b + s + 3d + 2c + e - L_{\mathrm{S}}(A, B, D) - L_{\mathrm{S}}(D, Q, Q^*) - |QR|,$$

where $L_{\mathrm{S}}(p, q, r)$ is the length of an SMT for $\{p, q, r\}$. If $\Psi > 0$, then $|DF| \leq \max(|DZ|, |DQ|)$.

Proof. If Q is a terminal, then $F = Q$. Otherwise the walk continues from Q to $Q_2 = Q + g$, $g > 0$. Let S' be the union of the edges AX, BX, XY, YD, YZ, ZR, ZQ, QQ_2, of length $1+b+s+d+c+e+f+g$. Since A, B, D are leaves and X, Y, Z have all their edges in S', the rest of S consists of trees attached to S' only at R, Q and Q_2. Hence S' is a shortest network containing $\{A, B, D, R, Q, Q_2\}$: replacing it by a shorter one would give a shorter network containing all terminals.

Put $x_Q = \mathrm{Re}(D-Q) = d+\frac{c-f}{2}$, so $Q^* = Q+2x_Q$, and suppose $g \geq 2x_Q$. Both Q_2 and Q^* lie on the horizontal line through Q, so $|Q^*Q_2| = g - 2x_Q$ and $L_{\mathrm{S}}(D, Q, Q_2) \leq L_{\mathrm{S}}(D, Q, Q^*) + g - 2x_Q$. The network formed by SMTs for $\{A, B, D\}$ and $\{D, Q, Q_2\}$ and the segment $[Q, R]$ contains the six points, and its length is at most

$$L_{\mathrm{S}}(A, B, D) + L_{\mathrm{S}}(D, Q, Q^*) + g - 2d - c + f + |QR| = L(S') - \Psi < L(S'),$$

a contradiction. Therefore $g < 2x_Q$; if $2x_Q \leq 0$, this shows that Q is a terminal.

If $f \le d$, the bound is Lemma 3.8. Let $f > d$. Then $g < 2x_Q < c + d$. For the image of the chain $D, Y, Z, Q, Q_2, \dots$ under $z \mapsto \overline{(D-z)}/d$, with edge lengths $1, \frac{c}{d}, \frac{f}{d}, \frac{g}{d}, \dots$, the hypothesis of Lemma 3.7 reads $g < d+c$, and its bound, mapped back, gives $|DF| \le \max(|DZ|, |DQ|, |DQ_2|)$. Finally $|DQ_2|^2 = |DQ|^2 - g(2x_Q - g) < |DQ|^2$, since $0 < g < 2x_Q$. □

The function Ψ is concave in the parameters: the points are affine functions of them, and the length of an SMT for three points, $\min_x(|xp| + |xq| + |xr|)$, is a convex function of (p, q, r). Consequently $\Psi > 0$ at the vertices of a bounded box implies $\Psi > 0$ on the box. The condition fails for large f; Lemma 3.12 is used only on bounded boxes.

4. Verification functions

4.1. **Splitting tables.** The local configuration is the tree $\mathcal{T}$ with the vertices A, B, X, Y, D, Z, R, Q (and U, V) and the edges of Figure 1. The tables of [14] list 158 splittings for d_regular and 1418 for d_steiner. Each consists of a set V^* of terminals among A, B, D (or A, B, U, V), a set S^- of edges of $\mathcal{T}$, a set of vertices of $\mathcal{T}$ outside V^* that the new network S^+ joins, and a set t^* of segments between terminals among A, B, D, U, V, E, F. We have checked by a program that each of them has the following properties, which make it a splitting in the sense of Definition 2.1 for every configuration: V^* is nonempty and contains only terminals of $\mathcal{T}$; after removing S^-, no point of V^* keeps an edge; S^+ avoids V^*; the terminals outside V^*, together with R and Q, which stand for $T_R \ni E$ and $T_Q \ni F$, are connected in $(\mathcal{T} \setminus S^-) \cup S^+$; and t^* is a forest on terminals which, after contracting $V \setminus V^*$ to one vertex, is a tree.

4.2. **The bound for a splitting.** Each record of the certificate names a box, a splitting τ and one of the nine lemma sets of [14], numbered $0, \dots, 8$. A lemma set supplies bounds for the edges of t^* at E and F and, for four-point reconnections, a pairing of the four points. For a parameter w we put

$$\tilde{F}_\tau(w) = \rho \sum_{\epsilon \in t^*} \ell_\epsilon(w) + N_\tau(w) - \sum_{\epsilon \in S^-} |\epsilon|,$$

where $|\epsilon|$ is 1 for the edge AX and the corresponding coordinate otherwise, and:

- $\ell_\epsilon(w)$ is the distance of the endpoints if both are among A, B, D, U, V. For an edge at E or F it is the bound of the lemma that applies: for AE, Lemma 3.4, 3.5, 3.11 or 3.7 according to the lemma set (0; 1 or 2; 5; 8); for DE, UE, VE (lemma set 5), Lemma 3.11; for DF and VF, Lemma 3.8 in the families with $f \le d$, and Lemma 3.12 in d_regular with $f \ge d$. Each bound is a maximum of distances from the terminal to points that are affine functions of w (for lemma set 8 the checker also includes the linear term $e + c - 1$ of [14, Lemma 36], which only increases it), and comes with the condition of its lemma.
- $N_\tau(w)$ is 0 if S^+ joins at most one point, the distance if it joins two, and the length of an SMT for the three points if it joins three. If it joins four points p_1, p_2, p_3, p_4, with the pairing $\{p_1, p_2\}, \{p_3, p_4\}$ given by the lemma set, then $N_\tau(w) = \min_{x,y \in \mathbb{C}} \big(|p_1 x| + |p_2 x| + |xy| + |p_3 y| + |p_4 y|\big)$.

