CHAPTER 6

Square Tilings

We can also represent planar graphs by squares, rather then circles, in the plane. There are in fact two quite different ways of doing this: the squares can correspond to the edges, a classic result [Brooks et al. 1940], or the squares can correspond to the nodes, a more recent result [Schramm 1993].

6.1. Electric current through a rectangle

The classical paper [Brooks et al. 1940] used a physical model of electrical currents to show how to relate square tilings to planar graphs. The ultimate goal was to construct tilings of a square with squares whose edge-lengths are all different. This will not be our concern (but see Exercise 6.2); we will allow squares of the same size and also the domain to be tiled can be any rectangle, not necessarily a square.

Consider a tiling $\mathcal{T}$ of a rectangle R with a finite number of squares, whose sides are parallel to the coordinate axes. We can associate a planar map with this tiling as follows. Let us call a maximal horizontal segment composed of edges of the squares a *long edge*. Represent every long edge by a single node. Each square "connects" two horizontal segments, and we can represent it by an edge connecting the two corresponding nodes, directed top-down. We get a directed graph $G_{\mathcal{T}}$ (Figure 6.1), with a single source s (representing the upper edge of the rectangle) and a single sink t (representing the lower edge).

It is easy to see that graph $G_{\mathcal{T}}$ is planar: it can be obtained by first considering the midpoints of the horizontal edges of the squares, connecting two of them horizontally if they are neighbors along a long edge, and vertically if they belong to opposite edges of the same square. This graph is clearly planar, and contracting the horizontal edges, we get $G_{\mathcal{T}}$.

A little attention must be paid to points where four squares meet; we call these points 4-*fold corners*. Suppose that squares A, B, C, D share a corner p, where A is the upper left, and B, C, D follow counterclockwise. In this case, we may consider the lower edges of A and B to belong to a single long edge, or to belong to different long edges. In the latter case, we may or may not imagine that there is an infinitesimally small square sitting at p, which may "connect" A with C or B with D (Figure 6.2). What this means is that we have to declare if the four edges of $G_{\mathcal{T}}$ corresponding to A, B, C and D are adjacent to the same node, two nonadjacent nodes, or two adjacent nodes. We can orient this horizontal edge arbitrarily. Deciding between these four possibilities will be called "resolving" the 4-fold corner p.

If we assign the edge length of each square to the corresponding edge, we get a flow f from s to t: If a node v represents a segment I, then the total flow into v is the sum of edge lengths of squares attached to I from the top, while the total flow

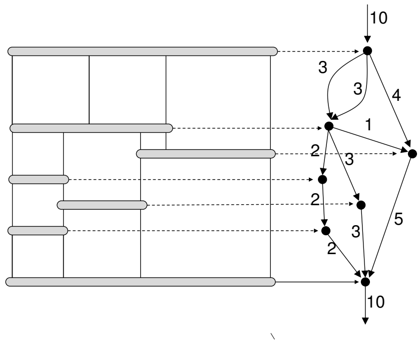

FIGURE 6.1. Constructing a graph with a harmonic function from a tiling of a rectangle by squares.

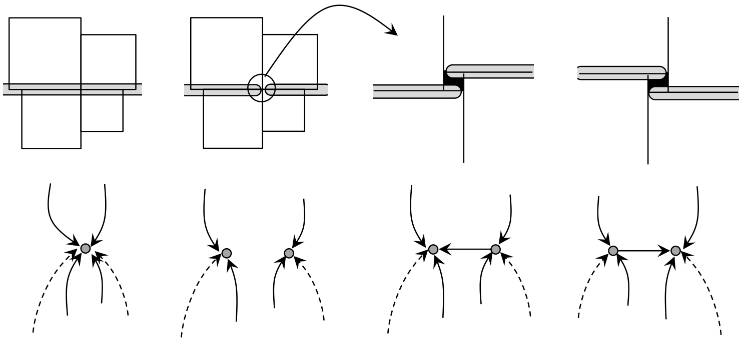

FIGURE 6.2. Possible resolutions of four squares meeting at a point.

out of v is the sum of edge length of squares attached to I from the bottom. Both of these sums are equal to the length of I.

Let $h(v)$ denote the distance of node v from the upper edge of R. Since the edge-length of a square is also the difference between the y-coordinates of its upper and lower edges, the function h is harmonic:

$$h(i) = \frac{1}{\deg(i)} \sum_{j \in N(i)} h(j)$$

for every node different from s and t (Figure 6.1).

