Finding minimal rectangulations

Michael R. Klug

February 16, 2023

Abstract

We give a polynomial-time algorithm to find a minimal rectangulation of a region in the plane that can be tiled with rectangles. This has been discovered several times (see, for example [6], [7], [3], [2]) – we were unaware of this and stumbled on it ourselves. After writing up the result, we became aware of the aforementioned results. This note simply our take on this problem, which we thought may be worth recording.

Let A be a compact subset of the plane whose boundary consists of a finite number of disjoint line segments all of whose endpoints are on the grid $\mathbb{Z}^2 \subset \mathbb{R}^2$ – see Figure 1. We call such regions *rectangulable* since such regions can be partitioned into rectangles that are disjoint except along their boundary as in Figure 1. We are interested in the problem of, given A, finding a rectangulation of A with the smallest number of rectangles. Such rectangulations will be called *minimal*.

Theorem 1. *There is a polynomial-time algorithm that, given a rectangulable region A, finds a minimal rectangulation of A.*

Before getting started on the proof of Theorem 1, we give an application which was our motivation for thinking about rectangulations:

Corollary 1. *There exists a polynomial-time algorithm that, given $f : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$, computes $\chi^{\mathrm{geom}}(f)$.*

Proof. (of Corollary) We will compute $\chi^{\mathrm{geom}}(f)$ by applying the algorithm in Theorem 1 to each of the f-monochromatic connected regions of $M(f)$ (of which there polynomially many). □

For the proof of Theorem 1, we introduce some terminology. For any pair of line segments we mention l and l' we will assume that they intersect transversely in their interior (i.e., they intersect in finitely many points). Given a rectangulable region in the plane A, we will refer to the *internal and external corners* of A as in Figure 1. We will refer to the line segments forming the boundary of A as the *walls* of A. We say that a line segment l *ends* at an internal corner c of A if one of its endpoints is at c. We also define l *pairing* the two internal corners c_1 and c_2 if it ends in c_1 and c_2. We say l *ends* in a wall of A or another line segment l' if an endpoint of l is on a non-corner part of A or the interior of l'. A *rectangulation* $\mathcal{R}$ of A is a collection of line segments such that the line segments cut A up into rectangles.

The *size* of $\mathcal{R}$, denoted $\mathrm{size}(R)$, is the number of rectangles that $\mathcal{R}$ cuts A into. A rectangulation $\mathcal{R}$ of A is *minimal* if it has the smallest size amongst all rectangulations of A.

The outline of our proof of Theorem 1 is as follows: We start by giving a method of describing particularly simple rectangulations of A and showing that a minimal rectangulation of this form can be found in polynomial time. Then we prove that there must exist a minimal rectangulation of this form.

We now give a method of describing certain simple rectangulations of A. A *corner pairing* is a set $P = \{P_1, ..., P_n\}$ where each $P_i = \{c_i^1, c_i^2\}$ where c_i^1 and c_i^2 are internal corners of A with a line segment l_i between them in A, $P_i \cap P_j = \emptyset$ for $i \neq j$, and the lines l_i and l_j are disjoint for $i \neq j$. A corner pairing P is *full* if no additional pair of internal corners P_{n+1} can be added to P to obtain a larger corner pairing – see Figure 2 for two different examples of full corner pairings and Figure 3 for all possible corner-pairing line segments. Given a corner pairing P, it determines rectangulation of A called the *realization* of P which we denote $\mathcal{R}(P)$ as in the right on Figure 1 (which is $\mathcal{R}(P_1)$ for the pairing P_1 in Figure 2) where we use the convention of drawing horizontal line segments from all of the internal corners in A that are not contained in any P_i (this is just a convention and we could have just as well taken vertical lines or some mix of vertical and horizontal lines). See Figure 4 for another example – namely, $\mathcal{R}(P_2)$ for P_2 in Figure 2.