Among the sets of four or more points joined in the tables, only $\{A, D, Q, R\}$ and $\{B, D, Q, R\}$ receive a pairing from a lemma set: $\{A, R\}, \{D, Q\}$ for lemma set 3, $\{A, D\}, \{Q, R\}$ for lemma sets 4 and 7, and $\{B, D\}, \{Q, R\}$ for lemma set 6. The checker rejects records whose splitting joins any other set of four or more points. In [14] these lemma sets carry conditions under which $N_\tau(w)$ is the length of a full Steiner tree with that pairing, given by a closed formula [6]; we do not use these conditions or formulas. Table 1 lists the correspondence.

Lemma 4.1. *Let B be a bounded box in the parameter space of a family, and let a record name B, a splitting τ and a lemma set. Suppose that at every vertex of B the conditions of the bounds*

lemma set	[14]	here	condition	used for
0	Lemmas 27, 28	Lemmas 3.4, 3.8	$c \le 1$; $f \le d$	AE; DF, VF
1, 2	Lemmas 29, 30	Lemma 3.5	linear	AE
5	Lemma 33	Lemma 3.11	$e < c$, $\Phi > 0$	AE, DE, UE, VE
8	Lemma 36	Lemma 3.7	$e < s + 1$	AE
3, 4, 6, 7	Lemmas 31, 32, 34, 35	pairings only	none	four-point S^+
all	—	Lemma 3.12	$\Psi > 0$	DF (d_regular, $f \ge d$)

TABLE 1. Lemma sets and the bounds they supply. In the families with $f \le d$, every lemma set supplies the bound of Lemma 3.8 for DF or VF; in d_steiner with $f \ge d$, no edge at F is used.

used by τ, other than the condition $f \le d$ of the family, hold and $\tilde{F}_\tau \le 0$. Then the conditions hold and $\tilde{F}_\tau \le 0$ on all of B, and every configuration of the family whose parameter w lies in B has a splitting of cost at most $\tilde{F}_\tau(w) \le 0$ at ρ.

Proof. The conditions are linear inequalities, or $\Phi > 0$ together with $e < c$, or $\Psi > 0$; each defines a convex set; thus they hold on all of B. Each term of $\tilde{F}_\tau$ is convex in w: distances between affine functions of w and their maxima are convex, the SMT length for three points and N_τ are partial minima of convex functions of (w, x) and (w, x, y), and the last sum is linear. Hence $\tilde{F}_\tau$ is at most its maximum over the vertices of B, which is at most 0. For a configuration with parameter $w \in B$, the bounds of Section 3 apply to the edges of t^* at E and F; in the families with $f \le d$ we use that $f \le d$ for the configuration, and in d_regular with $f \ge d$ Lemma 3.12 requires no condition on f. Take for S^+ an SMT of its points, or for four points a network attaining $N_\tau(w)$. Then the splitting τ has cost at most $\tilde{F}_\tau(w)$. □

4.3. **Unbounded boxes.** Boxes of the form $\prod_i [\mathrm{lo}_i, \mathrm{hi}_i]$ with some $\mathrm{hi}_i = \infty$ occur; their lower ends are finite. For a coordinate t and a splitting τ, let ϵ_t be the edge of $\mathcal{T}$ whose length is t. Increasing t with the other coordinates fixed translates the part of the configuration beyond ϵ_t. Let n_T be the number of edges of t^* whose term depends on t: an edge between two of A, B, D, U, V depends on t if ϵ_t separates its endpoints, and a bound from Lemma 3.4 or 3.8 depends on t if ϵ_t separates the terminal from one of the points in its maximum. Let $d_+ = 1$ if ϵ_t separates two points joined by S^+, and $d_- = 1$ if $\epsilon_t \in S^-$, and $d_+ = 0$, $d_- = 0$ otherwise. We call t *monotone* for τ if $n_T + d_+ = d_- = 0$, or if $n_T + d_+ \le 1$ and $d_- = 1$. This is the rule of [14, Algorithm 2 and Proposition 37], with the dependence of the bounds on t determined from the points they use.

