Complexity of Octagonal and Rectangular Cartograms

T. Biedl and B. Genc *

December 1, 2005

Abstract

In this paper, we study the complexity of rectangular cartograms, i.e., maps where every region is a rectangle, and which should be deformed such that given area requirements are satisfied. We study the closely related problem of cartograms with orthogonal octagons, and show that this problem is NP-hard. From our proof, it also follows that rectangular cartograms are NP-hard if we allow for the existence of a "sea", i.e., a region of arbitrarily high complexity on the outside of the drawing.

1 Introduction

A cartogram is a type of map used to visualize data. In a map regions are displayed in their true shapes and with their exact relations with the adjacent regions. However, such a map can only be used to demonstrate the actual area values of the regions. Sometimes, we need to display other data on a map, such as population, pollution, electoral votes, production rates, etc. One efficient way to do so is to modify the map such that the area of each shape corresponds to the data to be displayed. A map with given relationships between regions for which each region has pre-specified area is called a *cartogram*; see Section 2 for precise definitions. Cartograms are sometimes also called *diagrammatic maps* or *value-by-area maps*. See [1] for a web site with much information about cartograms.

There are two major cartogram types: contiguous area cartograms [2, 3, 6, 7, 12], where the regions are deformed but stay connected, and non-contiguous area cartograms [8], where regions preserve their shapes but may lose adjacency relationships. *Rectangular cartograms*, where every region is a rectangle is a specific type of contigous area cartograms which tries to preserve both the adjacency relations and the shape, but this does not exist for all area values. Kreveld and Speckmann [13] introduced the first automated algorithms for such cartograms. Heilmann et al. proposed RecMap [5] to approximate familiar land covering map region shapes by rectangles and to find a partition of the available screen space where the areas of these rectangular regions are proportional to given statistical values. Rahman

*School of Computer Science, University of Waterloo, Waterloo, ON N2L 3G1. {biedl,bgenc}@uwaterloo.ca. Research supported by NSERC.

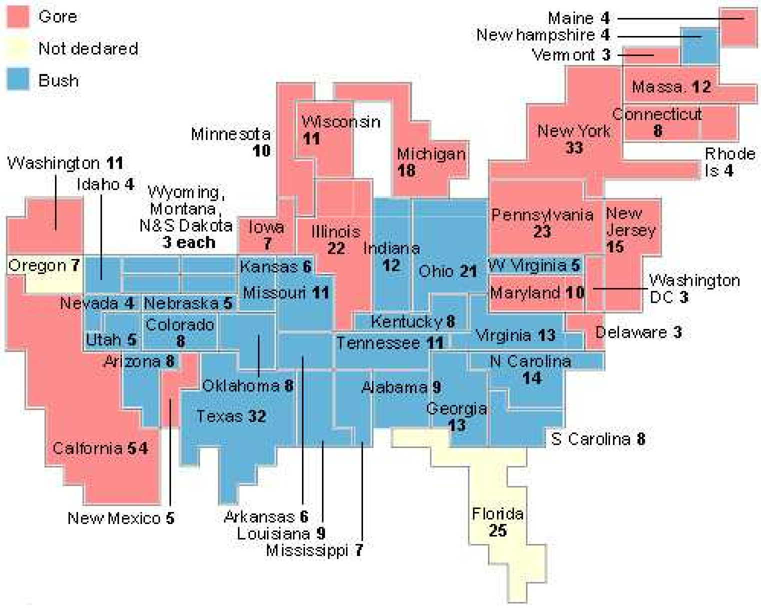

Figure 1: A cartogram of the 2000 US election. From Guardian Newspapers. (Image clipped for better fit.)

et al. studied slicing and good slicing graphs and their orthogonal drawings [9].They define an octagonal drawing as an orthogonal drawing of a plane graph such that each inner face is drawn as a rectilinear polygon of at most eight corners and the contour of the outer face is drawn as a rectangle. They show that any good slicing graph has an octagonal drawing with prescribed face areas.

It was left as an open problem whether testing the feasibility of a rectangular cartogram is NP-hard. In this paper, we make significant progress towards answering this question. We first study what we call *cartograms of orthogonal octagons*. These are cartograms where every region is an orthogonal polygon with at most 8 sides. We show that testing whether a cartogram of orthogonal octagons exist is NP-hard. In fact, in our reduction we use only two types of regions: rectangles and Z-shaped octagons.

We then use a very similar reduction to prove NP-hardness of a problem that is almost the same as rectangular cartograms: Here, all faces are rectangles, except for one face which is adjacent to the outside. Such a face exists in many real-life maps, corresponding to the "sea" around islands and peninsulas; see the examples in [13]. As in the paper by Kreveld and Speckmann, we also assume that the cartogram must be placed on a *canvas*, i.e., within a rectangle of fixed size.

