


On Primal-Dual Circle Representations

Stefan Felsner

Institut für Mathematik, Technische Universität Berlin, Germany


felsner@math.tu-berlin.de

 <https://orcid.org/0000-0002-6150-1998>

Günter Rote

Institut für Informatik, Freie Universität Berlin, Takustraße 9, 14195 Berlin, Germany

rote@inf.fu-berlin.de

 <https://orcid.org/0000-0002-0351-5945>

Abstract

The Koebe-Andreev-Thurston Circle Packing Theorem states that every triangulated planar graph has a contact representation by circles. The theorem has been generalized in various ways. The most prominent generalization assures the existence of a primal-dual circle representation for every 3-connected planar graph. We present a simple and elegant elementary proof of this result.

2012 ACM Subject Classification Human-centered computing → Graph drawings, Mathematics of computing → Graph algorithms

Keywords and phrases Disk packing, planar graphs, contact representation

Digital Object Identifier 10.4230/OASICS.SOSA.2019.8

Acknowledgements We thank Manfred Scheucher for implementing the algorithm and helping with the figures.

1 Introduction

For a 3-connected plane graph $G = (V, E)$ with face set F , a spherical *primal-dual disk representation* of G consists of two families of disks ($C_x: x \in V$) and ($D_y: y \in F$) on the sphere \mathbb{S}^2 with the following properties (see Figure 1).

- (i) The vertex-disks C_x have pairwise disjoint interiors.
- (ii) The face-disks D_y have pairwise disjoint interiors.

Moreover, for every edge $xx' \in E$ with dual edge yy' (i. e., y and y' are the two faces separated by xx'), the following holds:

- (iii) Circles C_x and $C_{x'}$ touch at a point p .
- (iv) Circles D_y and $D_{y'}$ touch at the same point p .
- (v) The common tangent of C_x and $C_{x'}$ in the point p is perpendicular to the common tangent of D_y and $D_{y'}$ in p .

► **Theorem 1.** *Every 3-connected plane graph G admits a primal-dual disk representation on the sphere. This representation is unique up to Möbius transformations.*

Given a primal-dual disk representation of a graph G , we can use stereographic projection to obtain a primal-dual circle representation in the plane. (In the plane, we stick to the more common terminology of *circle* packings, because a circle defines a unique disk; on the sphere, we have to specify which of the two parts bounded by a circle we mean, and therefore we speak of *disk* packings.) Changing the center of the stereographic projection leads to




© Stefan Felsner and Günter Rote;
licensed under Creative Commons License CC-BY
2nd Symposium on Simplicity in Algorithms (SOSA 2019).

Editors: Jeremy Fineman and Michael Mitzenmacher; Article No. 8; pp. 8:1–8:18

OpenAccess Series in Informatics



OASICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



■ **Figure 1** On the left, a primal-dual disk representation on the sphere, and its stereographic projection to the plane. The intersections of red primal disks with blue dual disks appear in purple. On the right, planes through the boundaries of the red disks define a polytope whose edges “cage” the sphere. The edge skeleton of this polytope is the dual graph (the touching graph of the blue disks).


different primal-dual circle representation in the plane. Figure 2 shows three primal-dual circle representations of K_4 in the plane where the projection center has been chosen as the center of one of the disks, the center of a digon formed by a primal-dual pair of intersecting disks, and the common point of four circles.

As a special case of the previous theorem we obtain the classical circle packing theorem:


► **Theorem 2.** *Every plane graph G admits a circle packing representation, i.e., it is the contact graph of a set of nonoverlapping disks in the plane.*

Our proof of Theorem 1 is constructive, in a sense: It computes a primal-dual circle representation in the plane by a limiting process. For the simplicity of the proof we choose the version where four of the circles are lines and all the other centers of circles are in the rectangle formed by these lines, see the right picture of Figure 2 and the larger example in Figure 3. The theorem then follows using an inverse stereographic projection.

Our proof combines ideas from an unpublished manuscript of Pulleyblank and Rote, from Brightwell and Scheinerman [6] and from Mohar [25]. All these proofs are based on an algorithm for iteratively improving estimates of the circle radii, whose idea goes back to Thurston [37, Section 13.7]. A distinguishing feature of our approach is the symmetric treatment of the primal and the dual family of circles. Four of the radii are already fixed at ∞ , and this helps to reduce the graph-theoretic argument in the proof of convergence to a simple statement about the number of edges of a plane bipartite graph (Lemma 4) and a connectivity argument. The core of the proof requires only 1.5 pages and four chains of equations and inequalities. The layout of the “kites” obtained from the limits of the radii is based on an auxiliary result of independent interest (Lemma 5): when polygonal shapes are



■ **Figure 2** If we project the symmetric primal-dual disk representation of K_4 on \mathbb{S}^2 to the plane by stereographic projection, we get different primal-dual circle representations, depending on the center of projection. In the right picture we see that circles may degenerate to lines.



■ **Figure 3** A larger example. Figure 1 includes a spacial image of this circle representation from a viewpoint on the left side and above the plane of Figure 3.

glued together along edges, local consistency conditions are sufficient to guarantee that these shapes form an overlap-free tiling.

Our simple and elementary proof of Koebe's Theorem, respectively its primal-dual version, is suited for a presentation in a class on Graph Theory, Discrete Geometry, Computational Geometry, or Graph Drawing.

In the next section we give a rather comprehensive account of the history of the theorem and mention some of its applications. The proof of the theorem is given in Section 3. Section 4 is devoted to the proof of Lemma 5.

2 History and Applications of the Theorem

In graph theory the study of circle contact representations can be traced back to the 1970's and 1980's; the term “coin representation” was used there. Wegner [39] and Jackson and Ringel [20] conjectured that every plane graph has a circle representation. The problem was popularized by Ringel [28], who also included it in a textbook from 1990 [19]. In a note written in 1991 [31], Sachs mentions that he found a proof of the circle packing theorem

which was based on conformal mappings. This eventually led him to the discovery that the theorem had been proved by Koebe as early as 1936 [21].

Thurston, in the context of the study of 3-manifolds, proved that any triangulation of the sphere has an associated “circle packing” which is unique up to Möbius transformations [37, Sections 13.6–7]. Thurston noted that this result was already present in earlier work of Andreev [2]. Nowadays the result is commonly referred to as the *Koebe-Andreev-Thurston Circle Packing Theorem*. At a conference talk in 1985, Thurston suggested connections between circle packings and the Riemann Mapping Theorem. A precise version was obtained by Rudin and Sullivan [29]. This line led to the study of discrete analytic functions and other aspects of discrete differential geometry, see to [35, 36, 5] for more on the topic.