Lemma 4.2. *Let B be a box whose lower ends are finite, and let B_0 be its face on which every unbounded coordinate equals its lower end. Suppose that every unbounded coordinate t of B is monotone for τ, that the lemma set is* 0*, so that the edges of t^* at E and F use only Lemmas 3.4 and 3.8 and S^+ joins at most three points, and that c is bounded on B if t^* contains AE. If the hypotheses of Lemma 4.1 hold at the vertices of B_0, then its conclusion holds on B.*

Proof. The condition $c \le 1$ of Lemma 3.4 holds at the vertices of B_0, which include $c = \mathrm{hi}_c$, and it does not involve the unbounded coordinates; the condition $f \le d$ of Lemma 3.8 is that of the family. Thus the bounds apply on B. Along a coordinate t, each distance term of $\tilde{F}_\tau$ changes at rate at most 1 in absolute value, and not at all if ϵ_t does not separate its two points; the same holds for a maximum of such terms. When t increases by h, some of the points joined by S^+, which are at most three, are translated by a vector of length h and the others stay fixed. The moved set or its complement has at most one element, so, up to a translation of all the points, a single point moves by h; the length of an SMT of the points then changes by at most h, since

a segment can be added. If all or none of the points move, the length does not change. The last sum changes at rate d_-. Therefore $\tilde{F}_\tau$ is non-increasing in t, because $\rho n_T + d_+ - d_- \le 0$ for $\rho < 1$. Hence $\tilde{F}_\tau(w) \le \tilde{F}_\tau(w_0)$, where $w_0 \in B_0$ has the unbounded coordinates of w replaced by their lower ends, and $\tilde{F}_\tau(w_0) \le 0$ by Lemma 4.1 applied to B_0. The rest of the proof of Lemma 4.1 applies unchanged. □

4.4. **The families with $f \ge d$.**

Lemma 4.3. *Let τ be a splitting without edges at F such that, if Q is one of the points joined by S^+, the edge ZQ is in S^-. Let $f \ge f_0 \ge 0$. If the conditions of the bounds used by τ hold at the parameter (w', f_0), they hold at (w', f), and the cost of τ for the configuration with parameter (w', f) is at most $\tilde{F}_\tau(w', f_0)$.*

Proof. Write $Q_0 = Z + f_0\bar{\omega}$ for the position of Q at (w', f_0). The terms of t^* do not involve f: the bounds for edges at E use only A, D, U, V, Y, Z, R, H, and of their conditions only $\Phi > 0$ involves f, and Φ is non-decreasing in f. The sum over S^- grows by $f - f_0$ if $ZQ \in S^-$ and is constant otherwise. If Q is joined by S^+, the network for (w', f_0) together with the segment $[Q_0, Q]$, which is part of the removed edge ZQ, joins the same points with Q in place of Q_0, and its length is larger by $f - f_0$. Otherwise S^+ does not depend on f. □

In d_steiner with $f \ge d$, the certificate uses only splittings as in Lemma 4.3 and checks $\tilde{F}_\tau(w', d)$ on boxes in the coordinates $w' = (b, c, d, s, u, v, e)$; by Lemma 4.3 with $f_0 = d$, this covers every configuration with $f \ge d$. This reduction of the case $f \ge d$ to $f = d$ is that of [14, Appendix H.5]. Then Q depends on d through $f = d$, and an unbounded d is accepted only if both d and f are monotone for τ and the lemma set is 0. The bounds for the rates of change along d and along f add up, by the triangle inequality, so the proof of Lemma 4.2 applies along the direction in which d and f increase together. In d_regular with $f \ge d$, f is a coordinate of the boxes, which also contain points with $f < d$; Lemma 3.12 is used only on bounded boxes, and a box with $\mathrm{hi}_f = \infty$ requires a splitting as in Lemma 4.3. For such a box, Lemma 4.3 with $f_0 = \mathrm{lo}_f$ reduces it to its face $f = \mathrm{lo}_f$, and if other coordinates are unbounded too, Lemma 4.2 is then applied to this face.

5. The certificate

5.1. **Families and records.** For each family, a binary tree of boxes is built. Its roots are the 2^n boxes in which each of the n coordinates ranges over $[0, 1]$ or $[1, \infty)$. A box is split into two along one coordinate, at the midpoint of a bounded side or at $2\,\mathrm{lo}$ for a side $[\mathrm{lo}, \infty)$. Each leaf is a record: the box, a splitting and a lemma set. Table 2 lists the families. A record whose box lies outside its family's domain is skipped: in the families with $f \le d$ if $\mathrm{lo}_f > \mathrm{hi}_d$, in d_regular with $f \ge d$ if $\mathrm{hi}_f < \mathrm{lo}_d$, and in d_steiner if $\max(\mathrm{lo}_u, \mathrm{lo}_v) > \max(\mathrm{hi}_b, 1)$, so that the box contains no point with $\max(u, v) \le \max(b, 1)$ (Lemma 2.5).