Our paper is structured as follows: In Section 2 we define the cartogram problem formally, using the language of planar graphs. In Section 3, we give the NP-hardness reduction, both for octagonal cartograms and rectangular cartograms with a sea. We conclude in Section 4 with open problems.

2 Definitions

In this section, we give a formal definition of our problem. Cartograms have in the past mostly been defined by giving a drawing of a graph and requesting to change the area of regions. Since for our types of cartograms, there exist drawings that look radically different, but have the same geometry for each region, we prefer not to use this approach, and instead define cartograms in the language of plane graphs with specified geometry.

We recall first some graph theory definitions. A graph $G = (V, E)$ is called *planar* if it can be drawn in the plane without crossing. Such a crossing defines a cyclic order of incident edges around each vertex; the collection of these cyclic orders is called the *planar embedding*. A planar drawing of a planar graph defines connected regions of the plane called *faces*; the unbounded region is called the *outer-face*. A planar embedding of a graph defines uniquely the faces, except for the choice of the outer-face. A planar graph where both a planar embedding and an outer-face have been specified is called a *plane graph*.

Given a plane graph, we define the *dual graph* by defining a vertex for every face. For every edge in the primal graph incident to faces f_1 and f_2, we define a *dual edge* in the dual graph incident to the vertices of the faces f_1 and f_2.

An *orthogonal cartogram dual* is a plane graph with one special vertex C (the *canvas*) where every incidence between a vertex and an edge is labeled with one of {N,S,E,W} such that the following holds:

- If edge (v, w) is labeled N at v, then it is labeled S at w.
- If edge (v, w) is labeled E at v, then it is labeled W at w.
- If we consider the labels of edges at a vertex v in cyclical order according to the planar embedding, then S never follows or precedes N, and W never follows or precedes E.

Thus, in effect, the labels around a vertex v describe the angles of an orthogonal polygon, see Figure 2, and the edges in the graph describe which polygons must be adjacent. In particular, from the graph labels we can read how many corners each polygon must have; this corresponds to the number of times the label changes while walking around the cyclic order at each vertex. We say that a vertex in the cartogram dual *has k corners* if the corresponding polygon has k corners.

It may not be straightforward to see that such a plane graph indeed gives rise to a valid drawing, but this can be shown using the technique of converting an orthogonal representation into an orthogonal drawing proposed by Tamassia [11].

All that is needed to specify a cartogram hence is to demand an area of each face. Thus an *orthogonal cartogram* is an orthogonal cartogram dual G, together with a positive integer area(v) for every vertex $v \neq C$ of G. An orthogonal cartogram can be *realized* if there exists a drawing of the dual graph of G such that

- If (v, w) is labeled N at v, then the dual edge of (v, w) is drawn horizontally with the face of v below the edge.

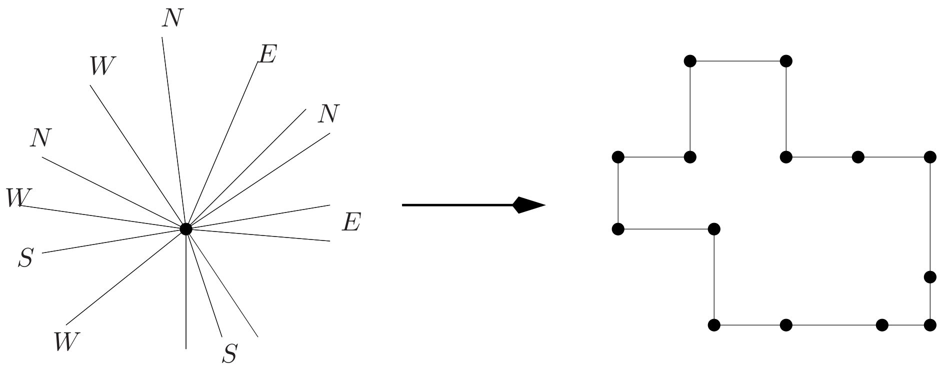

Figure 2: Reading the polygonal shape from the vertex labels. (We do not know the length of the edges of the polygon.)

- If (v, w) is labeled W at v, then the dual edge of (v, w) is drawn vertically with the face of v to the left of the edge.
- If $v \neq C$ is a vertex of G, then the face corresponding to v in the dual graph is an interior face in the drawing and has area equal to area(v).

We will sometimes additionally demand that the whole drawing fits inside a canvas. Thus we may specify a $w \times h$ rectangle R and demand that the drawing fit inside it. In particular, $w \cdot h$ must be at least the area of all other faces together, but it may be more, allowing for some "dead space" (also known as *the sea*) on the outer-face. Note that the aspect ratio of the rectangle for the canvas does not matter; if the drawing fits into any rectangle of area $w \cdot h$, then after suitable scaling it fits into all rectangles of area $w \cdot h$.