In the early 1990’s new proofs of the circle packing theorem were found. Colin de Verdière [7] gave an existential proof based on ‘invariance of domain’; this proof can also be found in [27, Chapter 8] and in the primal-dual setting in an early draft of a book manuscript by Lovász’s [22]. Colin de Verdière [8] gave another proof, which is based on the minimization of a convex function, and he extended circle packings to more general surfaces. Pulleyblank and Rote (unpublished) and Brightwell and Scheinerman [6] gave proofs of the primal-dual version (Theorem 1) based on an iterative algorithm, similar to the proof given in this note. Mohar [24] strengthened the result and proposed an iterative approach that obtains an ε -approximation for the radii and centers in time polynomial in the size of the graph and $\log(1/\varepsilon)$.

Primal-dual circle representations yield *simultaneous orthogonal drawings* of G and its dual G^* , i. e., straight-line drawings of G and G^* such that the outer vertex of G^* is at infinity and each pair of dual edges is orthogonal. The existence of such drawings was conjectured by Tutte [38]. In fact, it follows from Tutte’s “spider-web” embedding method via the Maxwell-Cremona correspondence, which produces a convex piecewise linear surface in \mathbb{R}^3 that vertically projects onto the drawing of G . Polarity will then yield a straight-line embedding of G^* with edges orthogonal to edges of G , see [26] or [30, Section 5]. However, unlike the embeddings implied by the circle theorem, primal-dual edge pairs in this embedding may not intersect.

Another consequence of primal-dual circle representations is known as the *Cage Theorem*. It says that every 3-connected planar graph is the skeleton of a convex 3-polytope such that every *edge* of the polytope is tangent to a given sphere. This strengthening of the Steinitz Theorem is easily derived from Theorem 1, see Figure 1. The Cage Theorem was generalized by Schramm [32], who showed that the sphere that is caged can be replaced by any smooth strictly convex body.

A stunning generalization of the Circle Packing Theorem is the *Monster Packing Theorem* of Schramm [34]. The statement (slightly simplified) is as follows: if each vertex v of a planar triangulation G has a prescribed convex prototype P_v , then there is a contact representation of G where each vertex is represented by a nonnegative homothet of its prototype. Some of these homothets may degenerate to points, but when the prototypes have a smooth boundary, such degeneracies are excluded. Contact representations of planar graphs with other shapes than circles have received quite some attention over the years, for example with triangles [9, 17, 1], rectangles and squares [13, 33], and pentagons and k -gons [16, 15].

The Circle Packing Theorem has been used to prove *separator theorems*. In particular, every planar graph with n vertices can be partitioned into components with at most $n/2$ vertices by removing $O(\sqrt{n})$ vertices. The approach was pioneered by Miller and Thurston and generalized to arbitrary dimensions by Miller, Teng, Thurston, and Vavasis [23]. The planar case is reviewed in [27, Chapter 8]. A slightly simpler proof was given by Har-Peled [18].




Figure 4 (a) A graph and its dual. (b) A cross-centered primal-dual circle packing for this graph. The areas of the primal disks are shaded. The two vertices ξ and ξ' are represented by “degenerate disks”: disjoint halfplanes bounded by C_ξ and $C_{\xi'}$, which “touch at infinity”. (c) The straight-line drawing of the two graphs induced by the circle packing; the centers ξ and ξ' of the degenerate disks lie infinitely far away to the left and to the right. The edge $\xi\xi'$ is not represented at all. The same holds for η and η' and the dual edge between them.

Bern and Eppstein [4, 11] relate circle packings to mesh generation techniques.

Not surprisingly, the theorem also has applications in Graph Drawing. Eppstein [12] used circle representations to prove that every planar graph with maximum degree 3 has a *Lombardi drawing*: a drawing in which the edges are drawn as circular arcs, meeting at equal angles at each vertex. Felsner, Igamberdiev, Kindermann, Klemz, Mchedlidze, and Scheucher [14] used circle representations to show that 3-connected planar graphs have planar *strongly monotone drawings*, i. e., straight-line drawings such that for any two vertices u, v there is a path which is monotone with respect to the connecting line of u and v .


3 Primal-Dual Circle Representation: The Proof

Let $G = (V, E)$ be a 3-connected plane graph with face set F and let $\xi\xi'$ be an edge in E with dual edge $\eta\eta'$, i. e., η and η' are the two faces on the sides of $\xi\xi'$. A *cross-centered primal-dual circle representation* of G with *central cross* $\xi\xi', \eta\eta'$ consists of two vertical lines C_ξ and $C_{\xi'}$, two horizontal lines D_η and $D_{\eta'}$, and two families of circles ($C_x: x \in V \setminus \{\xi, \xi'\}$) and ($D_y: y \in F \setminus \{\eta, \eta'\}$) with the following five properties, see Figures 3 and 4b for examples:

- (i) The vertex-circles C_x have pairwise disjoint interiors and are contained in the vertical strip between C_ξ and $C_{\xi'}$.
- (ii) The face-circles D_y have pairwise disjoint interiors and are contained in the horizontal strip between D_η and $D_{\eta'}$.

Moreover, for every edge $xx' \in E$ with $xx' \neq \xi\xi'$ and with dual edge yy' (i. e., y and y' are the two faces separated by xx'), the following holds:

- (iii) C_x and $C_{x'}$ are tangent at a point p with common tangent line $t_{xx'}$.
- (iv) D_y and $D_{y'}$ are tangent at the same point p with common tangent line $t_{yy'}$.
- (v) The lines $t_{xx'}$ and $t_{yy'}$ are orthogonal.



■ **Figure 5** Three primal-dual straight-line drawings of K_4 . They correspond to the primal-dual circle representations of Figure 2.

► **Theorem 3.** *Every 3-connected plane graph G admits a cross-centered primal-dual circle representation. Moreover, for a given central cross $\xi\xi', \eta\eta'$, this representation is unique up to scaling, translation, and horizontal or vertical reflections.*

Theorem 1 follows from Theorem 3 via inverse stereographic projection.

We give first an outline of the proof. A primal-dual circle representation of G induces a straight-line drawing of G and a straight-line drawing of the dual. Superimposing the two drawings yields a plane drawing whose faces are special quadrangles called kites, see Figures 6 and 7. After guessing radii for the circles, the shapes of the kites are determined. It is then checked whether the angles of kites meeting at a vertex sum up to 2π . If at some vertex the angle sum differs from 2π , the radii are changed to correct the situation. The process is designed to make the radii converge and to make the sum of angles meet the intended value at each vertex. The second part of the proof consists of showing that the kites corresponding to the final radii can be laid out to form a tessellation, thus giving the centers of a primal-dual circle representation of G .


Proof of Theorem 3. Given a cross-centered primal-dual circle representation of G we can use the centers of the circles C_x to obtain a planar straight-line drawing of G , see Figure 4c. Edges containing ξ or ξ' are represented by horizontal rays to the left and right respectively. The edge $\xi\xi'$ is missing. Similarly, the centers of the circles D_y yield a planar straight-line drawing of G^* with edges containing one of η and η' being represented by vertical rays.