Lemma 1. *Let A be a rectangulable region in the plane, P a side pairing of A and $\mathcal{R}(P)$ the realization of P. Then $\mathcal{R}(P)$ is a rectangulation of A. Let p be the size of P, u be the number of internal corners in A not contained in any pair in P, and k be the number of external corners of A. If P is maiximal, then*

$$4\,\mathrm{size}(\mathcal{R}(P)) = 2p + 3u + k$$

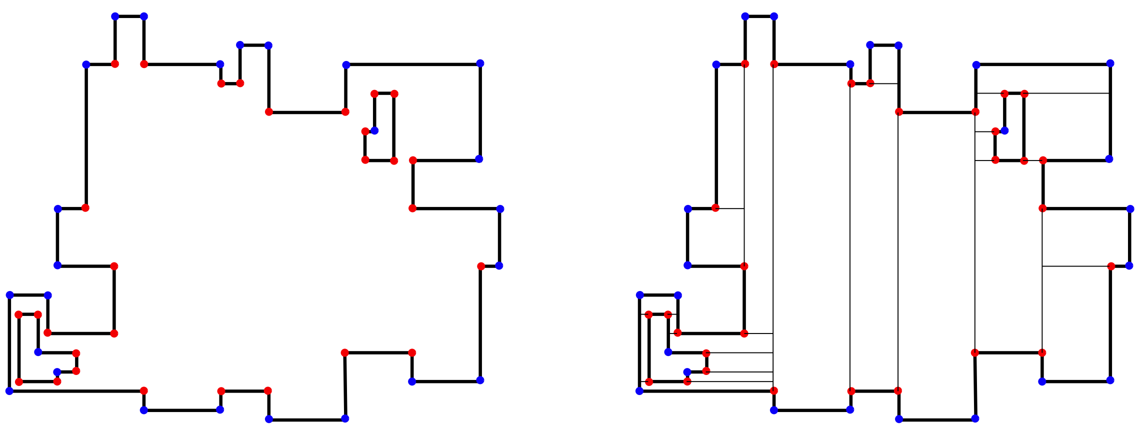

Figure 1: Example of a rectangulizable region in the plane and a rectangulation of it. On the left is an example of a connected but not simply-connected rectangulizable region in the plane. The internal corners are marked in red and the external corners are marked in blue. On the right is a minimal rectangulation of a rectangulizable region – here using 21 rectangles. This is in fact the rectangulation that results from the algorithm for finding such a minimal rectangulation that we present.

Proof. To see that $\mathcal{R}(P)$ is a rectangulation of A, we proceed by induction on the size of P – in fact we will want to strengthen the inductive hypothesis to allow all such regions A together with some number of internal line segments with at least one endpoint at an internal corner and the other endpoint at a wall or internal corner. As a base case, assume that P is empty, so the realization just consists of drawing horizontal lines from each internal corner of A to a wall of A. This is seen to be a rectangulation by induction on the number of internal corners of A. If there are no internal rectangles, then A itself is a rectangle, and if there are $n + 1$, then by adding the horizontal line

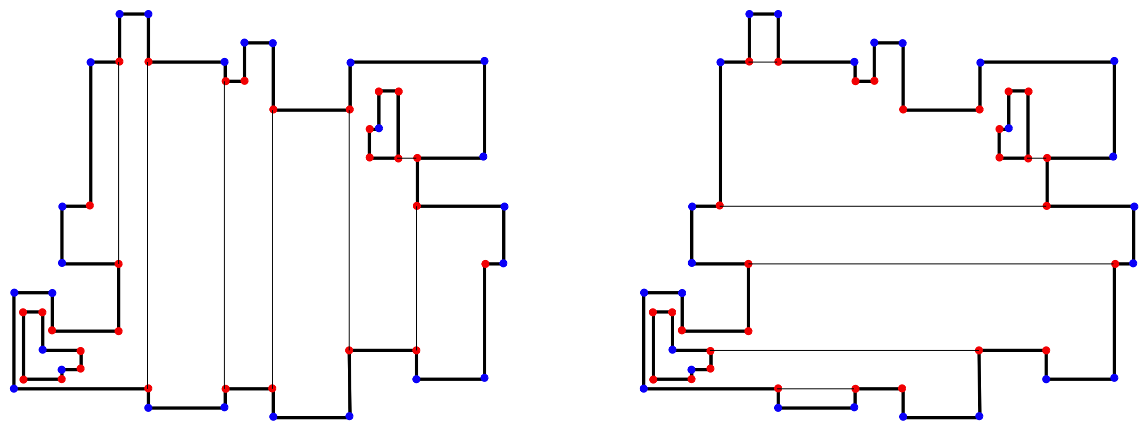

Figure 2: Different examples of full corner pairings. On the left is a full corner pairing P_1 that is also maximal. The result of applying the algorithm in Theorem 1 to P_1 is the minimal rectangulation $\mathcal{R}(P_1)$ shown on the right in in Figure 1. On the right is a full corner pairing P_2. Note that $|P_2| < |P_1|$.

