

Computing Cartograms with Optimal Complexity

Md. Jawaherul Alam

Department of Computer Science,
University of Arizona,
Tucson, AZ, USA
mjalam@email.arizona.edu

Stefan Felsner

Institut für Mathematik,
Technische Universität Berlin,
Berlin, Germany
felsner@math.tu-berlin.de

Stephen G. Kobourov

Department of Computer Science,
University of Arizona,
Tucson, AZ, USA
kobourov@cs.arizona.edu

Therese Biedl

School of Computer Science,
University of Waterloo,
Waterloo, ON N2L 3G1, Canada
biedl@uwaterloo.ca

Michael Kaufmann

Wilhelm-Schickhard-Institut für Informatik,
Eberhard Karls Universität Tübingen,
Tübingen, Germany
mk@informatik.uni-tuebingen.de

Torsten Ueckerdt

Institut für Mathematik,
Technische Universität Berlin,
Berlin, Germany
ueckerdt@math.tu-berlin.de

Abstract

We show how to compute cartograms with worst-case optimal polygonal complexity. Specifically we study *rectilinear duals* which are side-contact representations of a planar graph G with vertices represented by simple rectilinear polygons and adjacencies represented by a non-trivial contact between the corresponding polygons. A rectilinear dual is called a *cartogram* if the area of each region is equal to a pre-specified weight of the corresponding vertex. We show how to construct an area-universal rectilinear dual, i.e., one that can realize a cartogram with any pre-specified set of weights for the vertices of the graph. In a series of papers the polygonal complexity of such representations for maximal planar graphs was reduced in several steps from the initial 40 to 34, then to 12 and very recently to the currently best known 10, while it has been known that 8-sided polygons are sometimes necessary. Our construction uses only 8-sided polygons and hence is optimal in terms of polygonal complexity. The layout can be computed in linear time and has the appealing property that it is area-universal. The area-universality is inherited from a refinement which is an area-universal rectangular layout. The construction of the actual cartogram relies on a realization of the rectangular layout which requires numerical iteration.

For Hamiltonian maximal planar graphs we have an alternative construction which allows us to directly compute cartograms with 8-sided rectilinear polygons in linear time. Moreover, we prove that 8-sided rectilinear polygons are necessary by constructing a non-trivial lower bound example. If the Hamiltonian path has the extra property that it is one-sided, then we can reduce the polygonal complexity and realize cartograms with 6-sided polygons. This can be used to obtain 6-sided cartograms of maximal outer-planar graphs. Thus we have optimal (in terms of both polygonal complexity and running time) representations for Hamiltonian maximal planar and maximal outer-planar graphs.

1 Introduction

There is a large body of work about representing planar graphs as contact graphs, i.e., graphs whose vertices are represented by geometrical objects with edges corresponding to two objects touching in some specified fashion. Typical classes of objects might be curves, line segments, or polygons. An early result is Koebe’s 1936 theorem [12] that all planar graphs can be represented by touching disks.

In this paper, we consider contact representations of planar graphs, with vertices represented by simple polygons and adjacencies represented by a non-trivial contact between the corresponding polygons. We are specifically interested in the rectilinear weighted version where the vertices are represented by simple rectilinear polygons. This type of a representation is known as a *rectilinear dual* of the input planar graph.

In the weighted version the input is a planar graph $G = (V, E)$ along with a weight function $w : V(G) \rightarrow R^+$ that assigns a weight to each vertex of G . A rectilinear dual is called a *cartogram* if the area of each region is equal to the pre-specified weight of the corresponding vertex. Such representations have practical applications in cartography, geography, and sociology, but also in VLSI Layout, and floor-planning. Other applications can be found in visualization of relational data, where using the adjacency of regions to represent edges in a graph can lead to a more compelling visualization than drawing a line segment between two points [3].

For rectilinear duals (unweighted) and for cartograms (weighted) it is often desirable, for aesthetic, practical and cognitive reasons, to limit the *polygonal complexity* of the representation, measured by the number of sides (or by the number of corners). Similarly, it is also desirable to minimize the unused area in the representation, also known as “holes” in floorplanning and VLSI layouts. A given rectilinear dual is *area-universal* if it can realize a cartogram with any pre-specified set of weights for the vertices of the graph without disturbing the underlying adjacencies and without increasing the polygonal complexity.

With these considerations in mind, we study the problem of constructing area-universal rectilinear duals and show how to compute cartograms with worst-case optimal polygonal complexity and without any wasted area.

1.1 Related Work

In our paper and in most of the other papers cited here, “planar graph” refers to an inner-triangulated planar graph with a simple outer-face; this restriction is required if at most three rectilinear polygons are allowed to meet in a point and the union of all the polygons in the representation is a rectangle.

Rectilinear duals (unweighted) were first studied in graph theoretic context, and then with renewed interest in the context of VLSI layouts and floor planning. It is known that 8 sides are sometimes necessary and always sufficient [9, 15, 22].

The case when the rectilinear polygons are restricted to rectangles has been of particular interest and there are several (independent) characterizations of the class of planar graphs that allows such *rectangular duals* [19, 14, 13]. Buchsbaum *et al.* [3] give a historical overview and summarize the state of the art in the rectangle contact graphs literature.

In the above results on rectilinear duals and rectangular duals, the areas of the polygons is not considered; that is, these results deal with the unweighted version of the problem. The weighted version dates back to 1934 when Raisz described rectangular cartograms [16]. Algorithms by van

Kreveld and Speckman [20] and Heilmann *et al.* [10] yield representation with touching rectangles but the adjacencies may be disturbed and there can also be a small distortions on the weight. Recently, Eppstein *et al.* [8] characterized the class of planar graphs that have area-universal rectangular duals. Here area-universality of a rectangular dual implies that for any area-function for the vertices of the graph, there exists a combinatorially equivalent cartogram. The construction of the actual cartogram, given the area-universal rectilinear dual and the weight function, can be accomplished using a result by Wimer *et al.* [21] which in turn requires numerical iteration.

The result of Eppstein *et al.* above is restricted to planar graphs that have rectangular duals. Going back to the more general rectilinear duals, leads to a series of papers where the main goal has been to reduce the polygonal complexity while respecting all areas. De Berg *et al.* initially showed that 40 sides suffice [5]. This was later improved to 34 sides [11]. In a recent WADS [2] paper the polygonal complexity was reduced to 12 sides and an even more recently to 10 sides [1].

1.2 Our Results

Recall that the known lower bound on the polygonal complexity even for unweighted rectilinear duals is 8 [22] while the best known upper bound is 10. Here we present the first construction that matches the lower bound. Specifically, our construction produces 8-sided area-universal rectilinear duals in linear time, and is thus optimal in terms of the polygonal complexity. The realization of an actual cartogram with pre-specified set of weights given an area-universal rectilinear dual, can be accomplished using numerical iteration [21].

