Chapter 3
The Circle Packing Theorem

3.1 Planar Graphs, Maps and Embeddings

Definition 3.1 A graph $G = (V, E)$ is **planar** if it can be **properly drawn** in the plane, that is, if there exists a mapping sending the vertices to distinct points of $\mathbb{R}^2$ and edges to continuous curves between the corresponding vertices so that no two curves intersect, except at the vertices they share. We call such a mapping a **proper drawing** of G.

Remark 3.2 A single planar graph has infinitely many drawings. Intuitively, some may seem similar to one another, while others seem different. For example,

while

The following definition gives a precise sense to the above intuitive equivalence/non-equivalence of drawings.

Definition 3.3 A **planar map** is a graph endowed with a cyclic permutation of the edges incident to each vertex, such that there exists a proper drawing in which the clockwise order of the curves touching the image of a vertex respects that cyclic permutation.

The combinatorial structure of a planar map allows us to define faces directly (that is, without mentioning the drawing). Consider each edge of the graph as directed in both ways, and say that a directed edge $\vec{e}$ **precedes** $\vec{f}$ (or, equivalently,

A. Nachmias, *Planar Maps, Random Walks and Circle Packing*, Lecture Notes in Mathematics 2243, https://doi.org/10.1007/978-3-030-27968-4_3

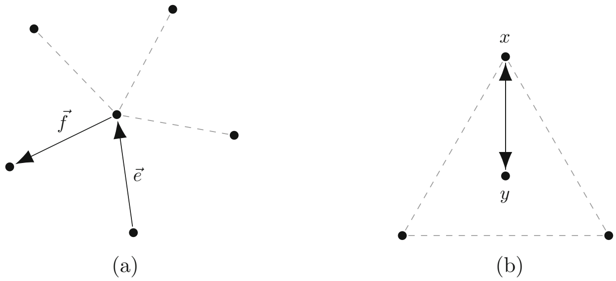

Fig. 3.1 Examples for the edge precedence relation. (**a**) $\vec{e}$ precedes $\vec{f}$. (**b**) (x, y) precedes (y, x)

$\vec{f}$ **succeeds** $\vec{e}$), if there exist vertices v, x, y such that $\vec{e} = (x, v)$, $\vec{f} = (v, y)$, and y is the successor of x in the cyclic permutation σ_v; see Fig. 3.1.

We say that $\vec{e}, \vec{f}$ **belong to the same face** if there exists a finite directed path $\vec{e}_1, \ldots, \vec{e}_m$ in the graph with $\vec{e}_i$ preceding $\vec{e}_{i+1}$ for $i = 1, \ldots, m-1$ and such that either $\vec{e} = \vec{e}_1$ and $\vec{f} = \vec{e}_m$, or $\vec{f} = \vec{e}_1$ and $\vec{e} = \vec{e}_m$. This is readily seen to be an equivalence relation and we call each equivalence class a **face**. Even though a face is a set of directed edges, we frequently ignore the orientations and consider a face as the set of corresponding undirected edges. Each (undirected) edge is henceforth incident to either one or two faces.

When the map is finite an equivalent definition of a face is the set of edges that bound a connected component of the complement of the drawing, that is, of $\mathbb{R}^2$ minus the images of the vertices and edges. This definition is not suitable for infinite planar maps since there may be a complicated set of accumulation points. Given a proper drawing of a finite planar map, there is a unique unbounded connected component of the complement of the drawing; the edges that bound it are called the **outer face** and all other faces are called **inner faces** . However, for any face in a finite map there is a drawing so that this face bounds the unique unbounded connected component, and because of this we shall henceforth refer to the outer face as an arbitrarily chosen face of the map.

We will use the following classical formula.

Theorem 3.4 (Euler's Formula) *Suppose G is a planar graph with n vertices, m edges and f faces. Then*

$$n - m + f = 2.$$

We now state the main theorem we will discuss and use throughout this course. Its proof is presented in the next section.

Theorem 3.5 (The Circle Packing Theorem [51]) *Given any finite simple planar map $G = (V, E)$, $V = \{v_1, \ldots, v_n\}$, there exist n circles in $\mathbb{R}^2$, $C_1, \ldots, C_n$, with disjoint interiors, such that C_i is tangent to C_j if and only if $\{i, j\} \in E$.*

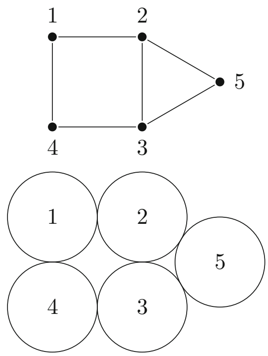

Fig. 3.2 Two distinct planar maps (of the same graph) with corresponding circle packings

Furthermore, for every vertex v_i, the clockwise order of the circles tangent to C_i agrees with the cyclic permutation of v_i's neighbors in the map.

Figure 3.2 gives examples for embeddings of maps which respect the cyclic orderings of neighbors, as guaranteed to exist according to the theorem.

First note that it suffices to prove the theorem for **triangulations**, that is, simple planar maps in which every face has precisely three edges. Indeed, in any planar map we may add a single vertex inside each face and connect it to all vertices bounding that face. The obtained map is a triangulation, and after applying the circle packing theorem for triangulations, we may remove the circles corresponding to the added vertices, obtaining a circle packing of the original map which respects its cyclic permutations. This is depicted in Fig. 3.3.