5.2. **Generation.** The trees were generated by a program derived from the generator of [14]. It processes the boxes depth first and, for each box, tries the splittings of the table with the lemma sets until it finds one whose bound is below -10^{-6} at every vertex in floating point; otherwise it splits the box along its widest side, where a side $[\mathrm{lo}, \infty)$ counts as having width $1/\mathrm{lo}$. We changed the order in which splittings are tried, the distribution of the work over threads and the size of the record identifiers, and we added Lemmas 3.11 and 3.12, the explicit four-point networks of Section 5.3, and the treatment of $f \ge d$ described in Section 4. The generator is not part of the proof: its output is checked by the two programs below, which share no code with it.

family	coordinates	n	records	checked	skipped	depth
d_regular, $f \le d$	b, c, d, s, e, f	6	95,569,010	63,854,665	31,714,345	51
d_regular, $f \ge d$	b, c, d, s, e, f	6	10,376,823	9,426,537	950,286	56
d_steiner, $f \le d$	b, c, d, s, u, v, e, f	8	436,460,311	341,304,562	95,155,749	61
d_steiner, $f \ge d$	b, c, d, s, u, v, e	7	12,696,761	12,690,883	5,878	75
total			555,102,905	427,276,647	127,826,258	

TABLE 2. The four families. "Checked" records were verified; "skipped" records lie outside the domain of their family (292,651 + 5,878 d_steiner records by the condition on u, v, the others by the condition on f); "depth" is the largest depth of a leaf in the tree.

5.3. **Checking.** The checker was written from the statements of [14] and from reading its verifier, whose record format it uses. It works in ball arithmetic with 128-bit precision (the Arb part of FLINT [9, 13]), with ρ computed as a ball from the exact rational 2141/2500, and it accepts a record only if the following hold.

(1) The splitting and lemma set exist and every lower end of the box is finite. If the box is unbounded, the lemma set is 0 and every unbounded coordinate is monotone for the splitting, as in Lemma 4.2; the one exception is an unbounded f in d_regular with $f \ge d$, which requires a splitting as in Lemma 4.3. If t^* contains AE, then c is bounded. A splitting with an edge at F is rejected on unbounded boxes in d_regular with $f \ge d$ and on all boxes in d_steiner with $f \ge d$; in the latter family f must be monotone for the splitting, which for a splitting without edges at F is the hypothesis of Lemma 4.3.

(2) At every vertex with finite coordinates, and there is at least one, the conditions of the bounds used hold for all points of the balls, and the computed upper bound for $\tilde{F}_\tau$ is at most 0. In d_steiner with $f \ge d$ the vertex is evaluated with $f = d$.

The checker computes upper bounds: distances and their maxima in ball arithmetic; the SMT length for three points by the formula $\left(\frac{1}{2}(|pq|^2 + |qr|^2 + |rp|^2) + 2\sqrt{3}\,\mathrm{area}\right)^{1/2}$ when every angle of the triangle is certainly at most 120°, and otherwise by the sum of the two sides at an angle that is not certainly at most 120°; and for four points, with the pairing $\{p_1, p_2\}, \{p_3, p_4\}$ of the lemma set, the length of the network p_1x, p_2x, xy, p_3y, p_4y, where x and y are double-precision points obtained by the Weiszfeld iteration [21]. Any x, y give an upper bound for N_τ. The lower bound for Ψ uses these upper bounds for the two SMT lengths.

A second program checks that the records of each family form the leaves of a tree as described: it rebuilds every box from the roots, recomputes the coordinate along which each internal box is split by the rule of the generator, and checks that every leaf is a record whose box equals the computed box, that every record is reached exactly once, and that every internal node has two distinct new children.

5.4. **Results.** All records of the four families pass the checker, with the numbers of Table 2, and the second program reports that the boxes of each family tile $[0, \infty)^n$.

Both programs also fail where they must: at slightly larger values of ρ the checker rejects many records of each family, it rejects altered records and records of small test certificates that violate conditions of the bounds, and removing one record makes the second program fail. The repository (see Data availability) describes these controls.

5.5. **Proof of Theorem 1.1.** Let $\rho = 2141/2500$. We prove $L_{\mathrm{S}}(V) \ge \rho L_{\mathrm{M}}(V)$ by induction on $n = |V|$. For $n \le 4$ this was shown in Section 2. Let $n \ge 5$, assume the inequality for all sets with

fewer points, and let S be an SMT for V. If some terminal has degree at least two, Lemma 2.3 applies. Otherwise S is full, and Lemmas 2.4 and 2.5 give a configuration with a parameter in the domain of one of the four families, with $\max(u, v) \le \max(b, 1)$ in d_steiner; in d_steiner with $f \ge d$ we take the parameter w' without f. The boxes of that family tile $[0, \infty)^n$; hence some record has a box containing the parameter, and this box is not skipped, since skipped boxes are disjoint from the domain. Since the record passed the checker, Lemmas 4.1, 4.2 and 4.3 show that the configuration has a splitting of cost at most 0 at ρ; by Section 4.1 it is a splitting of S. Lemma 2.2 gives $L_{\mathrm{S}}(V) \ge \rho L_{\mathrm{M}}(V)$. □

6. Limits of the published lemmas

In this section only the bounds of [14] are used, in the form stated there: for $|AE|$, those of Lemmas 3.4, 3.5, 3.6 and 3.7 (lemma sets 0, 1, 2, 5, 8), with the condition $e < \kappa c$ in Lemma 3.6; for $|DE|$, that of Lemma 3.6 (lemma set 5); for F, Lemma 3.8, which requires $f \le d$; and for four-point reconnections the closed formulas of [14] under their conditions. For a splitting τ and a lemma set, the *cost bound* at a parameter w is $\rho \sum_{\epsilon \in t^*} \ell_\epsilon(w) + L^+(w) - \sum_{\epsilon \in S^-} |\epsilon|$, with ℓ_ϵ as in Section 4 and L^+ the length of an SMT of the points joined by S^+, or the closed formula for four points. It is affine in ρ, and its *critical value* is the value of ρ at which it vanishes.