It is not hard to see that this construction can be reversed.

Theorem 6.1. *For every connected planar map G with two specified nodes s and t on the unbounded country, there is a unique tiling $\mathcal{T}$ of a rectangle such that (resolving the 4-fold corners appropriately) $G \cong G_{\mathcal{T}}$.*

Proof. Consider the harmonic function $f : V \to \mathbb{R}$ with $f(s) = 0$ and $f(t) = 1$ (obtained, say, as the 1-dimensional rubber band embedding with s and t nailed; in Figure 6.3, this is the vertical coordinate of each node). We assign a square to each edge uv with $f(v) > f(u)$, of side length $f(v) - f(u)$. This square will be placed so that its lower edge is at height $f(u)$, and its upper edge, at height $f(v)$. To find the horizontal position of these squares, we start from node s: we line up the

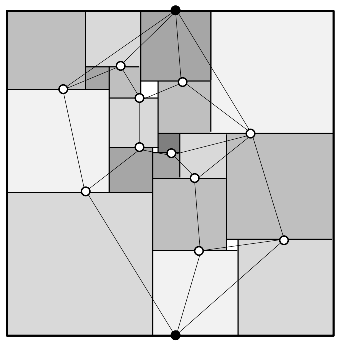

FIGURE 6.3. A planar graph and the tiling generated from it.

squares corresponding to the edges incident with s along the bottom line, in the order as these edges emanate from s. We go through the nodes in the increasing order of the values of f. Getting to a node v, those edges entering v from below have squares assigned to them whose top edges are at height $f(v)$, and these edges form a segment I_v of length

$$|I_v| = \sum_{\substack{u \in N(v) \\ f(u) < f(v)}} \big(f(v) - f(u)\big).$$

Since by the harmonic property of f we also have

$$|I_v| = \sum_{\substack{u \in N(v) \\ f(u) \geq f(v)}} \big(f(u) - f(v)\big),$$

so we can line up the squares corresponding to edges exiting v upwards, along I_v, in the order given by the embedding in the plane. When we get to t, we have filled up the rectangle. □

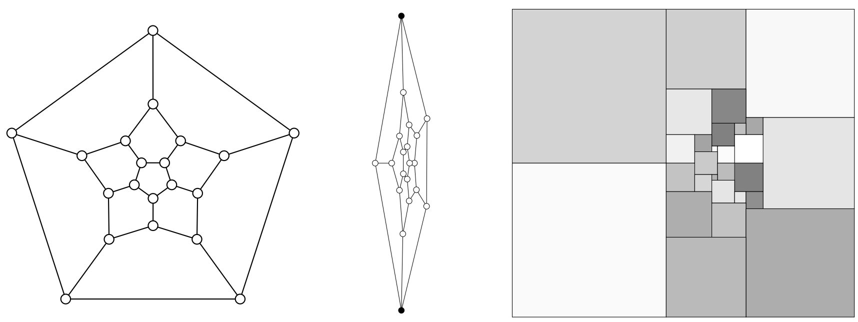

FIGURE 6.4. The dodecahedron graph, its rubber band embedding on the line (horizontally distorted to show the structure), and a square tiling generated from it.

6.2. Tangency graphs of square tilings

Let R be a rectangle in the plane, and consider a tiling of R by squares. Let us add four further squares attached to each edge of R from the outside, sharing the edge with R. We want to look at the tangency graph of this family of squares.

Since the squares do not have a smooth boundary, the conclusion of Exercise 5.8 does not apply, and the tangency graph of the squares may not be planar. Let us try to draw the tangency graph in the plane by representing each square by its center and connecting the centers of touching squares by a straight line segment. Similarly as in the preceding section, we get into trouble when four squares share a vertex. In this case we can specify arbitrarily one diametrically opposite pair as "infinitesimally overlapping", and connect the centers of these two square but not the other two centers. We call this a *resolved tangency graph* of the family of squares.

Every resolved tangency graph is planar, and it is easy to see that it has exactly one country that is a quadrilateral (namely, the unbounded country), and its other countries are triangles; briefly, it is a triangulation of a quadrilateral (Figure 6.5).

