

Computing Cartograms with Optimal Complexity

Md. Jawaherul Alam · Therese Biedl · Stefan Felsner ·
Michael Kaufmann · Stephen G. Kobourov · Torsten Ueckerdt

Received: 5 February 2013 / Revised: 28 May 2013 / Accepted: 3 June 2013 /
Published online: 4 July 2013
© Springer Science+Business Media New York 2013

Abstract In a rectilinear dual of a planar graph vertices are represented by simple rectilinear polygons, while edges are represented by side-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. The complexity of a cartogram is determined by the maximum number of corners (or sides) required for any polygon. In a series of papers the polygonal complexity of such representations for maximal planar graphs has been reduced from the initial 40 to 34, then to 12 and very recently to the currently best known 10. Here we describe a construction with 8-sided polygons, which is optimal in terms of polygonal complexity as 8-sided polygons are sometimes necessary.

This research was initiated at the Dagstuhl Seminar 10461 on Schematization.

Md. J. Alam (✉) · S. G. Kobourov
Department of Computer Science, University of Arizona, Tucson, AZ, USA
e-mail: mjalam@cs.arizona.edu

S. G. Kobourov
e-mail: kobourov@cs.arizona.edu

T. Biedl
David R. Cheriton School of Computer Science, University of Waterloo, Waterloo, ON, Canada
e-mail: biedl@uwaterloo.ca

S. Felsner
Institut für Mathematik, Technische Universität Berlin, Berlin, Germany
e-mail: felsner@math.tu-berlin.de

M. Kaufmann
Wilhelm-Schickhard-Institut für Informatik, Tübingen Universität, Tübingen, Germany
e-mail: mk@informatik.uni-tuebingen.de

T. Ueckerdt
Fakultät für Mathematik, Karlsruher Institut für Technologie, Karlsruhe, Germany
e-mail: torsten.ueckerdt@kit.edu

Specifically, we show how to compute the combinatorial structure and how to refine it into an area-universal rectangular layout in linear time. The exact cartogram can be computed from the area-universal layout with numerical iteration, or can be approximated with a hill-climbing heuristic. We also describe an alternative construction of cartograms for Hamiltonian maximal planar graphs, which allows us to directly compute the cartograms in linear time. Moreover, we prove that even for Hamiltonian graphs 8-sided rectilinear polygons are necessary, by constructing a non-trivial lower bound example. The complexity of the cartograms can be reduced to 6 if the Hamiltonian path has the extra property that it is one-legged, as in outer-planar graphs. Thus, we have optimal representations (in terms of both polygonal complexity and running time) for Hamiltonian maximal planar and maximal outer-planar graphs. Finally we address the problem of constructing small-complexity cartograms for 4-connected graphs (which are Hamiltonian). We first disprove a conjecture, posed by two set of authors, that any 4-connected maximal planar graph has a one-legged Hamiltonian cycle, thereby invalidating an attempt to achieve a polygonal complexity 6 in cartograms for this graph class. We also prove that it is NP-hard to decide whether a given 4-connected plane graph admits a cartogram with respect to a given weight function.

Keywords Planar graphs · Geometric representations · Contact graphs · Cartograms

1 Introduction

There is a large body of work about representing planar graphs as contact graphs, where 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 [23] theorem 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 interior-disjoint polygons and adjacencies represented by non-trivial contacts (shared boundary) between the corresponding polygons. We are specifically interested in rectilinear weighted version where the vertices are represented by simple (axis-aligned) rectilinear polygons. This type of a representation is known as a *rectilinear dual* of the input graph.

In the weighted version, the input is a planar graph $G = (V, E)$ 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 region is equal to the pre-specified weight of the corresponding vertex. Such representations have practical applications in cartography [30], geography [34], sociology [20], VLSI layout, floor-planning [27], etc. Other applications are 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 just drawing a line segment between two points [7].

For both rectilinear duals (unweighted) and 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 floor-planning 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 or increasing the polygonal complexity. With these considerations, 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 holes.

1.1 Related Work

In our paper and in most of the papers cited here, “planar graph” refers to an inner-triangulated planar graph with a simple outer-face; the former restriction is required if at most three rectilinear polygons are allowed to meet in a point and the latter restriction is customary to achieve that 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 [18,26,40]. 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* [24,25,36]. A summary of the history and the state of the art in the rectangle contact graphs literature can be found in [7].

In the above results on rectilinear and rectangular duals, the areas of the polygons are not considered; i.e., these results deal with the unweighted version of the problem. The weighted version dates back to 1934 when Raisz [30] described rectangular cartograms. Algorithms by van Kreveld and Speckmann [37] and Heilmann et al. [19] yield representations with touching rectangles but the adjacencies may be disturbed and there may also be a small distortion of the weights. Recently, Eppstein et al. [17] characterized the class of planar graphs with area-universal rectangular duals. Given an area-universal rectangular dual and a weight function the construction of the actual cartogram can be accomplished using a result by Wimer et al. [39], 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 and adjacencies. De Berg et al. [11] initially showed that 40 sides suffice. This was later improved to 34 sides [22]. In a recent paper [6] the polygonal complexity was reduced to 12 sides and even more recently to 10 sides [2]. But in practice for many maps the average region complexity is closer to 5 [12], so the natural question is to improve this further.

Side contact representations of planar graphs have also been studied without the restriction to rectilinear polygons. In the unweighted case 6-sided polygons are sometimes necessary and always sufficient [16]. The constructive upper bound relies on convex 6-sided polygons. In the weighted version, where the area of each polygon is prescribed, examples are known for which polygons with 7 sides are necessary [35]. This lower bound is matched by constructive upper bound of 7 sides if holes are allowed [3]. In the same paper it is shown that even allowing arbitrarily high

polygonal complexity and holes of arbitrary size, there exist examples with prescribed areas which cannot be represented with convex polygons. If holes are not allowed then the best previously known polygonal complexity is 10, and it is achieved with rectilinear polygons [2]. Recently Nöllenburg et al. [29] studied a slightly different variant of contact representations, where edges are weighted instead of vertices and a contact representation is sought where the edge weights are represented by the length of the common boundary between the two corresponding polygons.

1.2 Our Results

Recall that the known lower bound on the polygonal complexity even for unweighted rectilinear duals is 8 while the best known upper bound is 10. Here we present the first construction that matches the lower bound. The main difference to all previous papers is that we do not create the cartogram directly from the graph and the area function. Instead, we first create an 8-sided rectilinear dual layout of the graph with a very simple construction. Then we argue that this layout is area-universal, i.e., any assignment can be realized by deforming the rectilinear dual layout. With this approach, both the construction of the rectilinear dual layout and the argument that a cartogram exists are much simpler than in previous papers that achieve complexity 10 or more. The exact cartogram can be computed from the area-universal rectangular layout with numerical iteration, or can be approximated with a hill-climbing heuristic.

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-legged* (to be defined later), 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. Table 1 summarizes our contributions to upper and lower bounds on the complexity of cartograms for various graph classes.

Finally we address the problem of constructing small-complexity cartograms for 4-connected graphs (which is a subclass of Hamiltonian planar graphs). We first disprove a conjecture posed by two set of authors that any 4-connected maximal planar graph has a one-legged Hamiltonian cycle, thereby invalidating an attempt to achieve a polygonal complexity 6 in cartograms for this graph class. We also prove that it is NP-hard to decide whether a given 4-connected plane graph with triangular and

Table 1 Overview of results on the complexity of faces for cartograms

Complexity	Outer-planar	4-Connected	Hamiltonian	General
Upper	6 (Sect. 5.1)	8 (implied)	8 (Lemma 4)	8 (Theorem 1)
Lower	6 ^a [2]	6 (implied)	8 (Lemma 5)	8 [40]

A result marked “implied” is easily derived from the other results.

^a The result marked holds only if the outer-face is required to be a rectangle

quadrangular inner faces admits a cartogram with respect to a given weight function on the vertices.

The rest of the paper is organized as follows. In Sect. 2, we review some basic terminology. In Sect. 3, we show how to construct 8-sided area-universal rectilinear duals in linear time. In Sect. 4 we describe the alternative 8-sided construction for Hamiltonian graphs and prove that 8-sides is also a lower bound. In Sect. 5 we show that 6 sides are necessary and sufficient for one-legged Hamiltonian graphs and maximal outer-planar graphs. In Sect. 6 we address the problem of constructing cartograms for 4-connected graphs with small complexity. In Sect. 7 we summarize our contributions and consider several open problems. Parts of this paper appeared as preliminary versions in [1,4].

