The local weak limit of uniform random fullerenes

Myeongjin Kang. Prepared with AI assistance.

A proof of Conjecture 1 of Bille, Buchstaber, Coste, Kuriki and Spodarev, and its Lean certificate; MathDB record 358647.

mathdb:2611.00097v1 17 Nov 2026

Abstract

We prove Conjecture 1 of Bille, Buchstaber, Coste, Kuriki and Spodarev: the uniform random fullerene on n vertices converges locally weakly, in the sense of Eq. (4) of that paper, to the planar hexagonal lattice. The proof does not estimate the number A(k,l,p) of triangulations with a prescribed boundary, which the source names as its obstacle. It counts fullerenes with a short essential cycle by a bound independent of n, shows that every other fullerene has only O(1) vertices whose ball is not the lattice ball, and uses an elementary helical construction of many pairwise non-isomorphic fullerenes on each large even n. The statement is also formalized in Lean 4 with Mathlib; the endpoint Ful.endpoint is kernel checked and depends only on propext, Classical.choice and Quot.sound. The written proof was checked by hand by two AI reviewers of the same model family; its limitations and the limits of the formal audit are stated in Section 10.

1. Introduction

A fullerene is a finite, connected, 3-regular planar graph whose faces have degrees 5 and 6. This is Definition 2 of [1], quoted verbatim: ‘‘A fullerene (graph) is a finite, connected, 3-regular planar graph with faces of degrees both 5 and 6.’’ Euler’s formula forces exactly twelve pentagonal faces; all other faces are hexagons. If F has n vertices, m edges and f faces, then 3n = 2m, n - m + f = 2, n = 2f - 4 and m = 3f - 6.

The planar hexagonal lattice H of Eq. (3) of [1] has vertex set

$$V(H) = \{ (\sqrt{3}\, x + y\, \sqrt{3}/2,\ 3y/2 + c) \in \mathbb{R}^2 : x, y \in \mathbb{Z},\ c \in \{0,1\} \},$$

and every vertex is joined to its three nearest neighbours in the Euclidean distance. The ball $B_H(v,r)$ is the subgraph induced by the vertices at graph distance strictly smaller than r, as on p. 4 of [1].

Local weak convergence is Eq. (4) of [1]. A sequence of finite random graphs (G_n) converges locally weakly to a random rooted graph (G,o) if for every fixed integer r and every fixed rooted graph (H,o),

$$P(B_{G_n}(o_n, r) \cong B_H(o, r)) \to P(B_G(o, r) \cong B_H(o, r)) \quad (4)$$

as $n \to \infty$, where o_n is uniform on the vertices of G_n. Conjecture 1 of [1] reads, verbatim: ‘‘Let F_n be a random fullerene, chosen uniformly on the set of fullerenes with n vertices. The sequence (F_n) converges locally weakly towards the planar hexagonal lattice H.’’

Section 4 of [1] explains why its authors did not prove the conjecture. Their route reduces it to excluding more than one topological end, and the estimate they need is a bound on the number A(m,l,p) of planar graphs with m vertices, all faces triangles except one face of length l, and all interior vertices of degree 6 except exactly p of degree 5. They write in Section 4.3: ‘‘no efficient bound for A(k,l,p) is known, and we could not extract it from Thurston’s paper.’’ The proof below never estimates A(k,l,p). It replaces that estimate by two inputs: a bound, independent of n, on the number of fullerenes with a short essential cycle, and the fact that the number F(n) of fullerene isomorphism classes tends to infinity. Both are proved here elementarily.

We use the following reading, which is the reading of the source. F_n is uniform over isomorphism classes of graphs on n vertices, not over labelled graphs, and n runs over the values for which at least one fullerene exists. Every face has length 5 or 6; the phrase ‘‘faces of degrees both 5 and 6’’ in Definition 2 is read in the sense of the sentence immediately preceding it in [1], ‘‘one only allows faces with degree 5 or 6’’. Under this reading the dodecahedron, with n = 20, is a fullerene; the alternative reading excluding it changes only n = 20 and no limit. All fullerene orders are even, and no fullerene has 22 vertices.

The paper is organized as follows. Section 2 states the theorem in mathematical language. Section 3 proves that the dual of a fullerene is a simple triangulation, the ring identity and the bounded-disc lemma. Section 4 treats shortest essential cycles, clean rings, determinism and the n-independent count. Section 5 proves the local structure lemma away from pentagons. Section 6 proves that F(n) tends to infinity by an explicit helical family. Section 7 assembles the proof. Section 8 records where the Lean proof differs from the written one, Section 9 describes the Lean certificate, and Section 10 states the limits.

2. The theorem

Let F be a fullerene graph on n vertices, and let T be its plane dual. Since F is cubic, every face of T is a triangle; T is a triangulation of the sphere with exactly twelve vertices of degree 5 and all other vertices of degree 6. The vertices of F are the triangles of T, so F has n vertices if and only if T has n triangles. Write B_F(v,r) for the subgraph of F induced by the vertices at distance strictly smaller than r from v, and say that two rooted graphs are rooted-isomorphic if there is a graph isomorphism of the induced subgraphs taking the root to the root.

For n with at least one fullerene, let C(n) be the finite set of isomorphism classes of fullerenes on n vertices, and let F(n) = |C(n)|. Choose a class uniformly from C(n), choose a representative graph, and then choose a root o_n uniformly among its n vertices. For a rooted simple graph (K,o) and r ≥ 0, put

$$p_n(r,K,o) = P(B_{F_n}(o_n,r) \cong B_K(o,r)).$$

The limit is taken along the set of n for which C(n) is nonempty. The following is the theorem proved in this paper; it is the statement of the Lean theorem Ful.endpoint, whose exact text is quoted in Section 9.

Theorem 2.1. For every integer r ≥ 0 and every rooted simple graph (K,o) in a fixed universe,

$$p_n(r,K,o) \to 1 \text{ if } B_H(0,r) \cong B_K(o,r), \quad p_n(r,K,o) \to 0 \text{ otherwise,}$$

along the n with C(n) nonempty. Moreover C(n) is nonempty for infinitely many n, so the limit filter is proper and the statement is not vacuous.