For Hamiltonian maximal planar graphs we have an alternative construction which allows us to directly compute cartograms with 8-sided rectilinear polygons in linear time. Moreover, we prove that 8-sided rectilinear polygons are necessary by constructing a non-trivial lower bound example. If the Hamiltonian path has the extra property that it is one-sided, then we can reduce the polygonal complexity and realize cartograms with 6-sided polygons. This can be used to obtain 6-sided cartograms of maximal outer-planar graphs. Thus we have optimal (in terms of both polygonal complexity and running time) representations for Hamiltonian maximal planar and maximal outer-planar graphs.

The rest of the paper is organized as follows. In Section 2, we review basic terminology. In Section 3, we show how to construct 8-sided area-universal rectilinear duals in linear time. In Section 4 we describe the alternative 8-sided construction for Hamiltonian graphs and prove that 8-sides is also a lower bound. In Section 5 we show that 6-sides are necessary and sufficient for one-sided Hamiltonian graphs and maximal outer-planar graphs. In Section 6 we summarize our contributions and consider several open problems.

2 Preliminaries

A *planar graph*, $G = (V, E)$, is one that has a drawing without crossing in the plane along with an embedding, defined via the cyclic ordering of edges around each vertex. A *plane graph* is a fixed planar embedding of a planar graph. It splits the plane into connected regions called *faces*; the unbounded region is the *outer-face* and all other faces are called *interior faces*. A planar (plane) graph is *maximal* if no edge can be added to it without violating planarity. Thus each face of a maximal plane graph is a triangle. A *Hamiltonian cycle* in a graph G is a simple cycle containing all the vertices of G . A graph G is called *Hamiltonian* if it contains a hamiltonian cycle.

A set P of closed simple interior-disjoint polygons with an isomorphism $\mathcal{P} : V \rightarrow P$ is a polygonal contact representation of a graph if for any two vertices $u, v \in V$ the boundaries of $\mathcal{P}(u)$ and $\mathcal{P}(v)$ share a non-empty line-segment if and only if (u, v) is an edge. Such a representation is known as a *rectilinear dual* of the input graph if the vertices are represented by simple rectilinear polygons. In the weighted version the input is the graph G , along with a weight function $w : V(G) \rightarrow \mathbb{R}^+$ that assigns a weight to each vertex of G . A rectilinear dual is called a *cartogram* if the area of each polygon is equal to the pre-specified weight of the corresponding vertex. We define the *complexity of a polygon* as the number of sides it has. A common objective is to realize a given graph and a set of weights, using polygons with minimal complexity.

2.1 Canonical Orders and Schnyder Realizers

Next we briefly summarize the concepts of a “canonical order” of a planar graph [7] and that of a “Schnyder realizer” [18]. Let $G = (V, E)$ be a maximal plane graph with outer vertices u, v, w in clockwise order. Then we can compute in linear time [4] a *canonical order* or *shelling order* of the vertices $v_1 = u, v_2 = v, v_3, \dots, v_n = w$, which is defined as one that meets the following criteria for every $4 \leq i \leq n$.

- The subgraph $G_{i-1} \subseteq G$ induced by v_1, v_2, \dots, v_{i-1} is biconnected, and the boundary of its outer face is a cycle C_{i-1} containing the edge (u, v) .
- The vertex v_i is in the exterior face of G_{i-1} , and its neighbors in G_{i-1} form an (at least 2-element) subinterval of the path $C_{i-1} - (u, v)$.

A *Schnyder realizer* of a maximal plane graph G is a partition of the interior edges of G into three sets $\mathcal{S}_1, \mathcal{S}_2$ and \mathcal{S}_3 of directed edges such that for each interior vertex v , the following conditions hold:

- v has out-degree exactly one in each of $\mathcal{S}_1, \mathcal{S}_2$ and \mathcal{S}_3 ,
- the counterclockwise order of the edges incident to v is: entering \mathcal{S}_1 , leaving \mathcal{S}_2 , entering \mathcal{S}_3 , leaving \mathcal{S}_1 , entering \mathcal{S}_2 , leaving \mathcal{S}_3 .

Schnyder proved that any maximal plane graph has a Schnyder realizer and it can be computed in $O(n)$ time [18]. The first condition implies that \mathcal{S}_i , for $i = 1, 2, 3$ defines a tree rooted at exactly one exterior vertex and containing all the interior vertices such that the edges are directed towards the root.

The following well-known lemma shows a profound connection between canonical orders and Schnyder realizers.

Lemma 2.1 *Let G be a maximal plane graph. Then the following (a)–(b) hold.*

- A canonical order of the vertices of G defines a Schnyder realizer of G , where the outgoing edges of a vertex v are to its first and last predecessor (where “first” is w.r.t. the clockwise order around v), and to its highest-numbered successor.*
- A Schnyder realizer with the three trees $\mathcal{S}_1, \mathcal{S}_2, \mathcal{S}_3$ defines a canonical order, which is a topological order of the acyclic graph $\mathcal{S}_1^{-1} \cup \mathcal{S}_2^{-1} \cup \mathcal{S}_3$, where \mathcal{S}_k^{-1} is the tree \mathcal{S}_k with the direction of all its edges reversed.*

3 Cartograms with 8-Sided Polygons

In this section we show that 8-sided polygons are always sufficient and sometimes necessary for a cartogram of a maximal planar graph. We present a linear time algorithm that produces a rectilinear dual with 8-sided rectilinear polygons that is *area-universal*, i.e., any assignment of areas to the polygons can be realized by a combinatorially equivalent layout.

Our algorithm for constructing area-universal rectilinear duals has three main phases. In the first phase of the algorithm we create a contact representation of the graph G , where each vertex of G is represented by an (inverted) \mathbf{T} , which consists of a horizontal segment and a vertical segment. Figures 1(a)-(b) show a maximal planar graph and its contact representation using \mathbf{T} 's, where the three ends of each \mathbf{T} are marked with arrows. In the second phase of the algorithm we make both the horizontal and vertical segments of each \mathbf{T} into thin polygons with λ thickness for some $\lambda > 0$. We then have a contact representation of G with T -shaped polygons as illustrated in Figure 1(c). In the third phase of the algorithm we remove all the unused area in the representation by assigning each (rectangular) hole to one of the polygons adjacent to it, as illustrated in Figure 1(d). We show that the resulting representation is an area-universal rectilinear dual of G with polygonal complexity 8, as illustrated in Figure 1(e).