2 Preliminaries

A *planar graph* $G = (V, E)$ is one that has a drawing without crossings in the plane. A *plane graph* is a planar graph with a fixed planar embedding. It splits the plane into connected regions called *faces*; the unbounded region is the *outer-face* and all other faces are *interior faces*. Thus a plane graph is defined via a cyclic ordering of the edges around each vertex and a prespecified outerface. 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 is *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 $G = (V, E)$ 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 $(u, v) \in E$. ($\mathcal{P}(u)$ and $\mathcal{P}(v)$ may share multiple line-segments.) Such a representation is a *rectilinear dual* of the input graph if the polygons are rectilinear. In the weighted version the input is the graph G , along with a weight function $w : V(G) \rightarrow \mathbb{R}^+$ assigning a weight to each vertex. A rectilinear dual is a *cartogram* if the area of each polygon is equal to the assigned weight of the corresponding vertex. The *complexity of a polygon* is the number of sides in it. 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 [15] and that of a “Schnyder realizer” [33]. 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 [10] 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 neighbours in G_{i-1} form an (at least 2-element) subinterval of the path $C_{i-1} - (u, v)$.

A *Schnyder realizer* of G is a partition of the interior edges of G into the sets \mathcal{S}_1 , \mathcal{S}_2 and \mathcal{S}_3 of directed edges so 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 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 [33]. These two conditions imply that \mathcal{S}_i , for $i = 1, 2, 3$, defines a tree rooted at exactly one exterior vertex and containing all the interior vertices where the edges are directed towards the root. The following well-known lemma shows a profound connection between canonical orders and Schnyder realizers.

Lemma 1 *Let G be a maximal plane graph. Then the following conditions hold.*

- (a) *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 with respect to the clockwise order around v), and to its highest-numbered successor.*
- (b) *A Schnyder realizer with 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. Our algorithm for constructing 8-sided area-universal rectilinear duals has three main phases. In the first phase we create a contact representation of the graph G , where each vertex of G is represented by an upside-down **T**, i.e., a horizontal segment and a vertical segment. Figure 1a, b shows a maximal planar graph and its contact representation using **T**'s, where the three ends of each **T** are marked with arrows. In the second phase we make both the horizontal and vertical segments of each **T** into thin polygons with λ thickness for some $\lambda > 0$. We then have a contact representation of G with **T**-shaped polygons; see Fig. 1c. In the third phase we remove all the unused area in the representation by assigning each (rectangular) hole to one of the polygons adjacent to it; see Fig. 1d. We show that the resulting representation is an area-universal rectilinear dual of G with polygonal complexity at most 8, as illustrated in Fig. 1e.

3.1 Contact Representation with **T**'s

Our contact representation with **T**'s is similar to the approach described by de Fraysseix et al. [13].

Let G be a planar graph. As mentioned earlier, we assume that G is inner-triangulated with a simple outer-face. If need be, we can add two vertices (which we later choose as v_1 and v_2) and connect them to the outer-face to ensure that the

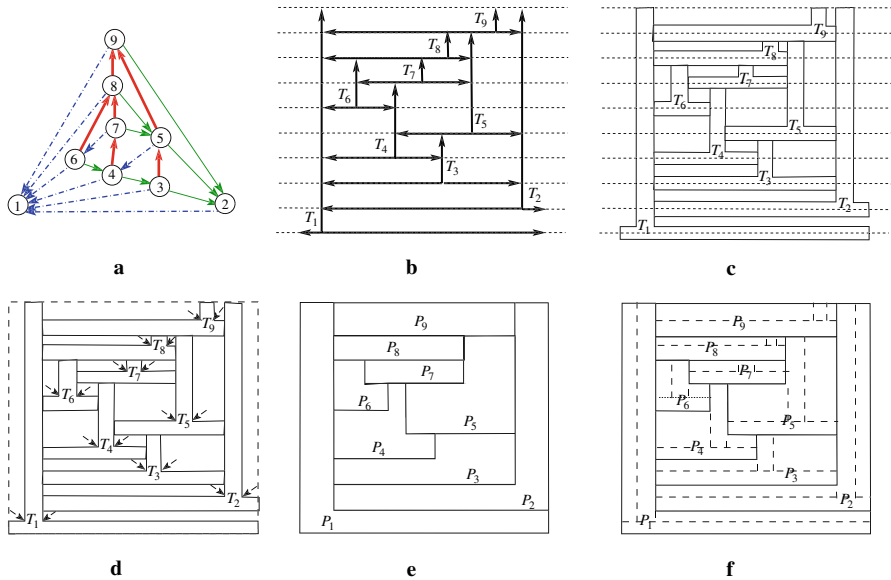


Fig. 1 Construction of a rectilinear dual of a maximal planar graph with 8-sided polygons

graph is maximal. Now let $v_1, v_2, v_3, \dots, v_n$ be a canonical order of the vertices in G with corresponding Schnyder trees $\mathcal{S}_1, \mathcal{S}_2$ and \mathcal{S}_3 rooted at v_1, v_2 and v_n . Add to \mathcal{S}_1 the edges $(v_2, v_1), (v_n, v_1)$ oriented towards v_1 and add to \mathcal{S}_2 the edge (v_n, v_2) oriented towards v_2 . In what follows, we sometimes identify vertex v_i with its canonical label i . We also denote by $\Phi_k(v)$ the parent of vertex v in tree \mathcal{S}_k .

We assign to each vertex v_i the **T-shape** T_i consisting horizontal and vertical segments h_i and b_i , respectively. Begin by placing T_1 and T_2 so 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 neighbours $v_{k_1}, v_{k_2}, \dots, v_{k_d}$ 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 $v_k = \Phi_3(v_{k_i})$ for $1 < i < d$ (Lemma 1), the topmost points of the corresponding vertical segments have y -coordinate k . On the other hand, the canonical labels of $\Phi_3(v_{k_1})$ and $\Phi_3(v_{k_d})$ are greater than k ; hence the topmost points of the corresponding vertical segments have y -coordinates greater than 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 these coordinates be 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 $(x_m, k), (x_m, \Phi_3(k))$




Fig. 2 **a** λ -Flattening of T , and **b** subdividing a T -shaped polygon into four rectangles

with $x_l < x_m < x_r$. Finally for $k = n$, we place T_n so that h_n touches b_1 to the left, b_2 to the right and the topmost point of b_n has y -coordinate $n + 1$.

We note that this representation can be computed in linear time in a way such that all coordinates are integers. We do this by pre-computing a topological order π of $\mathcal{S}_1^{-1} \cup \mathcal{S}_2$; then h_k is the segment $[\pi(\Phi_1(k)), \pi(\Phi_2(k))] \times k$ and b_k is the segment $\pi(k) \times [k, \Phi_3(k)]$.

3.2 λ -Fattening of T_i 's

Let Γ' be the contact representation of G using T 's obtained above. In this phase of the algorithm, we “fatten” 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 Fig. 2a. 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 is a contact representation Γ'' of G where each vertex v_i of G is represented by the T -shaped polygon $H_i \cup B_i$.

3.3 Removing Unused Area

In this step, we begin with the λ -fat T -shaped polygonal layout, Γ'' , from above and assign each (rectangular) hole to a polygon adjacent to it. We start by placing an axis-aligned rectangle of minimum size that encloses Γ'' . This creates five new bounded holes. Note that all these holes are rectangles, and each of them is bounded at the bottom by H_i for some vertex v_i . We assign each hole to this vertex. This assigns at most two holes to each vertex v_i : one hole L_i to the left of B_i , and one hole R_i to the right of B_i . Now 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 a rectilinear polygon with at most 8 sides 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 .




Fig. 3 The direct construction for 8-gons. The *left* picture shows the generic case (where i has incoming edges in both \mathcal{S}_1 and \mathcal{S}_2 and is represented by a T-shape.) The *right* picture shows the case that i has no incoming edges in \mathcal{S}_2 . Similarly it is also possible that the *left side* of the 8-gon is degenerate when i has no incoming edges in \mathcal{S}_1 ; the 8-gon may degenerate to a rectangle when i has no incoming edges in both \mathcal{S}_1 and \mathcal{S}_2

We preferred the above description for the computation of Γ since it gives the reader some intuition for the construction. However, we note here that the coordinates of P_i could be computed directly, without going through T-shapes and λ -fattening, using the values $\Phi_k(v_i)$ for $k = 1, 2, 3$ and a topological order π of $\mathcal{S}_1^{-1} \cup \mathcal{S}_2$. To this end, we take a topological ordering of the acyclic graph $\mathcal{S}_1^{-1} \cup \mathcal{S}_2$ and for each vertex $v_i \neq v_n$, we denote the index of v_i in this topological ordering by $\pi(i)$. Using $\lambda = 1$ and translating up by $\frac{1}{2}$, the coordinates for P_i are then as shown in Fig. 3 (after artificially defining $\Phi_1(1), \Phi_2(2), \Phi_3(n)$ to be suitable values.)