We are thus interested in the complexity of the following problems:

- Octagonal Orthogonal Cartogram: Given an orthogonal cartogram where every vertex has at most 8 corners, can it be realized within a given canvas?
- Rectangular Cartogram: Given an orthogonal cartogram where every vertex has four corners, can it be realized within a given canvas?
- Rectangular Cartogram with Sea: Given an orthogonal cartogram where every vertex except C has four corners, can it be realized within a given canvas?

We show that the first and third problem are NP-hard; the complexity status of the second problem remains open, but it seems likely that it is NP-hard as well because of the similarity to the other two problems.

We also note that for the first two problems, the requirement of fitting inside a canvas can be dropped, since this restriction can be simulated by adding more faces. On the other hand, the canvas restriction is crucial for the NP-hardness of the third problem.

3 NP-hardness

We show now that testing whether an octagonal cartogram can be realized is NP-hard. The proof is by reduction from PARTITION defined as follows. Assume that we are given a set A of positive integers $a_1 \dots a_n$ with $\sum_{i=1}^{n} a_i = 2S$ for some integer S. We want to find a subset I of A which satisfies $\sum_{a_i \in I} a_i = S$. It is known that this is NP-hard [4].

3.1 Construction

Given an instance of PARTITION $a_1, \dots, a_n$, we create the cartogram as follows. We have $2n+5$ faces $A_1, P_1, \dots, A_n, P_n$ and $M, B_1, \dots, B_4$, which are all rectangles except P_i, $i = 1, \dots, n$ is a Z-shaped octagon; see Figure 3(a). Figure 3(b) shows two other ways of representing this shape with the same angles; the choice between these two representations will be at the heart of our NP-hardness reduction.

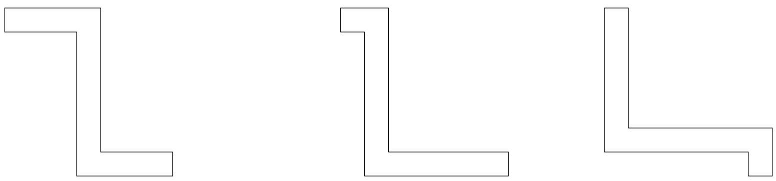

(a) (b)

Figure 3: (a) The shape for P_i. (b) Two other drawings with the same angles.

The adjacency relations between these faces are given in Figure 4, both as a drawing and by giving the cartogram dual. For easier visualization of the latter, we directed the edges and shown the label only for the tail of each edge; we also split the canvas C into multiple vertices. Furthermore, we show A_i and P_i only for the special cases $i = 1, n$ and for one generic i; the generic i needs to be repeated $n - 3$ times.

Now we explain the area requirements, depending on four parameters m, k, C and p. We will use $k = n$, $m = 2nS+4$, $C = 2n^2$ and $p = n+8$, but actually a wide range of parameters is possible and can be calculated from the proofs of the lemmas. The area requirements and purposes of faces are as follows:

- Each rectangle A_i corresponds to one number a_i of the PARTITION instance. We set area$(A_i) = C \cdot a_i$.
- Each Z-shaped octagon P_i acts as a "buffer" between rectangles A_i and A_{i+1} (or M); we set area$(P_i) = p$.
- M is a huge rectangle with area requirement m^2 that serves to split the rest of the canvas into essentially two parts.
- $F_1, \dots, F_4$ serve to build a frame that forces M to be an $m \times m$-square in any realization. We set area(F_1)=area$(F_3) = k$, area$(B_F) = m$ and area$(F_4) = m + 1$.