Since H is vertex-transitive, the root 0 of H is immaterial, and the two cases in Theorem 2.1 are the two possible values of the deterministic right-hand side of Eq. (4). The proof gives a rate: with F(n) the number of classes, the probability that the r-ball at a uniform root is not the H-ball is at most B(r)/n + A(r)/F(n), where A(r) and B(r) depend on r only. Since F(n) → ∞, this probability tends to 0, and Theorem 2.1 follows.

3. The dual triangulation

3.1. The dual is a simple triangulation.

Let F be a fullerene and T its plane dual: the vertices of T are the faces of F, and two vertices of T are joined by an edge for every edge of F that separates the two faces. Because F is cubic, every face of T is a triangle. Because F is connected and planar, T is a triangulation of the sphere. Euler gives that exactly twelve faces of F are pentagons and the rest are hexagons, so T has exactly twelve vertices of degree 5 and all other vertices of degree 6.

Proposition 3.1. T is a simple graph: it has no loop and no pair of parallel edges. Consequently T is a simple triangulation of the sphere.

Proof. A loop of T would be an edge of F with the same face on both sides, that is, a bridge of F. A bridge e of F separates F into two parts. On each side of e there is a face incident with e; the facial walk of that face starts with e, then runs along the boundary of one side, and returns along the boundary of the other side. Each side has a vertex of degree 2 in the side (the endpoint of e) and all other vertices of degree 3, so each side contains a cycle and its outer boundary walk has length at least 3. Hence the facial walk has length at least 2 + 3 + 3 = 8, which is impossible because every face has length 5 or 6.

mathdb:2611.00097v1 17 Nov 2026

It remains to exclude digons. Suppose that T has a digon D: two parallel edges between vertices x and y, bounding a disc whose interior contains k pentagons of T, all interior degrees being 5 or 6. Euler applied to the disc D, with d denoting degree within D, gives $d(x) + d(y) = 2 + k$. If $d(x) = 2$ then the two digon edges are the only edges at x, and the face of D at x would have degree 2, which is impossible because every face of T is a triangle; so $d(x) \geq 3$, and similarly $d(y) \geq 3$, hence $k \geq 4$.

If D has an interior edge joining x to y, that edge splits D into two strictly smaller digons; choose D minimal among all digons, so this case does not occur. Let $N1 = k - 4$. Then N1 is the length of the inner walk of D, the closed walk formed by the edges opposite the boundary edges. If $N1 = 0$ then the inner walk is a single vertex of degree 2, a hub, which is impossible because all degrees are at least 3. If $N1 = 1$ then the inner walk is a loop, which is impossible. Hence $k \geq 6$.

The two sides of a digon in the sphere have k1 and k2 interior pentagons with $k1 + k2 \leq 12$, because T has only twelve pentagons. If $k1 + k2 \leq 12$ and each side has at least 6 by the preceding paragraph, then $k1 = k2 = 6$. For such a digon $N1 = 2$, so the inner walk has length 2. It is either a single edge traversed twice, which would give interior vertices of degree 3, impossible, or two parallel edges bounding a strictly smaller digon D' with $k(D') = 6$. Repeating gives an infinite descent of digons, contradicting the finiteness of T. Hence T has no digon. This proves the proposition.

A simple triangulation of the sphere with at least four vertices is 3-connected, so its dual F is a polyhedral graph and, by Whitney's theorem, its embedding is unique. This identifies the literal class of Definition 2 with the classical class of polyhedral fullerenes. The Lean proof does not use this uniqueness; see Section 8.

3.2. The ring identity.

Let D be a triangulated closed disc in T with simple boundary cycle C of length $l \geq 3$. Assume that D has no chord, at least one interior vertex, and all interior degrees at most 6. Let $k = k(D)$ be the sum, over the interior vertices of D, of 6 minus the degree in T. Let a(D) be the number of triangles of D.

Write V_i for the number of interior vertices, E_i for the number of interior-interior edges, E_c for the number of edges joining a boundary vertex to an interior vertex, and $E_b = l$ for the number of boundary edges. Euler for the disc gives $V_i + l - (E_i + E_c + l) + a(D) = 1$, hence $E_i + E_c = V_i + a(D) - 1$. Counting triangle-edge incidences, each boundary edge lies in exactly one triangle of D (there is no chord), each cross edge lies in exactly two triangles, and each interior edge lies in exactly two triangles; so $3a(D) = l + 2E_c + 2E_i$. Substituting the Euler relation gives

$$a(D) = l + 2V_i - 2. \quad (5)$$

Let N1 be the number of triangles of D with exactly one vertex on C. The triangles with two vertices on C are exactly the l triangles sitting on the l boundary edges; no triangle has all three vertices on C, because that would be a chord triangle. Every triangle with a boundary vertex has exactly two cross edges, and every cross edge lies in exactly two such triangles, so $E_c = l + N1$. The sum of the degrees in D over the boundary vertices is $2l + E_c = 3l + N1$. The sum of the degrees in D over the interior vertices is $2E_i + E_c$, and this equals $6V_i - k$. Using (5) and $E_c = l + N1$,

$$2E_i + E_c = 3a(D) - l - E_c = 3a(D) - 2l - N1 = 3a(D) - 3l + 6 - k,$$

because $a(D) = l + 2V_i - 2$ gives $6V_i = 3a(D) - 3l + 6$. Comparing the two expressions gives

$$N1 = l + k(D) - 6. \quad (6)$$

Writing $t = N1$ for the number of triangles with exactly one boundary vertex, (6) reads $t = l + k - 6$.

This is the ring identity. It also gives that the sum of the degrees in D over the boundary vertices is $4l + k(D) - 6$.

3.3. Blocks and the bounded-disc lemma.

Let R be the closed union of the triangles of D that meet C. A block is the closure of a connected component of the complement of R in D. The following facts are used below. Every block X is a closed disc whose boundary is a sim-

mathdb:2611.00097v1 17 Nov 2026

ple cycle of length at least 3. Indeed, each triangle of R touches C, so R together with the outside of D is connected, and each complement component is an open disc. Its boundary walk lies in the N1 inner edges, each of which has a triangle of R on one side, so the total length of all block boundaries is at most N1. The walk is simple: a repeated vertex would produce a closed curve inside the component whose far side would have to contain triangles of R, but those touch C on the other side, or would contain nothing, which is impossible because every edge of the walk has a triangle of R on one side. Hence each block is a closed disc with a simple boundary cycle. Every interior vertex of D lies in some block or on the inner walk, and the sum of k(X) over the blocks is at most k(D).