Thus, it suffices to prove Theorem 3.5 for finite triangulations. In this case an important uniqueness statement also holds.

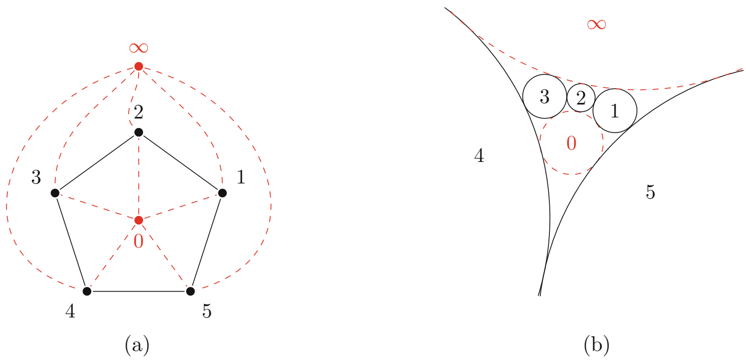

Fig. 3.3 Circle packing of a triangulation of a planar map. (**a**) A planar map and a triangulation. (**b**) A circle packing of the triangulation

Theorem 3.6 *Let $G = (V, E)$ be a finite triangulation on vertex set $V = \{v_1, \dots, v_n\}$ and assume that $\{v_1, v_2, v_3\}$ form a face. Then for any three positive numbers ρ_1, ρ_2, ρ_3, there exists a circle packing $C_1, \dots, C_n$ as in Theorem 3.5 with the additional property that C_1, C_2, C_3 are mutually tangent, form the outer face, and have radii ρ_1, ρ_2, ρ_3, respectively. Furthermore, this circle packing is unique, up to translations and rotations of the plane.*

3.2 Proof of the Circle Packing Theorem

We prove Theorem 3.6 which implies Theorem 3.5 as explained above. Therefore we assume from now on that our map is a triangulation. Denote by n, m and f the number of vertices, edges and faces of the map respectively, and observe that $3f = 2m$ since each edge is counted in exactly two faces, and each face is bounded by exactly three edges. Therefore, by Euler's formula (Theorem 3.4), we have that

$$2 = n - m + f = n - \frac{3}{2}f + f = n - \frac{1}{2}f,$$

thus

$$f = 2n - 4. \tag{3.1}$$

We assume the vertex set is $\{v_1, \dots, v_n\}$, that $\{v_1, v_2, v_3\}$ is the outer face and that ρ_1, ρ_2, ρ_3 are three positive numbers that will be the radii of the outer circles C_1, C_2, C_3 eventually. Denote by F° the set of inner faces of the map, and for a subset of vertices A let $F(A)$ be the set of inner faces with at least one vertex in A. We write $F(v)$ when we mean $F(\{v\})$.

Given a vector $\mathbf{r} = (r_1, \dots, r_n) \in (0, \infty)^n$, an inner face $f \in F^\circ$ bounded by the vertices v_i, v_j, v_k, and a distinguished vertex v_j, we associate a number $\alpha_f^{\mathbf{r}}(v_j) = \measuredangle v_i v_j v_k \in (0, \pi)$ which is the angle of v_j in the triangle $\triangle v_i v_j v_k$ created by connecting the centers of three mutually tangent circles C_i, C_j, C_k of radii r_i, r_j and r_k (that is, in a triangle with side lengths $r_i + r_j$, $r_j + r_k$ and $r_k + r_i$). This number can be calculated using the cosine formula

$$\cos(\measuredangle v_i v_j v_k) = 1 - \frac{2 r_i r_k}{(r_i + r_j)(r_j + r_k)},$$

however, we will not use this formula directly. For every $j \in \{1, \dots, n\}$ we define

$$\sigma_{\mathbf{r}}(v_j) = \sum_{f \in F(v_j)} \alpha_f^{\mathbf{r}}(v_j)$$

to be the *sum of angles* at v_i with respect to $\mathbf{r}$. Let $\theta_1, \theta_2, \theta_3$ be the angles formed at the centers of three mutually tangent circles C_1, C_2, C_3 of radii ρ_1, ρ_2, ρ_3. Equivalently, these are the angles of a triangle with edge lengths $r_1 + r_2$, $r_2 + r_3$ and $r_1 + r_3$. If the vector $\mathbf{r}$ was the vector of radii of a circle packing of the map satisfying Theorem 3.6, then it would hold that

$$\sigma_{\mathbf{r}}(v_i) = \begin{cases} \theta_i & i \in \{1, 2, 3\}\,, \\ 2\pi & \text{otherwise}\,, \end{cases} \tag{3.2}$$

and additionally $(r_1, r_2, r_3) = (\rho_1, \rho_2, \rho_3)$. The proof is split into three parts:

1. Show that there exists a vector $\mathbf{r} \in (0, \infty)^n$ satisfying (3.2);
2. Given such $\mathbf{r}$, show that a circle packing with these radii exists and that (r_1, r_2, r_3) is a positive multiple of (ρ_1, ρ_2, ρ_3); furthermore, this circle packing is unique up to translations and rotations.
3. Show that $\mathbf{r}$ is unique up to scaling all entries by a constant factor.