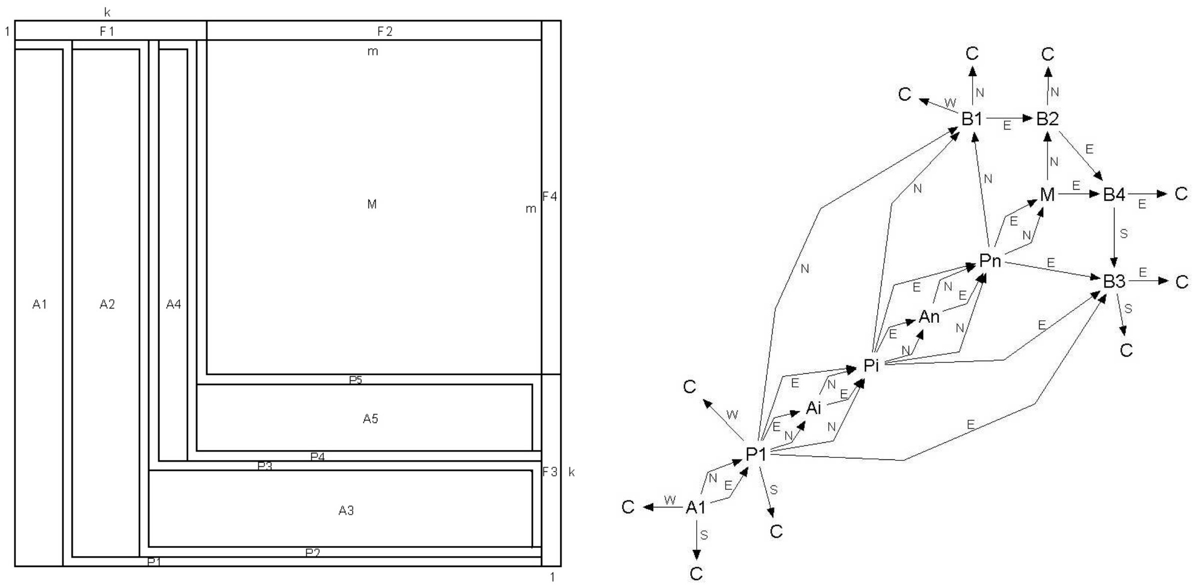

Figure 4: The drawing and the cartogram dual of the cartogram generated from the given Partition instance.

Some easy calculations show that with our choice of parameters we have $2CS + np = (m+k)^2 - m^2$; this shows that the area of all regions together is $(m+k+1)^2$ and by using an $(m+k+1) \times (m+k+1)$ as canvas, there is no empty space left for a sea.

To give an idea of the proportions of these rectangles, we show in Figure 5 the cartogram for the instance $\{1, 2, 4, 4, 5\}$ of Partition to scale, thus $k = n = 5$, $m = 2nS + 4 = 2 \cdot 5 \cdot 16 + 4 = 164$, $C = 2n^2 = 50$ and $p = n + 8 = 13$. Note that rectangle M is overwhelmingly large. The rest of the canvas therefore naturally splits into the area "left" and "below" M (and the tiny corner that is both); this will be crucial for our reduction.

3.2 Proof

We claim that our constructed cartogram is realizable iff the instance of Partition has a solution.

3.2.1 From cartogram to Partition

Assume first that we have a realization of the cartogram. We need some intermediary lemmas.

Lemma 3.1 *The widths and heights of M, F_1, F_2, F_3 and F_4 are as labeled in Figure 4.*

Proof Note that the left edges of F_3 and F_4 are collinear in any realization and touch the top and bottom of the canvas. Since F_3 and F_4 have a total area of $m + k + 1$ and the canvas has height $m + k + 1$, the widths of F_3 and F_4 has to be 1. Due to the individual

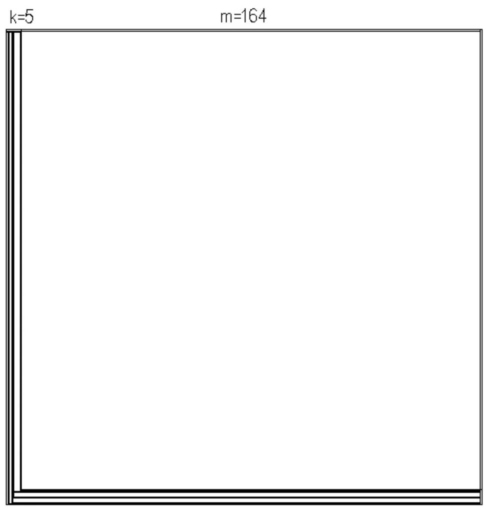

Figure 5: A drawing to scale for the instance $\{1, 2, 4, 4, 5\}$.

area requirements, this fixes the height of F_3 to k and the height of F_4 to $m + 1$. Similarly F_1 and F_2 will have a fixed height of 1 and F_1 will have a width of k while F_2 will have a width of m. This fixes the size of M to $m \times m$ in any realization. •

For the rest of the proof, we need two lines L and B. L marks the line through the left side of the rectangle M and B is the line through the bottom side of M. See also Figure 6. For each A_i, we now have two possible layouts: A_i may be placed to the left of line L or below line B. See Figure 6. One can show (though this is not crucial to our proof and hence will be omitted) that a rectangle A_i cannot be both below B and to the left of L, simply because there is not enough space for it. This choice between to the left of L or below B will be our main "decision gadget" in the reduction from PARTITION.

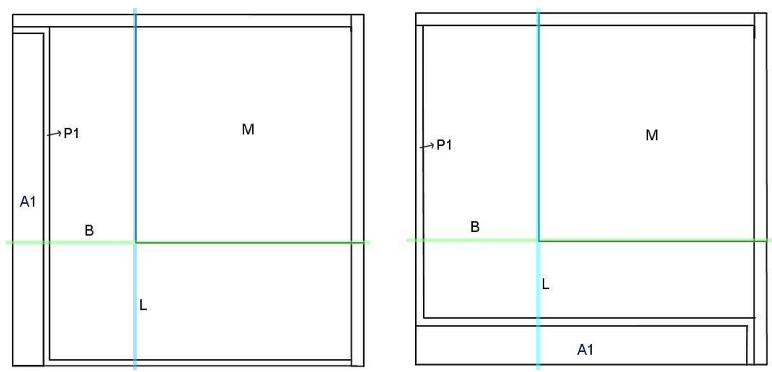

Figure 6: Two possible layouts for A_1. Drawing is not to scale.