For example, from the primal-dual circle representations of K_4 of Figure 2, we obtain plane straight-line drawings of K_4 and its dual that are displayed in Figure 5. The rightmost of these drawings corresponds to a cross-centered primal-dual circle representation of K_4 .


3.1 Kites

If we overlay the drawings of G and G^* , we get a partition of the plane into *kites*: quadrilaterals with right angles at two opposite vertices and a line of symmetry through the other two vertices. Figure 6 shows an example, and Figure 7 shows a generic kite. In the cross-centered case, there are *degenerate kites*: rectangular strips that are unbounded in one direction. They have a vertex with a 180° -angle in the midpoint of the only bounded edge. In addition, we have four quadrants, which can be regarded as *exceptional kites*. The bounded kites fill a rectangle between C_ξ , $C_{\xi'}$, D_η , and $D_{\eta'}$.

The kites are in bijection with the incident pairs (x, y) , where x is a primal vertex and y is a dual vertex. Since the involved circles or lines intersect orthogonally, the kite of x and y is completely determined by the radii r_x of C_x and r_y of D_y . (In the case of a line the radius



■ **Figure 6** The tessellation of the plane into kites obtained from the example in Figure 4c. Four kites are shaded, among them a degenerate kite (semi-infinite strip) and an exceptional kite (quadrant).



■ **Figure 7** The kite corresponding to the incident vertex-face pair x, y .

is ∞ .) For bounded kites, the angles at x and y are given by

$$\alpha_{xy} = 2 \arctan \frac{r_y}{r_x} \quad \text{and} \quad \alpha_{yx} = 2 \arctan \frac{r_x}{r_y}. \tag{1}$$

We extend these formulas to degenerate kites by taking the limits:

$$\alpha_{uw} = \begin{cases} 0, & \text{if } r_w \neq \infty \text{ and } r_u = \infty \\ \pi, & \text{if } r_w = \infty \text{ and } r_u \neq \infty \end{cases} \tag{2}$$

Then we have


$$\alpha_{uw} + \alpha_{wu} = \pi$$

for all pairs (u, w) forming a bounded or degenerate kite. We don't define the angles for the four exceptional kites because this would involve the undetermined expression $\frac{\infty}{\infty}$.

3.2 The Angle Graph

The number and combinatorial structure of the kites is captured by the angle graph. The *angle graph*, or *vertex-face incidence graph*, of a plane graph $G = (V, E)$ is the graph $G^\diamond = (U, K)$ whose node set $U = V \cup F$ represents both the vertices and faces of G , see Figure 8b. Its edges xy are the pairs with $x \in V$ and $y \in F$ that are incident in G , i. e., x is a vertex on the boundary of the face y . These edges are in bijection with the kites. The graph G^\diamond is plane and bipartite. Its faces corresponds to the edges of G , and they are 4-gons, i. e., G^\diamond is a quadrangulation. We choose the face $f_o = \xi\eta\xi'\eta'$ containing the four elements of the central cross as the outer face of G^\diamond . We denote its nodes by $U_o = \{\xi, \xi', \eta, \eta'\}$ and the remaining nodes by $U_{in} = U \setminus U_o$. We denote the four edges of the outer face by $K_o = \{\xi\eta, \eta\xi', \xi'\eta', \eta'\xi\}$.

An important property of the angle graph is that it cannot have a separating 4-cycle: If the nodes $xyx'y'$ with $x, x' \in V$ and $y, y' \in F$ would form some separating 4-cycle in G^\diamond , then x, x' would be a separating vertex pair in G , contradicting the 3-connectedness assumption for G .



■ **Figure 8** (a) The plane graph G from Figure 4a, (b) its angle graph G^\diamond , (c) the primal-dual completion $(G^\diamond)^\diamond$. The faces of this graph represent the kites, including the unbounded and exceptional kites.

We will need the following well-known basic fact about bipartite plane graphs, which is a consequence of Euler’s formula. For completeness, we include the detailed proof in Appendix A.

► **Lemma 4.** *A simple bipartite plane graph with $|S| \geq 4$ nodes has at most $|E| \leq 2|S| - 4$ edges, with equality if and only if the graph is connected and every face is a quadrilateral with four distinct vertices.*

In particular, G^\diamond contains $|K| = 2|U| - 4$ edges.

3.3 Angle Sums

We now come to the core of the argument. A hypothetical primal-dual circle representation of G contains a point for each $u \in U_{\text{in}}$. This point is fully surrounded by its incident kites. Hence, for every $u \in U_{\text{in}}$ we have:

$$\sum_{w: uw \in K} \alpha_{uw} = 2\pi \tag{3}$$

We now look at an arbitrary assignment $r: U_{\text{in}} \rightarrow \mathbb{R}_{>0}$ of radii. Additionally, we define $r_u = \infty$ for each $u \in U_o$. We can then form the corresponding kites and compute the angles according to (1) and (2). In particular, by (2), the degenerate kites have the correct angles:

$$\alpha_{uw} = \pi \text{ and } \alpha_{wu} = 0 \text{ whenever } u \in U_{\text{in}} \text{ and } w \in U_o. \tag{4}$$

Denote the angle sum at $u \in U_{\text{in}}$ by $\alpha_u = \alpha_u(r) = \sum_{w: uw \in K} \alpha_{uw}$. We want to find radii r such that $\alpha_u(r)$ becomes equal to the *target angle* 2π for all $u \in U_{\text{in}}$ in order to fulfill (3). Later we will show that a collection of radii with this property induces a primal-dual circle representation.

We first show that any choice of radii attains the correct target angles *on average*:

$$\sum_{u \in U_{\text{in}}} (\alpha_u(r) - 2\pi) = 0 \tag{5}$$

This follows from the following computation:

$$\begin{aligned} \sum_{u \in U_{\text{in}}} \alpha_u(r) &= \sum_{\substack{uw \in K \\ u, w \in U_{\text{in}}}} (\alpha_{uw} + \alpha_{wu}) + \sum_{\substack{uw \in K \\ u \in U_{\text{in}}, w \in U_o}} \alpha_{uw} = \sum_{\substack{uw \in K \\ u, w \in U_{\text{in}}}} \pi + \sum_{\substack{uw \in K \\ u \in U_{\text{in}}, w \in U_o}} \pi = \pi|K \setminus K_o| \\ &= \pi(|K| - 4) = \pi(2|U| - 8) = 2\pi(|U| - 4) = 2\pi|U_{\text{in}}| \end{aligned}$$

As a consequence, whenever $\alpha_u(r) \neq 2\pi$ for some u , the following two sets are both nonempty:

$$U_- = \{u \in U_{\text{in}} : \alpha_u(r) < 2\pi\} \quad \text{and} \quad U_+ = \{u \in U_{\text{in}} : \alpha_u(r) > 2\pi\}$$

If we increase the radius r_u of a node $u \in U_+$, leaving all remaining radii fixed, we observe from (1) that for every incident edge $uw \in K$, the angle α_{uw} decreases strictly to 0 as $r_u \rightarrow \infty$, with the possible exception of a single neighbor $w \in U_o$ with fixed angle $\alpha_{uw} = \pi$ according to (4). Hence, we can increase r_u to the unique value where $\alpha_u(r) = 2\pi$.