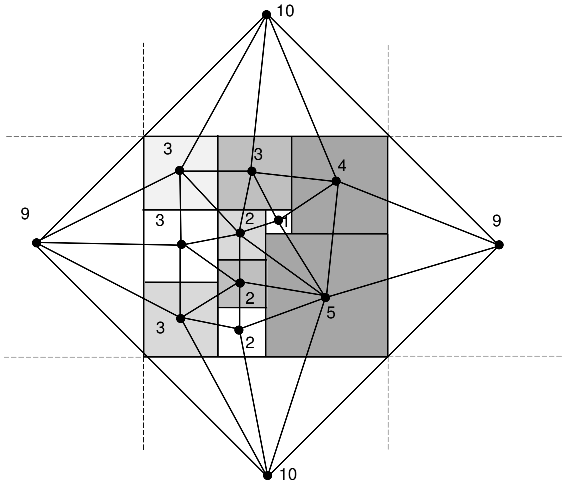

FIGURE 6.5. The resolved tangency graph of a tiling of a rectangle by squares. The numbers indicate the edge lengths of the squares.

Under some connectivity conditions, this fact has the following converse [Schramm 1993].

Theorem 6.2. *Every planar map in which the unbounded country is a quadrilateral, all other countries are triangles, and is not separated by a* 3*-cycle or* 4*-cycle, can be represented as a resolved tangency graph of a square tiling of a rectangle.*

Schramm proved a more general theorem, in which separating cycles were allowed; the prize to pay was that degenerate squares with edge-length 0 had to be allowed. It is easy to see that a separating triangle forces everything inside it to degenerate in this sense, and so we do not loose anything interesting by excluding these. Separating 4-cycles may or may not force degeneracy (Figure 6.6), and it does not seem easy to tell when they do.

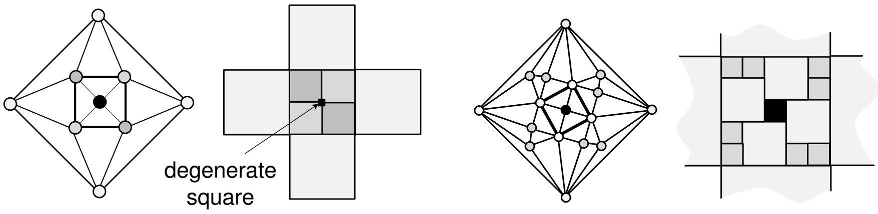

Figure 6.6. A separating 4-cycle that causes degeneracy and one that does not.

Before proving Theorem 6.2, let us make a detour to a well-known game. The following folklore fact is easy to prove:

Proposition 6.3. *Let G be a planar map in which the unbounded country is a quadrilateral $abcd$ and all other countries are triangles. Let us 2-color the nodes with black and white, so that a and c are black and b and d are white. Then either there is an all-black a-c path, or an all-white b-d path, but not both.*

Proof. From the planarity of G it is immediate that if an all-black a-c path exists, then it separates b and d, and so it excludes an all-white b-d path. To prove that one of these paths exists, let us start a walk through the map, avoiding the nodes. We enter through the edge ab, so that we have a black node to the left and a white node to the right, and leave through another edge that has differently colored endpoints (clearly again black on the left and white on the right). Going on similarly, we maintain that we cross only edges having a black endnode to our left and a white endnode to our right. It is easy to see that we never return to a triangle which we left earlier, and so we must return to the unbounded country through one of the edges bc or ad: Exiting through cd is impossible, since then we would have the wrong colors on our left and right (Figure 6.7, left).

Suppose (say) that we exit through bc; then the black-black edges of those triangles that we cross form a walk from a to c, which contains a black a-c path. $\square$

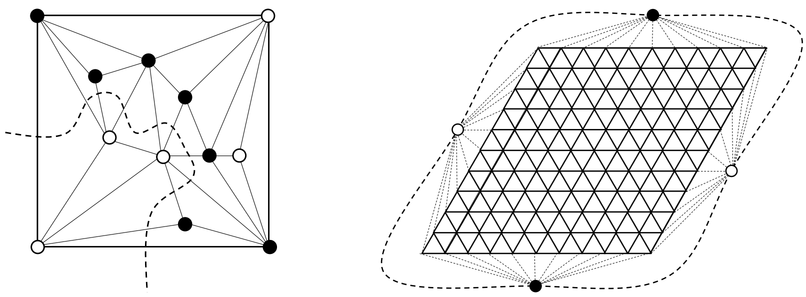

Figure 6.7. Left: Walking through a 2-colored triangulation of a square. Right: The game of Hex as a special case of the situation in Proposition 6.3.