on the internal corner(s) with say lowest y-coordinate(s), we obtain a region A' with fewer internal corners than A to which we can apply the inductive hypothesis. Therefore, in the case where P is empty, $\mathcal{R}(P)$ is a rectangulation of A. Now assume that P consists of $n+1$ pairs of internal corners. Given one such pair $\{c^1_{n+1}, c^2_{n+1}\} \in P$, let A' be A together with the line segment pairing c^1_{n+1} and c^2_{n+1}, and we consider A' together with a corner pairing $P - \{c^1_{n+1}, c^2_{n+1}\}$ which one fewer pair of internal corners, then the realization of $P - \{c^1_{n+1}, c^2_{n+1}\}$ is a rectangulation of A' and therefore, $\mathcal{R}(P)$ is a rectangulation of A.

To verify the rest of Lemma 1, we note that, since every rectangle has 4 external corners, $4\,\mathrm{size}(\mathcal{R}(P))$ is the number of external corners in $\mathcal{R}(P)$. Each element of P contributes 2 external corners, each of the u unparied corners contributes 3 external corners (and since P is assumes to be full, this is not redundant), and the k external corners of A are the remaining external corners. □

A vertex cover C of a graph G is a set of vertices such that each edge of G has at least one vertex in C. A vertex cover is minimal if it has the smallest size amongst all vertex covers of G. The problem of determining for a graph G and an integer k if G has a vertex cover of size less than or equal to k is NP-complete [4]. A matching M of G is a collection of edges of G no two of which share a vertex. A matching is maximal if it is as large as possible among all matchings of G. Maximal matchings of graphs can be constructed in polynomial time – see for example [1]. However, in the case of bipartite graph G, Kőnig's theorem says that the size of a maximal matching of G is the same as the size of a minimal vertex cover of G [5]. Furthermore, the proof constructs a minimal vertex cover given a maximal matching in polynomial time. A set of vertices of a graph is a vertex cover if and only if its complement is an independent set, and therefore the problem of finding a maximal independent set in a graph is equivalent to the problem of finding a minimal vertex cover. Therefore, by Kőnig's theorem in a bipartite graph, we can find a maximal independent set in polynomial time.

Let G_A be the graph whose vertices are the line segments in A with both ends at internal corners of A. The edges of G_A come in two types. For the first type, the *corner sharing edges*, we place an edge between two distinct vertices l_1 and l_2 when they share an internal corner. For the second type, the *intersection edges*, we put an edge between two line segments l_1 and l_2 if they

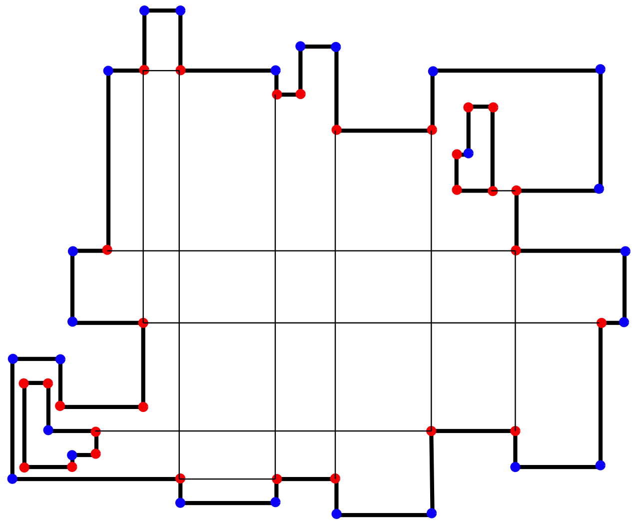

Figure 3: Rectangulable region with all corner pairing line segments drawn in. Note that many of them intersect and some corners have more than one such line segment ending at them.

intersect in their interiors. Note that G_A is biparite with the partition being into the vertices that are horizontal lines and the vertices that are vertical lines. Now, by construction, we have that corner pairings of A are in bijection with independent sets of vertices in G_A where the bijection takes the set of vertices to the pairs of the ends of the corresponding line segments. The resulting corner pairings are full if and only if no vertices can be added to the corresponding independent set in G_A. Therefore, we wish to find a maximal independent set in G_A and then the corresponding full corner pairing P will be as large as possible among all corner pairings and therefore $\mathcal{R}(P)$ will be a minimal rectangulation of A.