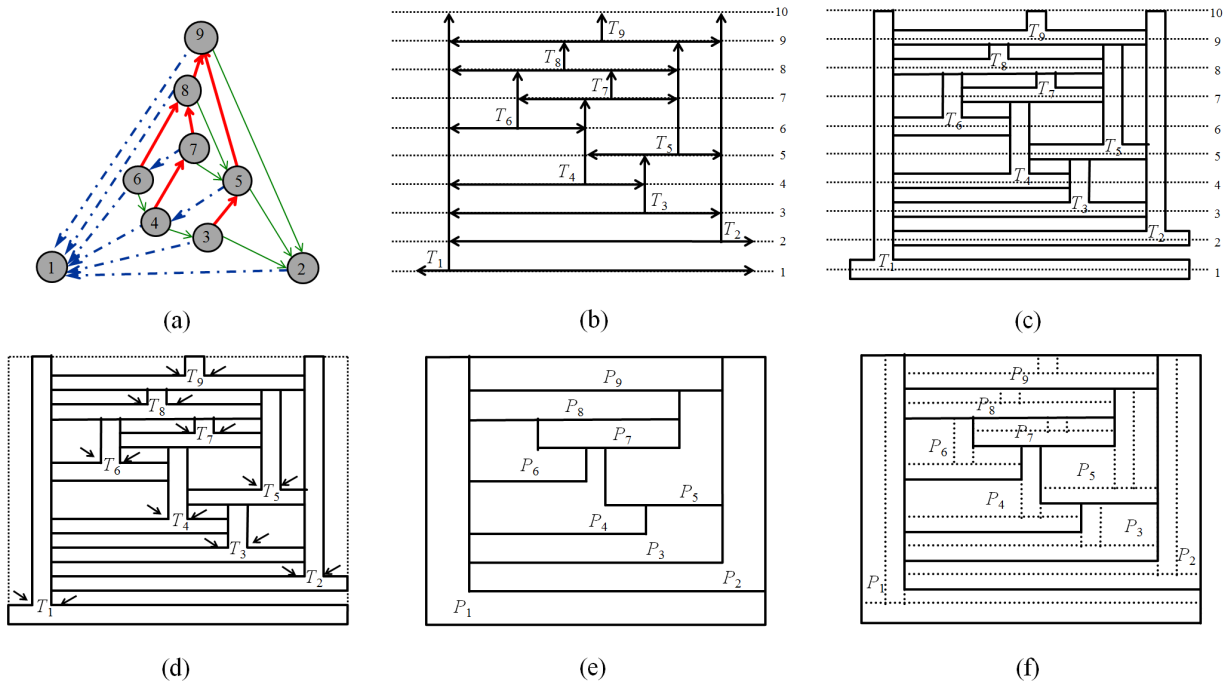


Figure 1: Illustration for the algorithm to obtain a rectilinear dual of a maximal planar graph with 8-sided polygons.

3.1 Constructing Contact Representation with \mathbf{T} 's

Our contact representation with \mathbf{T} 's is similar to the approach described by de Fraysseix *et al.* [6]. Let G be a maximal planar graph and let $v_1, v_2, v_3, \dots, v_n$ be a canonical order of the vertices in

G . Recall that G_k , $3 \leq k \leq n$, is the subgraph of G induced by the vertices v_1, v_2, \dots, v_k and C_k is the outer-cycle of G_k . The corresponding Schnyder trees $\mathcal{S}_1, \mathcal{S}_2$ and \mathcal{S}_3 rooted at v_1, v_2 and v_n , can be easily computed using Lemma 2.1. We also add to \mathcal{S}_1 the two edges (v_2, v_1) and (v_n, v_1) oriented towards v_1 and add to \mathcal{S}_2 the edge (v_n, v_2) oriented towards v_2 . Figure 1(a) illustrates an embedding of a maximal planar graph G , a canonical order of the vertices, and the corresponding Schnyder trees for G . Denote by $\Phi_i(k)$, $i = 1, 2, 3$ the canonical label of the parent of v_k in the tree \mathcal{S}_i . Let T_i be the \mathbf{T} representing v_i , with horizontal and vertical segments h_i and b_i .

The algorithm begins by placing T_1 and T_2 such that h_1 is placed at $y = 1$, h_2 is placed at $y = 2$, the topmost points of both b_1 and b_2 have y -coordinate $n + 1$ and the leftmost point of the h_2 touches b_1 . Next the algorithm iteratively constructs the contact representation by defining T_k so that h_k is placed at $y = k$ and the topmost point of b_k has y -coordinate $\Phi_3(k)$ for $3 \leq k < n$. After the k -th step of the algorithm we have a contact representation of G_k , and we maintain the invariant that the order of the vertical segments with non-empty parts in the half-plane $y > k$ corresponds to the same circular order of the vertices along $C_k - (v_1, v_2)$.

Consider inserting T_k for v_k . The neighbors $v_{k_1}, v_{k_2}, \dots, v_{k_d}$ of v_k in G_{k-1} form a subinterval of $C_{k-1} - (v_1, v_2)$ and hence the corresponding vertical segments are also in the same order in the half-plane $y > k - 1$ of the representation of G_{k-1} . Since v_k is the parent in \mathcal{S}_3 for all the vertices v_{k_i} , $1 < i < d$ (Lemma 2.1), the topmost points of the corresponding vertical segments have y -coordinate k . As v_{k_1} and v_{k_d} are the parents of v_k in \mathcal{S}_1 and \mathcal{S}_2 , the x -coordinates of b_{k_1} and b_{k_d} define the x -coordinates of the two endpoints of h_k . Let us assume that these coordinates are x_l and x_r ; then h_k is placed between the two points (x_l, k) , (x_r, k) and b_k is placed between the two points (x_m, k) , $(x_m, \Phi_3(k))$ with $x_l + 1 < x_m < x_r - 1$. Finally for $k = n$, we place T_n such that h_n touches b_1 to the left and b_2 to the right and the topmost point of b_n has y -coordinate $n + 1$. Figure 1(b) illustrates a contact representation of the maximal plane graph in Figure 1(a) computed using this algorithm.