Then the union of these n polygons define the rectilinear dual Γ of G which is contained inside the rectangle $[\frac{1}{2}, n + \frac{1}{2}] \times [1, n + 1]$. Thus we can compute the representation in linear time, and by scaling the representation by 2, we can make all coordinates to be integers of size $O(n)$.

3.4 Area-Universality

A rectilinear dual Γ is *area-universal* if any assignment of positive areas to its polygons can be realized by a combinatorially equivalent layout. Eppstein et al. [17] studied this concept 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* in a layout is the union of inner edges forming a consecutive part of a straight-line. A line-segment not contained in any other line-segment is *maximal*. A maximal line-segment s is called *one-sided* if it forms a full 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 2 [17] *A rectangular layout is area-universal if and only if each maximal segment is one-sided.*

No such characterization is known when some faces are not rectangles. Still we can use the characterization from Lemma 2 to show that the rectilinear dual obtained by the algorithm from the previous section is area-universal, with the following Lemma.

Lemma 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 Γ^* is area-universal. Specifically, recall that each polygon P_i consists of four rectangles H_i , B_i , L_i and R_i (as defined in the previous subsections), as illustrated in Fig. 2b. Any horizontal segment s not on the bounding box belongs to some H_i (either top or bottom), and by expanding it to both sides, it ends at $B_{\phi_1(v_i)}$ on the left and $B_{\phi_2(i)}$ on the right. So s is one-sided since it is a side of H_i . Any vertical segment s not on the bounding box belongs to some B_i (either left or right), and by expanding it to both sides it ends at H_i on the bottom and $H_{\phi_3(i)}$ on the top. So s is one-sided since it is a side of B_i .

Now given an assignment of areas $w : V \rightarrow \mathbb{R}^+$ to the vertices 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 exists a rectilinear dual of G that is combinatorially equivalent to Γ where these areas are realized. Figure 1f illustrates the rectangular layout obtained from the rectilinear dual in Fig. 1e. \square

So for any area-assignment, the rectilinear dual that we found can be turned into a combinatorially equivalent one that respects the area requirements. This proves our main result for maximal planar graphs. Omitting v_1 and v_2 from the drawing still results in a cartogram where the union of all polygons is a rectangle, so the result also holds for all inner-triangulated planar graphs with a simple outer-face.

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 [40]. The algorithm described in this section, thus 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 optimal polygonal complexity, we point out here a practical feature of the 8-sided area-universal rectilinear layout constructed with our algorithm. Earlier constructions, e.g., [11, 6], 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. In particular, we show in the followings that the *minimum feature size* (i.e., the smallest distance between non-adjacent edges) can be made as large as $1/2\sqrt{A} \cdot \min_{v \in V(G)} w(v)$ (where A is the sum of the weights), and that this is worst-case optimal.

Choose W and H such that $W \times H = A = \sum_{v \in V(G)} w(v)$. We are interested in cartograms within a rectangle of width W and height H . Define $w_{min} =$

$\min_{v \in V(G)} w(v)$. Recall that each vertex v_i is represented by the union of at most four rectangles $H_i \cup B_i \cup R_i \cup L_i$, with H_i and B_i non-empty. We can distribute the weight assigned to v_i arbitrarily among them. In particular, we can assign zero areas to the rectangles L_i and R_i and split the weight of v_i into two equal parts, assigning them to H_i and B_i , respectively. In this layout each original vertex is represented by rectangles H_i and B_i whose union is some fattened T or L , and all the necessary contacts remain. Note that the notion of area-universality for rectangular layouts are defined only for assignment of positive areas to the rectangles. However here we explain how we can assign zero areas to L_i and R_i for each vertex v_i in our layout. Consider the layout when we assign a small positive area $\epsilon > 0$ instead of zero to L_i and R_i for each vertex v_i . As we continuously decrease the value of ϵ towards zero, the coordinates of the points in the layout converges to a unique set of coordinates. These coordinates define a degenerate layout where some pairs of segments merge with each other as the area for L_i and R_i becomes zero for each vertex v_i . However since each edge incident to v_i is realized by a contact of H_i or B_i , this degenerate layout maintains all necessary contacts to realize each edge. Hence we can use this simplified layout to produce the cartogram.

The distribution of the weight of v_i in equal parts to H_i and B_i allows to bound the feature size. The height and width of each rectangle are bounded by H and W , respectively. Its weight is at least $w_{\min}/2$. Therefore, the height and width of each rectangle is at least $\frac{w_{\min}}{2 \max\{W, H\}}$. Thus the minimum feature size is at least $\frac{w_{\min}}{2 \max\{W, H\}}$. This is worst-case optimal for our layout, as the polygon with the smallest weight might need to reach from left to right and from top to bottom in the representation. We may choose $W = H = \sqrt{A}$ so that the minimum feature size is $\frac{w_{\min}}{2\sqrt{A}}$. Furthermore the rectangular layout based on only H_i and B_i yields a cartogram with at most $2n$ supporting lines, instead of the $3n$ supporting lines in the cartogram based on four rectangles per vertex.

3.6 Computing the Cartogram

The proof of Lemma 2 implies an algorithm for computing the final cartogram. Splitting the T -shaped polygons into four rectangles and distributing the weights yields an area-universal rectangular dual. This combinatorial structure has to be turned into an actual cartogram, i.e., into a layout respecting the given weights. Wimer et al. [39] gave a formulation of the problem combining flows and quadratic equations. Eppstein et al. [17] indicated that a solution can be found with a numerical iteration. Alternate methods also exist, based on non-linear programming [32], geometric programming [28], and convex programming [9]. Heuristic hill-climbing schemes converge much quicker and can be used in practice, at the expense of small errors [8, 21, 38].

3.7 Implementation and Experimental Results

We implemented the entire algorithm, along with a force-directed heuristic to compute the final cartogram. We treat each region as a rectilinear “room” containing an amount

of “air” equal to the weight assigned to the corresponding vertex. We then simulate the natural phenomenon of air pressure applied to the “walls”, which correspond to the line segment borders in our layout. At each iteration, we consider the segment that feels the maximum pressure and let it move in the appropriate direction.

For each vertex v_i of G , the polygon P_i contains air with volume $w(v_i)$. If the area of P_i is A_i , then the pressure applied to each of the walls surrounding P_i is given by $\mathcal{P}(v_i) = \frac{w(v_i)}{A_i}$. In Sect. 3.4, we saw that the maximal segments of the layout are the two horizontal and the two vertical segment associated with each polygon. For each polygon, the horizontal segment other than the base is entirely inside the polygon, hence it feels no “pressure” on it. For each of the other three segments s for the polygon P_i , the “inward force” it feels is given by

$$\mathcal{F}(s) = \sum_{v_j \in V(s)} [\mathcal{P}(v_j)l_j] - \mathcal{P}(v_i)l_i.$$

Here $V(s)$ is the set of vertices other than v_i whose corresponding polygon touches the segment s and l_i (resp. l_j) denotes the length of s that is shared with P_i (resp. P_j). At each iteration, we consider the segment that feels the maximum pressure and let it move in the appropriate direction. This process is guaranteed to converge with the desired cartogram according to [21], but in practice we stop when the cartographic errors fall below a certain threshold. Here, the *cartographic error* is defined as in [37] to be $\max_{v \in V} (|A(v) - w(v)|/w(v))$, where $A(v)$ denotes the area of the polygon representing v .

Some sample input–output pairs are shown in Fig. 4; more examples and movies showing the gradual transformation can be found at www.cs.arizona.edu/~mjalam/optocart.

We ran a few simple experiments to test the heuristic for time and accuracy. In the first experiment we generated 5 graphs on n vertices with each integer n in the range [10, 50] and assigned 5 random weight distributions with weights in the range [10, 100]. Next we ran the heuristic until the “cartographic error” dropped below 1% and recorded the average time. All the averages were below 50 ms, which confirms that good solutions can be found very quickly in practice; see Fig 5a. In the second experiment we fixed the time allowed and tested the quality of the cartograms obtained within the time limit. Specifically, we generated 5 graphs of n vertices with each integer n in the range [10, 50] and assigned 5 random weight distributions with weights in the range [10, 100]. We allowed the program to run for 1 millisecond and recorded the average cartographic error. Even with this small time limit, the average cartographic error was under 2.5%; see Fig. 5b. All of the experiments were run on an Intel Core i3 machine with a 2.2 GHz processor and 4 GB RAM.

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 *Hamiltonian* maximal planar graph. We first give a direct linear-time construction with 8-sided regions (without relying on numerical

iteration or heuristics as in the previous section.) We then prove that this is optimal by showing that 8 sides are necessary, with a non-trivial lower bound example.