Lemma 3.2 (bounded discs). Let D be a triangulated disc with simple boundary of length $l \geq 3$, possibly with chords, and suppose $k(D) \leq 5$. Then

$a(D) \leq f(l) := 2l(l-2)$.

Proof. Induct on a(D). If D has no interior vertex and no chord, then D is a single triangle, $a(D) = 1 \leq 6 = f(3)$. If D has a chord, the chord splits D into two discs D1 and D2 with boundary lengths $l1, l2 \geq 3$ satisfying $l1 + l2 = l + 2$, and $k(D1) + k(D2) = k(D) \leq 5$. By induction $a(Di) \leq f(li)$, and

$f(l) - f(l1) - f(l2) = 4(l1 - 2)(l2 - 2) \geq 0$,

so $a(D) \leq f(l)$.

Now suppose that D has no chord and at least one interior vertex. By (6), $N1 = l + k(D) - 6 \leq l - 1$. The blocks X_i of D have boundary lengths b_i with sum $b_i \leq N1$, and each satisfies $k(X_i) \leq k(D) \leq 5$, so by induction $a(X_i) \leq f(b_i)$. The function f is increasing on $[3,\infty)$ and superadditive, because $f(a+b) - f(a) - f(b) = 4ab \geq 0$. Hence

$a(D) = l + N1 + \Sigma_i\, a(X_i) \leq l + N1 + f(N1) \leq 2l - 1 + f(l-1)$.

Since $f(l-1) = 2(l-1)(l-3)$, the right-hand side is $2l^2 - 6l + 5$, which is at most $2l^2 - 4l = f(l)$ for $l \geq 3$. If D has no block at all, then $a(D) \leq l + N1 \leq 2l - 1 \leq f(l)$. This completes the induction.

Corollary 3.3. Let C be a simple cycle of T of length l. If one side of C has at most 5 interior pentagons, then that side has at most f(l) triangles. Consequently, if neither side of C has at most f(l) triangles, then each side has exactly 6 interior pentagons, all twelve pentagons of T lie strictly on the two sides, and C contains no pentagon.

Proof. The first sentence is Lemma 3.2 applied to the side, viewed as a disc with boundary C. If neither side has at most f(l) triangles, then by the first sentence neither side has at most 5 interior pentagons, so each side has at least 6. The two sides contain all pentagons not on C, and there are twelve, so each side has exactly 6 and none lies on C.

4. Short essential cycles

Call a simple cycle C of T essential if each of its two sides has at least 6 interior pentagons. By Corollary 3.3, an essential cycle has exactly 6 interior pentagons on each side and contains no pentagon. Let sys(T) be the minimum length of an essential cycle, or ∞ if there is none. The number sys(T) is an isomorphism invariant of T.

Lemma 4.1 (the shape of a shortest essential cycle). Let C be a shortest essential cycle of T, of length ell, and let A be one of its two sides, so $k(A) = 6$. Then either A is a clean ring, or $a(A) \leq G(ell) := f(ell+2) + 2ell$. A clean ring is a disc consisting of the ell triangles on the boundary edges and the ell triangles with one boundary vertex, together with a single block X whose boundary is again a shortest essential cycle; it has 2ell triangles outside X.

Proof. Suppose first that A has a chord. The chord splits A into discs A1 and A2 with $k(A1) + k(A2) = 6$. If $k(A1) = 6$, then the boundary cycle of A1, formed by the chord and one arc of C, has 6 pentagons on the A1 side and 6 on the other side: the A2 side contributes none, C contains no pentagon, and the part of T outside A contributes its 6 pentagons. So that boundary is a shorter essential cycle, contradicting the minimality of C. Hence $k(Ai) \leq 5$ for both i. By Lemma 3.2, $a(Ai) \leq f(li)$ with $l1 + l2 = ell + 2$, so $a(A) \leq f(l1) + f(l2) \leq f(ell+2)$.

Now suppose that A has no chord. By (6) with $l = ell$ and $k = 6$, $N1 = ell$. If some block X has $k(X) = 6$, then the boundary cycle of X is essential: it has 6 pentagons on the X side, and the other side has the remaining 6, because all 6 interior pentagons of A are in X and the other blocks have $k = 0$. Since $k(A) = 6$ is already attained by X, no

mathdb:2611.00097v1 17 Nov 2026

vertex of the inner walk is a pentagon either. Its length is at most the sum of all block boundary lengths, which is at most N1 = ell. By the minimality of C its length is exactly ell, and the sum of the block boundary lengths equals ell, so X is the only block and its boundary is the whole inner walk. Then A consists of the ell triangles on the boundary edges, the ell triangles with one boundary vertex, and X; it has 2ell + a(X) triangles and is a clean ring. If no block has k = 6, then every block has k ≤ 5, and

a(A) = ell + N1 + Σ_i a(X_i) ≤ 2ell + f(ell) ≤ G(ell).

This proves the lemma.

4.2. Determinism of clean rings.

Let D be a clean ring of length ell in T, with boundary cycle C and inner boundary cycle C'. Label the boundary vertices v_1,...,v_ell cyclically and define the ring word s(C) = (deg_D(v_1),...,deg_D(v_ell)), where D is the clean ring viewed as a disc with boundary C. At a boundary vertex v_i, the degree in D equals 2 plus the number of cross edges from v_i to the inner cycle, because two boundary edges meet at v_i. Hence s(C) determines, cyclically, the inner vertices and the number of ring triangles incident with each of them. Each inner vertex w has degree 6 in T, so its degree in the block X is 6 minus the number of ring triangles at w. Therefore s(C) determines the ring word s(X) of the inner boundary, and the whole ring D is determined by s(C) together with the rooted block X. There is a fixed combinatorial map Phi, independent of ell and of T, such that s(X) = Phi(s(C)).