3.4 Iteration and Convergence

The workhorse of the proof is the following infinite iteration.

$$\left. \begin{array}{l} \text{repeat forever:} \\ \quad \text{for each } u \in U_{\text{in}}: \\ \quad \quad \text{if } u \in U_+ \text{ then increase } r_u \text{ to reduce } \alpha_u(r) \text{ to } 2\pi \end{array} \right] \quad (6)$$

We will show that, for an arbitrary positive starting assignment, the radii converge to some limiting assignment \hat{r} , and this will imply that $\alpha_u(\hat{r}) = 2\pi$ for all $u \in U_{\text{in}}$.

Since radii can never decrease and every bounded monotone sequence is convergent, it is enough to show that the set of “divergent” nodes $D = \{u \in U : \lim r_u = \infty\}$ contains no other nodes than the four nodes of U_o . (The nodes $u \in U_o$ have $r_u = \infty$ fixed and are included in D by definition.)

The increase of r_u decreases the angle sum α_u , but not below 2π . It increases the angles at adjacent nodes, and it may hence cause some $w \in U_-$ to move to U_+ . A transition from U_+ to U_- , however, is impossible. It follows that some node u_0 must belong to U_- indefinitely unless the iteration comes to a halt with $U_- = U_+ = \emptyset$. Thus, as a consequence of the built-in behavior of the iteration, (a) U_- is disjoint from D from some time on, and (b) D is a proper subset of U .

Let us look at the subgraph $G^\circ[D]$ of G° induced by the divergent nodes. In order to apply Lemma 4, we will show that $G^\circ[D]$ has at least $2|D| - 4$ edges.

First, we wait for U_- to become disjoint from D . From that point onwards,

$$\sum_{u \in D \setminus U_o} \alpha_u(r) \geq \sum_{u \in D \setminus U_o} 2\pi = 2\pi|D \setminus U_o| = (2|D| - 8)\pi. \quad (7)$$

On the other hand, if $u \in D$ and $w \in U \setminus D$, then α_{uw} converges to 0 according to (1). Thus, in addition to (7), the inequality $\alpha_{uw} \leq 1/|U|^2$ will eventually hold for each such edge. Bounding these edges separately from the others, we get the following inequality at this point of the iteration:

$$\begin{aligned} \sum_{u \in D \setminus U_o} \alpha_u(r) &\leq |U|^2 \cdot \frac{1}{|U|^2} + \sum_{\substack{\text{kite with } x, y \in D \\ x \notin U_o \text{ or } y \notin U_o}} (\alpha_{xy} + \alpha_{yx}) = 1 + \sum_{\substack{xy \text{ edge of } G^\circ[D] \\ xy \notin K_o}} \pi \\ &= 1 + (|E(G^\circ[D])| - 4)\pi, \end{aligned} \quad (8)$$

where $E(G^\circ[D])$ is the edge set of $G^\circ[D]$. Comparing (7) and (8) gives $|E(G^\circ[D])| \geq 2|D| - 4 - 1/\pi$ and therefore $|E(G^\circ[D])| \geq 2|D| - 4$.

Since $U_o \subseteq D$ by definition and thus $|D| \geq 4$, we can apply Lemma 4. We conclude that $G^\circ[D]$ is connected and its faces are simple 4-cycles.

The outer face of $G^\circ[D]$ is the quadrilateral f_o formed by U_o . Our goal is to show that $D = U_o$ and $G^\circ[D]$ consists just of the single 4-cycle f_o . Since $G^\circ[D]$ is a proper subgraph of G° , $G^\circ[D]$ has some face f that is not a face of G° . This face f is an inner face of $G^\circ[D]$ because the outer face of $G^\circ[D]$ agrees with f_o . Suppose for contradiction that f does not coincide with the interior face bounded by f_o . Then it would form a separating 4-cycle in G° : it would contain nodes of U both in its interior (because it is not a face of G°) and in its exterior (because some nodes of f_o lie there). Since separating 4-cycles are excluded, we have shown that $D = U_o$.

This means that all radii r_u for $u \in U_{\text{in}}$ converge to some limits, which we denote by \hat{r}_u . It follows that all angles α_{uw} and all angle sums $\alpha_u(r)$ converge as well, and by the working of the iteration (6), their limits $\alpha_u(\hat{r})$ are bounded by $\alpha_u(\hat{r}) \leq 2\pi$. Since $\sum_{u \in U_{\text{in}}} (\alpha_u(\hat{r}) - 2\pi) = 0$ by (5), we must have $\alpha_u(\hat{r}) = 2\pi$ for all $u \in U_{\text{in}}$.

3.5 Uniqueness

We show that the radii are unique up to scaling. Let r and r' be two vectors of radii such that $\alpha_r(u) = \alpha_{r'}(u) = 2\pi$ for all $u \in U_{\text{in}}$. Scaling allows to assume that $r_{u_0} = r'_{u_0}$ for some $u_0 \in U_{\text{in}}$. Consider the set $S = \{u \in U_{\text{in}} : r_u > r'_u\}$ and observe that $u_0 \in \bar{S} = U_{\text{in}} \setminus S$.


$$\begin{aligned} |S| \cdot 2\pi &= \sum_{u \in S} \alpha_u(r) = \sum_{\substack{uw \in K \\ u, w \in S}} (\alpha_{uw}(r) + \alpha_{uw}(r)) + \sum_{\substack{uw \in K \\ u \in S, w \in U_o}} \alpha_{uw}(r) + \sum_{\substack{uw \in K \\ u \in S, w \in \bar{S}}} \alpha_{uw}(r) \\ &= \sum_{\substack{uw \in K \\ u, w \in S}} 2\pi + \sum_{\substack{uw \in K \\ u \in S, w \in U_o}} \pi + \sum_{\substack{uw \in K \\ u \in S, w \in \bar{S}}} \alpha_{uw}(r) \end{aligned}$$

Thus, the last sum has a constant value, independent of the radii r . However, if we change the radii from r to r' , then, by (1), every term $\alpha_{uw}(r)$ in the last sum increases, because $r_u > r'_u$ and $r_w \leq r'_w$. This means that the set of edges over which the sum is taken must be empty. In other words, if $w \in \bar{S}$, then every neighbor $u \in U_{\text{in}}$ of w must also belong to \bar{S} . Since $u_0 \in \bar{S}$ and $G^\circ[U_{\text{in}}]$ is connected, S must be empty. By a symmetric argument, the set $S' = \{u \in U_{\text{in}} : r'_u > r_u\}$ is empty as well, and this proves uniqueness of the radii r up to scaling.