Lemma 3.2 *The total area to the left of line L is less than $C(S + 1)$.*

Proof This holds by our choice of parameters. Note that it suffices to show $C(S+1) > k(m+k)$, since by Lemma 3.1, $k(m+k)$ is the area to the left of line L. For the proposed values of the parameters, and since we may assume that $n > 4$,

$$k(m+k) = n(2nS + 4 + n) = 2n^2S + n^2 + 4n < 2n^2S + 2n^2 = C(S+1).$$

as desired. •

An identical proof shows:

Lemma 3.3 *The total area below line B is less than $C(S+1)$.*

Lemma 3.4 *If I denotes the indices of rectangles A_i placed entirely to the left of L, then $\sum_{i\in I} a_i = S$.*

Proof By Lemma 3.2, we have $\sum_{i\in I} \text{area}(A_i) = \sum_{i\in I} Ca_i < C(S+1)$, so $\sum_{i\in I} a_i < S+1$, and hence $\sum_{i\in I} a_i \leq S$ since all numbers are integers. All rectangles with indices not in I are not entirely to the left of L and hence must be entirely below B. So $\sum_{i\notin I} \text{area}(A_i) < C(S+1)$ by Lemma 3.3, which similarly implies $\sum_{i\notin I} a_i \leq S$. Since $\sum_{i\in I} a_i + \sum_{i\notin I} a_i = 2S$, we must have equality for both sets. •

With this, a realization of the cartogram clearly gives a solution to the Partition instance.

3.2.2 From Partition to cartogram

Now we work on the other direction. Assume that the Partition instance has a solution I, i.e., $\sum_{i\in I} a_i = S$. Principally, the idea to construct the cartogram is easy: Let the width and height of $M, F_1, \ldots, F_4$ be as indicated in Figure 4, and position each A_i, P_i pair in one of the two fashions shown in Figure 6, depending on whether $i \in I$ or not.

The details of this are more complicated, because we need to choose the dimensions of A_i and P_i carefully such that all regions fit exactly into the region left by $M, F_1, \ldots, F_4$. We will only show that such coordinates must exist, by giving two layouts that don't quite work, and arguing that there exists a realization somewhere between them.

Let an *L-shape* be a 6-sided orthogonal polygon of which the third vertex on the boundary in clockwise order starting from the top leftmost vertex is the only reflex vertex of the shape boundary. Note that the region left inside the canvas after removing the rectangles M, F_1, F_2, F_3 and F_4 is an L-shape.

For the claims to come, we will need to analyze structures of various L-shapes, and hence introduce some notations. The *width* and *height* of an L-shape X are the width and height of its bounding box. Let $L(X)$ and $B(X)$ be the vertical and horizontal lines through the unique reflex vertex of X. We call the rectangle to the left of $L(X)$ the *left region*, and its width the *left width.* We call the region below B the *bottom region*, and its height the *bottom height.*

Let L_0 be the L-shape left of the canvas after placing $M, F_1, \ldots, F_4$. It has width and height $m+k$ and left width and bottom height k. Two other L-shapes are the areas occupied by the two layouts that we are going to define.

We also need some notations for the realizations of P_i. Consider Figure 6 again. In either method of realizing P_i, it is the union of three rectangles that overlap at the corners. One of these spans the height of the available area and is to the left of L; we call this the *left rectangle*. Another one spans the width of the available area and is below B; we call this the *bottom rectangle*. The third one is only adjacent to A_i and the outside, and usually quite small; we call this the *end rectangle*.

Now we are ready to define the layouts precisely; see Figure 7 for an illustration.