Consequently, if a cap A* with a marked boundary vertex is given, the rings are forced outward one after another: the first ring word is determined by the cap, and each subsequent ring word is Phi of the previous one. Hence a triangulation T whose shortest essential cycle has length ell is determined by the length ell, the two terminal caps A* and B* with marked boundary vertices, the number H of clean rings, and the relative rotation s ∈ Z/ell between the two caps. The number of triangles is

n = a(A*) + a(B*) + 2 ell H,

so H is determined by n once the two caps and ell are fixed.

Lemma 4.2 (the count of short-tube fullerenes). For every L ≥ 3, the number of triangulations T on n triangles with sys(T) ≤ L is at most

A(L) := Σ_{ell ≤ L} ell × |Cap(ell)|^2,

where Cap(ell) is the set of rooted triangulated discs with boundary length ell and at most G(ell) triangles. In particular this number does not depend on n.

Proof. Let T have an essential cycle of length at most L, and choose a shortest one, of length ell ≤ L. Its two sides are the two caps. By Lemma 4.1 each cap is either a clean ring, which is counted through its terminal disc, or has at most G(ell) triangles. In either case each cap is determined by a rooted disc in Cap(ell) with boundary ell: for a clean ring, the terminal disc together with the ring word determines the ring, and the ring word is Phi of the terminal word. The relative rotation between the two caps is an element of Z/ell. Since H is determined by n, the triangulation is determined by the two rooted caps and the rotation. Hence there are at most ell |Cap(ell)|^2 possibilities for a fixed ell, and summing over ell ≤ L gives the bound. Every T with sys(T) ≤ L is counted at least once, so the bound is valid.

This is the key point at which the proof avoids the source's A(k,l,p). A tube with given caps and circumference is rigid: once the two caps and the length are fixed, only the number of rings and one rotation remain, and the number of rings is fixed by n. So the count of fullerenes with a short essential cycle is bounded independently of n.

mathdb:2611.00097v1 17 Nov 2026

5. Local structure away from pentagons

Let Lambda be the triangular lattice, the dual of H. Fix r ≥ 0 and put

rho = r + 3, l0 = 2rho + 2 = 2r + 8, R = rho + 1 + f(l0).

Lemma 5.1 (lattice ball far from pentagons). Let t be a triangle of T such that no pentagon of T lies within distance R of any vertex of t, and suppose sys(T) > l0. Then the ball of T of radius rho around t is isomorphic, as a triangulated disc with root, to the ball of Lambda of radius rho; consequently the fullerene ball of radius r around the vertex of F corresponding to t is the ball of H.

Proof. Develop a map psi from the lattice ball X of radius rho + 1 around a lattice triangle t0 into T, crossing edges one at a time. Since X is simply connected and every interior vertex of X has degree 6, and every vertex of T within distance rho + 1 of t has degree 6 because no pentagon is within R > rho + 1, the development is a well-defined simplicial map, locally injective at every vertex.

Suppose that psi is not injective on the ball B(rho). Choose two distinct lattice points x, y in B(rho) with psi(x) = psi(y) at minimal lattice distance delta. Since psi is injective on stars, delta ≥ 3. Lattice balls are geodesically convex: a lattice geodesic uses at most two of the three lattice directions, and each of the three defining linear functionals is monotone along it. Take a lattice geodesic z_0 = x, z_1, ..., z_delta = y inside B(rho). Its image is a simple cycle C of length delta ≤ 2rho, and no vertex of C is a pentagon, because all its vertices are within distance rho < R of t.

If C is essential, then sys(T) ≤ delta ≤ 2rho < l0, contradicting the hypothesis. Hence C is not essential, so one side D of C has at most 5 interior pentagons. By Lemma 3.2, a(D) ≤ f(delta) ≤ f(l0), since delta ≤ 2rho < l0 and f is increasing. If D contained a pentagon, then that pentagon would be within distance a(D) ≤ f(l0) of C, hence within distance f(l0) + rho < R of t, contradiction. Therefore D has no pentagon, so all its interior vertices have degree 6 and D is flat.

Develop D into Lambda. There is a developing map psi_D from D to Lambda which is the inverse of psi along the path, normalized so that psi_D(psi(t0)) = t0. Unique continuation across edges shows that psi_D composed with psi is the identity on the strip of lattice triangles on the D side of the path. Hence

z_0 = psi_D(psi(z_0)) = psi_D(psi(z_delta)) = z_delta,

contradicting x ≠ y. Therefore psi is injective on B(rho).

It remains to exclude extra adjacencies. If x and y are non-adjacent lattice points in B(rho) whose images under psi are adjacent in T, apply the same argument to the path from x to y together with that edge. The resulting cycle has length at most 2rho + 1 ≤ l0, so it is not essential; one side has at most 5 pentagons, and the flat development maps the extra edge to a lattice edge, so x and y are adjacent in Lambda, contradiction. Hence psi preserves and reflects adjacency on B(rho).

Thus psi is an isomorphism of triangulated discs from the lattice ball of radius rho onto its image. Every triangle of T within dual distance r of t has all its vertices within distance r + 1 of t0, and the choice rho = r + 3 gives the required margin: the three triangles at a vertex of F are all inside the developed region. Therefore the fullerene ball of radius r around the vertex corresponding to t is the H-ball. This proves the lemma.

A triangle t is called bad if some pentagon lies within distance R of one of its vertices. Call a vertex of T bad if some pentagon lies within distance R of it. Each pentagon has at most 6ˆ(R+1) vertices within distance R, because every vertex of T has degree at most 6. There are twelve pentagons, so there are at most $12 \times 6^{(R+1)}$ bad vertices of T. Each bad vertex of T is a corner of at most 6 triangles, so there are at most $72 \times 6^{(R+1)}$ bad triangles. In the non-exceptional case, every bad fullerene vertex is a bad triangle: if its triangle had no vertex within R of a pentagon, Lemma 5.1 would make its ball good. Hence the number of bad fullerene vertices is at most

B(r) := $72 \times 6^{(R+1)}$.

mathdb:2611.00097v1 17 Nov 2026

6. The number of fullerenes tends to infinity

This section gives the construction in corrected form; an earlier hand-written version contained two immaterial slips. The construction is a helical family of fullerenes with an elementary verification; it replaces the deep input $F(n) \to \infty$ by a self-contained count. The two corrections are these: the maximum change of the height functional D under one lattice step is L, not L - 1, attained in the direction w^2; and the thresholds actually stated give only $F(n) \geq c\, n^{(1/3)}$, not $F(n) \geq c\, \sqrt{(n)}$. Neither correction affects the argument, because only $F(n) \to \infty$ is used.