The reader familiar with basic algebraic topology will notice that the proof above is very similar to one of the (algorithmic) proofs of Sperner's Lemma. In fact it would be easy to derive Proposition 6.3 from Sperner's Lemma.

A special case of this proposition is the fact that in every game of Hex, one or the other player wins, but not both (Figure 6.7, right).

Proof of Theorem 6.2. Let G be the planar map with unbounded country $abcd$. It will be convenient to set $V' = V \setminus \{a,b,c,d\}$. Recall that $V'(P)$ denotes the set of inner nodes of the path P. For a path P and two nodes u and v on the path, let us denote by $P(u,v)$ the subpath of P consisting of all nodes strictly between u and v. Later on we will also need the notation $P[u,v)$ for the subpath formed by the nodes between u and v, with u included but not v, and $P[u,v]$ for the subpath with both included.

We consider the node-path and node-cut polyhedra, where a and c are the special nodes and b and d are deleted, or vice versa (see Appendix C.2.1). Proposition 6.3 implies that the a-c node-cut polyhedron of $G \setminus \{b,d\}$ is the same as the b-d node-path polyhedron of $G \setminus \{a,c\}$, and vice versa. Another way of saying this is that the b-d node-path polyhedron of $G \setminus \{a,c\}$ is the blocker of the a-c node-path polyhedron of $G \setminus \{b,d\}$.

Let $\mathcal{P}$ be the set of a–c paths in $G \setminus \{b,d\}$ and $\mathcal{Q}$, the set of b–d paths in $G \setminus \{a,c\}$. The a-c node-cut polyhedron is described by the linear constraints

$$x_u \ge 0 \qquad (u \in V') \tag{6.1}$$

$$x\big(V'(P)\big) \ge 1 \qquad (P \in \mathcal{P}). \tag{6.2}$$

Consider the solution $\overline{x}$ of (6.1)–(6.2) minimizing the objective function $\sum_u x_u^2$, and let R^2 be the minimum value. By Proposition C.2, $\overline{y} = (1/R^2)\overline{x}$ minimizes the same objective function over the blocker (which is the b-d node-path polyhedron of $G \setminus \{a,c\}$), and the optimum value is $1/R^2$. Let us rescale the optimizers to get $z = \frac{1}{R}\overline{x} = R\overline{y}$. Then we have

$$z\big(V'(P)\big) \ge \frac{1}{R} \qquad (P \in \mathcal{P}) \tag{6.3}$$

$$z\big(V'(Q)\big) \ge R \qquad (Q \in \mathcal{Q}) \tag{6.4}$$

$$\sum_{u \in V'} z_u^2 = 1. \tag{6.5}$$

It will be convenient to define $z_a = z_b = z_c = z_d = 0$.

We assign the *length* $\widehat{z}_{ij} = \frac{1}{2}(z_i + z_j)$ to every edge ij. Using this, we can define (as usual) the length $\widehat{z}(P)$ of path as the sum of lengths of its edges, and the *distance* $d_z(u,v)$ of two nodes as the minimum length of any path connecting them. Inequalities (6.3) and (6.4) say that $d_z(a,c) \ge 1/R$ and $d_z(b,d) \ge R$, and the minimality of $\overline{x}$ and $\overline{y}$ implies that we have equality here. It is easy to see that $d_z(i,j) = \widehat{z}_{ij}$ for every edge ij, so the edge is a shortest path between its endpoints.

We know that $\overline{y}$ is in the a-c node-path polyhedron, and it is a minimal vector in there, so it is in the a-c node-path polytope, and it can be written as a convex combination of indicator vectors $\mathbb{1}_P$ of sets $V'(P)$, where P is an a–c path. It follows that z can be written as

$$z = \sum_{P \in \mathcal{P}} \lambda_P \mathbb{1}_P, \qquad \sum_P \lambda_P = R, \qquad \lambda_P \ge 0. \tag{6.6}$$

Similarly, we have a decomposition

$$z = \sum_{Q \in \mathcal{Q}} \mu_Q \mathbb{1}_Q, \qquad \sum_Q \mu_Q = \frac{1}{R}, \qquad \mu_Q \ge 0. \tag{6.7}$$

Let $\mathcal{P}' = \{P \in \mathcal{P} : \ \lambda_P > 0\}$, and define $\mathcal{Q}'$ analogously. Obviously, a node $u \in V'$ has $z_u > 0$ if and only if it is contained in one of the paths in $\mathcal{P}'$, which must then be equivalent to be contained in one of the paths in $\mathcal{Q}'$ (we'll prove later that all nodes have $z_u > 0$).