3.2 λ -Fattening of T_i 's

Let Γ' be the contact representation of G using \mathbf{T} 's obtained by the previous algorithm. In this phase of the algorithm, we “fatten” \mathbf{T} 's so that each vertex is represented by a T -shaped polygon. We replace each horizontal segment h_i by an axis-aligned rectangle H_i which has the same width as h_i , and whose top (bottom) side is $\lambda/2$ above (below) h_i , for some $0 < \lambda$, as illustrated in Figure 2(a). Similarly, we replace each vertical segment b_i by an axis-aligned rectangle B_i which has the same height as b_i and whose left (right) side is $\lambda/2$ to the left (right) of b_i . We call this process λ -fattening of T_i . Note that this process creates intersections of H_i with B_i , $B_{\Phi_1(i)}$ and $B_{\Phi_2(i)}$ and intersection of B_i with $H_{\Phi_3(i)}$. We remove these intersections by replacing H_i by $H_i - B_{\Phi_1(i)} - B_{\Phi_2(i)}$ and replacing B_i by $B_i - H_i - H_{\Phi_3(i)}$. The resulting layout would then be a contact representation Γ'' of G where each vertex v_i of G is represented by the T -shaped polygon $H_i \cup B_i$. Figure 1(c) illustrates such a representation of maximal plane graph in Figure 1(a).

3.3 Removing unused area

In this step, we begin with the λ -fat T -shaped polygonal layout, Γ'' , from above and remove all of the unused area, to obtain a rectilinear dual of G with eight-sided polygons. We do this by assigning each (rectangular) hole to one of the polygons adjacent to it. We start by placing an axis-aligned rectangle of minimum size that encloses Γ'' . This creates five new bounded holes out

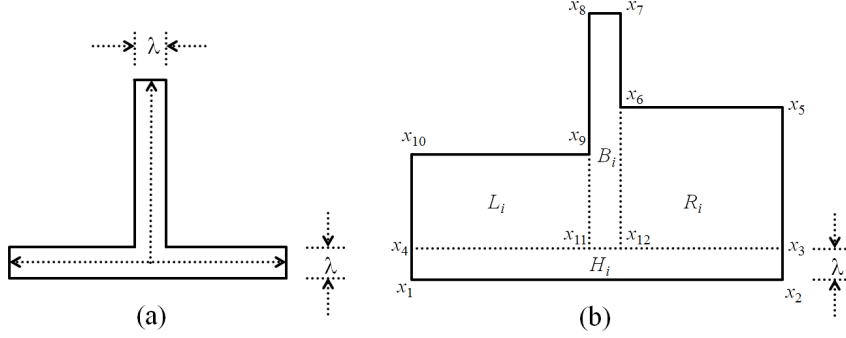


Figure 2: (a) λ -fattening of T , and (b) subdividing a T -shaped polygon into four rectangles.

of the unbounded region outside Γ'' : (i) to the left of B_1 and above H_1 , (ii) to the right of B_1 and between H_1 and H_2 , (iii) to the right of B_2 and above H_2 , (iv) to the left of B_n and above H_n , (v) to the right of B_n and above H_n . Call these L_1, R_1, R_2, L_n and R_n , respectively. Now we associate each T -shaped polygon, representing v_i in Γ'' , with two holes: one above H_i and to the left of B_i , which we call the *left hole* of v_i ; and the other to the top of H_i and to the right of B_i , which we call the *right hole* of v_i , as illustrated in Figure 1(d). Note that the left hole of v_i is a rectangle bounded by the top of H_i , the left of B_i , the right of $B_{\Phi_1(i)}$ and the bottom of H_j where v_j is either the parent of v_i in \mathcal{S}_3 or the child of v_i in \mathcal{S}_2 with the lowest canonical label. Similarly the right hole of v_i is a rectangle bounded by the top of H_i , the right of B_i , the left of $B_{\Phi_2(i)}$ and the bottom of H_j , where v_j is either the parent of v_i in \mathcal{S}_3 or the child of v_i in \mathcal{S}_1 with the lowest canonical label. (The left hole of v_1 is L_1 , the right hole of v_2 is R_2 and the left and right holes of v_n are L_n and R_n , respectively.) For $1 \leq i \leq n$, denote by L_i and R_i the left hole and the right hole of v_i .

Lemma 3.1 *Each hole in Γ'' is uniquely associated with exactly one polygon by the above procedure.*

Proof: First note that each hole in Γ'' (other than the five newly added ones) corresponds to a triangular inner face of G and is surrounded by the three polygons representing the vertices of the face. Thus each hole contains exactly one reflex corner of one of the three polygons. Each of the five newly added holes also contains exactly one reflex corner of a polygon. With $2n - 5$ inner faces, there are exactly $2n$ holes while each T -shape polygon for each of the n vertex gives two reflex corners. Since no hole shares a reflex corners with any other hole, each hole is uniquely assigned to the polygons. Therefore each hole in Γ'' is either the left hole or the right hole of v_i for some $1 \leq i \leq n$. #

We now can combine each T -shape region with its two associated holes to obtain an 8-sided rectilinear polygon. Specifically, for each vertex v_i , define $P_i = T_i \cup L_i \cup R_i$. It is easy to see that P_i is an 8-sided rectilinear polygon since the left side of L_i has the same x -coordinate as the left side of H_i and the right side of R_i has the same x -coordinate as the right side of H_i . Thus we have a rectilinear dual, Γ , of G where each vertex v_i is represented by P_i . Figure 1(e) illustrates such a dual for the graph in Figure 1(a).

3.4 Area-Universality

A rectilinear dual Γ is *area-universal* if any assignment of areas to its polygons can be realized by a combinatorially equivalent layout. Eppstein *et al.* [8] introduced the concept of “area-universality” for the case when all the polygons are rectangles and the outer-face boundary is also a rectangle (which they call a *rectangular layout*). They gave a characterization of area-universal rectangular layouts using the concept of “maximal line-segment”. A *line-segment* of a layout is the union of inner edges of the layout forming a consecutive part of a straight-line. A line-segment that is not contained in any other line-segment is maximal. A maximal line-segment s is called *one-sided* if it is part of the side of at least one rectangular face, or in other words, if the perpendicular line segments that attach to its interior are all on one side of s .

Lemma 3.2 [8]. *A rectangular layout is area-universal if and only if each maximal segment in the layout is one-sided.*

No such characterization is known when some faces are not rectangles, but we can use the characterization in Lemma 3.2 to show that the rectilinear dual obtained by the algorithm from the previous section is area-universal, with the following Lemma.

Lemma 3.3 *Let Γ be the rectilinear dual obtained by the above algorithm. Then Γ is area-universal.*

Proof: To show the area-universality of Γ , we divide all the polygons in Γ into a set of rectangles such that the resulting rectangular layout is area-universal. Specifically, we divide each polygon P_i into four rectangles H_i , B_i , L_i and R_i (as defined in the previous subsection) by adding three auxiliary segments: one horizontal and two vertical, as illustrated in Figure 2(b).