The radii determine shape and size of the kites. Below we show that the kites can be laid out to form a tessellation of the plane. The line C_ξ is vertical, hence, the tessellation is unique up to scaling, translation, and horizontal or vertical reflection. Since the tessellation determines the circles, uniqueness carries over to the cross-centered primal-dual circle representation with fixed central cross.

3.6 Laying out the Kites

We now show that the kites defined by the limiting radii \hat{r} can be laid out in the plane with the intended side-to-side contacts. Figure 9 illustrates this task. We will use Lemma 5 below, which warrants the existence of such a layout if certain local matching conditions are fulfilled. We invite the reader to skip forward and read the statement of Lemma 5 in Section 4. We apply this lemma to the graph H of the vertices and edges of the *bounded* kites, see Figures 6 and 9. This graph is a subgraph of the *primal-dual completion* of $G = (V, E)$ (which, by the



■ **Figure 9** Laying out the kites.

way, is nothing but the angle graph $(G^\diamond)^\diamond$ of the angle graph of G , see Figure 8c). The nodes of H are vertices, faces and edges of G . Specifically, $V_H = (V \setminus \{\xi, \xi'\}) \cup (F \setminus \{\eta, \eta'\}) \cup (E \setminus \{\xi\xi'\})$, and the edges of H are the pairs $(z, e) \in ((V \setminus \{\xi, \xi'\}) \cup (F \setminus \{\eta, \eta'\})) \times (E \setminus \{\xi\xi'\})$ with z incident to the edge $e \in E$ in G . Each bounded face of H is a quadrilateral representing a bounded kite. It contains one node from V , one node from F , and two nodes from E .

The 3-connectivity of G and of G^* easily implies that H is 2-connected, as required for Lemma 5. (In fact, the first proof of Lemma 5 shows that connectedness of H is sufficient, provided that the outer face is a simple cycle.) We know that two adjacent kites fit together locally because they have the same edge lengths by construction: these lengths are defined by the same radius r_u . This is condition (iii) of Lemma 5.

Moreover, as we have shown, the kites around a vertex $u \in U_{\text{in}}$ form a complete angle of $\alpha_u(\hat{r}) = 2\pi$. Every right angle of a kite, if it is an interior node of H , is complemented by the right angles of three other kites to again form a complete angle of 2π . This is condition (i) of Lemma 5.

The vertices of H incident to the outer face of H are either points where one or two right angles of kites meet, forming an angle of 90° or 180° , or they are nodes $u \in U_{\text{in}}$ which are adjacent in G^\diamond to some node $w \in U_o$, forming an angle $\alpha_{uw} = \pi$ by (4). Since this angle is not part of H , the incident angles in H around u sum up to $\alpha_u(\hat{r}) - \alpha_{uw} = 2\pi - \pi = \pi$. In summary, the angle sums of nodes incident to the outer face of H are either 90° or 180° , and thus condition (ii) for Lemma 5 is fulfilled, and moreover, the layout of the bounded kites must form a rectangle R . The unbounded kites can be attached edge by edge along the boundary of R . This yields the claimed cover of the whole plane.

3.7 Constructing the Circle Representation

Finally, we derive a cross-centered primal-dual circle representation from the layout of the kites. The kites induce a straight-line drawing of G and a straight-line drawing of the dual G^* with the edges incident to one node of U_o being rays and edges induced by U_o omitted. For every primal-dual pair xx', yy' of edges the point p where xx' and yy' meet is a right angle in each of the four involved kites. This implies (v).

For a node $u \in U_{\text{in}}$, consider the set of kites containing u . These kites can be put together in the cyclic order given by the rotation of u in G^\diamond to form a polygon P_u surrounding u ,

because $\alpha_u(\hat{r}) = 2\pi$. By the geometry of the kites, all edges incident to $u \in V \cap U_{\text{in}}$ have the same length \hat{r}_u , and the circle C_u of radius \hat{r}_u centered at u is inscribed in P_u and touches P_u at the common corners of neighboring kites. For $u \in \{\xi, \xi'\}$, the polygon P_u obtained by gluing the corresponding unbounded kites is a halfplane and the vertical line C_u goes through the right-angle corners of the involved kites. From the incidences of the kites, and since the polygons P_u for $u \in V$ are pairwise disjoint, we obtain that the family $(C_x : x \in V)$ satisfies (i) and (iii).

Dually, the polygons P_u corresponding to $u \in F$ also tile the plane, and the family $(D_y : y \in F)$ satisfies Properties (ii) and (iv). This concludes the proof of Theorem 3. ◀

4 Tiling a Convex Polygon

The following lemma says that certain local consistency conditions around each vertex and along each edge are sufficient to guarantee a global nonoverlapping layout of faces with prescribed shapes.

► **Lemma 5.** *Let H be a 2-connected plane graph (possibly drawn with curved edges). For each bounded face f of H , a simple polygon P_f is given whose corners are labeled with the vertices from the boundary of f in the same cyclic order. Denote the corner of P_f labeled with v by p_{fv} and the angle of P_f at this corner by β_{fv} . For each vertex v , let F_v denote the set of incident bounded faces. We assume the following conditions:*

- (i) $\sum_{f \in F_v} \beta_{fv} = 2\pi$ for every inner vertex v .
 - (ii) $\sum_{f \in F_v} \beta_{fv} \leq \pi$ for every vertex v on the outer face.
 - (iii) $\|p_{fv} - p_{fw}\| = \|p_{gv} - p_{gw}\|$ for every inner edge vw of H with incident faces f and g .
- Then there is a crossing-free straight-line drawing of H in which every bounded face f can be obtained from P_f by a rigid motion, i. e., translation and rotation.*


Lemma 5 or similar statements have been used explicitly or implicitly in other situations, beyond the context of circle packings. For example in [16, 15] it is used in the context of contact representations with pentagons and k -gons. In fact, our second proof of Lemma 5 slightly generalizes a proof from [16].

We give two proofs of Lemma 5, a more geometric one and a more combinatorial one.

Proof 1. We proceed by induction on the number of interior vertices. The tool that we need is that every simple polygon can be subdivided into convex pieces, or if we want, into triangles, by inserting diagonals between its vertices.

To make the induction go through, we have to strengthen the assumption of the lemma and require that each polygon P_f is convex. This can be achieved by inserting diagonals and subdividing it into convex pieces. On the other hand, we don't require H to be 2-connected, and we even allow H to have multiple edges. (Showing beforehand that separating vertices or multiple edges cannot actually occur would be more tedious.) We do however maintain the requirement that H is connected and that the outer face is a simple cycle.