From conditions (6.3)-(6.5) we can derive some simple but powerful properties of the paths in $\mathcal{P}'$ and $\mathcal{Q}'$. It is trivial from the topology of G that $|V(P) \cap V(Q)| \ge 1$ for $P \in \mathcal{P}'$ and $Q \in \mathcal{Q}'$. On the other hand, (6.5) implies that

$$1 = \sum_{u \in V'} z_u^2 = \Big(\sum_P \lambda_P \mathbb{1}_P\Big)^{\mathsf T}\Big(\sum_Q \mu_Q \mathbb{1}_Q\Big) = \sum_{P,Q} \lambda_P \mu_Q |V(P) \cap V(Q)|$$
$$\ge \sum_{P,Q} \lambda_P \mu_Q = \Big(\sum_P \lambda_P\Big)\Big(\sum_Q \mu_Q\Big) = 1.$$

We must have equality here, which implies that

$$|V(P) \cap V(Q)| = 1 \qquad (P \in \mathcal{P}',\, Q \in \mathcal{Q}'). \tag{6.8}$$

For any path $Q \in \mathcal{Q}'$, we have

$$\widehat{z}(Q) = z^{\mathsf T}\mathbb{1}_Q = \sum_{P \in \mathcal{P}'} \lambda_P \mathbb{1}_P^{\mathsf T}\mathbb{1}_Q = \sum_{P \in \mathcal{P}'} \lambda_P |V'(P) \cap V'(Q)| = \sum_{P \in \mathcal{P}'} \lambda_P = R. \tag{6.9}$$

Similarly, for every $P \in \mathcal{P}'$, we have

$$\widehat{z}(P) = \frac{1}{R}. \tag{6.10}$$

So the paths in $\mathcal{P}'$ and $\mathcal{Q}'$ are shortest paths with respect to the metric d_z. It follows that they are chordless, and also for every path $P \in \mathcal{P}'$ and $u, v \in V(P)$, the subpath $P[u, v]$ is a shortest path between its endpoints. We get by the same kind of argument that the common nodes of any two paths $P, P' \in \mathcal{P}'$ are encountered by the two paths in the same order.

A consequence of (6.9) is that a b-d path Q satisfies $\widehat{z}(Q) = R$ (so it is a shortest b-d path) if and only if $|V(P) \cap V(Q)| = 1$ for every $P \in \mathcal{P}'$.

We say that a $P' \in \mathcal{P}'$ is *to the right* of $P \in \mathcal{P}'$, if every node of P' is either a common point of P and P', or is separated from b by P. Clearly this defines a partial ordering of $\mathcal{P}'$. In this case we say that P and P' *do not cross*. We can similarly define noncrossing paths in $\mathcal{Q}'$, and a partial ordering.

We continue with two somewhat more elaborate properties of the paths in $\mathcal{P}'$ and $\mathcal{Q}'$. The first of these arguments could be omitted at the cost of making the arguments later less transparent. However, under the name of "uncrossing", this method is a standard and powerful step in many proofs in graph theory, and it is worth describing.

Claim 1. *We can choose the decompositions in* (6.6) *and* (6.7) *so that the families* $\mathcal{P}'$ *and* $\mathcal{Q}'$ *are pairwise non-crossing.*

In other words, "to the right" defines a total order on $\mathcal{P}'$. Similarly, we get a total order on $\mathcal{Q}'$.

To prove the Claim, let $P, P' \in \mathcal{P}'$ be a pair of paths that are crossing. Using that their common nodes are in the same order along both, we can construct two other a-c paths P_0 and P_0', consisting of the common nodes and those nodes of each path that are to the left (resp. to the right) of the other path (Figure 6.8, left). Clearly $\widehat{z}(P) + \widehat{z}(P') = \widehat{z}(P_0) + \widehat{z}(P_0')$, and since P and P' satisfy (6.3) with equality, so do P_0 and P_0'. Let $\varepsilon = \min\{\lambda_P, \lambda_{P'}\}$, and let us decrease λ_P and $\lambda_{P'}$ by ε and increase λ_{P_0} and $\lambda_{P_0'}$ by ε. This gives another decomposition in (6.6). So one of P and P' (say, P') drops out of $\mathcal{P}'$, and P_0 and P_0' enter it (they may have been there already, in which case only their weight increases).