Proof of Theorem 3.6, Step 1: Finding the Radii Vector r

Observation 3.7 *For every* $\mathbf{r}$*,*

$$\sum_{i=1}^{n} \sigma_{\mathbf{r}}(v_i) = \left|F^\circ\right| \pi = (2n - 5)\pi.$$

Proof Follows immediately since each inner face f bounded by the vertices v_i, v_j, v_k contributes the three angles $\alpha_f^{\mathbf{r}}(v_i)$, $\alpha_f^{\mathbf{r}}(v_j)$ and $\alpha_f^{\mathbf{r}}(v_k)$ which sum to π. By (3.1), there are $2n - 5$ inner faces. □

We now set

$$\delta_{\mathbf{r}}(v_i) = \begin{cases} \sigma_{\mathbf{r}}(v_i) - \theta_i & j \in \{1, 2, 3\}\,, \\ \sigma_{\mathbf{r}}(v_j) - 2\pi & \text{otherwise.} \end{cases} \tag{3.3}$$

Using this notation, our goal is to find $\mathbf{r}$ for which $\delta_{\mathbf{r}} \equiv 0$. It follows from Observation 3.7 that for every $\mathbf{r}$,

$$\sum_{i=1}^{n} \delta_{\mathbf{r}}(v_i) = \sum_{i=1}^{n} \sigma_{\mathbf{r}}(v_i) - \theta_1 - \theta_2 - \theta_3 - (n - 3) \cdot 2\pi = 0\,. \tag{3.4}$$

We define

$$\mathcal{E}_{\mathbf{r}} = \sum_{i=1}^{n} \delta_{\mathbf{r}}(v_i)^2.$$

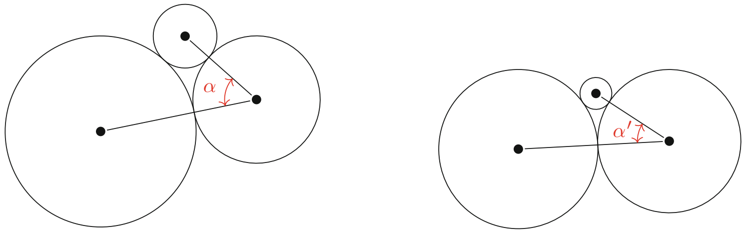

Fig. 3.4 When the radius of a circle corresponding to a vertex increases, while the radii of the circles corresponding to its two neighbors in a given face decrease, the vertex's angle in the corresponding triangle decreases (see Observation 3.8)

We would like to find $\mathbf{r}$ for which $\mathcal{E}_{\mathbf{r}} = 0$. We will use the following geometric observation; see Fig. 3.4.

Observation 3.8 *Let* $\mathbf{r} = (r_1, \ldots, r_n)$ *and* $\mathbf{r}' = (r'_1, \ldots, r'_n)$*, and let* $f \in F^\circ$ *be bounded by* v_i, v_j, v_k*.*

- *If* $r'_i \le r_i$, $r'_k \le r_k$ *and* $r'_j \ge r_j$*, then* $\alpha_f^{\mathbf{r}'}(v_j) \le \alpha_f^{\mathbf{r}}(v_j)$.
- *If* $r'_i \ge r_i$, $r'_k \ge r_k$ *and* $r'_j \le r_j$*, then* $\alpha_f^{\mathbf{r}'}(v_j) \ge \alpha_f^{\mathbf{r}}(v_j)$.
- $\alpha_f^{\mathbf{r}}(v_j)$ *is continuous in* $\mathbf{r}$.

Proof A proof using the cosine formula is routine and is omitted. □

We now define an iterative algorithm, whose input and output are both vectors of radii normalized to have ℓ_1 norm 1. We start with the vector $\mathbf{r}^{(0)} = \left(\frac{1}{n}, \ldots, \frac{1}{n}\right)$, and given $\mathbf{r} = \mathbf{r}^{(t)}$ we construct $\mathbf{r}' = \mathbf{r}^{(t+1)}$. Write $\delta = \delta_{\mathbf{r}}$ and $\delta' = \delta_{\mathbf{r}'}$, and similarly $\mathcal{E} = \mathcal{E}_{\mathbf{r}}$ and $\mathcal{E}' = \mathcal{E}_{\mathbf{r}'}$. We begin by ordering the set of reals $\{\delta(v_i) \mid 1 \le i \le n\}$. If $\delta \equiv 0$ we are done; otherwise, we may choose $s \in \mathbb{R}$ such that the set $S = \{v \mid \delta(v) > s\} \ne \varnothing$ and its complement $V \setminus S$ are non-empty and such that the gap

$$\operatorname{gap}_\delta(S) := \min_{v \in S} \delta(v) - \max_{v \notin S} \delta(v) > 0$$

is maximal over all such s. See Fig. 3.5 for illustration.

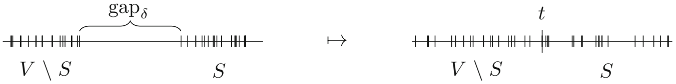

Fig. 3.5 Left: finding the maximum gap between two consecutive values of δ, and splitting the set of values into S and its complement. Right: moving from $\mathbf{r}$ to $\mathbf{r}'$ closes the gap between S and $V \setminus S$

Once we choose S, a step of the algorithm consists of two steps:

1. For some $\lambda \in (0, 1)$ to be chosen later, we set

$$(\mathbf{r}_\lambda)_i = \begin{cases} r_i & v_i \in S, \\ \lambda r_i & v_i \notin S. \end{cases}$$

2. We normalize $\mathbf{r}_\lambda$ so that the sum of entries is 1, letting $\bar{\mathbf{r}}_\lambda$ be the normalized vector. Note that this step does not change the vector δ.