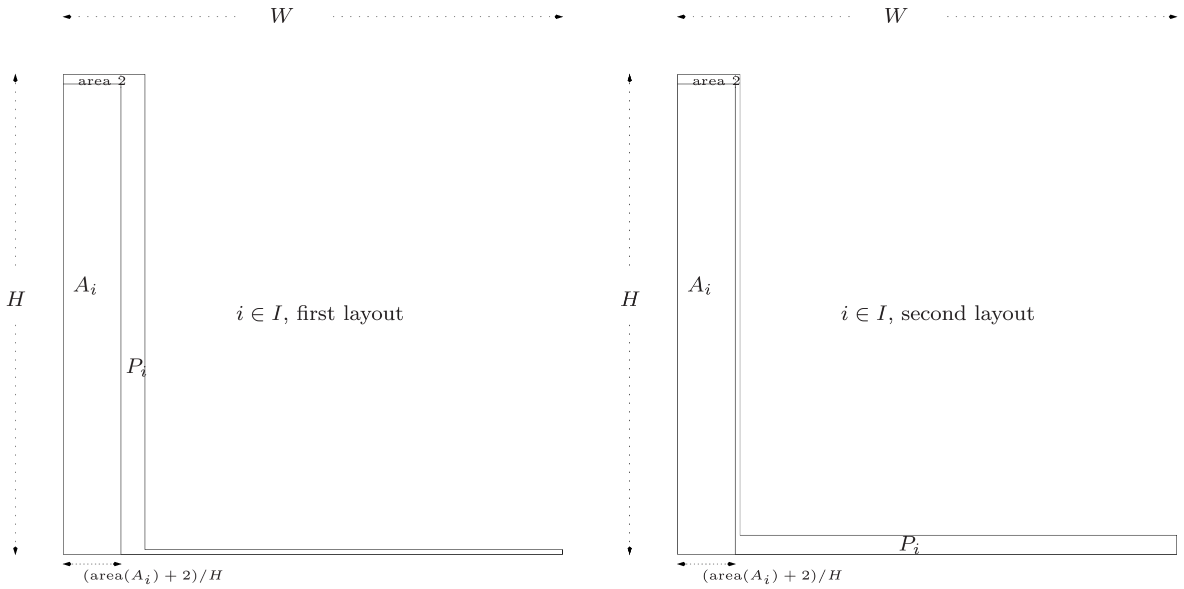

Figure 7: The two layouts for an A_i, P_i pair if $i \in I$. The layouts for $i \notin I$ can be obtained by flipping A_i and P_i along the diagonal.

- Set $H = m + k$ and $W = m + k$; these keep track of the bounding box of the L-shape into which we place all the remaining faces.
- For $i = 1, \ldots, n$:
  - Let A_i be placed at the bottom left corner of the remaining free region.
  - If $i \in I$, make A_i (almost) as tall as possible. More precisely, set the width of A_i to be $(\text{area}(A_i) + 2)/H$, and set the height such that the area is correct. Notice that the end rectangle of P_i then will have area at least 2.
  - If $i \notin I$, make A_i (almost) as wide as possible. More precisely, set the height of A_i to be $(\text{area}(A_i) + 2)/W$, and set the width such that the area is correct. Notice that the end rectangle of P_i then will have area at least 2.

- Choose as shape for P_i the one that corresponds to whether $i \in I$ as in Figure 6.
- The dimensions for the rectangles of P_i depend on which layout we are creating:
  - In the first layout, define dimensions of P_i such that the left rectangle has area at least $n+4$ and the bottom rectangle has area at least 2. Recall that the total area of P_i is $n+8$, and that the three rectangles of P_i overlap, so these constraints can be satisfied simultaneously.
  - In the second layout, define dimensions of P_i such that the bottom rectangle has area at least $n+4$ and the left rectangle has area at least 2.
- The union of A_i and P_i is an L-shape. Decrease W by the left width of that L-shape, and H by the bottom height of that L-shape.

Let L_1 be the union of $A_1, P_1, \dots, A_n, P_n$ in the first layout, and L_2 be the union of $A_1, P_1, \dots, A_n, P_n$ in the second layout. Note that both L_1 and L_2 are L-shapes with width and height $m+k$, and their area is the same. Also, both L_1 and L_2 contain all A_i's with $i \in I$ in their left region, and all A_i's with $i \notin I$ in their bottom region. L_1 contains almost all area of the P_i's in its left region, while L_2 contains almost all area of the P_i's in the bottom region.

Unfortunately neither L_1 nor L_2 is a realization of the cartogram, since they don't fit into L_0. Our goal is to show that some layout "between" L_1 and L_2 does fit. To do so, we need to show that L_1 is too wide and L_2 too slim.

Lemma 3.5 *The left width of L_1 is not smaller than the left width of L_0, and the left width of L_2 is not bigger than the left width of L_0.*

Proof Note that L_0, L_1 and L_2 all have the same height and width. Therefore the left width determines the area of the left region. For L_0, this left region has area

$$k(m+k) = n(2nS + 4 + n) = 2n^2 S + n^2 + 4n = C \cdot S + n^2 + 4n.$$

L_1 was defined such that its left region contains all A_i with $i \in I$. (It also contains some of the areas of the other A_i's in the corner, but this is irrelevant here.) Furthermore, it contains the left rectangle for *all* P_i's by definition; each of these has area at least $n+4$ by construction. Therefore, the left region of L_1 has area at least

$$\sum_{i \in I} \text{area}(A_i) + \sum_{i=1}^{n} (n+4) = C \cdot S + n^2 + 4n,$$

which is at least as much as the area of the left region of L_0. So the left width of L_1 cannot be smaller than the left width of L_0.