Let Lambda = Z[w] with $w = \exp(i\,\pi/3)$, so w^2 = w - 1, and let the unit steps be ±1, ±w, ±w^2. For $s \geq 1$ define a convex lattice polygon P_s with boundary traversed counterclockwise through six marked vertices

V0 = (0,0), V1 = (1,0), V2 = (-1,2), V3 = (-1-s,2), V4 = (-1-s,1), V5 = (-s,0).

Its side vectors are S0 = 1, S1 = 2w^2, S2 = s w^3, S3 = w^4, S4 = w^5, S5 = s, and they sum to 0. The cap P_s has 2(s+1) + (2s+1) = 4s + 3 triangles and boundary length L = 2s + 5. At the six marked vertices the numbers of cap triangles are 3, 1, 2, 2, 2, 2; every unmarked boundary vertex is straight and has 3 cap triangles.

Rotate the side S_j by w^(-j). The resulting closed path gamma on the cylinder Y = C/<c> has one step in direction 1 and then 4 + 2s steps in direction w, so its chiral vector is c = 1 + (2s+4)w. This vector is primitive because its coefficient of 1 is 1. Its length is L = 2s + 5, equal to the boundary length of P_s. Glue P_s on the left of gamma. Each marked vertex loses exactly 60 degrees of turn: the cap contributes 3, 1, 2, 2, 2, 2 triangles and the cylinder side contributes 2, 4, 3, 3, 3, 3, so the total at each marked vertex is 5, while unmarked vertices have 3 + 3 = 6. Thus the marked vertices become the twelve degree-5 vertices of a fullerene dual, and all other vertices have degree 6.

The helical model makes the gluing explicit. Define D(x + y w) = y - (L-1)x. This is a surjection from Lambda to Z with kernel <c>, so the vertices of the cylinder Y are exactly the integers. In these labels the neighbours of a vertex differ by ±1, ±(L-1) or ±L, and the triangles of Y are, for each integer j,

$U_j = \{j, j+L-1, j+L\}, \quad D_j = \{j, j+1, j+L\}.$

The curve gamma has labels g_0 = 0 and g_i = i - L for $1 \leq i \leq L-1$, so gamma is the path -L+1, ..., 0 closed by the edge from 0 to -L+1. Simplicity of gamma and of Y is immediate from these labels. The band between the two caps has vertices [-m,0] and triangles U_j and D_j for j in [-m,-L], so it has exactly A(m) = m - L + 1 rhombi, with A(m+1) = A(m) + 1. The second cap is the image of P_s under $x \to -m - x$ on the labels.

At a vertex x of gamma the number of tube triangles is 2, plus 1 if $x \leq -1$, plus 1 if x = -L+1. Thus x = 0 has 2, x = -L+1 has 4, and all other vertices have 3, matching the tube counts 2, 4, 3, 3, 3, 3 at V0,...,V5. The total degree is 5 at the six marked vertices and 6 elsewhere. The number of vertices is m + 1 + 2s and the number of triangles is

$n = 8s + 6 + 2(m - L + 1) = 4s + 2m - 2.$

Since m is free, every even $n \geq N(s)$ is realised by some T_s(m), where N(s) is the least even number for which some m > 2L - 2 gives n = 4s + 2m - 2.

Simplicity is checked as follows. A band edge joining two non-consecutive gamma labels would need a difference in {1, L-1, L} other than a consecutive pair; solving the unit equations, whose x-part forces the coefficient to lie in {0, ±1}, leaves only consecutive pairs. Hence no chord of a cap coincides with a band edge, and the single chord of each cap is not doubled. Requiring m > 2L - 2 makes the two gamma curves disjoint, so the two caps do not meet. Thus T_s(m) is a simple sphere triangulation with exactly twelve degree-5 vertices, that is, the dual of a fullerene.

For m large in terms of s, sys(T_s(m)) = L. The upper bound is a translate of gamma deep in the band: it is a pentagon-free simple cycle with six pentagons on each side, hence essential, of length L. For the lower bound, let C be a shortest essential cycle, of length ell ≤ L. By Lemma 4.1 the caps have at most G(L) triangles, so a vertex v at distance greater than G(L) + L from all pentagons lies on a clean ring of length ell whose cycle lies in the ball of radius L about v, which is a patch of Y. The cycle has pentagons on both sides, so it is the generator of the annulus and lifts to a path from 0 to ±c. Hence its length is at least the hexagonal distance from 0 to c, which is 1 + (2s+4) = L. Such a vertex exists for m large, because the pentagons are confined to bounded neighbourhoods of the caps while the band is long.

mathdb:2611.00097v1 17 Nov 2026

Since sys is an isomorphism invariant, the T_s are pairwise non-isomorphic for distinct s. Therefore, for every M, every even n ≥ max_{s ≤ M} N(s) has F(n) ≥ M, and F(n) → ∞ along the even orders. The argument gives only about F(n) ≥ c n^(1/3); the true order n^9 is not addressed and is not needed. This is the elementary input that removes the only non-formalizable step of the written proof.

7. Proof of the theorem

Fix r ≥ 0. Keep the notation of Section 5:

rho = r + 3, l0 = 2r + 8, R = rho + 1 + f(l0), B(r) = 72 × 6^(R+1),

and put A(r) = A(l0), the count of Lemma 4.2.

Let T be the dual of a fullerene on n vertices. Call T exceptional if sys(T) ≤ l0. By Lemma 4.2 there are at most A(r) exceptional triangulations on n triangles. If T is not exceptional, then sys(T) > l0, and by Lemma 5.1 every triangle at distance greater than R from all pentagons has the lattice ball of radius r. Hence every bad fullerene vertex lies within distance R of a pentagon, and by the count of Section 5 there are at most B(r) bad vertices. Therefore, for a uniform class and a uniform root,

P(the r-ball is not the H-ball) ≤ B(r)/n + A(r)/F(n). (7)

In (7) the first term accounts for bad roots inside a non-exceptional class, and the second for the exceptional classes.