This describes the main step, but we want to repeat this procedure to get rid of all crossings; to this end, we have to justify that we are making progress. This is not quite obvious, since we have replaced two paths by (possibly) three, creating new crossings with the other paths. Looking more carefully, we see that

(a) the remaining path P and the new paths P_0 and P_0' are mutually noncrossing, and

(b) if a further path in $\mathcal{P}'$ crosses either one of P_0 and P_0', then it crosses at least one of P and P'; and if it crosses both P_0 and P_0', then it crosses both P and P'.

From these observations it follows that the sum

$$\sum_{P_1,P_2\in\mathcal{P}' \text{ crossing}} \lambda_{P_1}\lambda_{P_2}$$

decreases at the above operation, and so if we start with a decomposition minimizing this sum, then the family $\mathcal{P}'$ will consist of non-crossing paths. We argue for $\mathcal{Q}'$ similarly. This proves the Claim.

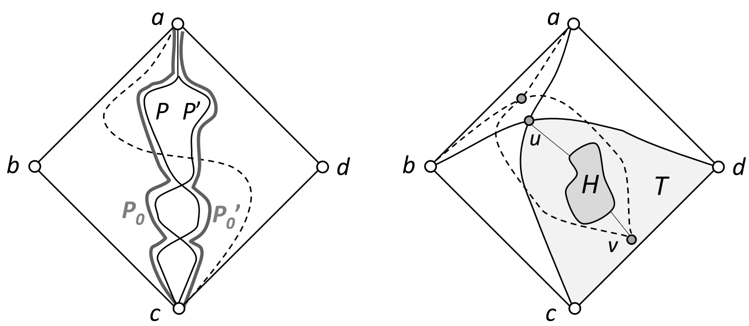

FIGURE 6.8. Left: uncrossing two paths. Right: every node has positive weight.

Let $u \in V'$. Every path in $\mathcal{P}'$ either goes through u or it separates u from exactly one of b or d in the topology of the plane. Let $\mathcal{P}_u^-$, $\mathcal{P}_u$ and $\mathcal{P}_u^+$ denote the sets of paths in $\mathcal{P}'$ separating u from b, passing through u, and separating u from d, respectively. Clearly the sets $\mathcal{P}_u^-$, $\mathcal{P}_u$ and $\mathcal{P}_u^+$ form intervals in the ordering of $\mathcal{P}'$. We define the partition $\mathcal{Q}' = \mathcal{Q}_u^- \cup \mathcal{Q}_u \cup \mathcal{Q}_u^+$ analogously: the paths in $\mathcal{Q}^-$ separate u from a etc. If $Q \in \mathcal{Q}'$ is any path through u, then it is easy to tell which paths in $\mathcal{P}'$ separate u from b: exactly those whose unique common node with Q lies on the subpath $Q[b,u)$.

Clearly $\lambda(\mathcal{P}_u) = \mu(\mathcal{Q}_u) = z_u$. We can express the distance of a node $u \in V'$ with $z_u > 0$ from b as

$$d_z(b,u) = \frac{z_u}{2} + \lambda(\mathcal{P}_u^-). \tag{6.11}$$

Indeed, choosing any $Q \in \mathcal{Q}_u$, a path $P \in \mathcal{P}'$ belongs to $\mathcal{P}_u^-$ if and only if it intersects $Q(b,u)$, and such paths contribute λ_P to z_v for a single node v on $Q(b,u)$.

Claim 2. *Every node $i \in V'$ satisfies $z_i > 0$.*

Suppose that there are nodes with $z_i = 0$, and let H be a connected component of the subgraph induced by these nodes. There are no separating 3- and 4-cycles by the hypothesis of the theorem, hence H has at least five neighbors in $V \setminus V(H)$. Each neighbor v of H has weight $z_v > 0$, and so it is contained in a path $P_v \in \mathcal{P}'$ as well as in a path $Q_v \in \mathcal{Q}'$. The path P_v is disjoint from H and separates H either from b or from d, and similarly, the path Q_v separates H from either a or c. Thus there are two neighbors u and v of H so that both P_u and P_v separate H from (say) b, and both Q_u and Q_v separate h from (say) a. We choose the same path for P_u and P_v if possible, and similarly for Q_u and Q_v, but (6.8) implies that we cannot succeed in both cases. So we may assume that $Q_u \neq Q_v$.