We will choose an appropriate λ that will decrease all values of $\delta(v)$ for $v \in S$, increase all values of $\delta(v)$ for $v \notin S$, and will close the gap. This is made formal in the following two claims.

Claim 3.9 For every $\lambda \in (0, 1)$, setting $\mathbf{r}' = \bar{\mathbf{r}}_\lambda$, we have that $\delta'(v) \leq \delta(v)$ for any $v \in S$, and $\delta'(v) \geq \delta(v)$ for any $v \notin S$.

Claim 3.10 There exists $\lambda \in (0, 1)$ such that setting $\mathbf{r}' = \bar{\mathbf{r}}_\lambda$ gives that $\mathrm{gap}_{\delta'}(S) = 0$.

Proof of Claim 3.9 Consider $v_j \notin S$ and an inner face v_i, v_j, v_k.

Case I $v_i, v_k \notin S$. In this case, the radii of C_i, C_j, C_k are all multiplied by the same number λ, so $\alpha_f^{\mathbf{r}'}(v_j) = \alpha_f^{\mathbf{r}}(v_j)$.

Case II $v_i, v_k \in S$. In this case, the radii of C_i, C_k remain unchanged and the radius of C_j decreases, thus by Observation 3.8, $\alpha_f^{\mathbf{r}'}(v_j) \geq \alpha_f^{\mathbf{r}}(v_j)$.

Case III $v_i \notin S$, $v_k \in S$. In this case the radii of C_i, C_j are multiplied by λ and the radius of C_k is unchanged. The angles of $\triangle v_i v_j v_k$ remain unchanged if we multiply all radii by λ^{-1}, thus we could just as easily have left C_i, C_j unchanged and increased the radius of C_k. By Observation 3.8, we get that $\alpha_f^{\mathbf{r}'}(v_j) \geq \alpha_f^{\mathbf{r}}(v_j)$.

It follows that $\delta'(v) \geq \delta(v)$ for any $v \notin S$. An identical argument shows that $\delta'(v) \leq \delta(v)$ for all $v \in S$. □

In order to prove Claim 3.10, we present another claim.

Claim 3.11

$$\lim_{\lambda \searrow 0} \sum_{v \notin S} \delta_{\mathbf{r}_\lambda}(v) > 0.$$

Proof of Claim 3.10 Using Claim 3.11 The function $\lambda \mapsto \mathrm{gap}_{\delta_{\mathbf{r}_\lambda}}(S)$ is continuous on $(0, 1]$ by the third bullet of Observation 3.8, and its value at $\lambda = 1$ is $\mathrm{gap}_\delta(S) > 0$. Claim 3.11 says that if $\mu > 0$ is small enough, then

$$\sum_{v \notin S} \delta_{\mathbf{r}_\mu}(v) > 0\,,$$

from which it follows that $\max_{v \notin S} \delta_{\mathbf{r}_\mu}(v) > 0$. By (3.4), we also have

$$\sum_{v \in S} \delta_{\mathbf{r}_\mu}(v) < 0,$$

meaning that $\min_{v \in S} \delta_{\mathbf{r}_\mu}(v) < 0$ and therefore $\mathrm{gap}_{\delta_{\mathbf{r}_\mu}}(S) < 0$. By continuity, there exists $\lambda \in (\mu, 1)$ such that $\mathrm{gap}_{\delta_{\mathbf{r}_\lambda}}(S) = 0$. $\square$

Proof of Claim 3.11 We first show that for each face $f \in F(V \setminus S)$ bounded by v_i, v_j, v_k, the sum of angles at the vertices belonging to $V \setminus S$ converges to π as $\lambda \searrow 0$. We show this by the following case analysis. The statements in cases II and III can be justified by drawing a picture or appealing to the cosine formula.

Case I If $v_i, v_j, v_k \notin S$ then since the face is a triangle, $\alpha_f^{\mathbf{r}_\lambda}(v_i) + \alpha_f^{\mathbf{r}_\lambda}(v_j) + \alpha_f^{\mathbf{r}_\lambda}(v_k) = \pi$ for all $\lambda \in (0, 1)$.

Case II If $v_i, v_j \notin S$ but $v_k \in S$ then $\lim_{\lambda \searrow 0} \alpha_f^{\mathbf{r}_\lambda}(v_k) = 0$, hence $\lim_{\lambda \searrow 0} \alpha_f^{\mathbf{r}_\lambda}(v_i) + \alpha_f^{\mathbf{r}_\lambda}(v_j) = \pi$.

Case III If $v_i \notin S$ but $v_j, v_k \in S$ then $\lim_{\lambda \searrow 0} \alpha_f^{\mathbf{r}_\lambda}(v_j) + \alpha_f^{\mathbf{r}_\lambda}(v_k) = 0$, hence $\lim_{\lambda \searrow 0} \alpha_f^{\mathbf{r}_\lambda}(v_i) = \pi$.