Now fix a rooted simple graph (K,o). If the H-ball of radius r is rooted-isomorphic to the K-ball, then every good root has a ball isomorphic to the K-ball, so p_n(r,K,o) ≥ 1 - P(bad). If the two balls are not rooted-isomorphic, then every root with a K-ball is bad, so p_n(r,K,o) ≤ P(bad). In either case

|p_n(r,K,o) - 1_{B_H(0,r) ≅ B_K(o,r)}| ≤ B(r)/n + A(r)/F(n).

By Section 6, F(n) → ∞ along the n with at least one fullerene, so the right-hand side tends to 0. This proves Theorem 2.1 and Eq. (4), with the limit taken along the fullerene orders. The non-vacuity is Section 6: fullerenes exist on every sufficiently large even n, so the filter is proper.

The proof gives the rate in (7) but no lower bound on F(n) beyond F(n) → ∞. No step uses the exact order n^9 of Thurston, Engel-Smillie or Rukhovich; that order is quoted in [1] only for context. The constants A(r) and B(r) are enormous but finite and depend on r only.

8. Where the Lean proof differs from the written proof

The Lean development proves the same endpoint but organizes several steps differently. The differences are recorded here so that the formal constants are not mistaken for the constants of Sections 3 to 7.

First, the Lean proof replaces the flat-disc step of Section 5 by a corner-count argument. Instead of developing the flat disc D into the lattice and inverting the developing map, it counts corners and bounds the total turning along closed sequences obtained from the development (the modules FulLocal and FulInj), and rules out a collision along a lattice geodesic by that turning bound. The written proof instead uses the flat developing map directly. The two arguments prove the same injectivity statement, but they are not literally the same proof.

Second, the Lean constants are looser. The Lean endpoint takes L = 4(r+2) + 4 = 4r + 12, whereas the written proof uses l0 = 2r + 8. The Lean bad-vertex count is 72 × 7^(Rad r), whereas the written count is 72 × 6^(R+1) with R = r + 4 + f(2r+8). The Lean proof does not claim the constants of the written proof, and the constants of the written proof are not those of the Lean proof.

Third, the Lean proof does not use the systole to prove non-isomorphism in the helical family. It uses two graph-distance bounds on the explicit family, a Lipschitz bound dist_F ≤ 6 dist_T + 5 and a lower bound dist_F ≥ (label range)/(2L), and chooses s_j = 30^(j+1) so that the diameter ranges are pairwise disjoint for large n. This avoids essential-cycle reasoning and Whitney uniqueness for the family. The written proof uses sys(T_s(m)) = L and the fact that sys is an isomorphism invariant.

mathdb:2611.00097v1 17 Nov 2026

Fourth, the Lean proof fixes one embedding per isomorphism class by Classical.choose and uses only that embedding. The fullerene property asserts that some genus-0 rotation system with faces of length 5 or 6 exists; a graph admitting several such embeddings is one class. The exceptional set is defined through the chosen embedding, and a class is exceptional as soon as its chosen embedding has a short essential cycle. No Whitney uniqueness of embeddings is proved or used. The written proof mentions Whitney only to identify the literal class with the polyhedral class, and that identification is not needed for the endpoint.

Fifth, the Lean digon descent is shorter than the written block argument. It uses first-return exponents at a minimal digon side, both at least 2 and summing to 6, to produce a strictly smaller digon, rather than the block and ring identity of Section 3. The Lean helical family is given by an explicit rotation table on integer labels in [-s, m+s], with circumference L = 2s + 5, valid for s ≥ 4 and m ≥ 3(2s+5), and with 2m + 4s - 2 vertices; the written family is valid for s ≥ 1 and m > 2L - 2. The Lean table proof checks local laws by case analysis and forces genus by total curvature, and it proves simplicity by construction rather than by the chord argument of Section 6.

9. The Lean certificate

The formal development uses Lean 4.34.1 and Mathlib at revision d13f23b723b8a846827a245b89c10fc7d3f11612. The statement lives in the module FulStatement and the endpoint in the module FulEndpoint. The definitions below are quoted from FulStatement. Lean's Unicode notation is typeset with the corresponding mathematical glyphs where the typesetting fonts have them. Where a glyph is unavailable, the following identifiers are used: the blackboard-bold N is written Nat, the blackboard-bold Z is written Int, the blackboard-bold R is written Real, the blackboard-bold N with an infinity superscript is written ENat, the neighbourhood-filter notation is written nhds, the powerset notation is written principal, the lattice meet is written inf, the graph-isomorphism notation is written ≅g; Lean's token there is the Unicode character U+2243 followed by g; the summation sign is typeset as the Greek capital Sigma (U+03A3) and the angle brackets as U+2329 and U+232A, while Lean's tokens are the n-ary summation sign U+2211 and the angle brackets U+27E8 and U+27E9. Apart from these substitutions, the tokens are those of the modules.

Rotation system and fullerene.

```
structure RotationSystem {V : Type u} (G : SimpleGraph V) where
  σ : Equiv.Perm G.Dart
  tail_eq : ∀ d, (σ d).fst = d.fst
  transitive : ∀ d d' : G.Dart, d.fst = d'.fst → σ.SameCycle d d'

structure IsFullerene {V : Type u} [Fintype V] (G : SimpleGraph V) : Prop where
  connected : G.Connected
  cubic : ∀ v, Nat.card (G.neighborSet v) = 3
  embedding : ∃ ρ : RotationSystem G,
    Nat.card V + ρ.numFaces = Nat.card G.edgeSet + 2 ∧ ∀ d, ρ.faceLen d = 5 ∨ ρ.faceLen d = 6
```

Balls and rooted isomorphism.

```
def ballSet {V : Type u} (G : SimpleGraph V) (o : V) (r : Nat) : Set V :=
  {v | G.edist o v < (r : ENat)}

def RootedBallIso {V : Type u} {W : Type v} (G : SimpleGraph V) (o : V) (K : SimpleGraph W) (o' : W)
    (r : Nat) : Prop :=
  ∃ e : G.induce (ballSet G o r) ≅g K.induce (ballSet K o' r),
    ∀ (h : o ∈ ballSet G o r) (h' : o' ∈ ballSet K o' r), e ⟨o, h⟩ = ⟨o', h'⟩
```