Now we prove that all maximal line-segments are one-sided. Consider an arbitrary maximal horizontal line-segment s of Γ^* and say it has y -coordinate $i + \lambda/2$ (the case for $i - \lambda/2$ is symmetric) for some integer i . Then s is the top side of H_i , which is maximal since its left end is on the right side of $B_{\Phi_1(i)}$ and its right end is on the left side of $B_{\Phi_2(i)}$. Then s is one-sided since it is a side of the rectangle H_i .

Now consider an arbitrary maximal vertical line-segment s . It is then either the left or the right side of B_i for some i , which is maximal in either case since it has its bottom end on the top side of H_i and its top end on the bottom side of $H_{\Phi_3(i)}$. Again s is symmetric in this case since it is a side of B_i .

Now given any assignment of areas $w : V \rightarrow R^+$ to the vertices V of G , we split $w(v_i)$ arbitrarily into four parts and assign the four values to its four associated rectangles. Since Γ^* is area-universal, there exist a rectilinear dual of G that is combinatorially equivalent to Γ for which these areas are realized. Figure 1(f) illustrates the rectangular layout obtained from the rectilinear dual in Figure 1(e).

#

Recall that the lower bound on the complexity of polygons in any rectilinear dual (and hence in any cartogram) is 8, as proven by Yeap and Sarrafzadeh [22]. The algorithm described in this section, along with this lower bound leads to our main theorem.

Theorem 1 *Eight-sided polygons are always sufficient and sometimes necessary for a cartogram of an inner triangulated planar graph with a simple outer-face.*

3.5 Feature Size and Supporting Line Set

In addition to the polygonal complexity optimality, we point out here a practical feature of the 8-sided area-universal rectilinear layout constructed with our algorithm. Earlier constructions, e.g., [5, 2], often rely on “thin connectors” to maintain adjacencies, whereas our construction does not. Moreover, we have the freedom to choose how to divide the area assigned to any vertex v_i among the four rectangles associated with it. This flexibility makes it possible to achieve other desired properties, such as a large *minimum feature size*. Specifically, let G be a maximal planar graph and Γ be a cartogram of G , for some weight function $w : V(G) \rightarrow \mathbb{R}^+$, computed by the above algorithm. Choose W and H such that $W \times H = A = \sum_{v \in V(G)} w(v)$. W.l.o.g. we can assume that W and H are the width and the height of Γ , respectively. Define $w_{min} = \min_{v \in V(G)} w(v)$. For each vertex v_i of G , we can assign zero areas to the rectangles L_i and R_i and split its weight into two equal parts to H_i and B_i .

Since the height of each of the rectangles H_i and B_i is at most $\max(W, H)$ and their weights are at least $w_{min}/2$, the width is at least $\frac{w_{min}}{2\max(W, H)}$, and the same holds for its height. Thus the minimum feature size of Γ is at least $\frac{w_{min}}{2\max(W, H)}$ and it is worst-case optimal, as the polygon with the smallest weight might need to reach from the leftmost to the rightmost (or topmost to bottommost) polygon in the representation. We may choose $W = H = \sqrt{A}$. Then the minimum feature size is $\frac{w_{min}}{2\sqrt{A}}$. Furthermore it is not difficult to show that assigning zero areas to the rectangles L_i and R_i yields a cartogram with at most $2n$ supporting lines, instead of the $3n$ supporting lines in the original cartogram Γ .

4 Cartograms for Hamiltonian Graphs

In this section we show that 8-sided polygons are always sufficient and sometimes necessary for a cartogram of a maximal planar *Hamiltonian* graph. We prove that 8 sides are necessary by constructing a non-trivial lower bound example.

Let v_1, \dots, v_n be a Hamiltonian cycle of a maximal planar graph G . Consider a plane embedding of G with the edge (v_1, v_n) on the triangular outer-face. The Hamiltonian cycle splits the plane graph G into two outer-planar graphs which we call the *left graph* G_l and *right graph* G_r . Edges on the Hamiltonian cycle belong to both graphs. The naming is with respect to a planar drawing of G in which the vertices v_1, \dots, v_n are placed in increasing order along a vertical line, and the edges are drawn with y -monotone curves with leftmost edge (v_1, v_n) . Call this drawing Γ ; see Figure 3(a).

For each vertex v_i in the graph we define several indices as follows. First $b_l(i)$ is the index of the bottommost neighbor of v_i in G_l , i.e., $b_l(i)$ is the smallest index j such that (v_j, v_i) is an edge in G_l . Next $b_l^{-1}(i)$ is the (possibly empty) set of indices j such that $i = b_l(j)$. The other two indices, $b_r(i)$ and $b_r^{-1}(i)$, as defined analogously, using G_r in place of G_l . It is easy to see that these four indices $b_l(i)$, $b_l^{-1}(i)$, $b_r(i)$, $b_r^{-1}(i)$ can be computed for all the vertices v_i of G in linear time.

Here we define a special property of the embedded Hamiltonian cycle, called “one-sidedness”. We call a Hamiltonian cycle v_1, \dots, v_n , in an embedded maximal planar graph with (v_1, v_n) on the outer-face, *one-sided* if for every $i = 2, \dots, n$ either $b_l(i) = i - 1$ or $b_r(i) = i - 1$. In other words, a Hamiltonian cycle is one-sided, if for every vertex in the drawing Γ it has no edges going downwards on the left or on the right.

We also define two special types of edges associated with every vertex v_i , $1 < i < n$: a *left i -chord* in G is an edge $(v_k, v_{b_l(k)})$ where $k > i > b_l(k)$, and a *right i -chord* in G is analogously

defined.

4.1 Sufficiency of 8-sided Polygons

While the upper bound of 8 is not an improvement over the results of the previous section, the algorithm described here directly computes the desired cartogram in linear time, rather than first constructing area-universal rectilinear dual and then applying numerical iteration to realize the desired weights. We argue sufficiency with the following theorem.

Lemma 4.1 *Let $G = (V, E)$ be a maximal planar Hamiltonian graph and let $w : V \rightarrow R^+$ be a weight function. Then a cartogram with 8-sided polygons can be computed in linear time.*

Proof: Let v_1, \dots, v_n be a Hamiltonian cycle and Γ be its drawing defined above with (v_1, v_n) on the outer-face. Suppose R is a rectangle of width W and height H where $W \times H = \sum_{v \in V} w(v)$. We now describe our algorithm to compute a cartogram of G inside R , such that each vertex of G is represented by an 8-sided polygon; see Figure 3(b). For each vertex v_i in the graph, the corresponding polygon is subdivided into three rectangles, which we call the *left leg*, *body*, and *right leg* of v_i , and the legs have width $\lambda_i = w(v_i)/(2H + W)$. This number λ_i is the minimum feature size of the polygon for the vertex v_i .