It follows that

$$\lim_{\lambda \searrow 0} \sum_{v \notin S} \sigma_{\mathbf{r}_\lambda}(v) = |F(V \setminus S)|\pi. \tag{3.5}$$

For convenience, set

$$\theta(v_i) = \begin{cases} \theta_i & 1 \le i \le 3, \\ 2\pi & \text{otherwise,} \end{cases}$$

so that $\delta_{\mathbf{r}}(v) = \sigma_{\mathbf{r}}(v) - \theta(v)$ for all $v \in V$. Then

$$\lim_{\lambda \searrow 0} \sum_{v \notin S} \delta_{\mathbf{r}_\lambda}(v) = |F(V \setminus S)|\pi - \sum_{v \notin S} \theta(v). \tag{3.6}$$

Let $\bar{F} = F^\circ \setminus F(V \setminus S)$, so every face in $\bar{F}$ contains only vertices of S. We will show that

$$|\bar{F}|\pi < \sum_{v \in S} \theta(v). \tag{3.7}$$

If (3.7) holds, then we can add the negative quantity $|\bar{F}|\pi - \sum_{v \in S} \theta(v)$ to the right side of (3.6), obtaining $|F^\circ|\pi - \sum_{v \in V} \theta(v) = (2n-5)\pi - (2n-5)\pi = 0$. It follows that (3.6) is strictly positive, proving the claim. Thus it suffices to show (3.7).

In the rest of the proof, we fix an embedding of G in the plane with (v_1, v_2, v_3) as the outer face. Let $G[S]$ be the subgraph of G induced by S. Partition S into equivalence classes, $S = S_1 \cup \cdots \cup S_k$, where two vertices are equivalent if they are in the same connected component of $G[S]$. Then $G[S] = G[S_1] \cup \cdots \cup G[S_k]$. Let $\bar{F}_j$ be the set of faces in $\bar{F}$ that appear as faces of $G[S_j]$, so that we have the disjoint union $\bar{F} = \bar{F}_1 \cup \cdots \cup \bar{F}_k$.

Since S is nonempty, it is enough to show that for all $1 \le j \le k$,

$$|\bar{F}_j|\pi < \sum_{v \in S_j} \theta(v). \tag{3.8}$$

Let m_j and f_j denote the number of edges and faces, respectively, of $G[S_j]$. Observe that $|\bar{F}_j| \le f_j - 1$. If $|\bar{F}_j| = 0$, then (3.8) is trivial. If $|\bar{F}_j| \ge 1$, then $G[S_j]$ has at least one inner face, and since it is a simple graph, every face must have degree at least 3. (The degree of a face is the number of directed edges that make up its boundary.) Because the sum of the degrees of all the faces equals twice the number of edges, we have $2m_j \ge 3f_j$. Euler's formula now gives

$$|S_j| + f_j - 2 = m_j \ge \frac{3}{2} f_j,$$

and hence $f_j \le 2|S_j| - 4$. Thus, the left side of (3.8) satisfies

$$|\bar{F}_j|\pi \le (2|S_j| - 5)\pi.$$

If S_j contains all of v_1, v_2, v_3, then the right side of (3.8) is

$$\theta_1 + \theta_2 + \theta_3 + (|S_j| - 3) \cdot 2\pi = (2|S_j| - 5)\pi.$$

Otherwise, at least one of the θ_i is replaced by 2π and so the right side of (3.8) is strictly greater than the left side. In fact, (3.8) holds except when $v_1, v_2, v_3 \in S_j$ and $|\bar{F}_j| = f_j - 1 = 2|S_j| - 5$. We now show that this situation cannot occur.

The equality $|\bar{F}_j| = f_j - 1$ means that every inner face of $G[S_j]$ is an element of $\bar{F}_j$ and therefore a face of G. Since $v_1, v_2, v_3 \in S_j$, the outer face of $G[S_j]$ is (v_1, v_2, v_3), which is the same as the outer face of G. So, every face of $G[S_j]$ is also a face of G. But this is impossible: if we choose any $v \in V \setminus S$, then v must lie in some face of $G[S_j]$, which then cannot be a face of G. Therefore, it cannot be true that $v_1, v_2, v_3 \in S_j$ and also $|\bar{F}_j| = f_j - 1$, so we conclude that (3.8) always holds. □

We now analyse the algorithm. Let $\lambda \in (0, 1)$ be the one guaranteed by Claim 3.10, and set $\mathbf{r}' = \bar{\mathbf{r}}_\lambda$.

Claim 3.12 $\mathcal{E}' \le \mathcal{E}\left(1 - \frac{1}{2n^3}\right)$.

Proof As depicted in Fig. 3.5, define

$$t = \min_{v \in S} \delta'(v) = \max_{v \notin S} \delta'(v).$$

By (3.4) we have that $\sum_{i=1}^{n} \delta(v_i) = \sum_{i=1}^{n} \delta'(v_i) = 0$, hence

$$\mathcal{E} - \mathcal{E}' = \sum_{i=1}^{n} \delta(v_i)^2 - \sum_{i=1}^{n} \delta'(v_i)^2 = \sum_{i=1}^{n} (\delta(v_i) - \delta'(v_i))^2 + 2\sum_{i=1}^{n} (t - \delta'(v_i))(\delta'(v_i) - \delta(v_i)).$$