The hexagonal lattice.

```
def hexLattice : SimpleGraph (Int × Int × Bool) where
  Adj p q := hexDist4 p q = 4
  symm := ⟨fun p q h => by
    have : hexDist4 q p = hexDist4 p q := by
      unfold hexDist4
      ring
    rw [this]
```

mathdb:2611.00097v1 17 Nov 2026

```
    exact h⟩
  loopless := ⟨fun p h => by
    have h' : hexDist4 p p = 0 := by
      unfold hexDist4
      ring
    rw [h'] at h
    norm_num at h⟩
```

The uniform random fullerene and the limit.

```
def fullereneBallProb (n : Nat) (r : Nat) {W : Type v} (K : SimpleGraph W) (o : W) : Real :=
  (Σ c : FullereneClass n, ballFraction (Quotient.out c).1 r K o) / (numFullerenes n : Real)

def FullereneOrders : Set Nat := {n | Nonempty (FullereneClass n)}

def fullereneLimit : Filter Nat := atTop inf principal FullereneOrders

def NonVacuous : Prop := Filter.NeBot fullereneLimit

def LocalWeakConvergence : Prop :=
  ∀ (r : Nat) (W : Type v) (K : SimpleGraph W) (o : W),
    Tendsto (fun n => fullereneBallProb n r K o) fullereneLimit
      (nhds (if RootedBallIso hexLattice hexRoot K o r then 1 else 0))
```

In the module FulStatement, FullereneClass n is the quotient of labelled fullerenes on Fin n by graph isomorphism, numFullerenes n is its cardinality, and ballFraction averages over roots. Thus fullereneBallProb is the probability of Eq. (4): first a uniform isomorphism class, then a uniform root. The theorem proved in FulEndpoint is, verbatim,

```
theorem endpoint : LocalWeakConvergence.{v} ∧ NonVacuous := ⟨localWeakConvergence, Helix.nonVacuous⟩
```

Its statement is exactly Theorem 2.1 together with the non-vacuity of the limit filter. The module FulEndpoint proves it from the definitions in FulStatement with no additional hypothesis. The audit output for the endpoint and the nine named theorems is quoted verbatim from the audit record:

```
'Ful.endpoint' depends on axioms: [propext, Classical.choice, Quot.sound]
'Ful.localWeakConvergence' depends on axioms: [propext, Classical.choice, Quot.sound]
'Ful.Helix.nonVacuous' depends on axioms: [propext, Classical.choice, Quot.sound]
'Ful.Helix.numFullerenes_ge' depends on axioms: [propext, Classical.choice, Quot.sound]
'Ful.bad_bound' depends on axioms: [propext, Classical.choice, Quot.sound]
'Ful.exc_card_le' depends on axioms: [propext, Classical.choice, Quot.sound]
'Ful.transfer_ball' depends on axioms: [propext, Classical.choice, Quot.sound]
'Ful.ball_of_flat' depends on axioms: [propext, Classical.choice, Quot.sound]
'Ful.graphIso_of_mapIso' depends on axioms: [propext, Classical.choice, Quot.sound]
'Ful.mapOf_isFullerene' depends on axioms: [propext, Classical.choice, Quot.sound]
```

The independent audit replayed a fresh extracted build: 50 of 50 modules compiled and kernel-checked with exit 0, with no warnings or extra output. The axiom walk covered 27,127 constants reachable from the named theorems, with no missing constants; the only axioms in the closure are propext, Classical.choice and Quot.sound. Token scans over all 50 files found no sorry, admit, axiom, unsafe, opaque, native_decide, ofReduceBool, notation, syntax, elab, run_cmd or attribute priority declarations; implemented_by occurs 5 times and extern 57 times, all in core declarations (Lean.Name.*, Nat.repr and List.attachWith for implemented_by, and core arithmetic and arrays for extern), and partial occurs once, in a comment in FulChart. There are no declaration redefinitions of Mathlib names, and every declaration sits inside namespace Ful. The endpoint has no hypotheses; internal theorems with hypotheses, such as transfer_ball and local_ball with a systole hypothesis, discharge them from the non-exceptional condition, and exceptional classes are counted by exc_card_le.

The audit's reading depth must be stated plainly. The auditor read 17 statement and bridge modules in full, every line: FulStatement, FulMap, FulBallIso, FulLattice, FulLatticeAut, FulGraphOf, FulBridge, FulMapOf, FulMapFin, FulFmap, FulBad, FulTransfer, FulTransfer2, FulPsi, FulEndpoint, FulEuler and FulCycle. The other 33 modules were read at declaration level: every definition in full and every theorem statement with its hypotheses, but not the proof scripts. For those 33 proof scripts the auditor relied on the kernel, the token scans and the axiom walk, which cover the complete text or the compiled constants. The auditor did not read those proof scripts line by line. Upstream Mathlib oleans were trusted; the Mathlib checkout was not hashed or inspected by the auditor. The kernel

mathdb:2611.00097v1 17 Nov 2026

checker replays each named module, but that is a check of the compiled objects, not a line-by-line reading of every proof.

Three modelling decisions matter for fidelity, and all are recorded in the fidelity note of the formal development. First, planarity is rendered as a rotation system: a permutation of the darts that preserves tails, whose cycles at each vertex are the darts with that tail, with V + F = E + 2, where faces are the orbits of the face permutation. For a connected graph this is the combinatorial definition of a genus-zero cellular embedding; the equivalence with topological planarity is classical and is not proved in Lean. Second, ‘‘faces of degrees both 5 and 6’’ is read as ‘‘every face has length 5 or 6’’, following the sentence before Definition 2 in [1]. The stricter reading, requiring at least one face of each degree, excludes only the dodecahedron at n = 20 and changes no limit. Third, the uniform law is over isomorphism classes of graphs, not labelled graphs. The labelled reading is not formalized. The two models differ by bounded factors if one has a uniform bound on the automorphism group; the audit records that the bound $|\mathrm{Aut}| \leq 120$ is not proved in Lean, so the labelled reading is outside the formal certificate.