The rest of the argument consists of analyzing the topology around H. If $P_u = P_v$, then we may assume that u comes before v along the path P_u. The path Q_v separates u from a, so P_u meets Q_v before u, and then meets it again at v, contradicting (6.8).

Suppose that $P_u \neq P_v$. Then, as before, it follows that $P_u[a,u)$ meets Q_v. Since P_u and Q_v have only one intersection point, it follows that Q_v is disjoint from $P_u[u,c]$. Similarly, P_v is disjoint from $Q_u[u,d]$.

The paths P_u and Q_u cross each other at u and split the interior of $abcd$ into four open parts (Figure 6.8, right). The subgraph H is contained in the part T incident with the edge cd, and since v is a neighbor of H and not on the paths P_u and Q_u, it must also be in T. Now the path $P_v[a,v]$ must enter T through a node of $P_u(u,c]$, since it is disjoint from $Q_u[u,d]$. Similarly, the path $Q_v[b,v]$ must enter T though $Q_u(u,d]$. But then $P_v[a,v]$ and $Q_v[b,v]$ must have an intersection point outside T, different from v, contradicting (6.8).

After all this preparation, we can describe the squares representing G. We start with the rectangle $\mathfrak{R}$ with one vertex at $(0,0)$ and the opposite vertex at $(R, 1/R)$. Every node $u \in V'$ will be represented by a square S_u centered at $\mathbf{p}_u = (d_z(a,u), d_z(b,u))^\mathsf{T}$ and having edge length z_u. We represent node a by the square S_a of side length R attached to the top of $\mathfrak{R}$. We represent b, c and d similarly by squares attached to the other edges of $\mathfrak{R}$.

Claim 3. *Let $i,j \in V$. If i and j are adjacent, then S_i and S_j are tangent. If i and j are nonadjacent, then S_i and S_j are either disjoint or tangent along a single vertex.*

The Claim is easily checked when i and/or j belong to $\{a,b,c,d\}$, so suppose that $i,j \in V'$.

First, let $ij \in E(P)$, where (say) $P \in \mathcal{P}'$. Assume that P encounters i before j. Then $d_z(a,j) = d_z(a,i) + \widehat{z}_{ij}$ (since P is a shortest path). On the other hand, let (say) $d_z(b,j) \geq d_z(b,i)$, and let $Q \in \mathcal{Q}'$ go through i and $Q' \in \mathcal{Q}'$ go through j. Then $Q'' = Q[b,i] \cup \{ij\} \cup Q'[j,d]$ is a b-d path, and since Q'' intersects P in at least two nodes, it is not a shortest b-d path. Thus $z(Q'') > R$, which implies that

$$d_z(b,i) = \widehat{z}(Q[b,i]) > R - \widehat{z}(Q'[j,d]) - \widehat{z}_{ij} = \widehat{z}(Q'[b,j]) - \widehat{z}_{ij} = d_z(b,j) - \widehat{z}_{ij}.$$

So $|d_z(b,j) - d_z(b,i)| < \widehat{z}_{ij}$. This proves that the squares S_i and S_j are tangent along horizontal edges.

Second, let $ij \in E(G)$, and assume that ij is an not an edge of any path in $\mathcal{P}' \cup \mathcal{Q}'$. It is clear that $\mathcal{P}_i \cap \mathcal{P}_j = \emptyset$, since for any path in $\mathcal{P}'$ passing through both i and j, the edge ij would be a chord. Since $\mathcal{P}_i$ and $\mathcal{P}_j$ are intervals in the ordering of $\mathcal{P}'$, it follows that (say) $\mathcal{P}_i$ is completely to the left of $\mathcal{P}_j$. Since i and j are adjacent, no path in $\mathcal{P}'$ can separate i and j, and hence $\cap\mathcal{P}_j^- = \mathcal{P}_i^- \cup \mathcal{P}_i$. Applying (6.11) for both i and j, we see that

$$d_z(b,j) - d_z(b,i) = \frac{z_j}{2} - \frac{z_i}{2} + \lambda(\mathcal{P}_j^-) - \lambda(\mathcal{P}_i^-) = \frac{z_j}{2} - \frac{z_i}{2} + \lambda(\mathcal{P}_i) = \frac{1}{2}(z_i + z_j). \tag{6.12}$$