Our algorithm is a line-sweep, where a horizontal line L sweeps Γ from bottom to top, while the algorithm iteratively computes the cartogram of G inside R . Unlike the algorithm in the previous section, this algorithm does not rely on numerical iteration, as it computes the cartogram directly. We start the construction with the polygon for v_1 , which is simply a rectangle with area $w(v_1)$ at the bottom of R ; call this rectangle R_1 . After the i -th step of our algorithm, the horizontal line L is at vertex v_i in Γ , and the algorithm has constructed the polygons for all the vertices v_1, \dots, v_i , while completely filling up a rectangle R_i at the bottom of R . The top side of R_i would contain the followings from left to right:

1. A set $S_l(i)$ of horizontal lines of width λ_i , for all the left i -chords $(v_i, v_{b_l(i)})$ with visibility up to the top of the polygon for $v_{b_l(i)}$, in the order of the intersection of these edges with L ;
2. The top side of the polygon for v_i ;
3. A set $S_r(i)$ of horizontal lines of width λ_i , for all the right i -chords $(v_i, v_{b_r(i)})$ with visibility up to the top of the polygon for $v_{b_r(i)}$, in the order of the intersection of these edges with L .

Suppose we are at the i -th step of our algorithm. We compute R_i from R_{i-1} as follows. If $b_l(i) = k < i - 1$, then due to planarity, $b_l^{-1}(i - 1)$ is empty and the rightmost left i -chord is (v_i, v_k) . Then we use the horizontal line segment of width λ_i on top of R_{i-1} with visibility up to the top of the polygon for v_k to draw the left leg of v_i . We also now have $S_l(i) = S_l(i - 1) - h$; see Figure 3(c). Otherwise, if $b_l(i) = i - 1$, then we add to the set $S_l(i - 1)$ horizontal line-segments with widths λ_j for all the indices $j \in b_l^{-1}(i - 1)$ to obtain the set $S_l(i)$. We then draw the left leg of i with width λ_i on top of the polygon for v_{i-1} ; see Figure 3(d). We compute $S_r(i)$ and the right leg of v_i in a similar fashion. The top of the body of v_i spans from the right of its left leg to the left of its right leg. We shift the top of R_i to the appropriate y -coordinate such that the polygon for v_i (the union of the left, the right leg, and the body) has area $w(v_i)$.

From this construction it is easy to see that, after the n -step of the algorithm, the polygon for each vertex v_i inside R has the correct area. We now show that the correct adjacencies are also




Figure 3: (a) A maximal planar Hamiltonian graph G , (b) an 8-sided polygon for vertex i , (c)–(d) illustration for the algorithm to construct a cartogram of G , (e) a cartogram of G with 8-sided polygons.

preserved. Consider an edge $e = (v_i, v_j)$ of G , where $i > j$. If e is a Hamiltonian edge, i.e., $i = j + 1$, then clearly the polygons for v_i and v_j are adjacent. Next consider the case that e is in G_l (the case that e is in G_r is analogous). If $j = b_l(i)$, then this adjacency is maintained by the left leg of v_i . Otherwise suppose $b_l(i) = k$; then (v_i, v_k) is the rightmost left j -chord and hence the left leg of v_j is adjacent to the left leg of v_i . Therefore, the algorithm has computed the desired cartogram of G . It is also not difficult to argue that this algorithm runs in linear time. $\#$

The proof of Lemma 4.1 gives an algorithm for constructing a cartogram of a maximal planar Hamiltonian graph with 8-sided polygons. Figure 3(e) illustrates a cartogram for the graph in Figure 3(a) computed by the above algorithm.

Let G be a maximal planar Hamiltonian graph and let Γ be the cartogram of G obtained by the above algorithm. It is easy to show that Γ is also an area-universal rectilinear dual (using a similar argument to that in the previous section). However, instead of using the area-universality of the rectilinear dual and then numerical iteration to realize a desired weight function, we described a linear-time algorithm for the direct computation of the cartogram. This algorithm also guarantees a minimum feature size for the cartogram equal to $\min_{v_i \in V} \lambda_i = \frac{w_{\min}}{2H+W}$, where $w_{\min} = \min_{v \in V} w(v)$. We may choose $W = H = \sqrt{A}$, where $A = \sum_{v \in V} w(v)$. Then the minimum feature size is $\frac{w_{\min}}{3\sqrt{A}}$.

4.2 Necessity of 8-sided Polygons

While it was known that 8-sided rectilinear polygons are necessary for general planar graphs [17], here we show that 8-sided rectilinear polygons are necessary even for maximal planar Hamiltonian graphs. We argue necessity with the following lemma.

Lemma 4.2 *Consider the maximal planar Hamiltonian graph $G = (V, E)$ in Figure 4(a). Define $w(i) = w(j) = D$ and $w(v) = \delta$ for $v \in V \setminus \{i, j\}$, where $D \gg \delta$. Then any cartogram of G with weight function w requires at least one 8-sided polygon.*

Proof Sketch Assume for a contradiction that G admits a cartogram Γ with respect to w such that all the polygons used in Γ have complexity at most 6. It is easy to see that if $\{u, v, w\}$ is some separating triangle in G , i.e., three mutually adjacent vertices whose removal disconnect the graph,




Figure 4: (a) A maximal planar Hamiltonian graph with a weight function that requires at least one 8-sided polygon in any cartogram, (b) illustration for the proof of Lemma 4.2.

then the region used for the inside of the separating triangle contains at least one reflex corner of the polygon of u , v , or w . Note that the 5-vertex set $\{a, c, e, g, i\}$ in G is the union of the five separating triangles $\{a, c, g\}$, $\{a, c, e\}$, $\{c, e, g\}$, $\{a, e, i\}$, and $\{e, g, i\}$ with disjoint interiors. Since all the polygons in Γ are either 4-sided or 6-sided, the union of the polygons for these five vertices has at most five reflex corners and hence each of the five separating triangles above contains the only reflex corner of the polygon for a , c , e , g , or i . A similar argument shows that the separating triangles $\{a, i, m\}$, $\{i, k, m\}$, $\{k, a, m\}$ contain the only reflex corners of k , m , and exactly one of $\{a, i\}$, and therefore j, l, n are rectangles. With the help of some simple case analysis it can be shown that this is possible only when the sum of the areas of j and l is less than the sum of the areas of n, m and k . Assigning large areas to j and l and small areas to n, m , and k leads to a contradiction; therefore, at least one of the polygons in Γ must contain 8 sides. $\#$

Lemma 4.1 together with Lemma 4.2 yield the following theorem.