If $v \in S$, then $t \le \delta'(v) \le \delta(v)$ and if $v \notin S$, then $t \ge \delta'(v) \ge \delta(v)$. Thus, in both cases $(t - \delta'(v))(\delta'(v) - \delta(v)) \ge 0$. Taking $u \in S$ and $v \notin S$ with $\delta'(u) = \delta'(v) = t$, we have that

$$\mathcal{E} - \mathcal{E}' \ge (\delta(u) - t)^2 + (\delta(v) - t)^2 \ge \frac{(\delta(u) - \delta(v))^2}{2} \ge \frac{\mathrm{gap}_\delta(S)^2}{2}.$$

Since $\mathrm{gap}_\delta(S)$ was chosen to be the maximal gap we may bound,

$$\mathrm{gap}_\delta(S) \ge \frac{1}{n}\left(\max_{v \in V} \delta(v) - \min_{v \in V} \delta(v)\right).$$

For every $v \in V$,

$$\max_{w \in V} \delta(w) - \min_{w \in V} \delta(w) \ge |\delta(v)|,$$

and thus

$$n\left(\max_{v \in V} \delta(v) - \min_{v \in V} \delta(v)\right)^2 \ge \sum_{i=1}^{n} \delta(v_i)^2 = \mathcal{E}.$$

Hence

$$\mathcal{E} - \mathcal{E}' \ge \frac{1}{2n^2}\left(\max_{v \in V} \delta(v) - \min_{v \in V} \delta(v)\right)^2 \ge \frac{1}{2n^2} \cdot \frac{\mathcal{E}}{n},$$

and we conclude that

$$\mathcal{E}' \le \mathcal{E}\left(1 - \frac{1}{2n^3}\right).$$

□

Write $\mathcal{E}^{(t)} = \mathcal{E}_{\mathbf{r}^{(t)}}$. By iterating the described algorithm, we obtain from Claim 3.12 that

$$\mathcal{E}^{(t)} \leq \mathcal{E}^{(0)} \left(1 - \frac{1}{2n^3}\right)^t \longrightarrow 0 \qquad \text{as } t \to \infty\,.$$

By our normalization $\left\|\mathbf{r}^{(t)}\right\|_{\ell_1} = 1$. Thus, by compactness, there exists a subsequence $\{t_k\}$ and a vector $\mathbf{r}^{\infty}$ such that $\mathbf{r}^{(t_k)} \to \mathbf{r}^{\infty}$ as $k \to \infty$. From continuity of $\mathcal{E}$ we have that $\mathcal{E}\left(\mathbf{r}^{\infty}\right) = 0$, meaning that (3.2) is satisfied. For $\mathbf{r}^{\infty}$ to be feasible as a vector of radii, we also have to argue that it is positive (the fact that no coordinates are ∞ follows since $||\mathbf{r}^{\infty}||_{\ell_1} = 1$).

Claim 3.13 $\mathbf{r}_i^{\infty} > 0$ for every i.

Proof Let $S = \{v_i \in V : \mathbf{r}_i^{\infty} > 0\}$. Because of the normalization of $\mathbf{r}$, we know that S is nonempty. Assume for contradiction that $S \subsetneq V$. We repeat the exact same argument used in the proof of Claim 3.11 showing first by case analysis that

$$\lim_{t\to\infty} \sum_{v \notin S} \sigma_{\mathbf{r}^{(t)}}(v) = |F(V \setminus S)|\pi$$

and then deducing that

$$\lim_{t\to\infty} \sum_{v \notin S} \delta_{\mathbf{r}^{(t)}}(v) > 0.$$

This contradicts that $\lim_{t\to\infty} \mathcal{E}^{(t)} = 0$, so we conclude that $S = V$. $\square$

Proof of Theorem 3.6, Step 2: Drawing the Circle Packing Described by $\mathbf{r}^{\infty}$

Given the vector of radii $\mathbf{r}^{\infty}$ satisfying (3.2), we now show that the corresponding circle packing can be drawn uniquely up to translations and rotations. In fact, we provide a slightly more general statement which is due to Ori Gurel-Gurevich and Ohad Feldheim [personal communications, 2018].

Let $G = (V, E)$ be a finite planar triangulation on vertex set $\{v_1, \ldots, v_n\}$ and assume that $\{v_1, v_2, v_3\}$ is the outer face. A vector of positive real numbers $\boldsymbol{\ell} = \{\ell_e\}_{e\in E}$ indexed by the edge set E is called *feasible* if for any face enclosed by edges e_1, e_2, e_3, the lengths $\ell_{e_1}, \ell_{e_2}, \ell_{e_3}$ can be made to form a triangle. In other words, these lengths satisfy three triangle inequalities,

$$\ell_{e_i} + \ell_{e_j} > \ell_{e_k} \qquad \{i, j, k\} = \{1, 2, 3\}\,.$$

Given a feasible edge length vector $\boldsymbol{\ell}$ we may again use the cosine formula to compute, for each face f, the angle at a vertex of the triangle formed by the three corresponding edge lengths. We denote these angles, as before, by $\alpha_f^{\boldsymbol{\ell}}(v)$ where v is a vertex of f. Similarly, we define

$$\sigma_{\boldsymbol{\ell}}(v) = \sum_{f \in F(v)} \alpha_f^{\boldsymbol{\ell}}(v)$$

to be the sum of angles at a vertex v.