Proposition 6.1. *Consider the d_regular configuration with*

$$w_* = (b, c, d, s, e) = (1.3656,\ 1.10927,\ 0.365604,\ 1,\ 0.171605)$$

and $f = d$.

(a) *For each of the* 158 *splittings of the d_regular table of* [14] *that has no edge at F, and each lemma set whose conditions hold at w_*, the cost bound is positive at $\rho = 0.856$. The largest critical value is* $0.8559867412\ldots$*; it is attained by the splitting with $V^* = \{A, B\}$, S^- the seven edges of $\mathcal{T}$, S^+ joining D, Q, R, and $t^* = \{AB, BD\}$.*

(b) *The same holds for every combination of a lemma set, a nonempty $V^* \subseteq \{A, B, D\}$, a set S^- of edges of $\mathcal{T}$ containing those at V^*, and a set of at most four vertices of $\mathcal{T}$ outside V^* that, joined by S^+, connects the terminals outside V^* together with R and Q, when t^* is a shortest forest joining every point of V^* to the other terminals among A, B, D, E, with edges at E priced by the bound of the lemma set, and the length of S^+ is replaced by the length of an SMT of its points; there are* 42,303 *such combinations. Their critical values are at most* 0.8559867413*.*

At $w_1 = (1.34514,\ 1.07421,\ 0.347656,\ 1.00391,\ 0.167969)$, with $f = d$, both statements hold with 0.8572 *in place of* 0.856*; the largest critical value in* (a) *is* $0.8571349212\ldots$*, and the bound in* (b) *is* 0.8571349213*.*

Proof. The statements are finite computations at points with rational coordinates, done by three programs. The first evaluates in ball arithmetic, from a transcription of the verifier of [14], all $158 \cdot 9$ pairs of a splitting and a lemma set; at both points the conditions of 527 pairs hold, and each of these has positive cost bound at 0.856 (at w_*) or 0.8572 (at w_1). The result is the same if the conditions of lemma sets 3 and 7, which are not linear, are ignored. The second evaluates the formulas generated by the code of [14] in floating point and finds the same largest critical values. The third enumerates the 42,303 combinations of (b) in ball arithmetic with lower bounds for $L(S^+)$: the distance, the length of an SMT for three points, and, for four points $p_1, \ldots, p_4$, the minimum over the three pairings of a lower bound for $\min_{x,y}(|xp_i| + |xp_j| + |xy| + |yp_k| + |yp_l|)$ obtained from vectors $u_1, \ldots, u_5$ of norm at most 1 with $u_1 + u_2 + u_3 = 0$ and $u_4 + u_5 = u_3$, since then the sum is at least $-\langle u_1, p_i\rangle - \langle u_2, p_j\rangle - \langle u_4, p_k\rangle - \langle u_5, p_l\rangle$; the program also computes, for each pairing, the bound $|U_{ij}U_{kl}|$ given by points U_{ij}, U_{kl} that form equilateral triangles with p_i, p_j and with p_k, p_l, and the largest bound for three of the points, and keeps the largest of

these valid bounds. The shortest t^* is computed as a minimum spanning tree on lower bounds of the edge lengths. Every tree that joins four points with at most two Steiner points is a network of one of the three pairings, possibly with coinciding vertices, so the bound applies to every S^+. □

In (b), S^+ is one network, as in the tables of [14]. Joining at most four vertices of $\mathcal{T}$ is no restriction for such a network: after removing S^-, the terminals outside V^* and the trees T_R, T_Q lie in at most four components, one vertex from each suffices, and the length of an SMT does not decrease when points are added.

The cost bounds in (b) are continuous in f, and there are finitely many of them; hence (b) also holds at (w_*, f) for $f \in (d, d+\delta]$ and some $\delta > 0$; for these f no bound for F from [14] applies. Let $\rho \geq 0.856$, and consider a certificate that uses only the lemmas of [14] and splittings of the kind in Proposition 6.1(b), and a record of it that is valid on its whole box, where the box contains (w_*, f) for such an f. Since $f > d$ there, the splitting has no edge at F. The conditions of the bounds for edges at E (lemma sets 0, 1, 2, 5, 8) are linear, so if they hold at the vertices of the box, they hold at (w_*, f). The function that the record checks is $\rho \sum_{t^*} \ell_\epsilon + N - \sum_{S^-} |\epsilon|$, where N is the length of an SMT of the points joined by S^+ if there are at most three, and otherwise N_τ for the pairing of the lemma set, or any other convex function of w that is at least the length of an SMT of the points. It is convex in w (proof of Lemma 4.1), and at (w_*, f) it is at least the cost bound of (b), which is positive. Consequently it is positive at a vertex of the box, and the record fails. In the formulation of [14], which checks this family on boxes in (b, c, d, s, e) with $f = d$ and uses the closed formulas for four points, part (a) gives the same conclusion at w_*, for lemma sets 3 and 7 even without their conditions, which are not linear; the closed formulas are convex. Hence, with the lemmas of [14] alone and with any partition, neither kind of certificate reaches $\rho = 0.856$, and a complete certificate of either kind for any value at least 0.856 requires further lemmas. Lemmas 3.11 and 3.12 both apply at $(w_*, f = d)$, with $\Phi > 0.06$ and $\Psi > 0.26$.