The inductive step proceeds as follows: If there is an interior vertex w , we take the k faces f_1, \dots, f_k incident to w and place the corresponding polygons P_{f_1}, \dots, P_{f_k} successively around the origin, see Figure 10. By condition i, they completely surround the origin. By convexity of the faces, each face is confined within its own sector, disjoint from the other sectors. Thus the faces don't overlap, and their union forms a simple polygon P that contains w in its interior. (It is star-shaped around w .) We triangulate P geometrically. We remove w from the graph H , and we insert the appropriate new edges into H , replacing the faces f_1, \dots, f_k by the new triangular faces, with the triangles as the corresponding polygons.



■ **Figure 10** Triangulating the union of the faces surrounding a vertex w .

(Here is the point where multiple edges could conceivably be created.) When performing this replacement, by construction, the angle sum $\sum_{f \in F_v} \beta_{fv}$ remains the same around every vertex $v \neq w$ (conditions i and ii), and the new polygons have matching edge lengths, both among themselves and with the previously existing faces (condition iii). The resulting graph H' has one interior vertex less, and it is still connected, because the boundary vertices of P , which include all neighbors of w , are connected through the boundary edges of P . By induction, its faces can be laid out in the plane without overlap and with adjacent faces touching along their common edges. The triangular faces that were added form a polygon that is congruent to P . Cutting the polygon P into the faces P_{f_1}, \dots, P_{f_k} from which it was originally formed, we obtain the position for w and a drawing of the original graph H .

In the base case of the induction, there are no interior vertices. We simply merge adjacent polygons pairwise along their common interior edge. By condition ii, the two new resulting interior angles are $\leq \pi$. Hence the polygon resulting from each merge is again a convex polygon. In the end, we have a single convex polygonal face, and there is nothing left to prove. ◀

Proof 2. The proof proceeds in four steps. (A) In the first step, we define positions for every face. Let H^* be the dual graph of H without the vertex corresponding to the outer face of H . Let S be a spanning tree of H^* . Then by (iii) we can glue the polygons P_f of all bounded faces f of H together along the edges of S . This determines a unique position for every polygon, up to a global motion.

(B) Since a vertex belongs to several faces, this layout might prescribe several inconsistent positions for the same vertex. In the second step, we show that such contradictory constraints do not arise, and each vertex has a unique position. For the edges of S we already know that the polygons of the two incident faces touch in such a way that corners corresponding to the same vertex coincide. For the edges of the complement \bar{S} of S we still need to show this. The set \bar{S} , considered as a subset of the edges of H , forms a forest. Let v be a leaf of this forest that is an inner vertex of H , and let e be the edge of \bar{S} incident to v . Then for all incident edges $e' \neq e$ of v we already know that the polygons of the two incident faces of e' touch in the right way. But then also the two polygons of the two incident faces of e touch in the right way because v fulfills property (i). Since the set of edges we still have to check remains a forest, we can iterate this process until all inner edges of H are checked. After gluing all the polygons P_f , every vertex v has an unambiguous position.

(C) Let P_o be the cycle formed by the boundary edges of H in this drawing. As the third step, we will show that P_o forms a convex polygon. We know from property (ii) that when we traverse P_o with the interior on its left, we make only left turns, but it is conceivable

that P_0 makes several loops and intersects itself. We show that this is not the case. Let $H = (V, E)$ with face set F , and let V_o be the set of outer vertices of H . Denoting $d_o = |V_o|$, we claim that

$$\sum_{v \in V_o} \sum_{f \in F_v} \beta_{fv} = (d_o - 2)\pi. \tag{9}$$

To see this, we express the angle sum B over all polygons P_f in two different ways. Property (i) gives

$$B = \sum_{v \in V \setminus V_o} \sum_{f \in F_v} \beta_{fv} + \sum_{v \in V_o} \sum_{f \in F_v} \beta_{fv} = (|V| - d_o)2\pi + \sum_{v \in V_o} \sum_{f \in F_v} \beta_{fv}. \tag{10}$$

On the other hand, let us denote the degree of each bounded face f by d_f . Then the angle sum of P_f is $(d_f - 2)\pi$. Summing this over all bounded faces gives

$$B = \sum_f (d_f - 2)\pi = [(2|E| - d_o) - 2(|F| - 1)]\pi = (|E| - |F| + 2 - d_o)2\pi + (d_o - 2)\pi. \tag{11}$$

Comparing the right-hand sides of (10) and (11), Euler's Formula gives the claim (9).

Thus, the sum of angles at the outer vertices has just the right value for a d_o -gon. Hence, the image P_o of the boundary edges is a convex polygon and therefore nonintersecting.

(D) We finally prove that the glued polygons P_f tile the interior of P_o without holes or overlap. Since we will refer to this argument later, we formulate it as a separate lemma:

► **Lemma 6.** *Let H be a 2-connected plane graph (possibly drawn with curved edges). Let H' be a straight-line drawing of H in the plane (possibly with crossings), with the following properties:*


- (a) *For each bounded or unbounded face f of H , the edges of the face cycle in H' form a simple polygon P_f .*
- (b) *For each inner edge e with incident faces f and f' , the interior regions of P_f and $P_{f'}$ lie on different sides of e .*

Then H' contains no crossings.

In our case, the assumptions of this lemma are fulfilled: For the bounded faces, Property (a) holds by assumption, and for the unbounded face, it has just been established in Step (C). Property (b) has been established in Step (B). Thus, our second proof of Lemma 5 is complete once we prove Lemma 6:

Proof of Lemma 6. Let P_o denote the outer boundary, which is a simple polygon by assumption (a). We prove that the polygons P_f tile the interior of P_o without holes or overlap, using a covering number argument. Consider a point p on the plane which does not lie on an edge of one of the polygons. We can move this point to infinity along a straight ray which avoids all polygon vertices. We keep track of the number $X(p)$ of polygons in which p is contained. Whenever we cross an edge e of some polygon, we leave one polygon and enter another polygon, keeping $X(p)$ constant, unless e is an edge of P_o . In the last case, $X(p)$ changes by ± 1 in the correct way. This argument remains valid if we cross several edges simultaneously (but we are about to show that this situation never occurs). Since $X(p) = 0$ when p is far away outside all polygons, it follows that all points p , except those on polygon boundaries, have $X(p) = 1$ if they lie inside P_o , and $X(p) = 0$ if they lie outside P_o . Consequently, the union of the polygons P_f is the polygon bounded by P_o , and the polygons cover it without overlap. ◀

◀



■ **Figure 11** After placing a few faces, the outer cycle P might cross itself.