For the other claim, we can show similarly that the bottom height of L_2 is not smaller than the bottom height of L_0. Since L_2 and L_0 have the same area, same height and same width, this implies that the left width of L_2 cannot be bigger than the left width of L_0. •

Lemma 3.6 *There exists a realization of the cartogram such that all rectangles of all P_i's have area at least 2.*

Proof L_1 was defined by having at least area $n + 4$ in the left rectangle of each P_i, and at least 2 units area in the bottom rectangle, whereas for L_2 the area distribution was vice versa. We can now define intermediary layouts between L_1 and L_2, where we gradually shift area from the left rectangle of each P_i to the bottom rectangle. This can be done in such a way that the function of layouts is continuous. In the beginning we thus have L_1, where the left width is at least k (by Lemma 3.5), and at the end we have L_2 where the left width is at most k. By the mean value theorem, therefore at some point we have a layout L^* with left width k, and its height and width (as for all layouts we create) is $m + k$. Since also all layouts have total area $(m + k)^2 - m^2$ (which is the total area of all A_i's and P_i's together), therefore L^* has exactly the shape of L_0. Combining it with the layout for $M, F_1, \dots, F_4$ gives therefore the desired realization. Finally note that all rectangles of all P_i's in both L_1 and L_2 have area at least 2, so this also holds for all intermediary drawings, and hence for our cartogram as well. •

Thus given a solution to PARTITION, we can obtain a cartogram, which proves our main result:

Theorem 3.7 *Testing whether an orthogonal octagonal cartogram is realizable is NP-hard.*

3.3 Rectangular cartograms with sea

We now show that essentially the same construction also leads to an NP-hardness result for rectangular cartograms if a sea is allowed. In our construction we eliminate octagon P_i and replace it with a rectangle R_i that connects to A_i at the bottom and to B_1 at the top. Rectangle R_i has area requirement 1. The canvas is unchanged, i.e., a square of side length $m+k+1$; note that this leaves some empty space for the sea since R_i requires less area than P_i. The resulting sea will take on the role of a "buffer" between rectangles. See Figure 8.

With exactly the same proof as before one shows that if this cartogram can be realized, then the PARTITION instance has a solution; note that nowhere in this part of the proof did we make use of the octagons P_i.

On the other hand, if PARTITION has a solution, then we create a cartogram as before. Now we can place R_i inside the end rectangle of P_i (if $i \in I$) or inside the left rectangle of P_i (if $i \notin I$); we know that these rectangles have area at least 2 and hence there is sufficient space for R_i. We thus obtain the following theorem.

Theorem 3.8 *Testing whether a rectangular cartogram with a sea is realizable is NP-hard (with the canvas restriction).*

Note that the use of a canvas is mandatory in this case, because the adjacency relations do not guarantee the exact shape of the cartogram and without canvas there is a trivial solution as in Figure 9.

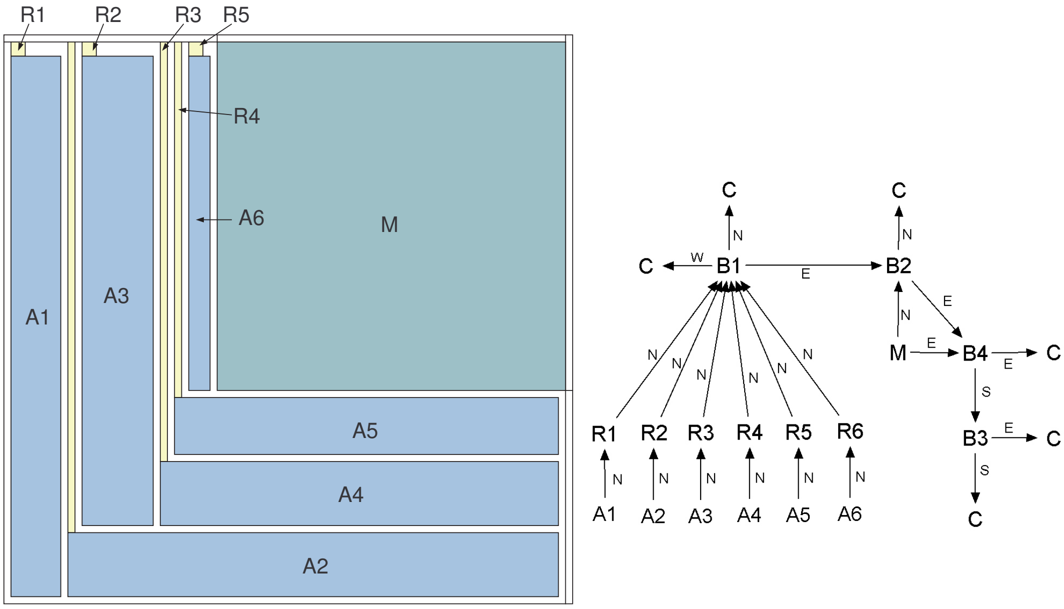

Figure 8: NP-hardness with only rectangular regions and a sea. Cartogram dual on the right; we have omitted (numerous) arcs from A_i and R_i to the sea.