Remark 6.2. The partitions of [19] were generated with the generator, the splitting tables and the lemmas of [14], which at the commit pinned in [19] are identical to those at the commit we use, and [19, Section 7] takes their coverage of each case domain from the recursive subdivision of the generator. In the family d_regular with $f \geq d$, which [14] checks on boxes in (b, c, d, s, e) with $f = d$, the argument above shows that none of the records of [19] that pass at the values 0.856 and 0.860 reported there can have a box containing w_*. Consequently, with or without the proofs of the lemmas of [14] given here, these partitions are not complete certificates at these values.

7. Formal proofs

Both formal proofs of Theorem 1.1, in Lean 4 [15] with Mathlib [16], take the statements (H1)–(H3) below as hypotheses. In the formal statements a point of the plane is a complex number. A finite set of segments connects a set V if any two points of V are joined by a path in the graph whose edges are these segments, and $L_{\mathrm{M}}(V)$ and $L_{\mathrm{S}}(V)$ are the infima of the total lengths of the finite sets of segments that connect V, with endpoints in V for L_{M} and with arbitrary endpoints for L_{S}. These are the lengths of Section 1: the edges of a minimum spanning tree or of an SMT form such a set, and the segments of such a set form a network that contains V. The configurations have the coordinates (2.1), and the splittings are the entries of the tables of [14] in the numeric form used by our checker, which a program of the repository produces from the tables.

The theorem `S43.Final.steiner_ratio_ge` states that $\frac{2141}{2500} L_{\mathrm{M}}(V) \leq L_{\mathrm{S}}(V)$ for every finite set V of points, for every relation $\mathcal{R}$ that satisfies the three hypotheses below. For Theorem 1.1,

$\mathcal{R}(w, E, F)$ means that some full SMT with at least five terminals, placed as in Lemma 2.5, contains a configuration of one of the two types with parameter w whose hidden terminals (Section 3.1) are E and F; in d_regular the unused coordinates u and v are 0. Let $\rho = 2141/2500$.

(H1) If $\mathcal{R}(w, E, F)$, then each of the following bounds of Section 4.2 holds whenever its condition holds at w: the bounds of the lemma sets 0, 1, 2 and 8 for $|AE|$; the bounds of lemma set 5 for $|AE|$, for $|DE|$ in d_regular and for $|UE|$ and $|VE|$ in d_steiner, under the condition of Lemma 3.6 or that of Lemma 3.11; $|DF| \le \max(|DZ|, |DQ|)$ in d_regular, under $f \le d$ or $\Psi > 0$; and $|VF| \le \max(|VY|, |VZ|, |VQ|)$ in d_steiner, under $f \le d$.

(H2) If $\mathcal{R}(w, E, F)$, then w lies in the domain of its configuration type: all coordinates are nonnegative, $u = v = 0$ in d_regular, and $\max(u, v) \le \max(b, 1)$ in d_steiner.

(H3) If for every (w, E, F) with $\mathcal{R}(w, E, F)$ some entry of the table of its configuration type satisfies $\rho L(t^*) + L_{\mathrm{S}}(S^+) - L(S^-) \le 0$, where $L_{\mathrm{S}}(S^+)$ is the length of an SMT for the points that S^+ joins and every edge of t^* at E or F is one of the edges in Table 1 and is counted with its true length, then $\rho L_{\mathrm{M}}(V) \le L_{\mathrm{S}}(V)$ for every finite set V.

In the Lean files these hypotheses are called `HTrap`, `HDomain` and `HDuHwang`. For the relation above, (H1) holds by Lemmas 3.4–3.8, 3.11 and 3.12, and (H2) by Lemma 2.5. Hypothesis (H3) rests on the splitting lemma of Du and Hwang (Lemma 2.2), the reduction to full trees and to the two configurations with their coordinates (Lemmas 2.3–2.5), the definition of the hidden terminals (Section 3.1), the case of sets of at most four points (Section 2.6, which uses Pollak's theorem), the check of Section 4.1 that every entry of the tables is a splitting together with the transcription of the tables into their numeric form, and the induction in the proof of Theorem 1.1. In the formal proof all of these are hypotheses; here they are proved on paper, except the check and the transcription of the tables, which are done by programs.

The theorem `S43.Final.certificate` uses none of (H1)–(H3). For each of the two configuration types and each point w of its domain, it gives an entry of the table and a lemma set whose conditions for the edges of t^* at E and F hold at w and for which $\rho L(t^*) + L_{\mathrm{S}}(S^+) - L(S^-) \le 0$, where these edges are counted with the bounds of the lemma set. Together with (H1)–(H3) it implies `steiner_ratio_ge`.