Our second proof of Lemma 5 generalizes to shapes P_f with curved edges: The matching condition (iii) of the lemma must then be strengthened in an appropriate way. The angle conditions corresponding to (i) and (ii) are not so straightforward to formulate, depending on the generality of the allowed boundaries, and an additional constraint is required to guarantee an overall convex shape (or at least a shape without self-overlap).

4.1 Comparison with Other Proofs

4.1.1 Lemma 5

We are aware of only one other proof of a statement like Lemma 5 in the literature: Brightwell and Scheinerman [6] (who did not formulate it as a separate lemma) gave a proof that is similar in spirit to our second proof. They successively place the polygons in some appropriate order, such that the boundary P of the placed polygons is always a simple cycle in the graph. In this way, what done in two separate Steps (A) and (B) in our proof, the placement of the faces and ensuring the consistency of the vertex positions, is achieved together. Step (C) is not necessary in their case, because the outer face is a triangle, and therefore it is automatically non-intersecting. The same statement actually applies to the application of Lemma 5 in the proof of our main theorem in Section 3.6, because the outer face is a rectangle in this case.

Step (D) is omitted in [6]. However, some argument like Lemma 6 is necessary, as illustrated by a hypothetical situation in Figure 11: The shaded faces have already been drawn, in a locally consistent way. While the outer boundary P forms a simple cycle in the graph, it self-intersects in the plane. It is conceivable that such a boundary can be completed with the remaining polygons to a locally consistent where the outer boundary becomes, say, the rectangular outline. It requires a proof that this cannot occur.

The existence of an appropriate face order for the face placement is assumed without justification in [6]. It is not hard to show that such an order can be chosen greedily: A proper subset of bounded faces enclosed by a simple cycle P can always be extended by an additional face f , so that $f \cup P$ is a connected curve, and the boundary remains a simple cycle.

(Alternatively, one can choose an edge st on the outer face and use a “bipolar orientation” (or an “ s - t -numbering”), which is known to exist for any 2-connected graph. This results in an acyclic orientation of the dual graph, and any linear extension of this acyclic orientation is a suitable face order. We are grateful to Therese Biedl (private communication) for this observation.)

4.1.2 Lemma 6

Instead of Condition (b) of Lemma 6, we can stipulate that all faces are oriented consistently:

(b') *Every bounded face cycle of H is oriented in the same way in H and H' .*

Up to reflecting the drawing and reversing the orientation of *every* face, this is equivalent to Condition (b): Condition (b') clearly implies Condition (b). On the other hand, Condition (b) implies that adjacent bounded faces must be oriented consistently. (When the interior of the face is on the left, they must both be ordered clockwise or both counterclockwise.) Since the graph is 2-connected, the dual graph of the bounded faces is connected, and hence all bounded faces have to be ordered consistently.

A slightly different condition has been used by Devillers, Liotta, Preparata and Tamassia [10, Lemma 16]. Their lemma states that the following condition, in conjunction with Property (a), is sufficient to guarantee a non-crossing drawing:

(b'') *The cyclic order of the edges around every vertex is the same in H and H' .*

In contrast to Lemma 6, where the outer face of the initial drawing H is fixed and has to remain unchanged in H' , this variation gives up the a-priori distinction between inner faces and the outer face. From the cyclic order in (b''), one can infer the face structure by walking around each face boundary, keeping the area of the face always to the left. The area of the face might turn out to be the inner (bounded) or the outer (unbounded) region bounded by the face cycle, depending on the orientation (counterclockwise or clockwise).

Our proof of Lemma 6 can be adapted to this situation: The regions P_f denote the (bounded or unbounded) face areas, and the goal is to show that these regions tile the whole plane, i. e., $X(p) = 1$ everywhere. To prove this, one has to establish that there is exactly one unbounded face. This can be shown by an account of the angle sums, like in Step (C) of our second proof of Lemma 5.

Lemma 16 of Devillers et al. [10] is stated for connected graphs and not just 2-connected graphs. In this case, face cycles are no longer simple polygons. The proof in [10] is sketchy, as acknowledged by one of the authors (private communication), and we could not fill all gaps. It is fortunate that Lemma 6 offers an alternate approach.

Di Battista and Vismara [3, Lemma 4.5] have previously proved another variation of the lemma where all interior faces are triangles. In this special case, condition (a) becomes trivial for the interior faces. Instead of condition (a), the only requirement in addition to (b'') is that and the boundary of the outer face turns only in one direction (cf. condition (ii) of Lemma 5 and the discussion in Step (C) of our second proof of Lemma 5). Their proof is by induction on the number of interior vertices, and the main argument proceeds by retriangulating the hole that is left after removing an interior vertex, like our first proof of Lemma 5.

References

- 1 Nieké Aerts and Stefan Felsner. Straight Line Triangle Representations. *Discr. and Comput. Geom.*, 57:257–280, 2017. doi:10.1007/s00454-016-9850-y.
- 2 E. M. Andreev. Convex polyhedra in Lobačevskiĭ spaces. *Mat. Sb. (N.S.)*, 81 (123):445–478, 1970. English: Math. USSR, Sb. 10, 413–440 (1971).
- 3 Giuseppe Di Battista and Luca Vismara. Angles of planar triangular graphs. *SIAM J. Discrete Math.*, 9(3):349–359, August 1996.
- 4 Marshall W. Bern and David Eppstein. Quadrilateral Meshing by Circle Packing. *Int. J. Comput. Geometry Appl.*, 10:347–360, 2000.