4 Conclusion and open problems

In this paper, we studied the complexity of realizing a rectangular cartogram that is bounded by a canvas. The main question (is this NP-hard?) remains open, but we showed that two closely related results are indeed NP-hard. In particular, the small (and realistic) step of adding a sea bounded by a canvas to a rectangular cartogram makes the problem NP-hard.

The most pressing open problem is to resolve the complexity of rectangular cartograms. Can we do away with the sea?

Another very interesting problem is whether this problem is actually NP-complete (i.e., whether it is in NP). To show this, one would have to show that any realization can be specified with coordinates that are polynomial in the input. Since the input has area requirements, but the realization would (presumably) be specified with x-coordinates and y-coordinates, this is far from trivial (the coordinates might well be radicals.)

Finally, we are interested in exploring cartograms with orthogonal octagons (or k-gons for some small number of k) further. Note that k-gons are more flexible than rectangles, and thus more cartograms will be realizable. In particular, Speckmann et al. showed that for any vertex-weighted plane triangulated graph, G, there exists a cartogram with at most 60 corners per face that has G as dual graph and the areas according to the vertex weights of G [10]

Also, a number of existing heuristics seem to rely on using k-gons for small k. What are good heuristics for cartograms with orthogonal octagons?

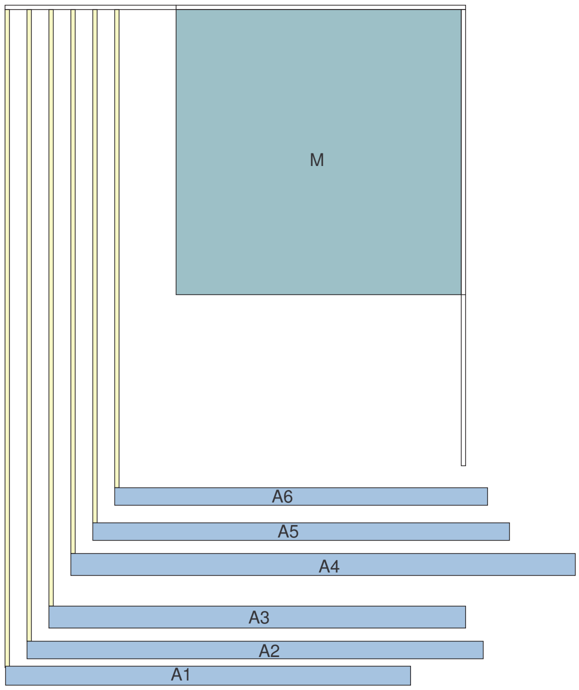

Figure 9: Layout is trivial without the canvas.

Acknowledgments

The authors would like to thank Marc van Kreveld and Bettina Speckmann for helpful input.

References

[1] Cartogram central. http://www.ncgia.ucsb.edu/projects/Cartogram_Central/, Web site at University of California at Santa Barbara, last accessed November 26, 2005.

[2] J. Dougenik, N. Chrisman, and D. Riemeyer. An algorithm to construct continuous area cartograms. *Professional Geographer*, 37:75–81, 1985.

[3] H. Edelsbrunner and R. Waupotitsch. A combinatorial approach to cartograms. In *SCG '95: Proceedings of the eleventh annual symposium on Computational geometry*, pages 98–108. ACM Press, 1995.

[4] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. Freeman, 1979.

[5] R. Heilmann, D. A. Keim, C. Panse, and M. Sips. Recmap: Rectangular map approximations. In *INFOVIS '04: Proceedings of the IEEE Symposium on Information Visualization (INFOVIS'04)*, pages 33–40. IEEE Computer Society, 2004.

[6] D. H. House and C. J. Kocmoud. Continuous cartogram construction. In *VIS '98: Proceedings of the conference on Visualization '98*, pages 197–204. IEEE Computer Society Press, 1998.

[7] D. A. Keim, S. C. North, and C. Panse. Cartodraw: A fast algorithm for generating contiguous cartograms. *IEEE Transactions on Visualization and Computer Graphics*, 10(1):95–110, 2004.

[8] J. Olson. Noncontiguous area cartograms. *Professional Geographer*, 28:371–380, 1976.

[9] M. S. Rahman, K. Miura, and T. Nishizeki. Octagonal drawings of plane graphs with prescribed face areas. In *Graph-Theoretic Concepts in Computer Science: 30th International Workshop*, volume 3353 of *Lecture Notes in Computer Science*, pages 320–331, 2004.

[10] B. Speckmann. Personal communication.

[11] R. Tamassia. On embedding a graph in the grid with the minimum number of bends. *SIAM J. Computing*, 16(3):421–444, 1987.

[12] W. Tobler. Pseudo-cartograms. *The American Cartographer*, 13:43–50, 1986.

[13] M. v. Kreveld and B. Speckmann. On rectangular cartograms. In *European Symposium on Algorithms (ESA'04)*, volume 3221 of *Lecture Notes in Computer Science*, pages 724–735, 2004.