Theorem 2 *Eight-sided polygons are always sufficient and sometimes necessary for a cartogram of a maximal planar Hamiltonian graph.*

5 Cartograms with 6-sided polygons

In this section we give an algorithm for constructing cartograms of a subclass of maximal planar Hamiltonian graph with rectilinear 6-sided polygons. We then use this result to obtain cartograms of maximal outer-planar graphs with rectilinear 6-sided polygons.

5.1 One-Sided Hamiltonian Cycles

Recall the special property of a Hamiltonian cycle, called “one-sidedness” defined in the previous section. Also recall that the algorithm from Lemma 4.1 assigns an 8-sided polygon to a vertex v_i only if neither its bottommost neighbor in G_l , nor its bottommost neighbor in G_r , is v_{i-1} . Since the Hamiltonian cycle v_1, \dots, v_n is one-sided if $b_l(i) = i - 1$ or $b_r(i) = i - 1$ for every $i = 1, \dots, n$ the following lemma is a direct consequence of Lemma 4.1.

Lemma 5.1 *Let $G = (V, E)$ be a maximal planar graph with a one-sided Hamiltonian cycle and let $w : V \rightarrow \mathbb{R}^+$ be a weight function. Then a cartogram with 6-sided polygons can be computed in linear time.*

We will use this result for constructing cartograms of maximal outer-planar graphs with 6-sided polygons. To this end, we first derive conditions which are equivalent to the definition of one-sidedness. Given a Hamiltonian cycle v_1, \dots, v_n we fix a plane embedding of G with outer triangle $\{v_1, v_k, v_n\}$. Then for every $i = 2, \dots, n$ we consider the plane subgraph $G_i \subseteq G$ induced by v_1, v_2, \dots, v_i with the embedding inherited from G . Moreover, define $w_i := v_{n-i+1}$ and let $\tilde{G}_i \subseteq G$ be the subgraph of G induced by w_1, w_2, \dots, w_i , i.e., $\tilde{G}_i = G \setminus G_{n-i}$ and $\tilde{G}_{n-i} = G \setminus G_i$

Lemma 5.2 *Let v_1, \dots, v_n be a Hamiltonian cycle in a maximal plane graph G with (v_1, v_n) on the outer triangle. Then the following are equivalent:*

- (a) $b_l(i) = i - 1$ or $b_r(i) = i - 1$ for $i = 2, \dots, n$, i.e., the cycle is one-sided.
- (b) The edge (v_{i-1}, v_i) is an outer edge of G_i for $i = 2, \dots, n$.
- (c) v_{n-1} is an outer vertex and the vertex v_i has at least two neighbors with a larger index for $i = 1, \dots, n - 2$.
- (d) w_1, \dots, w_n is a canonical ordering for G .
- (e) G admits a Schnyder realizer in which every inner vertex is a leaf in the tree rooted at w_1 or w_2 .

Figure 5 shows an example of a maximal plane graph with a one-sided Hamiltonian cycle, the corresponding canonical ordering, and Schnyder realizer.




Figure 5: A maximal plane graph with a one-sided Hamiltonian cycle (left) and the corresponding Schnyder realizer (right).

Lemma 5.2 has an elementary proof using the concepts of involved. Because we refer to (b) in the next section we prove the first equivalence here.

Proof of (a) \iff (b): For $i = 2, \dots, n$ we argue that (a) holds for i if and only if (b) holds for i . Indeed, (v_{i-1}, v_i) is an inner edge in G_i if and only if there are boundary edges (v_i, v_j) and (v_i, v_k) with $j, k < i - 1$ in G_l and G_r , respectively. But this is equivalent to $b_l(i) = j < i - 1$ and $b_r(i) = k < i - 1$. #

Once we have a one-sided Hamiltonian cycle, we can build a 6-sided cartogram via Lemma 5.1 in linear time. Alternately we could obtain from it a Schnyder wood, rooted such that every vertex is a leaf in \mathcal{S}_1 or \mathcal{S}_2 , and hence obtain a 6-sided cartogram via the algorithm in Section 3. In the next section we use a one-sided Hamiltonian cycle to obtain a cartogram with 6-gons for maximally outer-planar graphs. In this particular case, Lemma 5.1 and the algorithm in Section 3 give combinatorially equivalent cartograms, however, we use Lemma 5.1 due to its linear runtime.

Note that not every Hamiltonian maximally planar graph admits an one-sided Hamiltonian cycle; e.g., the graph in Figure 4 does not even admit a cartogram with 6-gons. However, we suspect that some non-trivial subclasses of Hamiltonian maximally planar graphs are also one-sided Hamiltonian. In particular, we conjecture that every 4-connected maximal planar graph is one-sided Hamiltonian.

5.2 Cartograms for Maximal Outer-planar Graphs

An *outer-planar graph* is a graph that has a planar embedding with all its vertices on the outer face. Such an embedding is called an *outer-planar embedding* of the graph. A *maximal outer-planar graph* is one to which no edge can be added without violating outer-planarity. It is trivial to see that each internal face in an outer-planar embedding of a maximal outer-planar graph is a triangle, and for $n \geq 3$ the outer-face is a simple cycle containing all vertices. We first show sufficiency and then necessity of polygonal complexity 6 in cartograms of maximal outer-planar graphs.

Lemma 5.3 *A cartogram of a maximal outer-planar graph using 6-gons can be computed in linear time. Moreover, the cartogram has a rectangular outer boundary.*

Proof: Choose an outer-planar embedding of G . Let v_1, \dots, v_n be the Hamiltonian cycle in G consisting of all outer edges. Add a vertex v_{n+1} into the outer face connect it by an edge to every v_i , so that (v_1, v_n) remains an outer edge. Then v_1, \dots, v_n, v_{n+1} is a Hamiltonian cycle of the resulting maximal plane graph G' . Indeed this cycle is one-sided since for $i = 2, \dots, n$ the vertex v_i has only three neighbors in G_l , namely v_{i-1} , v_{i+1} and v_{n+1} , which implies $b_l(i) = i - 1$ here. The new vertex v_{n+1} has only one neighbor in G_r , namely v_n , i.e., $b_r(n+1) = n$. With Lemma 5.2 we conclude that v_1, \dots, v_n, v_{n+1} is one-sided, and with Lemma 5.1 we get a cartogram of G' with 6-sided rectilinear polygons (We give v_{n+1} any positive weight.)

Finally note that removing the polygon for v_{n+1} from the cartogram leaves a cartogram of the original graph G with 6-sided rectilinear polygons. This representation fits inside a rectangular area since the removed polygon is 6-sided occupying two consecutive sides of the outer-boundary.