4.1 Sufficiency of 8-Sided Polygons

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 Γ 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 as y -monotone circular arcs with leftmost edge (v_1, v_n) ; see Fig. 6a.

Lemma 4 *Let $G = (V, E)$ be a Hamiltonian maximal planar graph with a weight function $w : V \rightarrow \mathbb{R}^+$. 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 the drawing defined above with (v_1, v_n) on the outer-face. Suppose R is a rectangle of width W and height H where $W \times H = A = \sum_{v \in V} w(v)$. Each vertex v_i will be represented as the union of three rectangles, the *left leg*, the *body* B_i , and *right leg* of v_i . We set the width of the legs to $\lambda_i = w(v_i)/(2H + W)$; see Fig. 6b.




Fig. 4 Input–output pairs: weights are assigned at random in the range $[10, 100]$, and the cartographic error in the output is less than 1%. The colors indicate air-pressure: a gray region has almost the desired area; the lighter a region is, the more it needs to shrink; the *darker* a regions is, the more it needs to grow




Fig. 5 Experimental results for graphs with 10–50 vertices. Each sample point corresponds to 25 graphs. **a** Plotting the average time it takes to reach cartographic error of 1 %. **b** Plotting the average cartographic error achieved in 1 ms




Fig. 6 **a** A Hamiltonian maximal planar graph G , **b** an 8-sided polygon for vertex i , **c** illustration for the algorithm to construct a cartogram of G , **d** a cartogram of G with 8-sided polygons

Our algorithm places vertices v_1, \dots, v_n in this order, and also reserves vertical strips for legs of all vertices that have earlier neighbours. More precisely, let \mathcal{L}_j be all vertices v_k with an edge (v_i, v_k) in G_l for which $i \leq j < k$. Similarly define \mathcal{R}_j with respect to edges in G_r . In the drawing Γ , \mathcal{L}_j are those vertices above v_j for which the horizontal ray left from v_j crosses an incident edge.

We place vertices v_1, \dots, v_j with the following invariant: The horizontal line through the top of B_j intersects, from left to right: (a) a vertical strip of width λ_k for each $v_k \in \mathcal{L}_j$ (which will be used for the left leg of v_k), in descending order, (b) a non-empty part of the top of B_j , and (c) a vertical strip of width λ_k for each $v_k \in \mathcal{R}_j$ (which will be used for the right leg of v_k), in ascending order. We start by placing B_1 as a rectangle that spans the bottom of R . At the left and right end of the top of B_1 , we reserve vertical strips of width λ_k for each vertex in \mathcal{L}_1 and \mathcal{R}_1 , respectively.

To place B_i , $i > 1$, first locate the vertical strips reserved for v_i in previous steps (since $v_i \in \mathcal{L}_{i-1}$ and $i \in \mathcal{R}_{i-1}$, there always are such strips, though they may have started only at the top of B_{i-1}). Since vertical strips are in descending/ascending order, the strips for v_i are the innermost ones. Let B_i be a rectangle just above B_{i-1} connecting these strips. Choose the height of B_i so large that it, together with the

left and right leg inside the strips, has area $w(v_i)$; we discuss the appropriate height below.

Finally, at the top left of the polygon of v_j we reserve a new vertical strip of width λ_k for each vertex k in $\mathcal{L}_i - \mathcal{L}_{i-1}$. Similarly reserve strips for vertices in $\mathcal{R}_i - \mathcal{R}_{i-1}$. By planarity, the vertices in $\mathcal{L}_i - \mathcal{L}_{i-1}$ must have smaller indices than vertices in \mathcal{L}_{i-1} , and so this can be done such that the order required for the invariant is respected, see Fig. 6c.

Clearly this algorithm takes linear time and constructs 8-gons of the correct area. To see that it creates contacts for all edges, consider an edge (v_i, v_k) with $i < k$ in G_l (edges in G_r are similar.) By definition $k \in \mathcal{L}_i$. If $v_k \in \mathcal{L}_i - \mathcal{L}_{i-1}$, then we reserved a vertical strip for v_k when placing v_i . This vertical strip is used for the left leg of v_k , which touches v_i . Otherwise ($v_k \notin \mathcal{L}_i - \mathcal{L}_{i-1}$) we have $v_k \in \mathcal{L}_{i-1}$. When v_{i-1} was placed, there was a vertical strip for v_k . There was also a vertical strip for $v_i \in \mathcal{L}_{i-1}$. These two strips must be adjacent, because by planarity (and edge (v_i, v_k)) there can be no vertex v_j with $i < j < k$ in \mathcal{L}_{i-1} . So these strips create a contact between the two left legs of v_i and v_k .

We now discuss the choice of $\lambda_i = w(v_i)/(2H + W)$. Each leg of v_i has height $\leq H$ and width λ_i , hence area $\leq H\lambda_i$. Then B_i has area $\geq w(v_i) - 2H\lambda_i$ and width $\leq W$, hence height $\geq \frac{w(v_i) - 2H\lambda_i}{W} = \lambda_i$. It follows that B_i has positive height. Also all vertical strips fit: after placing B_i , we have a strip of width λ_k for each vertex $v_k \in (\mathcal{L}_i \cup \mathcal{R}_i)$, and these strips use width

$$\sum_{v_k \in (\mathcal{L}_i \cup \mathcal{R}_i)} \frac{w(v_k)}{2H+W} \leq \frac{2 \sum_{v_k \in V - \{v_i\}} w(v_k)}{2H+W} < \frac{2A}{2H} - \frac{2w(v_i)}{2H+W} = W - 2\lambda_i.$$

Hence B_i has width $> 2\lambda_i$ and the polygon of v_i has minimum feature size λ_i . □

Figure 6d illustrates a cartogram for graph in Fig. 6a, obtained by the above algorithm. This algorithm also gives a minimum feature size for the cartogram: $\min_{v_i \in V} \{\lambda_i\} = \frac{w_{\min}}{2H+W}$, where $w_{\min} = \min_{v \in V} \{w(v)\}$. Choosing $W = \sqrt{2}A$, $H = \sqrt{A/2}$, yields the minimum feature size $\frac{w_{\min}}{2\sqrt{2}\sqrt{A}}$.

4.2 Necessity of 8-Sided Polygons

While it was known that 8-sided polygons are necessary for general planar graphs [31], the graph used for this was not Hamiltonian. Here we show that this necessity holds even for Hamiltonian maximal planar graphs.

Lemma 5 *Consider the Hamiltonian maximal planar graph $G = (V, E)$ in Fig. 7a. Define $w(j) = w(l) = D$ and $w(v) = \delta$ for $v \in V \setminus \{j, l\}$, where $D \gg \delta$. Then any cartogram of G with respect to w contains a polygon of at least 8 sides.*

Proof Assume for a contradiction that G admits a cartogram Γ with respect to w such that each polygon $\{P(v)\}$ in Γ representing the vertex v has complexity at most 6. Observe that if $\{u, v, x\}$ is a *separating triangle* in G , i.e., three mutually




Fig. 7 **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 5

adjacent vertices whose removal disconnect the graph, then the region R_{uvw} used for the graph inside the separating triangle contains at least one reflex corner of the polygon $P(u)$, $P(v)$, or $P(x)$. 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 . In particular, the outer boundary of R_{aei} contains exactly one reflex corner from one of $P(a)$, $P(e)$ and $P(i)$, hence it is a rectangle, say 1234. By symmetry, assume that the reflex corner of $P(i)$ is *not* used for R_{aei} .

The 4-vertex set $\{a, i, k, m\}$ is the disjoint union of three separating triangles $\{a, k, m\}$, $\{k, i, m\}$, $\{i, a, m\}$, containing l, j and n , respectively, in the interior. The reflex corner of $P(i)$ is not used for R_{aei} , so, it cannot be used for any of these triangles. Thus each of $P(j)$, $P(l)$, $P(n)$ contains exactly one reflex corner from $P(a)$, $P(k)$ and $P(m)$. In particular, $R_{aei} = 1234$ must contain the reflex corner of $P(a)$. We also conclude that $P(j)$, $P(l)$, $P(n)$ are rectangles, since there are no additional reflex corners to accommodate additional convex corners from $P(j)$, $P(l)$ and $P(n)$.

Assume the naming in Fig. 7b is such that edge 12 belongs to $P(i)$, edges 23 and 34 belong to $P(a)$ and edge 41 belongs to $P(e)$. By the adjacencies, $P(k)$ must occupy corners 1 and 4 and $P(n)$ must occupy corner 2, while corner 3 (which is the reflex corner of $P(a)$) could belong to n or l .