Finally, assume that i and j are nonadjacent nodes. If there is a path $Q \in \mathcal{Q}'$ going through both i and j, encountering (say) i before j, then

$$d_z(b,j) = d_z(b,i) + d_z(i,j) > d_z(b,i) + \frac{z_i + z_j}{2}$$

and hence S_i and S_j are disjoint. We conclude similarly is there is a path in $\mathcal{P}'$ through both i and j. We are left with the case when $\mathcal{P}_i \cap \mathcal{P}_j = \mathcal{Q}_i \cap \mathcal{Q}_j = \emptyset$. As in the previous proof, we may assume that $\mathcal{P}_i$ is completely to the left of $\mathcal{P}_j$, but we can only conclude that $\cap\mathcal{P}_j^- \subseteq \mathcal{P}_i^- \cup \mathcal{P}_i$, and so the computation in (6.12) gives an inequality only: $d_z(b,j) - d_z(b,i) \geq \frac{1}{2}(z_i + z_j)$. Similarly, $|d_z(a,j) - d_z(a,i)| \geq \frac{1}{2}(z_i + z_j)$. If strict inequality holds in both directions, then S_i and S_j are disjoint; if equality holds in both cases, then S_i and S_j have a vertex in common. This proves Claim 3.

Consider G as drawn in the plane so that node i is at the center $\mathbf{p}_i$ of the square S_i, and every edge ij is a straight segment connecting $\mathbf{p}_i$ and $\mathbf{p}_j$. This gives a planar embedding of G. Indeed, we know from Claim 3 that every edge ij is covered by the squares S_i and S_j. This implies that edges do not cross, except possibly in the degenerate case when four squares S_i, S_j, S_k, S_l share a vertex (in this clockwise order around the vertex), and $ik, jl \in E$. Since the squares S_i and S_j are touching along their vertical edges, we have $ij \in E$. Similarly, $jk, kl, li \in E$, and hence i, j, k, l form a complete 4-graph. But this is impossible in a triangulation of a quadrilateral that has no separating triangles.

Next, we argue that the squares S_i ($i \in V \setminus \{a,b,c,d\}$) tile the rectangle $\mathfrak{R} = [0,R] \times [0,1/R]$. It is easy to see that all these squares are contained in the rectangle $\mathfrak{R}$, and they are non-overlapping by Claim 3. The total area covered by the squares is $\sum_{i \in V'} z_i^2 = 1$, which is just the area of $\mathfrak{R}$, so they must cover $\mathfrak{R}$.

Finally, we show that an appropriately resolved tangency graph of the squares S_i is equal to G. By the above, it contains G (where for edges of type (iii), the 4-corner is resolved so as to get the edge of G). Since G is a triangulation of the outside quadrilateral, the containment cannot be proper, so G is the whole tangency graph. □

Exercise 6.1. Figure out how to resolve the common points of four squares in the tiling in Figure 6.2 in order to get the dodecahedron graph.

Exercise 6.2. Verify that the graph in Figure 6.9 gives rise to a tiling of a square by different size squares. (This construction provides the minimum number of different squares tiling a square.)

Exercise 6.3. Prove that if G is a resolved tangency graph of a square tiling of a rectangle, then every triangle in G is a country.

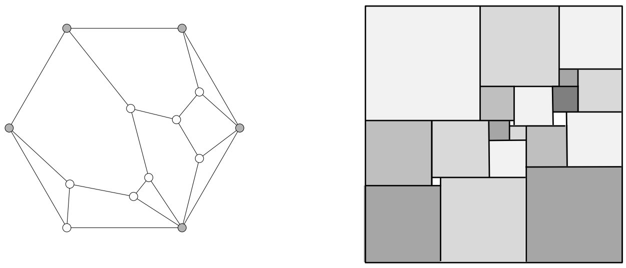

FIGURE 6.9. Tiling a square with the smallest number of different squares.

Exercise 6.4. Construct a resolved tangency graph of a square tiling of a rectangle that contains a quadrilateral with a further node in its interior.

Exercise 6.5. Let G be a maximal planar graph with at least 4 nodes, with unbounded country abc. Let us 2-color the nodes different from a, b and c with black and white. Prove that either there is a black node connected by black paths to each of a, b and c, or there is a white node connected by white paths to each of a, b and c, but not both.