Lemma 2. *Let A be a rectangulable region in the plane. Among the rectangulations of A in the set $\{\mathcal{R}(P) : P$ is a full corner pairing of A $\}$, we can find such a rectangulation of smallest size amongst rectangulations of the form $\mathcal{R}(P)$ in polynomial time.*

Proof. Consider the graph G_A whose vertices are the corner pairing line segments in A and whose edges come in two forms: (1) two vertices v and u are adjacent if the corresponding line segments intersect, and (2) two vertices v and u are adjacent if the corresponding corner pairings share an internal corner. Note that G_A is biparite with a biparition being given by separating the vertices into those representing vertical line segments and those representing horizontal line segments.

By construction, corner pairings P correspond to independent sets in G_A. Furthermore, by Lemma 1, we see that $\text{size}(\mathcal{R}(P))$ is minimized when the size of P is maximized. Therefore, we want to find a maximal independent set in G_A. By Kőnig's theorem as in the preceding discussion, this can be done in polynomial time. $\square$

Now we prove that we can restrict attention to rectangulations of the form $\mathcal{R}(P)$ for P a full corner pairing. Before doing this, we need one technical lemma. Given two line segments l and l' in the plane, the set of *internal intersections* of l and l' is the set of points in the intersection of the interiors of l and l' (note that intersections at the endpoints are not included).

In what follows, we use the notation 1_{claim} for the function with $1_{\text{claim}}(x) = 1$ if the claim holds true for x and 0 otherwise.

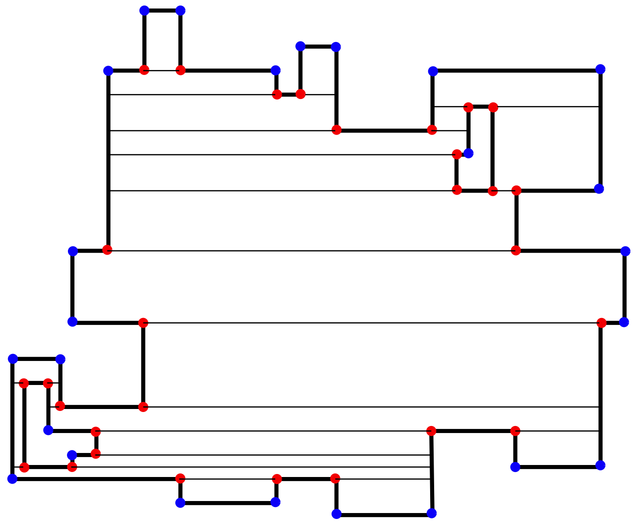

Figure 4: A non-minimal rectangulation of a rectangulable region. This is the rectangulation $\mathcal{R}(P_2)$ where P_2 is the full but not maximal corner pairing on the right in Figure 2. Note that $\mathcal{R}(P_2)$ consists of 22 rectangles whereas the rectangulation of the same region shown in Figure 3 has 21 rectangles.

Lemma 3. *Let A be a rectangulable region in the plane together with several disjoint corner-pairing line segments $w_1, ..., w_n$. Let $l_1, .., l_m$ be line segments in A that are disjoint from A and all of the w_i, except possibly at their endpoints, such that each l_i ends at a unique internal corner of A that is not an endpoint of any w_j. Suppose that there are k external corners of A. Then the number of external corners in $A \cup \{w_1, ..., w_n\} \cup \{l_1, ..., l_m\}$ is*

$$2n + k + 3m + 4\sum_{i \neq j} 1_{l_i \text{ and } l_j \text{ have an interior intersection}} + \sum_{i \neq j} 1_{l_i \text{ and } l_j \text{ share an endpoint}}$$

Proof. We proceed by induction on m the number of lines $l_1, ..., l_m$. The base case, $m = 0$ is immediate since each wall contributes 2 external corners to $A \cup \{w_1, ..., w_n\}$.