The certificate theorem is proved by running a checking program written in Lean on the certificate, as compiled code in the first proof and in Lean's kernel in the second; Appendix A describes the two programs and their trusted bases.

Concluding remarks

Section 6 shows how the estimates for the positions of the hidden terminals limit the method; with Lemmas 3.11 and 3.12 it reaches $2141/2500$. We do not know the largest value that these lemmas allow, nor whether estimates of this kind can bring the method close to $\sqrt{3}/2$.

Appendix A. The checking programs of the formal proofs

A.1. **Checks run as compiled code.** The checking program of the first proof follows the checker of Section 5.3, with intervals whose endpoints are dyadic rationals with 100 fractional bits, rounded outward, in place of balls. It bounds $L_{\mathrm{S}}(S^+)$ by the length of an explicit network: the segment for two points; for three points, a path of two edges or a star whose centre is computed in floating point and rounded to a point whose coordinates are multiples of 2^{-40}; and for four points, the network of Section 5.3, with its two Steiner points obtained in the same way. The same networks give the lower bound for Ψ. As for the checker (Section 5.3), these networks connect their points; thus their lengths are upper bounds for $L_{\mathrm{S}}(S^+)$. The data of a family describe its tree (Section 5.1) as a sequence of codes in depth-first order: the coordinate along which each internal box is split, and the splitting and the lemma set of each leaf. The program

recomputes every box from the roots and the codes and checks every leaf, so the covering of the domains by the checked boxes is part of the formal proof. A proof in Lean shows that if the program accepts the data of the four families, then every point of the domain of each configuration type lies in an accepted box and the inequality holds there. This proof contains the enclosure properties of the interval operations, the convexity argument of Lemma 4.1, the arguments of Lemmas 4.2 and 4.3 for unbounded boxes and for $f \geq d$, and the covering of the domains by the boxes of the trees.

The data enter the proof only as string constants, read from the data files at build time; the correctness proof holds for every string; the data files therefore need not be trusted. Theorems stating that the program accepts the leaves are proved by `native_decide`, which evaluates the program as compiled code. The trusted base therefore includes the Lean compiler and its interpreter for compiled code, the C compiler that it uses, the Lean runtime and the compiled functions of Lean's core library that the program calls. Besides the standard axioms `propext`, `Classical.choice` and `Quot.sound`, the theorem `steiner_ratio_ge` depends on further axioms, one for each compiled check, each stating that the check returns true. None of the definitions that the compiled checks evaluate is declared with `implemented_by`, `extern`, `unsafe`, `opaque` or `partial`.

The modules without data and a small test case, the complete certificate at $\rho = 4/5$ of d_regular with $f \geq d$ (427 records), can be built on a single machine, and deliberately altered versions of this test case are rejected. The repository describes an independent rebuild of this proof and replays of its exported proofs by three kernel checkers. The Lean files are in the directory `lean/` of the repository. The data files are in the archive, and the programs in `lean/tools/` regenerate them from the certificate files.

A.2. **The kernel-checked proof.** The second proof assumes (H1)–(H3) in the same form: its final theorem `S43K.Final.steiner_ratio_ge` has the same statement as `steiner_ratio_ge`, and its certificate theorem `S43K.Final.certificate` has the same statement as `certificate`. It reuses the files of the first proof that define the networks, the boxes, the tables and the four families, together with the proofs of the convexity argument, of the enclosure properties of the networks, of the arguments for unbounded boxes and of the covering of the domains. The checking program and the proof that its acceptance implies the certificate theorem are new.

The new program checks the same inequalities on the same boxes with the same intervals, whose endpoints are multiples of 2^{-100}, but it stores each endpoint as a natural number shifted by 2^{256}, so that the kernel can evaluate it with its built-in arithmetic on natural numbers. Lean proofs show that each operation returns the same interval as in the first program, except the square root, for which a separate proof shows that its interval contains the exact value. These proofs require that no subtraction of natural numbers be cut off at zero, and to ensure this the program rejects every record with a finite box end of 2^{26} or more, which bounds all intermediate values. The points of the explicit networks are computed by the program at each vertex of a box, with the constructions of Torricelli and of Melzak in exact arithmetic instead of floating-point iterations, and are rounded to multiples of 2^{-40}. As in the first proof, the correctness of this computation does not have to be proved.

For every record, a theorem states that the program accepts it, and it is proved by `decide +kernel`, so that the kernel evaluates the program. The boxes are recomputed from the codes of the trees, as in the first proof, and the records of each family are taken in the order of its tree. Each generated module holds a run of consecutive records and exports a single theorem: that its records carry the walk through the tree from one list of pending boxes to the next. Further generated modules chain these theorems for each family, and the final module derives the two theorems from the four chains. The monotone coordinates of the splittings and the reading of

the two tables, which the first proof also checks as compiled code, are decided in the kernel as well. Both theorems use only the axioms `propext`, `Classical.choice` and `Quot.sound`, and the trusted base of this proof is Lean's kernel, including the kernel's built-in arithmetic on natural numbers. The repository describes a rebuild of this proof on other machines. The Lean files of this proof, for the same versions of Lean and Mathlib, and the generator of its modules, which produces them from the data files of the first proof, are in the directory `lean_kernel/` of the repository.