Now consider the rectangles $P(j)$ and $P(l)$. If D is sufficiently big, then these two rectangles each occupy almost half of rectangle 1234. Therefore, either their x -range or their y -range must overlap. Assume their y -range overlaps, the other case is similar. Which polygon should occupy the area horizontally between $P(j)$ and $P(l)$? It cannot be k , because $P(k)$ contains corners 1 and 2 and hence would obtain 2 reflex angles from $P(j)$ and $P(l)$. So it must be $P(m)$, since n is not adjacent to j and l . But $P(m)$ must also separate $P(n)$ from both $P(j)$ and $P(l)$. Regardless of whether n or l occupies corner 3, this is not possible without two reflex vertices for m . Thus some polygon must have 8 sides. □

Lemmas 4 and 5 yield the following theorem.

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

5 Cartograms with 6-Sided Polygons

Here we study cartograms with rectilinear 6-gons. We first note that these are easily constructed for outer-planar graphs. Then we generalize this technique to a class of maximal planar Hamiltonian graphs.

5.1 Maximal Outer-planar Graphs

Our algorithm from Lemma 4 naturally gives drawings of maximal outer-planar graphs that use 6-sided polygons. Another linear-time algorithm for constructing a cartogram of a maximal outer-planar graph with 6-sided rectilinear polygons is also described in [2], however, our construction based on Lemma 4 is much simpler. Any maximal outer-planar graph G can be made into a maximal Hamiltonian graph by duplicating G and gluing the copies together at the outer-face such that $G_l = G = G_r$. (This graph has double edges, but the algorithm in Lemma 4 can handle double edges as long as one copy is in the left and one in the right graph.) Create the drawing based on Lemma 4 with all vertices having double the weight, and cut it in half with a vertical line. This gives a drawing of G with 6-sided rectilinear polygons as desired.

5.2 One-Legged Hamiltonian Cycles

We now aim to find more maximal Hamiltonian graphs which have cartograms with 6-sided polygons. In a Hamiltonian cycle v_1, \dots, v_n , call vertex v_j *two-legged* if it has a neighbor v_{j^l} in G_l with $i^l < j - 1$ and also a neighbor v_{j^r} in G_r with $i_r < j - 1$. Call a Hamiltonian cycle *one-legged* if none of its vertices is two-legged. In the construction from Lemma 4, the polygon of v_j obtains a reflex vertex on both sides only if it has a neighbor below v_{j-1} on both sides, or in other words, if it is two-legged. Hence we have:

Lemma 6 *Let $G = (V, E)$ be a maximal planar graph with a one-legged 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.*

It is a natural question to characterize graphs with such Hamiltonian cycles. The following lemma gives a nice relation between a one-legged Hamiltonian cycle, a canonical ordering and a Schnyder realizer of a graph.

Lemma 7 *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. Define $w_i := v_{n-i+1}$. Then these are equivalent:*

- (a) *The Hamiltonian cycle is one-legged.*
- (b) *For $i = 2, \dots, n$, edge (v_{i-1}, v_i) is an outer edge of the embedded graph G_i induced by v_1, v_2, \dots, v_i .*

- (c) v_{n-1} is an outer vertex and 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 $(\mathcal{S}_1, \mathcal{S}_2, \mathcal{S}_3)$ in which w_1, w_2 and w_n are the roots of $\mathcal{S}_1, \mathcal{S}_2$ and \mathcal{S}_3 , respectively and every inner vertex is a leaf in \mathcal{S}_1 or \mathcal{S}_2 .

Proof (a) \iff (b): For $i = 2, \dots, n$ we argue that a vertex v_i is one-legged 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 outer edges (v_i, v_j) and (v_i, v_k) with $j, k < i - 1$ in G_l and G_r , respectively. But then v_i is two-legged by definition.

(b) \iff (c): Since v_n is an outer vertex and $G_n = G$, (b) holds for $i = n$ if and only if v_{n-1} is an outer vertex. For $i = 2, \dots, n - 1$ we argue that (b) holds for i if and only if (c) holds for $i - 1$. Let v_i^l , respectively v_i^r , denote the third vertex in the inner facial triangle containing the edge (v_{i-1}, v_i) in G_l , respectively G_r . Now (v_{i-1}, v_i) is an inner edge in G_i if and only if both, v_i^l and v_i^r , have a smaller index than v_{i-1} , which in turn holds if and only if the index of every neighbor of v_{i-1} , other than v_i , is less than $i - 1$.

(c) \implies (d): By (c) $\{w_1, w_2, w_n\} = \{v_n, v_{n-1}, v_1\}$ is the outer triangle of G . Moreover, \tilde{G}_3 , induced by v_n, v_{n-1}, v_{n-2} , is a triangle, hence its outer boundary is a simple cycle C_3 with the edge (w_1, w_2) . Thus the first condition of canonical ordering is met for $i = 4$. Assuming (c) and the first condition for $i = 4, \dots, n - 1$, we show that the second and first condition hold for i and $i + 1$, respectively. Finally, the second condition holds for $i = n$ since w_n is an outer vertex.

Note that w_i is in the exterior face of \tilde{G}_{i-1} since w_n lies in the exterior face and the path w_i, \dots, w_n is disjoint from vertices in \tilde{G}_{i-1} and the embedding is planar. By (c) w_i has at least two neighbours in \tilde{G}_{i-1} . If the neighbours do not form a subinterval of the path $C_{i-1} \setminus (w_1, w_2)$, there is a non-triangular inner face in \tilde{G}_i , containing a vertex w_j with $j > i$ in its interior. Then the path w_j, \dots, w_n , which is disjoint from \tilde{G}_i , starts and ends in an interior and the exterior face of \tilde{G}_i , respectively, contradicting planarity. Thus the second condition of canonical ordering is satisfied for i . Moreover \tilde{G}_i is inner-triangulated, has a simple outer cycle C_i containing the edge (w_1, w_2) . Thus, the first condition holds for $i + 1$.

(d) \implies (c): Since w_1, \dots, w_n is a canonical ordering, (w_1, w_2) is an outer edge. In particular, $w_2 = v_{n-1}$ is an outer vertex. Clearly v_1 has at least two neighbours and every neighbor has a larger index, i.e., (c) holds for $i = 1$. Moreover, by the second condition of a canonical ordering every vertex $v_i = w_{n-i+1}$, for $i = 2, \dots, n - 2$, has at least two neighbours in $\tilde{G}_{n-i} = G \setminus G_i$, which is the subgraph induced by v_n, \dots, v_{i+1} .

(d) \implies (e): Consider the Schnyder realizer $(\mathcal{S}_1, \mathcal{S}_2, \mathcal{S}_3)$ of G defined by the canonical order w_1, \dots, w_n according to Lemma 1. One can easily show by induction that for $i = 3, \dots, n - 1$ the outer cycle C_i of \tilde{G}_i consists of the edge (w_1, w_2) , the $w_i w_1$ -path P_1 in \mathcal{S}_1 , and the $w_i w_2$ -path P_2 in \mathcal{S}_2 . Due to the counterclockwise order of edges in a Schnyder realizer, no vertex on P_1 , respectively P_2 , has an incoming inner edge in \tilde{G}_i in \mathcal{S}_2 , respectively \mathcal{S}_1 . Thus considering only edges in \tilde{G}_i every outer vertex in \tilde{G}_i , different from w_1, w_2 , is a leaf in \mathcal{S}_1 or \mathcal{S}_2 . When in the canonical ordering vertex w_{i+1} is attached to \tilde{G}_i , some vertices on C_i become inner vertices of




Fig. 8 A graph with a one-legged Hamiltonian cycle and the corresponding Schnyder realizer

\tilde{G}_{i+1} . Every inner edge in \tilde{G}_{i+1} , which was not an edge in \tilde{G}_i is in \mathcal{S}_3 . Thus every inner vertex in \tilde{G}_i is a leaf in either \mathcal{S}_1 or \mathcal{S}_2 .

(e) \implies (d): Consider a canonical ordering x_1, x_2, \dots, x_n of G defined by the Schnyder realizer $(\mathcal{S}_1, \mathcal{S}_2, \mathcal{S}_3)$ according to Lemma 1. We must show that x_1, \dots, x_n gives a Hamiltonian cycle. We know that $\{x_1, x_2, x_3\}$ is a triangle, hence C_3 consists of the edge (x_1, x_2) , the x_3x_1 -path P_1 in \mathcal{S}_1 , and the x_3x_2 -path P_2 in \mathcal{S}_2 . We will show that for all i the outer-face of \tilde{G}_i consists of the edge (x_1, x_2) , the x_ix_1 -path P_1 in \mathcal{S}_1 , and the x_ix_2 -path P_2 in \mathcal{S}_2 . For $i = 3, \dots, n$, x_{i+1} is attached to \tilde{G}_i . If x_i is not a predecessor of x_{i+1} , then (say) the $x_{i+1}x_2$ path in \mathcal{S}_2 meets the x_ix_1 path in \mathcal{S}_1 at some vertex x_ℓ , $\ell \neq i, i + 1$. This means that x_ℓ has an incoming edge in both \mathcal{S}_1 and \mathcal{S}_2 —a contradiction. So x_i is adjacent to x_{i+1} , which can also be used to show that the shape of the outer-face is as desired. \square

Figure 8 shows a one-legged Hamiltonian cycle, corresponding canonical ordering, and Schnyder realizer in a graph.