Theorem 3.14 *Let G be a finite triangulation and $\boldsymbol{\ell}$ a feasible vector of edge lengths. Assume that $\sigma_{\boldsymbol{\ell}}(v) = 2\pi$ for any internal vertex v. Then there is a drawing of G in the plane so that each edge e is drawn as a straight line segment of length ℓ_e and no two edges cross. Furthermore, this drawing is unique up to translations and rotations.*

It is easy to use the theorem above to draw the circle packing given the radii vector $\mathbf{r}^\infty$ satisfying (3.2). Indeed, given $\mathbf{r}^\infty$ we set $\boldsymbol{\ell}$ by putting $\ell_e = \mathbf{r}_i^\infty + \mathbf{r}_j^\infty$ for any edge $e = \{v_i, v_j\}$ of the graph. Condition (3.2) implies that $\boldsymbol{\ell}$ is feasible. We now apply Theorem 3.14 and obtain the guaranteed drawing and draw a circle C_i of radii $\mathbf{r}_i^\infty$ around v_i for all i. Theorem 3.14 guarantees that for any edge $\{v_i, v_k\}$ the distance between v_i, v_j is precisely $\mathbf{r}_i^\infty + \mathbf{r}_j^\infty$ and thus C_i and C_j are tangent. Conversely, assume that v_i, v_j do not form an edge. To each vertex v let A_v be the union of triangles touching v, each triangle is the space bounded by a face touching v in the drawing of Theorem 3.14. Since G is a triangulation and v_i and v_j are not adjacent we learn that A_{v_i} and A_{v_j} have disjoint interiors. Furthermore, $C_i \subset \mathrm{Int}(A_{v_i})$ since the straight lines emanating from v_i have length larger than $\mathbf{r}_i^\infty$. By the same token $C_j \subset \mathrm{Int}(A_{v_j})$ and we conclude that C_i and C_j are not tangent.

Lastly, we note that by (3.2) the outer boundary of the polygon we drew is a triangle with angles $\theta_1, \theta_2, \theta_3$ and hence (r_1, r_2, r_3) is a positive multiple of (ρ_1, ρ_2, ρ_3). Step 2 of the proof of Theorem 3.5 is now concluded.

Proof of Theorem 3.14 We prove this by induction on the number of vertices n. The base case $n = 3$ is trivial since the feasibility of $\boldsymbol{\ell}$ guarantees that the edge lengths of the three edges of the outer face can form a triangle. Any two triangles with the same edge lengths can be rotated and translated to be identical, so the uniqueness statement holds for $n = 3$.

Assume now that $n > 3$ so that there exists an internal vertex v. Denote by $v_1, \ldots, v_m$ the neighbors of v ordered clockwise. We begin by placing v at the origin and drawing all the faces to which v belongs, see Fig. 3.6, left. That is, we draw the edge $\{v, v_1\}$ as a straight line interval of length $\ell_{\{v,v_1\}}$ on the positive x-axis emanating from the origin and proceed iteratively: for each $1 < i \le m$ we draw the edge $\{v, v_i\}$ as a straight line interval of length $\ell_{\{v,v_i\}}$ emanating from the origin (v) at a clockwise angle of $\alpha_f^{\boldsymbol{\ell}}(v)$ from the previous drawn line segment of $\{v, v_{i-1}\}$, where $f = \{v, v_{i-1}, v_i\}$. This determines the location of $v_1, \ldots, v_m$ in the plane and allows us to "complete" the triangles by drawing the straight line

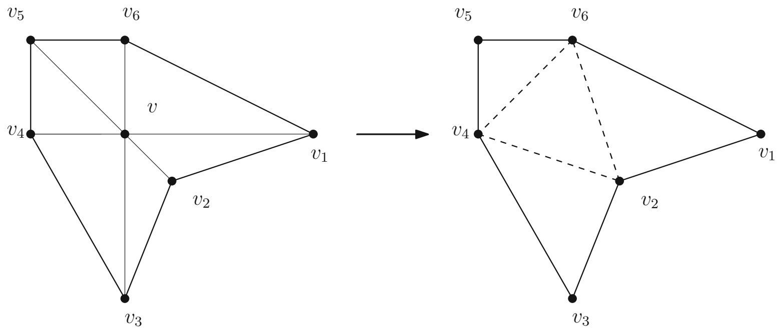

Fig. 3.6 On the left, we first draw the polygon surrounding v. On the right, we then erase v and the edge emanating from it, replacing it with diagonals that triangulate the polygon while recording the lengths of the diagonals in $\boldsymbol{\ell}'$. The latter is the input to the induction hypothesis

segments connecting v_i to v_{i+1}, each of length $\ell_{\{v_i,v_{i+1}\}}$ where $1 \leq i \leq m$ (where $v_{m+1} = v_1$). Denote these edges by $e_1, \ldots, e_m$.

Since $\sigma_{\boldsymbol{\ell}}(v) = 2\pi$ we learn that these m triangles have disjoint interiors and that the edges $e_1, \ldots, e_m$ form a closed polygon containing the origin in its interior. It is a classical fact [64] that every closed polygon can be triangulated by drawing some diagonals as straight line segments in the interior of the polygon. We fix such a choice of diagonals and use it to form a new graph G' on $n-1$ vertices and $|E(G)| - 3$ edges by erasing v and the m edges emanating from it and adding the new $m-3$ edges corresponding to the diagonals we added. Furthermore, we generate a new edge length vector $\boldsymbol{\ell}'$ corresponding to G' by assigning the new edges lengths corresponding to the Euclidean length of the drawn diagonals and leaving the other edge lengths unchanged. See Fig. 3.6, right.