Data availability

The certificate files (the records and the trees of the four families, 94,104,688,256 bytes uncompressed), compressed with xz, the data files of the formal proofs, and the SHA-256 checksums of the uncompressed files are archived at doi:10.5281/zenodo.23716011. The checker, the coverage program, the splitting tables, the generator, the logs of all runs, the control runs, the build commands, the programs of Section 6 and the Lean files of Section 7, for Lean `v4.35.0-rc3` and Mathlib commit `c55e6e786f49`, are available at `https://github.com/tejasvi/steiner-ratio-bound/tree/v4`, and a copy of this repository is in the archive. The modules generated for the kernel-checked proof are not archived, since a program in `lean_kernel/` produces them from the data files; the archive `lean_kernel_lists.tar.xz`, of 63,906,192 bytes, holds the SHA-256 hashes of these modules and of the files compiled from them, so that a rebuild can be compared file by file. The generator, the splitting tables and the formulas of the splittings are taken or derived from the code of [14], published under the MIT License, Copyright (c) 2025–2026 Yisi Ke and Yankai Shu; the license text accompanies these files.

References

[1] F. R. K. Chung and R. L. Graham, *A new bound for Euclidean Steiner minimal trees*, Ann. New York Acad. Sci. **440** (1985), 328–346.

[2] F. R. K. Chung and F. K. Hwang, *A lower bound for the Steiner tree problem*, SIAM J. Appl. Math. **34** (1978), 27–36.

[3] D.-Z. Du and F. K. Hwang, *A new bound for the Steiner ratio*, Trans. Amer. Math. Soc. **278** (1983), 137–148.

[4] D.-Z. Du and F. K. Hwang, *The Steiner ratio conjecture of Gilbert and Pollak is true*, Proc. Natl. Acad. Sci. USA **87** (1990), 9464–9466.

[5] D.-Z. Du and F. K. Hwang, *A proof of the Gilbert–Pollak conjecture on the Steiner ratio*, Algorithmica **7** (1992), 121–135.

[6] D.-Z. Du, F. K. Hwang, G. D. Song and G. Y. Ting, *Steiner minimal trees on sets of four points*, Discrete Comput. Geom. **2** (1987), 401–414.

[7] D.-Z. Du, F. K. Hwang and E. Y. Yao, *The Steiner ratio conjecture is true for five points*, J. Combin. Theory Ser. A **38** (1985), 230–240.

[8] D.-Z. Du, E. Y. Yao and F. K. Hwang, *A short proof of a result of Pollak on Steiner minimal trees*, J. Combin. Theory Ser. A **32** (1982), 396–400.

[9] The FLINT team, *FLINT: Fast Library for Number Theory*, `https://flintlib.org`.

[10] R. L. Graham and F. K. Hwang, *Remarks on Steiner minimal trees*, Bull. Inst. Math. Acad. Sinica **4** (1976), 177–182.

[11] E. N. Gilbert and H. O. Pollak, *Steiner minimal trees*, SIAM J. Appl. Math. **16** (1968), 1–29.

[12] A. O. Ivanov and A. A. Tuzhilin, *The Steiner ratio Gilbert–Pollak conjecture is still open*, Algorithmica **62** (2012), 630–632.

[13] F. Johansson, *Arb: efficient arbitrary-precision midpoint-radius interval arithmetic*, IEEE Trans. Comput. **66** (2017), 1281–1292.

[14] Y. Ke, T. Huang, Y. Shu, D. He, J. Gai and L. Wang, *Towards solving the Gilbert–Pollak conjecture via large language models*, in: Proceedings of the 43rd International Conference on Machine Learning, PMLR **306** (2026), 56442–56498; we cite the numbering of arXiv:2601.22365v3. Code: `https://github.com/keyisi2006/Steiner-Ratio`, commit 8f19604.

[15] L. de Moura and S. Ullrich, *The Lean 4 theorem prover and programming language*, in: Automated Deduction – CADE 28, Lecture Notes in Comput. Sci. **12699**, Springer, 2021, 625–635.
[16] The mathlib Community, *The Lean mathematical library*, in: Proc. 9th ACM SIGPLAN Int. Conf. Certified Programs and Proofs (CPP 2020), 2020, 367–381.
[17] H. O. Pollak, *Some remarks on the Steiner problem*, J. Combin. Theory Ser. A **24** (1978), 278–295.
[18] J. H. Rubinstein and D. A. Thomas, *The Steiner ratio conjecture for six points*, J. Combin. Theory Ser. A **58** (1991), 54–77.
[19] J. Savva, *Steiner_Ratio: independent verification of the Gilbert-Pollak certificate layer and a certificate-layer bound of 0.860*, Zenodo, 2026, doi:10.5281/zenodo.22223485.
[20] T. Tao et al., *Optimization constants in mathematics*, online repository, `https://teorth.github.io/optimizationproblems/`, 2026.
[21] E. Weiszfeld, *Sur le point pour lequel la somme des distances de n points donnés est minimum*, Tôhoku Math. J. **43** (1937), 355–386.

ORCID: `https://orcid.org/0000-0003-4668-3639`