Once we have a one-legged Hamiltonian cycle, we can build a 6-sided cartogram via Lemma 6 in linear time. Alternately we could obtain a Schnyder wood, where every vertex is a leaf in \mathcal{S}_1 or \mathcal{S}_2 , and hence obtain a 6-sided cartogram via the algorithm in Sect. 3. However, we prefer the construction of Lemma 6 due to its linear runtime.

Not every Hamiltonian maximal planar graph admits a one-legged Hamiltonian cycle; for example, the graph in Fig. 7 does not even admit a cartogram with 6-gons. However, we believe that some non-trivial subclasses of Hamiltonian maximal planar graphs are also one-legged Hamiltonian.

6 Cartograms for 4-Connected Graphs

Since 4-connected planar graphs are a strict subset of Hamiltonian planar graphs, the following conjecture, which we raised in a preliminary version of this paper [4] (also raised in a different context a decade ago by de Fraysseix et al. [14]), would be natural:

Conjecture 1 [4, 14] *Every 4-connected maximal planar graph has a one-legged Hamiltonian cycle.*

We disprove this conjecture in Sect. 6.1. In Sect. 6.2, we address the issue of existence of a cartogram with rectangles for a given 4-connected graph (not necessarily maximal) along with a given weight function and show that this problem is NP-hard.

6.1 Hamiltonian Canonical Cycle

Recall that a one-legged Hamiltonian cycle is by Lemma 7 (after reversing the order of vertices along the Hamiltonian cycle) the same as a Hamiltonian cycle that induces a canonical ordering. We call such a cycle a Hamiltonian canonical cycle. More precisely, a *Hamiltonian canonical cycle* in a maximal planar graph G is a canonical order v_1, v_2, \dots, v_n of the vertices of G such that $v_1v_2 \dots v_n$ is also a Hamiltonian cycle of G . We now disprove the above conjecture by constructing a 4-connected maximal planar graph with no Hamiltonian canonical cycle.

Theorem 3 *There exist 4-connected maximal planar graphs that do not have any Hamiltonian canonical cycle in any embedding.*

Proof We construct a 4-connected planar graph G as follows. Let K be the graph in Fig. 9. Take two copies of K and identify their vertices $\{a, b, c, d\}$. The resulting graph is 4-connected and triangulated. In what follows, we will always consider the one copy of K that does not contain the outer-face, and hence is embedded as in Fig 9.

For any embedded subgraph H of G bounded by a separating 4-cycle C_H , we call the subgraph induced by the vertices of $H - C_H$ the *interior* of H . Assume for a contradiction that G has a Hamiltonian canonical cycle \mathcal{C} . We consider the Hamiltonian cycle to be visiting the vertices in order. Since it is a Hamiltonian canonical cycle, from the third vertex on every vertex must have at least two neighbours that were visited earlier. Since the first two vertices form an edge on the outer-face, we cannot reach the interior of K before the third vertex of \mathcal{C} .

The first vertex of \mathcal{C} that is in the interior of K must be one of $\{e, h, g, f\}$, because no other vertices have two neighbours on the boundary of K , and only those can have been visited earlier. Without loss of generality we may assume e is the first vertex visited in the interior of K , which means that a and b were visited earlier. Next we




Fig. 9 The graph K used in Theorem 3

must visit *all vertices* in the interior of K , for K is bounded by a 4-cycle, and there are not enough unvisited vertices left on the boundary to leave the interior of K and return to it again later. Say after visiting the interior of K , we exit the 4-cycle $\{a, b, c, d\}$ at vertex d (the case of ending at c is similar.)

Consider the subgraph $L_3 \cup \{o\}$, which is bounded by a 4-cycle $\{o, h, d, g\}$. The first vertex in the interior of $L_3 \cup \{o\}$ to be visited must necessarily be r , for d is visited after all other vertices of L_3 , and no other vertex in the interior has two neighbours $\neq d$ on the 4-cycle $\{o, h, d, g\}$. By the above argument for separating 4-cycles, we must visit the entire interior of $L_3 \cup \{o\}$ in one contiguous subpath that uses only vertices in the interior or on the boundary of $L_3 \cup \{o\}$. Let P be the minimal such path. Then P begins at r and does not visit d .

Since the first vertex of P is r , we must have visited o and at least one of $\{g, h\}$ before P . By symmetry, assume that we visited h already (we may or may not have visited g .) In summary, we know the following about P :

- (i) P begins at r and at this time h and o were visited already.
- (ii) P visits all vertices in the interior of L_3 , and it may or may not visit g .
- (iii) P ends with d still unvisited.
- (iv) Every vertex of P has at least two neighbours that were visited earlier by P .

By analyzing the structure of L_3 , we will show that such a path P cannot exist, which gives the desired contradiction. Consider the close-up of graph L_3 in Fig. 10.

We claim that without loss of generality the next vertex after r on P is A_1 . It can only be A_1 or B_1 , for no other vertex in the interior of L_3 has two neighbours that were already visited. If it were B_1 , then g was already visited. But in this case L_3 is entirely symmetric with respect to who has been visited already, and a symmetric argument will lead to a contradiction. So we may assume that the next vertex after r on P is A_1 . Consider vertex A_2 , which has degree 4. Neighbor h was visited already, and neighbor d will be visited much later. So the adjacent vertices of A_2 in P must be A_1 and A_3 . Therefore the next two vertices on P after A_1 are A_2 and A_3 . This part of P is marked with a solid black in Fig. 10 and corresponds to the beginning of P .

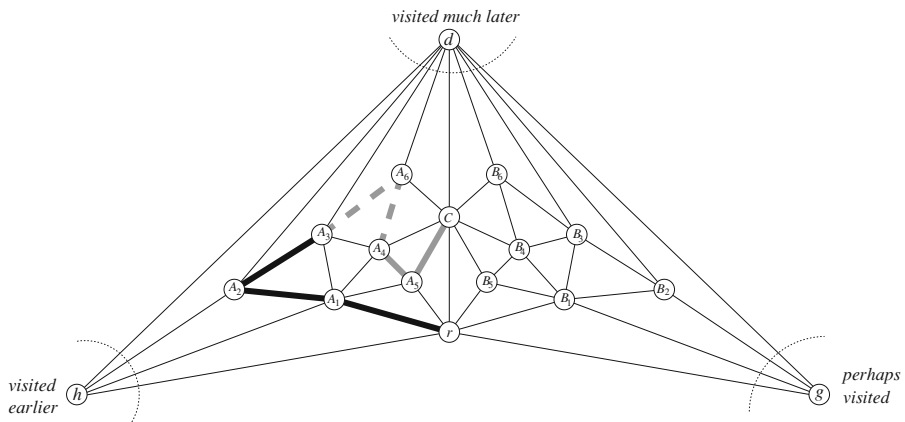


Fig. 10 The graph L_3 used for constructing K

Consider the separating 4-cycle $\{h, r, C, d\}$. When we reach A_1 , we have visited h and r already. We will not visit d with P . So the last vertex of P that is on or inside $\{h, r, C, d\}$ must be C . Consider vertex A_5 , which has degree 4. Its neighbours r and A_1 were already visited, so the adjacent vertices of A_5 in P must be A_4 and C . Therefore the two vertices on P before C are A_5 and A_4 . This part of P is marked with a solid gray in Fig. 10 and corresponds to the last part of P on or inside $\{h, r, C, d\}$.

Consider vertex A_6 , which has degree 4. Neither C nor d can be its adjacent vertices on P because the part of P near C is already fixed, and d comes much later. So A_6 must be between A_3 and A_4 on P . This part of P is marked with a dashed line in Fig. 10. But now we have a problem: A_6 has only one neighbor A_3 that was visited earlier by C ; all the other neighbours of A_3 are visited later. This contradicts that C is a Hamiltonian canonical cycle. \square

6.2 Cartograms with Rectangles

Another question one could ask is: Given a plane weighted graph, how easy is it to test whether it has a cartogram of a given complexity?

For rectilinear layouts, we can answer this question in polynomial time. Specifically, the question whether a graph has a rectilinear dual with an upper bound on the number of corners can be phrased as a maximum flow problem and hence be solved in linear time.