For the induction step, assume we have line segments $l_1, ..., l_{m+1}$ as in the statement of the lemma. Then l_{m+1} contributes 3 or 4 external corners from its endpoints. In particular, from the endpoint at the unique internal corner that l_{m+1} ends at that is not an endpoint of one of the w_i, l_{m+1} contributes 2 external corners if one of the l_i for $i < m + 1$ has this as an endpoint and 1 external corner otherwise The other 2 external corners are from the other endpoint of l_{m+1}. In addition, for each internal intersection of l_{m+1} with another line l_i, we obtain 4 more external corners. The result then follows by induction. $\square$

Lemma 4. *Let A be a rectangulable region in the plane. There exists a minimal rectangulation of A of the form $\mathcal{R}(P)$ for P a full corner pairing of A.*

Proof. Let R be a minimal rectangulation of A. Let P be a maximal set of endpoints of corner-pairing line segments in R so that no two such sets of endpoints intersect. We show that

$$\text{size}(R) = \text{size}(\mathcal{R}(P))$$

thus proving the result.

Note first that every line segment l in R must have at least one end as a internal corner of A, since otherwise both of the ends of l are either other line segments in R or walls of A and by removing l from R, we would obtain a new rectangulation of A with one fewer rectangle than R, thus contradicting the minimality of R.

Let $w_1, ..., w_n$ denote the set corner-pairing line segments with endpoints in P and let $l_1, .., l_m$ denote the other line segments in R so that each l_i ends at an internal corner of A - note then that we meet the conditions of Lemma 3 with these choices of w_i and l_j.

Let c denote the number of internal corners of A that are not contained in any set in P. By Lemma 1, $\mathcal{R}(P)$ is a rectangulation of A. By Lemma 1, we see that $\mathcal{R}(P)$ has $2n + k + 3c$ external corners since all of the lines in $\mathcal{R}(P)$ other than those whose endpoints are in P are horizontal and hence disjoint and there is exactly one line segment coming from each of the c internal corners of A that are not contained in any set in P. Therefore

$$4\,\mathrm{size}(\mathcal{R}(P)) = 2n + k + 3c$$

since each rectangle contributes 4 external corners.

Now, R may differ from $\mathcal{R}(P)$, but we claim that it has the same size. This is because in a rectangulation of A, every internal corner of A must have at least one line segment of the rectangulation ending at that corner. Therefore, $m \geq c$ and by Lemma 3, we have

$$4\,\mathrm{size}(R) \geq 2n + k + 3c$$

Therefore by minimality of R,

$$4\,\mathrm{size}(R) = 2n + k + 3c = 4\,\mathrm{size}(\mathcal{R}(P))$$

□

We now have all of the pieces in place for the proof of Theorem 1:

Proof. (of Theorem 1) Given the region A, using Lemma 2 we can find a minimal rectangulation of A in polynomial time since Lemma 4 guarantees that there is a minimal rectangulation of the form $\mathcal{R}(P)$ for P a full maximal corner pairing. □

References

[1] Jack Edmonds. Paths, trees, and flowers. *Canadian Journal of mathematics*, 17:449–467, 1965.

[2] David Eppstein. Graph-theoretic solutions to computational geometry problems. In *International Workshop on Graph-Theoretic Concepts in Computer Science*, pages 1–16. Springer, 2009.

[3] L Ferrari, Pathamadi V Sankar, and Jack Sklansky. Minimal rectangular partitions of digitized blobs. *Computer vision, graphics, and image processing*, 28(1):58–71, 1984.

[4] Richard M Karp. Reducibility among combinatorial problems. In *Complexity of computer computations*, pages 85–103. Springer, 1972.

[5] Dénes Kőnig. Gráfok és mátrixok. *Matematikai és Fizikai Lapok*, 38:116–119, 1931.

[6] Witold Lipski Jr, Elena Lodi, Fabrizio Luccio, Cristina Mugnai, and Linda Pagli. On two-dimensional data organization ii. *Fundamenta Informaticae*, 2(1):245–260, 1978.

[7] Tatsuo Ohtsuki. Minimum dissection of rectilinear regions. In *Proc. 1982 IEEE Symp. on Circuits and Systems, Rome*, pages 1210–1213, 1982.