It is clear that $\boldsymbol{\ell}'$ is feasible and that the angle sum at each internal vertex of G' is 2π. Therefore we may apply the induction hypothesis and draw the graph G' according to the edge lengths $\boldsymbol{\ell}'$. This drawing is unique up to translations and rotations by induction. Note that in this drawing of G', the polygon corresponding to $e_1, \ldots, e_m$ must be the exact same polygon as before, up to translations and rotations, since it has the same edge lengths and the same angles between its edges. Since it is the same polygon, we can now erase the diagonals in this drawing and place a new vertex in the same relative location where we drew v previously, along with the straight line segments connecting it to $v_1, \ldots, v_m$. Thus we have obtained the desired drawing of G. The uniqueness up to translations and rotations of this drawing follows from the uniqueness of the drawing of G' and the fact that the location of v is uniquely determined in that drawing. $\square$

Proof of Theorem 3.6, Step 3: Uniqueness

Theorem 3.15 (Uniqueness of Circle Packing) *Given a simple finite triangulation with outer face v_1, v_2, v_3 and three radii ρ_1, ρ_2, ρ_3, the circle packing with $C_{v_1}, C_{v_2}, C_{v_3}$ having radii ρ_1, ρ_2, ρ_3 is unique up to translations and rotations.*

Proof We have already seen in step 2 that given the radii vector $\mathbf{r}$ the drawing we obtain is unique up to translations and rotations. Thus, we only need to show the uniqueness of $\mathbf{r}$ given ρ_1, ρ_2, ρ_3.

To that aim, suppose that $\mathbf{r}^a$ and $\mathbf{r}^b$ are two vectors satisfying (3.2). Since the outer face in both vectors correspond to a triangle of angles $\theta_1, \theta_2, \theta_3$ we may rescale so that $\mathbf{r}_i^a = \mathbf{r}_i^b = \rho_i$ for $i = 1, 2, 3$. After this rescaling, assume by contradiction that $\mathbf{r}^a \neq \mathbf{r}^b$ and let v be the interior vertex which maximizes $\mathbf{r}_v^a/\mathbf{r}_v^b$. We can assume without loss of generality that this quantity is strictly larger than 1, as otherwise we can swap $\mathbf{r}^a$ and $\mathbf{r}^b$.

Now we claim that for each $f = (v, u_1, u_2) \in F(v)$, we have $\alpha_f^{\mathbf{r}^a}(v) \leq \alpha_f^{\mathbf{r}^b}(v)$, with equality if and only if the ratios $\mathbf{r}_{u_i}^a/\mathbf{r}_{u_i}^b$, for $i = 1, 2$, are both equal to $\mathbf{r}_v^a/\mathbf{r}_v^b$. This is a direct consequence of Observation 3.8. Indeed, scale all the radii in $\mathbf{r}^b$ by a factor of $\mathbf{r}_v^a/\mathbf{r}_v^b$ to get a new vector $\mathbf{r}'$ such that $\mathbf{r}_v^a = \mathbf{r}'_v$ and $\mathbf{r}_u^a \leq \mathbf{r}'_u$ for all $u \neq v$. The second bullet point in Observation 3.8 implies that $\alpha_f^{\mathbf{r}^a}(v) \leq \alpha_f^{\mathbf{r}'}(v) = \alpha_f^{\mathbf{r}^b}(v)$. As well, if either $\mathbf{r}_{u_1}^a < \mathbf{r}'_{u_1}$ or $\mathbf{r}_{u_2}^a < \mathbf{r}'_{u_2}$, then the cosine formula yields the strict inequality $\alpha_f^{\mathbf{r}^a}(v) < \alpha_f^{\mathbf{r}'}(v)$. Thus, $\alpha_f^{\mathbf{r}^a}(v) = \alpha_f^{\mathbf{r}^b}(v)$ only if $\mathbf{r}_{u_i}^a/\mathbf{r}_{u_i}^b = \mathbf{r}_v^a/\mathbf{r}_v^b$ for $i = 1, 2$.

Now, since $\alpha_f^{\mathbf{r}^a}(v) \leq \alpha_f^{\mathbf{r}^b}(v)$ for each $f \in F(v)$, while $\sigma_{\mathbf{r}^a}(v) = \sigma_{\mathbf{r}^b}(v) = 2\pi$, the equality $\alpha_f^{\mathbf{r}^a}(v) = \alpha_f^{\mathbf{r}^b}(v)$ must hold for each f. Therefore, each neighbor u of v satisfies $\mathbf{r}_u^a/\mathbf{r}_u^b = \mathbf{r}_v^a/\mathbf{r}_v^b$. Because the graph is connected, the ratio $\mathbf{r}_u^a/\mathbf{r}_u^b$ must be constant for all vertices $u \in V(G)$. But this contradicts that $\mathbf{r}_v^a/\mathbf{r}_v^b > 1$ while $\mathbf{r}_{v_i}^a/\mathbf{r}_{v_i}^b = 1$ for $i = 1, 2, 3$. We conclude that $\mathbf{r}^a = \mathbf{r}^b$. □