- 5 Alexander I. Bobenko and Boris A. Springborn. Variational principles for circle patterns and Koebe's theorem. *Trans. Amer. Math. Soc.*, 356:659–689, 2004.
- 6 G. R. Brightwell and E. R. Scheinerman. Representations of Planar Graphs. *SIAM J. Discr. Math.*, 6(2):214–229, 1993.
- 7 Yves Colin de Verdière. Empilements de cercles: convergence d'une méthode de point fixe. *Forum Math.*, 1:395–402, 1989.
- 8 Yves Colin de Verdière. Un principe variationnel pour les empilements de cercles. *Invent. Math.*, 104:655–669, 1991.
- 9 Hubert de Fraysseix, Patrice Ossona de Mendez, and Pierre Rosenstiehl. On Triangle Contact Graphs. *Comb., Probab. and Comput.*, 3(02):233–246, 1994.
- 10 Olivier Devillers, Giuseppe Liotta, Franco P. Preparata, and Roberto Tamassia. Checking the convexity of polytopes and the planarity of subdivisions. *Comput. Geom.*, 11(3-4):187–208, 1998. doi:10.1016/S0925-7721(98)00039-X.
- 11 David Eppstein. Diamond-Kite Meshes: Adaptive Quadrilateral Meshing and Orthogonal Circle Packing. In *Proc. Mesh. Roundtable*, pages 261–277. Springer, 2012.
- 12 David Eppstein. A Möbius-invariant power diagram and its applications to soap bubbles and planar Lombardi drawing. *Discrete Comput. Geom.*, 52:515–550, 2014.
- 13 Stefan Felsner. Rectangle and Square Representations of Planar Graphs. In J. Pach, editor, *Thirty Essays in Geometric Graph Theory*, pages 213–248. Springer, 2013. doi:10.1007/978-1-4614-0110-0_12.
- 14 Stefan Felsner, Alexander Igamberdiev, Philipp Kindermann, Boris Klemz, Tamara Mchedlidze, and Manfred Scheucher. Strongly Monotone Drawings of Planar Graphs. In *Proc. 32nd Intern. Symp. Comput. Geom. (SoCG 2016)*, volume 51 of *LIPICs*, pages 37:1–15, 2016.
- 15 Stefan Felsner, Hendrik Schrezenmaier, and Raphael Steiner. Equiangular polygon contact representations, 2017. page.math.tu-berlin.de/~felsner/Paper/kgons.pdf.
- 16 Stefan Felsner, Hendrik Schrezenmaier, and Raphael Steiner. Pentagon Contact Representations. *Electronic J. Combin.*, 25(3):article #P3.39, 38 pp., 2018.
- 17 Daniel Gonçalves, Benjamin Lévêque, and Alexandre Pinlou. Triangle Contact Representations and Duality. *Discr. and Comput. Geom.*, 48(1):239–254, 2012.
- 18 S. Har-Peled. A Simple Proof of the Existence of a Planar Separator. arXiv:1105.0103, 2011.
- 19 Nora Hartsfield and Gerhard Ringel. *Pearls in Graph Theory: A Comprehensive Introduction*. Academic Press, 1990.
- 20 Brad Jackson and Gerhard Ringel. Colorings of circles. *Amer. Math. Monthly*, 91:42–49, 1984.
- 21 Paul Koebe. Kontaktprobleme der konformen Abbildung. *Ber. Verh. Sächs. Akad. Leipzig, Math.-Phys. Klasse*, 88:141–164, 1936.
- 22 László Lovász. Geometric Representations of Graphs. web.cs.elte.hu/~lovasz/geomrep.pdf, December 2009. Draft version.
- 23 Gary L. Miller, Shang-Hua Teng, William Thurston, and Stephen A. Vavasis. Separators for sphere-packings and nearest neighbor graphs. *J. ACM*, 44(1):1–29, 1997.
- 24 B. Mohar. Circle packings of maps in polynomial time. *Europ. J. Comb.*, 18:785–805, 1997.
- 25 Bojan Mohar. Circle packings of maps—the Euclidean case. *Rend. Sem. Mat. Fis. Milano*, 67:191–206, 2000.
- 26 David Orden, Günter Rote, Francisco Santos, Brigitte Servatius, Herman Servatius, and Walter Whiteley. Non-crossing frameworks with non-crossing reciprocals. *Discrete and Computational Geometry*, 32:567–600, 2004. doi:10.1007/s00454-004-1139-x.
- 27 János Pach and Pankaj K. Agarwal. *Combinatorial Geometry*. John Wiley & Sons, 1995.

- 28 Gerhard Ringel. 250 Jahre Graphentheorie. In *Graphs in research and teaching*, pages 136–152. Franzbecker, 1985.
- 29 Burt Rodin and Dennis Sullivan. The convergence of circle packings to the Riemann mapping. *J. Differential Geom.*, 26:349–360, 1987.
- 30 Günter Rote, Francisco Santos, and Ileana Streinu. Pseudo-triangulations — a survey. In Jacob E. Goodman, János Pach, and Richard Pollack, editors, *Surveys on Discrete and Computational Geometry—Twenty Years Later*, volume 453 of *Contemporary Mathematics*, pages 343–410. American Mathematical Society, 2008. [arXiv:math/0612672](https://arxiv.org/abs/math/0612672).
- 31 H. Sachs. Coin Graphs, Polyhedra, and Conformal Mapping. *Discr. Math.*, 134:133–138, 1994.
- 32 Oded Schramm. How to Cage an Egg. *Invent. Math.*, 107:534–560, 1992.
- 33 Oded Schramm. Square tilings with prescribed combinatorics. *Isr. J. Math.*, 84:97–118, 1993.
- 34 Oded Schramm. Combinatorically Prescribed Packings and Applications to Conformal and Quasiconformal Maps. [arXiv:0709.0710](https://arxiv.org/abs/0709.0710), 2007. Modified version of PhD thesis from 1990.
- 35 Kenneth Stephenson. Circle packing: a mathematical tale. *Notices Amer. Math. Soc.*, 50:1376–1388, 2003.
- 36 Kenneth Stephenson. *Introduction to Circle Packing: The Theory of Discrete Analytic Functions*. Cambridge Univ. Press, 2005.
- 37 William Thurston. *Geometry and Topology of Three-Manifolds*. Princeton Univ., 1980. Lecture Notes. Electronic edition: library.msri.org/books/gt3m/.
- 38 W. T. Tutte. How to draw a graph. *Proc. London Math. Soc. (Ser. 3)*, 13:743–767, 1963.
- 39 G. Wegner. Problems in geometric convexity: Problem 86. In J. Tölke and J. M. Wills, editors, *Contributions to Geometry. Proceedings of the Geometry Symposium in Siegen 1978*, page 274. Birkhäuser, 1979.

A Proof of Lemma 4

► **Lemma 4.** *A simple bipartite plane graph with $|S| \geq 4$ nodes has at most $|E| \leq 2|S| - 4$ edges, with equality if and only if the graph is connected and every face is a quadrilateral with four distinct vertices.*

Proof. If the graph is not connected, we add a minimal set of edges to make it connected while keeping it plane and bipartite, resulting in a larger edge set E' .

Since the graph is bipartite, every face cycle has even length. Moreover, every face cycle contains at least 4 edges (possibly visiting both sides of a single edge). To see this, note that the only possible exception, a “digonal” face cycle, would have to be the two sides of a single isolated edge, or two parallel edges. Since $|S| \geq 3$ and the graph is connected and has no multiple edges this cannot happen.

Denoting the set of faces by F , standard double-counting gives the relation $4|F| \leq 2|E'|$, because every edge has 2 sides, and every face cycle goes through at least 4 sides of edges. Euler’s formula gives then $|E'| + 2 = |S| + |F| \leq |S| + |E'|/2$ and therefore $|E'| \leq 2|S| - 4$, with equality if and only all face cycles have length 4. Together, in the chain $|E| \leq |E'| \leq 2|S| - 4$, equality cannot hold if the original graph with edge set E was disconnected ($|E| < |E'|$).

We still have to exclude face cycles of length 4 that are not quadrilaterals (i. e., with 4 distinct vertices). Such a cycle could only be the face surrounding a path with two edges. This is excluded because $|S| \geq 4$ and the graph is connected. ◀