One further convention is worth recording. The ball is defined with the strict inequality edist o v ≺ r, verbatim the source’s ‘‘smaller than r’’. With this convention the source’s r corresponds to the usual radius r - 1. At r = 0 both balls are empty and the root condition is vacuous, which is harmless. The limit is quantified over all rooted simple graphs (K,o) in any universe, which is formally stronger than the source’s restriction to locally finite rooted graphs. The root of H is (0,0,false); the lattice is proved vertex-transitive, so the root is immaterial.

10. Limits

This section states plainly what was not read and what the objections to significance are. No claim of peer review, priority or worldwide novelty is made.

Prior art. The prior-art search was bounded and had no citation index. It could not check MathSciNet, zbMATH or any citation index, and it did not reach a journal final. The journal version of [1] is J. Phys. A: Math. Theor. 58(2) (2025) 025212, 31 pages, DOI 10.1088/1751-8121/ad9efb, as listed by a Russian mathematics portal in a search summary. That journal version was not read, and whether it keeps Conjecture 1 as an open conjecture is unknown. The arXiv version read here is v2 of 17 January 2025; the MathDB record’s status note dates from June 2023 and therefore predates the journal version. No proof or disproof of Conjecture 1 was found in the bounded search, and a later talk abstract by the same group is reported as listing the conjecture as open. This is bounded evidence, not a novelty certificate.

Works not read. Thurston [4] was not read. Brinkmann and coauthors on nanotube caps were not read, nor was Brinkmann, Goedgebeur and McKay [5] beyond its bibliographic record. Graver’s patch papers were not read. Kardos and Skrekovski 2008, the 2009 preprint’s reference [4] for the classification of cuts, was not read. The Kardos habilitation could not be delivered, and the MATCH 71(1) 99-114 item could not be delivered. Angel, Benjamini and Horesh [6] were not read. Kutnar and Marusic were consulted only for the size-5 case, and a 2009 preprint of Kardos, Krnc, Luzar and Skrekovski on cyclic 7-edge-cuts was consulted at figure and theorem level, not in full. The cached texts of Engel-Smillie [2] and Rukhovich [3] were available, but were not read in full; only opening pages and keyword searches were consulted for this draft. The source [1] was read as extracted text, in passages, not as page images. The literature on cap and nanotube enumeration was not systematically searched.

The rigid-tube picture. The statement that a short essential cycle forces a rigid tube with two caps is close to folklore in the cap and nanotube literature. The printed results seen stop at cut size 7 and proceed by case analysis; no printed statement of the rigid-tube picture uniformly in the circumference was found. The clean-ring structure is the flat cylinder picture at total curvature 2π, and one AI reviewer reads the bounded-disc lemma as a discrete form of the Alexandrov-Fiala-Bol inequality for discs with curvature below 2π, those classical inequalities were not read here. A specialist route through Thurston of perhaps a few pages may exist, but no complete bridge was found. The source authors, who cite Thurston, say they could not extract the needed bound.

Significance. The objections to significance are these. The result is the expected picture: fullerenes are lattice points in a cone over CHˆ9/Gamma, short 6|6 curves correspond to cusps, and thin parts have small volume, so rar-

mathdb:2611.00097v1 17 Nov 2026

ity of tubes is expected. Each step is a discrete shadow of classical non-negative curvature, and the core lemmas use only the fact that degrees are at most 6, so the argument is a discrete non-negative-curvature argument rather than a fullerene-specific one. A Thurston-literate specialist could plausibly write a short proof. The elementary count F(n) $\to \infty$ is not new; Thurston, Engel-Smillie and Rukhovich already give order nˆ9. What is new here, if anything, is the removal of the source's A(k,l,p) obstacle by a count independent of n, together with the local-to-global step. Significance is modest, and no worldwide novelty is claimed.

What was machine checked, and what was not. The written proof of Sections 3 to 7 was checked by hand by two AI reviewers of the same model family; it was not machine checked, and two immaterial slips in an earlier version were corrected. The Lean endpoint was kernel checked, but the auditor read 17 modules line by line and the other 33 at declaration level, relying on the kernel for their proof scripts. The Lean constants are not the constants of the written proof. The formal certificate does not cover the labelled model, and it does not prove the Whitney uniqueness of embeddings. Upstream Mathlib oleans were trusted. The journal version, Thurston, the Kardos habilitation, the MATCH item and any citation index were not read. The source was read as extracted text, not as page images. These limits bound what this draft establishes.

AI assistance. This draft was prepared with AI assistance; the author is Myeongjin Kang. The mathematical proof, the hand reviews and the Lean development were produced with AI assistance. No invented author, affiliation, date, acknowledgement, citation, page number or result appears here. All bibliographic details are those printed in the inputs, and all Lean statements are quoted from the named modules.

References

[1] A. Bille, V. Buchstaber, S. Coste, S. Kuriki and E. Spodarev, ‘‘Random eigenvalues of graphenes and the triangulation of plane’’, arXiv:2306.01462v2, 17 January 2025. Definition 2, Eq. (3), Eq. (4), Conjecture 1, Remark 4 and Sections 4.2-4.3.

[2] P. Engel and P. Smillie, ‘‘The number of convex tilings of the sphere by triangles, squares, or hexagons’’, Geometry & Topology 22 (2018) 2839-2864.

[3] A. D. Rukhovich, ‘‘On the growth rate of the number of fullerenes’’, Russian Mathematical Surveys 73:4 (2018) 734-736, DOI 10.1070/RM9837.

[4] W. P. Thurston, ‘‘Shapes of polyhedra and triangulations of the sphere’’, Geometry & Topology Monographs 1 (1998) 511-549.

[5] G. Brinkmann, J. Goedgebeur and B. McKay, ‘‘The generation of fullerenes’’, Journal of Chemical Information and Modeling 52 (2012) 2910-2918.

[6] O. Angel, I. Benjamini and N. Horesh, ‘‘An isoperimetric inequality for planar triangulations’’, Discrete & Computational Geometry 59 (2018) 802-809.

mathdb:2611.00097v1 17 Nov 2026