In contrast to this, we now show that testing whether a graph has a cartogram with an upper bound on the number of corners is NP-hard. Note that the graph in our construction has some faces that are quadrangular (rather than the preferred triangular), but it is 4-connected. We show that it is NP-hard to test whether it has a rectangular cartogram, i.e., where all faces are rectangles. Let us call the problem **RectangleCartogram (RC)**.

We show that this problem is NP-hard by a reduction from the well-known NP-hard problem **Partition** defined as follows. Given a (multi-)set of positive integers $S = \{x_1, \dots, x_n\}$ with $\sum_{i=1}^n x_i = 2A$ for some integer A , we want to find a subset I of S such that $\sum_{x_i \in I} x_i = A$.

Given an instance of **Partition**, we construct an instance of **RC**.¹ We first construct a 4-connected graph G as follows. For each integer x_i of S , we have a subgraph with eight vertices: $X_i, p_i, p_{i+1}, q_i, q_{i+1}, a_i, b_i, c_i$. We highlight such a subgraph for x_i in Fig. 11. We also have four vertices L, T, R and B that forms the outerface. Call these four vertices *boundary vertices* of G . The constructed graph G is then 4-connected with a quadrangular outerface and each of its internal faces is a triangle or quadrangle.

We define the weight function as follows. For each vertex X_i , we define $w(X_i) = x_i$. We give a very small weight $\delta < 1/(5n + 2)$ to each vertex a_i, b_i, c_i, p_i, q_i for $1 \leq i \leq n$ and to each vertex L, R, T and B . We give a very large weight $W(M)$ to

¹ The construction is quite similar to the one by Biedl and Genc for testing whether a given rectilinear drawing can be used to satisfy area-constraints [5].

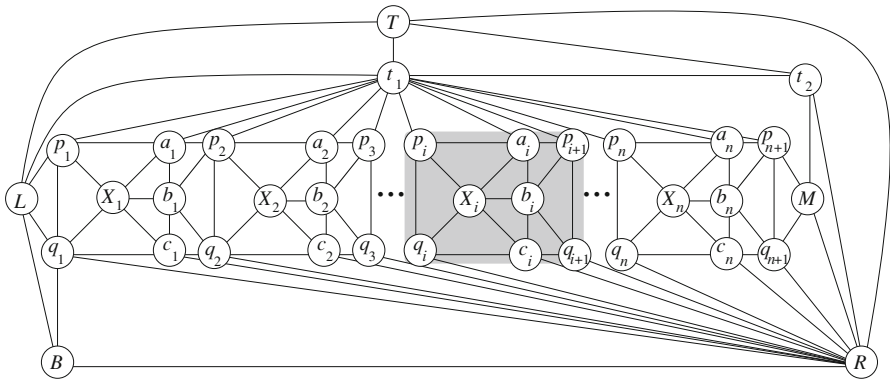


Fig. 11 The graph constructed from an instance of a **Partition** problem

the vertex M such that $\sqrt{W(M) + 2A + (5n + 2)\delta} - \sqrt{W(M)} < 1$. Finally we give weights to the vertices t_1 and t_2 such that

$$w(t_1) : w(t_2) = \sqrt{W(M) + 2A + (5n + 2)\delta} - \sqrt{W(M)} : \sqrt{W(M)}.$$

We now have the following lemma.

Lemma 8 *There exists a subset I of S such that $\sum_{x_i \in I} x_i = A$ if and only if there is a rectangular cartogram Γ of G with respect to the weight function w where the rectangles for L, T, R and B forms the outer-boundary of Γ .*

Proof Suppose first that there is a rectangular cartogram Γ of G with respect to w where the rectangles for the four vertices L, T, R and B define a rectangular region F' which consists of the rectangles for the remaining vertices. Without loss of generality, we assume that the rectangles L, T, R and B are to the left, top, right and bottom of F' , respectively; see Fig 12. Since the face formed by t_1, t_2, M and p_{n+1} is a rectangle, this face will correspond to a point of degree four in Γ incident to rectangles for the four vertices. Thus the rectangles $R(t_1)$ and $R(t_2)$ for the vertices t_1 and t_2 , respectively, have the same height and together they span the top of the region F' . Define the rectangular region $F = F' - R(t_1) - R(t_2)$. Scale Γ in the x and y direction by a reciprocal factor such that the area of Γ remains unchanged but F becomes a square. Then from the weight distribution of t_1 and t_2 , the rectangle $R(M)$ for M is also a square. The left and bottom of this square define two lines l_v and l_h , respectively. Now the rectangle for each vertex X_i lies either to the left of l_v or to the bottom of l_h , but not both since

$$\sqrt{W(M) + 2A + (5n + 2)\delta} - \sqrt{W(M)} < 1 \leq \sqrt{x_i}.$$

Thus each subgraph for x_i has a drawing in one of two configurations; see Fig. 12a, b, respectively. Then the vertices corresponding to the rectangles X_i lying to the left of l_v form the subset I of S . The reason is as follows. Define the four rectangles R_1, R_2, R_3, R_4 that partitions F in this way: $R_1 = R(M)$, R_2, R_3 are the rectangles





Fig. 12 Illustration of the proof of Lemma 8

Fig. 13 Construction of a cartogram from a solution of Problem **Partition**



to the left of l_v ; R_2 to the top and R_3 to the bottom of l_h and R_4 is the rectangle to the right of l_v and to the bottom of l_h . Then the area of $R_2 + R_3$ and $R_4 + R_3$ are equal and neither can contain rectangles with total weights $A + 1$ since the area of R_4 is less than 1. However since the weights for each X_i is a positive integer, it must be the case that the rectangles X_i lying in $R_2 + R_3$ has a total area of A .

Conversely if we are given a subset I of S such that $\sum_{x_i \in I} x_i = A$, then we construct a cartogram of G as follows. We draw the rectangles for L, T, R, B, t_1 and t_2 such that they enclose a square region F of size $W(M) + 2A + (5n + 2)\delta$, and the rectangles L, T, R and B are to the left, top, right, and bottom of the drawing, respectively. We also draw the square $R(M)$ for M to the right-top corner of F . From the weight distribution, the top-left corner of $R(M)$ will be a point of degree four. We then draw the subgraphs for x_i in the way illustrated in Fig. 12a if $x_i \in I$; otherwise we draw this subgraph in the way illustrated in Fig. 12b. Thus we obtain a cartogram of G , as in Fig. 13. Here $x_1, x_2 \in I, x_3 \notin I$ and so on. \square

We can thus reduce an instance S of Problem **Partition** to an instance (H, w) of Problem **RectangleCartogram** as follows. We construct the graph H by taking two

copies G_1, G_2 of G with boundary vertices L_i, T_i, R_i and B_i for $i = 1, 2$ and then identifying T_1 with T_2 , B_1 with B_2 and R_1 with L_2 . This graph H is thus 4-connected planar graph with four vertices on the outerface and each of its internal faces is either a triangle or a quadrangle. The weight function w is as described in the paragraph before Lemma 8 for each copy of G . Then if H has a cartogram with weight function w , then this gives an embedding of H and we consider one of the two copies of G that does not contain the outerface. Then by Lemma 8, we get a subset I of S with $\sum_{x_i \in I} x_i = A$. Conversely, if we are given a subset I of S with $\sum_{x_i \in I} x_i = A$ we can construct a cartogram for both copies of G using the algorithm in Lemma 8 and then place these two copies beside each other to find a cartogram for H (more precisely, we identify the rectangles R_2 and L_1 and reduce the height of $T = T_1 \cup T_2$ and $B = B_1 \cup B_2$ by half). Thus S has a solution if and only if G has a rectangular cartogram with respect to w . This yields the following theorem.

Theorem 4 *Problem RectangleCartogram is NP-hard.*

7 Conclusion and Open Problems

We presented a cartogram construction for maximal planar graphs with optimal polygonal complexity. For the precise realization of the actual cartogram this approach requires numerical iteration. Even though the simple heuristic works well in practice, a natural open problem is whether everything can be computed with an entirely combinatorial linear-time approach. We also presented an entirely different linear-time construction for Hamiltonian maximal planar graphs and showed that the resulting 8-sided cartograms are optimal. Finally, we showed that if the graph admits a one-legged Hamiltonian cycle, only 6 sides are needed. Maximal outer-planar graphs contain one-legged Hamiltonian cycles but not all 4-connected maximal planar graphs do. It remains to identify larger classes of planar graphs which are one-legged Hamiltonian and thus have 6-sided cartograms.

All of the constructions in this paper yield area-universal rectilinear duals with optimal polygonal complexity. While Eppstein et al. [17] characterized area-universal *rectangular* layouts, a similar characterization for general area-universal *rectilinear* layouts is a natural open problem. Another open problem related to area-universal layout is whether given such area-universal rectangular layout one can efficiently (or at least more efficiently than the non-linear programming [32], geometric programming [28], and convex programming [9]) realize a given weight function.