#

A linear-time algorithm for constructing a cartogram of a maximal outer-planar graph with 6-sided rectilinear polygons is also described in [1], however, the construction above is simpler. A lower bound of 6 on the complexity of such a cartogram with the additional constraint that the outer-boundary is a rectangle was proved in [17]. (If we do not insist on the additional property, then maximal outer-planar graphs admit cartograms that use only rectangles [1].)

We thus have the following theorem.

Theorem 3 *Six-sided polygons are sufficient and sometimes necessary for a cartogram of a maximal outer-planar graph.*

6 Conclusion and Open Problems

For maximal planar graphs we presented a cartogram construction with optimal polygonal complexity. For the realization of the actual cartogram this approach requires numerical iteration. A natural open problem is whether the same result can be accomplished with an entirely combinatorial linear time approach.

We also presented such an entirely combinatorial linear-time construction for Hamiltonian maximal planar graphs and showed that the resulting 8-sided cartograms are optimal. We also showed that if the graph admits a one-sided Hamiltonian cycle, only 6 sides are needed. This condition is met by outer-planar graphs. It remains to identify larger classes of planar graphs which are one-sided Hamiltonian and thus have 6-sided cartograms. We conjecture that 4-connected maximal planar graphs have this property.

All of the constructions in this paper yield area-universal rectilinear duals with optimal polygonal complexity. While Eppstein *et al.* [8] characterized area-universal *rectangular* layouts, this remains to be done for general area-universal *rectilinear* layouts.

For some classes of graphs the unweighted and weighted versions of the problem have the same polygonal complexity, as in the case of general planar graphs where we have shown that the tight bound of 8-sided for weighted graphs matches the tight bound for unweighted graphs. On the other hand, maximal planar Hamiltonian graphs have a tight bound of 6 in the unweighted case, while we have shown that the tight bound is 8 in the weighted case. It would be interesting to study when the weighted version of the problem increases the polygonal complexity.

In a similar vein, rectilinear representations are often desirable for practical and technical reasons (e.g., for VLSI layout or floor-planning). Sometimes, insisting on rectilinear representation increases the underlying polygonal complexity. For example, general (unweighted) planar graphs can be represented by 6-sided polygons (tight bound) while 8 are needed in the rectilinear case. For the weighted version, we also now know that 8 sided are sufficient in the rectilinear case, but can we get away with fewer sides if we do not insist on rectilinear layouts?

Acknowledgments This work was initiated at the Dagstuhl Seminar 10461 on Schematization. We also thank Marcus Krug, Ignaz Rutter, Henk Meijer, Emilio Di Giacomo, and Andreas Gerasch for useful discussions.

References

- [1] M. J. Alam, T. Biedl, S. Felsner, A. Gerasch, M. Kaufmann, and S. G. Kobourov. Linear-time algorithms for proportional contact graph representations. Technical Report CS-2011-19, 2011.
- [2] T. Biedl and L. E. Ruiz Velázquez. Orthogonal cartograms with few corners per face. In *Proc. Data Structures and Algorithms Symposium (WADS'11)*, volume 6844 of *Lecture Notes in Computer Science*, pages 98–109. Springer, 2011.
- [3] A. L. Buchsbaum, E. R. Gansner, C. M. Procopiuc, and S. Venkatasubramanian. Rectangular layouts and contact graphs. *ACM Transactions on Algorithms*, 4(1), 2008.
- [4] M. Chrobak and T. Payne. A linear-time algorithm for drawing planar graphs. *Inform. Process. Lett.*, 54:241–246, 1995.
- [5] M. de Berg, E. Mumford, and B. Speckmann. On rectilinear duals for vertex-weighted plane graphs. *Discrete Mathematics*, 309(7):1794–1812, 2009.
- [6] H. de Fraysseix, P. O. de Mendez, and P. Rosenstiehl. On triangle contact graphs. *Combinatorics, Probability and Computing*, 3:233–246, 1994.
- [7] H. de Fraysseix, J. Pach, and R. Pollack. How to draw a planar graph on a grid. *Combinatorica*, 10(1):41–51, 1990.
- [8] D. Eppstein, E. Mumford, B. Speckmann, and K. Verbeek. Area-universal rectangular layouts. In *Proceedings of the 25th ACM Symposium on Computational Geometry*, pages 267–276. ACM, 2009.
- [9] X. He. On floor-plan of plane graphs. *SIAM Journal of Computing*, 28(6):2150–2167, 1999.
- [10] R. Heilmann, D. A. Keim, C. Panse, and M. Sips. Recmap: Rectangular map approximations. In *10th IEEE Symp. on Information Visualization (InfoVis 2004)*, pages 33–40, 2004.
- [11] A. Kawaguchi and H. Nagamochi. Orthogonal drawings for plane graphs with specified face areas. In *4th Conf. on Theory and Applications of Models of Comp.*, pages 584–594, 2007.
- [12] P. Koebe. Kontaktprobleme der konformen Abbildung. *Berichte über die Verhandlungen der Sächsischen Akademie der Wissenschaften zu Leipzig. Math.-Phys. Klasse*, 88:141–164, 1936.
- [13] K. Koźmiński and E. Kinnen. Rectangular duals of planar graphs. *Networks*, 15:145–157, 1985.
- [14] S. M. Leinwand and Y.-T. Lai. An algorithm for building rectangular floor-plans. In *21st Design Automation Conference*, pages 663–664. IEEE Press, 1984.
- [15] C.-C. Liao, H.-I. Lu, and H.-C. Yen. Compact floor-planning via orderly spanning trees. *Journal of Algorithms*, 48:441–451, 2003.
- [16] E. Raisz. The rectangular statistical cartogram. *Geographical Review*, 24(3):292–296, 1934.
- [17] I. Rinsma. Nonexistence of a certain rectangular floorplan with specified area and adjacency. *Environment and Planning B: Planning and Design*, 14:163–166, 1987.
- [18] W. Schnyder. Embedding planar graphs on the grid. In *Proceedings of the 1st ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 138–148, 1990.
- [19] P. Ungar. On diagrams representing graphs. *J. London Math. Soc.*, 28:336–342, 1953.
- [20] M. J. van Kreveld and B. Speckmann. On rectangular cartograms. *Computational Geometry*, 37(3):175–187, 2007.
- [21] S. Wimer, I. Koren, and I. Cederbaum. Floorplans, planar graphs, and layouts. *IEEE Transactions on Circuits and Systems*, 35(3):267–278.
- [22] K.-H. Yeap and M. Sarrafzadeh. Floor-planning by graph dualization: 2-concave rectilinear modules. *SIAM Journal on Computing*, 22:500–526, 1993.