Is it NP-hard to test whether a graph with a given area-assignment has a k -sided cartogram? Here we show that for 4-connected graphs with triangular and quadrangular faces and for cartograms containing only rectangles the problem is NP-hard. It would be interesting to investigate the complexity of the problem when all the inner faces are triangles. On the other hand, if a rectilinear layout of the graph is given, then it is again known to be NP-hard to test whether the layout can be deformed into a cartogram, even if all faces have at most 8 sides [5]. Does the freedom to choose the rectilinear layout help?

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 for weighted graphs matches the tight bound for unweighted graphs. On the other hand, Hamiltonian maximal planar 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 improve this to 7 sides if we do not insist on rectilinear layouts?

Acknowledgments Alam and Kobourov’s research was funded in part by NSF Grants CCF-0545743 and CCF-1115971. Biedl’s research was supported by NSERC. Felsner’s and Ueckerdt’s research was partially supported by EUROGIGA project GraDR and DFG Fe 340/7-2. Kaufmann’s research was partially supported by EUROGIGA project GraDR. The work was done when Ueckerdt was at Technische Universität Berlin.

References

1. Alam, M.J., Kobourov, S.G.: Proportional contact representations of 4-connected planar graphs. In: Graph Drawing (GD’12), pp. 211–223. Springer (2013)
2. Alam, M.J., Biedl, T.C., Felsner, S., Gerasch, A., Kaufmann, M., Kobourov, S.G.: Linear-time algorithms for hole-free rectilinear proportional contact graph representations. In: International Symposium on Algorithms and Computation (ISAAC’11), pp. 281–291. Springer, Berlin (2011)
3. Alam, M.J., Biedl, T.C., Felsner, S., Kaufmann, M., Kobourov, S.G.: Proportional contact representations of planar graphs. In: Graph Drawing (GD’11), pp. 26–38. Springer, Heidelberg (2012)
4. Alam, M.J., Biedl, T.C., Felsner, S., Kaufmann, M., Kobourov, S.G., Ueckerdt, T.: Computing cartograms with optimal complexity. In: Symposium on Computational Geometry (SoCG’12), pp. 21–30. Bethesda (2012)
5. Biedl, T.C., Genc, B.: Complexity of octagonal and rectangular cartograms. In: Canadian Conference on Computational Geometry (CCCG’05), pp. 117–120. Ontario (2005)
6. Biedl, T.C., Velázquez, L.E.R.: Orthogonal cartograms with few corners per face. In: Algorithms and Data Structures, Symposium (WADS’11), pp. 98–109. Springer, Berlin (2011)
7. Buchsbaum, A.L., Gansner, E.R., Procopiuc, C.M., Venkatasubramanian, S.: Rectangular layouts and contact graphs. *ACM Trans. Algorithms* **4**, 28 (2008)
8. Cederbaum, I.: Analogy between VLSI floorplanning problems and realisation of a resistive network. *IEE Circuits Devices Syst.* **139**(1), 99–103 (1992)
9. Chen, T., Fan, M.K.H.: On convex formulation of the floorplan area minimization problem. In: Symposium on Physical Design, pp. 124–128. (1998)
10. Chrobak, M., Payne, T.H.: A linear-time algorithm for drawing a planar graph on a grid. *Inf. Process. Lett.* **54**(4), 241–246 (1995)
11. de Berg, M., Mumford, E., Speckmann, B.: On rectilinear duals for vertex-weighted plane graphs. *Discrete Math.* **309**(7), 1794–1812 (2009)
12. de Berg, M., Mumford, E., Speckmann, B.: Optimal BSPs and rectilinear cartograms. *Int. J. Comput. Geom. Appl.* **20**(2), 203–222 (2010)
13. de Fraysseix, H., Ossona de Mendez, P.: On triangle contact graphs. *Comb. Probab. Comput.* **3**, 233–246 (1994)
14. de Fraysseix, H., Ossona de Mendez, P.: On topological aspects of orientations. *Discrete Math.* **229**(1–3), 57–72 (2001)
15. de Fraysseix, H., Pach, J., Pollack, R.: How to draw a planar graph on a grid. *Combinatorica* **10**(1), 41–51 (1990)

16. Duncan, C.A., Gansner, E.R., Hu, Y.F., Kaufmann, M., Kobourov, S.G.: Optimal polygonal representation of planar graphs. *Algorithmica* **63**(3), 672–691 (2012)
17. Eppstein, D., Mumford, E., Speckmann, B., Verbeek, K.: Area-universal and constrained rectangular layouts. *SIAM J. Comput.* **41**(3), 537–564 (2012)
18. He, X.: On floor-plan of plane graphs. *SIAM J. Comput.* **28**(6), 2150–2167 (1999)
19. Heilmann, R., Keim, D.A., Panse, C., Sips, M.: Recmap: Rectangular map approximations. In: IEEE Symposium on Information Visualization (INFOVIS'04), pp. 33–40. Minneapolis (2004)
20. House, D., Kocmoud, C.: Continuous cartogram construction. In: Proceedings of IEEE Visualization, pp. 197–204. Triangle Park (1998)
21. Izumi, T., Takahashi, A., Kajitani, Y.: Air-pressure model and fast algorithms for zero-wasted-area layout of general floorplan. *IEICE Trans. Fundam. Electron. Commun. Comput. Sci.* **E81–A**, 857–865 (1998)
22. Kawaguchi, A., Nagamochi, H.: Orthogonal drawings for plane graphs with specified face areas. In: Theory and Applications of Model of Computation (TAMC'07), pp. 584–594. Springer, Berlin (2007)
23. Koebe, P.: Kontaktp Probleme der konformen Abbildung. Berichte über die Verhandlungen der Sächsischen Akademie der Wissenschaften zu Leipzig. Math. Phys. Klasse **88**, 141–164 (1936)
24. Koźmiński, K., Kinnen, E.: Rectangular duals of planar graphs. *Networks* **15**, 145–157 (1985)
25. Leinwand, S.M., Lai, Y.T.: An algorithm for building rectangular floor-plans. In: Design Automation Conference, pp. 663–664. IEEE Press, Piscataway (1984)
26. Liao, C.C., Lu, H.I., Yen, H.C.: Compact floor-planning via orderly spanning trees. *J. Algorithms* **48**, 441–451 (2003)
27. Michalek, J., Choudhary, R., Papalambros, P.: Architectural layout design optimization. *Eng. Optim.* **34**(5), 461–484 (2002)
28. Moh, T.S., Chang, T.S., Hakimi, S.L.: Globally optimal floorplanning for a layout problem. *IEEE Trans. Circuits Syst.* **43**(9), 713–720 (1996)
29. Nöllenburg, M., Prutkin, R., Rutter, I.: Edge-weighted contact representations of planar graphs. In: Graph Drawing (GD'12), pp. 224–235. Springer (2013)
30. Raisz, E.: The rectangular statistical cartogram. *Geogr. Rev.* **24**(3), 292–296 (1934)
31. Rinsma, I.: Nonexistence of a certain rectangular floorplan with specified area and adjacency. *Environ. Plan.* **14**, 163–166 (1987)
32. Rosenberg, E.: Optimal module sizing in VLSI floorplanning by nonlinear programming. *Methods Models Oper. Res.* **33**, 131–143 (1989)
33. Schnyder, W.: Embedding planar graphs on the grid. In: Symposium on Discrete Algorithms (SODA'90), pp. 138–148. San Francisco (1990)
34. Tobler, W.: Thirty five years of computer cartograms. *Ann. Assoc. Am. Geogr.* **94**, 58–73 (2004)
35. Ueckerdt, T.: Geometric representations of graphs with low polygonal complexity. Ph.D. Thesis, Technische Universität, Berlin (2011)
36. Ungar, P.: On diagrams representing maps. *J. Lond. Math. Soc.* **28**, 336–342 (1953)
37. van Kreveld, M.J., Speckmann, B.: On rectangular cartograms. *Comput. Geom.* **37**(3), 175–187 (2007)
38. Wang, K., Chen, W.K.: Floorplan area optimization using network analogous approach. In: IEEE Transactions Circuits and Systems, pp. 167–170 (1995)
39. Wimer, S., Koren, I., Cederbaum, I.: Floorplans, planar graphs, and layouts. *IEEE Trans. Circuits Syst.* **35**(3), 267–278 (1988)
40. Yeap, K.H., Sarrafzadeh, M.: Floor-planning by graph dualization: 2-concave rectilinear modules. *SIAM J. Comput.* **22**, 500–526 (1993)