

Complexity measures for mosaic drawings

Citation for published version (APA):

Bouts, Q. W., Speckmann, B., & Verbeek, K. A. B. (2017). Complexity measures for mosaic drawings. In *WALCOM: Algorithms and Computation - 11th International Conference and Workshops, WALCOM 2017, Proceedings* (pp. 149-160). (Lecture Notes in Computer Science; Vol. 10167). Springer.
https://doi.org/10.1007/978-3-319-53925-6_12

Document license:

Unspecified

DOI:

[10.1007/978-3-319-53925-6_12](https://doi.org/10.1007/978-3-319-53925-6_12)

Document status and date:

Published: 01/01/2017

Document Version:

Accepted manuscript including changes made at the peer-review stage

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

Visual Complexity Measures for Mosaic Drawings

Quirijn W. Bouts, Bettina Speckmann, and Kevin Verbeek

Department of Mathematics and Computer Science, TU Eindhoven, The Netherlands
[q.w.bouts|b.speckmann|k.a.b.verbeek]@tue.nl

Abstract. Graph Drawing uses a well established set of complexity measures to determine the quality of a drawing, most notably the area of the drawing and the complexity of the edges. For contact representations the complexity of the shapes representing vertices also clearly contributes to the complexity of the drawing. Furthermore, if a contact representation does not fill its bounding shape completely, then also the complexity of its complement is visually salient.

We study the complexity of contact representations with variable shapes, specifically mosaic drawings. Mosaic drawings are drawn on a tiling of the plane and represent vertices by *configurations*: simply-connected sets of tiles. The complement of a mosaic drawing with respect to its bounding rectangle is also a set of simply-connected tiles, the *channels*.

We prove that simple mosaic drawings without channels necessarily require $\Omega(n^2)$ area. This bound is tight. If we use only straight channels, then every outerplanar graph with k ears requires at least $\Omega(nk)$ area. Also this bound is tight: we show how to draw outerplanar graphs with k ears in $O(nk)$ area with L-shaped vertex configurations and straight channels. Finally, we argue that L-shaped channels are strictly more powerful than straight channels, but may still require $\Omega(n^{7/6})$ area.

1 Introduction

Graph Drawing uses a well established set of complexity measures to determine the quality of a drawing (see, for example, the overview by Purchase [10]). For node-link drawings of planar graphs the arguably most prominent measures are the area of the drawing and the complexity (number of bends) of the edges. In addition to the classic node-link drawings, there are also other well-established drawings styles, most notably *contact representations*. Here the vertices of a graph are represented by a variety of possible shapes and the edges are implied by point or side contacts between these shapes. In certain settings the shapes are fixed: for example circles in Koebe's theorem [8] or rectangles in rectangular duals and rectangular cartograms [9, 11, 12]. In other scenarios there is a certain variability in the shapes: for example, when using rectilinear polygons for rectilinear cartograms [1, 5] or mosaic tiles and pixels [2, 3] to compose shapes.

Contact representations do not draw edges explicitly and hence the complexity of the edges is not a valid quality measure. However, in scenarios where

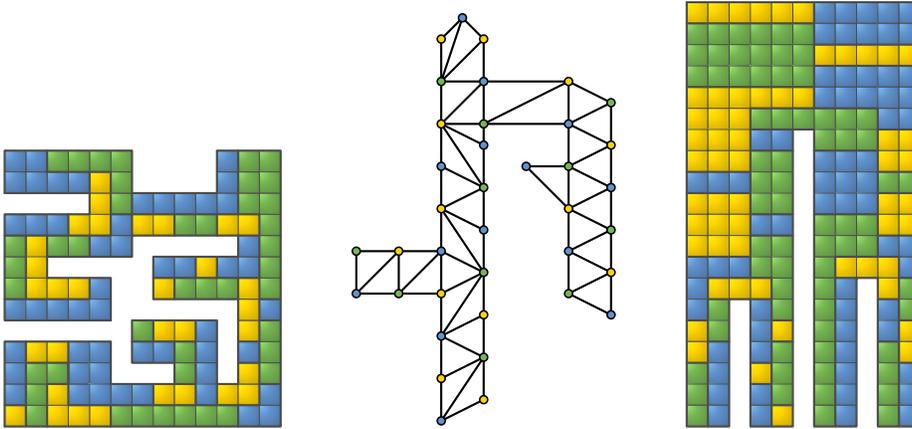


Fig. 1: A mosaic drawing with complex channels (left) of an outerplanar graph (middle). The same graph drawn with straight channels (right).

the shapes representing vertices are not fixed, the complexity of these shapes clearly contributes to the visual quality of the drawing. For example, there are several sequences of papers which strive to represent (weighted) planar triangulated graphs with rectilinear polygons of the lowest possible complexity, which is 8 in both the weighted [1] and in the unweighted case [4].

Contact representations are said to be *proper* if the shapes corresponding to vertices form a partition of the bounding shape. For example, a rectangular dual is a proper contact representation, since it consists of a partition of a rectangle into rectangles. Similarly, certain outerplanar graphs have a proper touching triangle representation [7], which consists of a partition of a triangle into triangles. If a contact representation is not proper, then the complexity of its complement with respect to an appropriate bounding shape is also visually salient.

We propose to study the complexity of contact representations with variable shapes. Specifically, we focus on mosaic drawings, which were recently introduced by Cano *et al.* [3]. The same drawing style was independently described by Alam *et al.* [2] as pixel and voxel (in 3D) drawings. Mosaic drawings are drawn on a tiling of the plane and represent vertices by so-called *configurations*: simply-connected sets of tiles (see Fig. 1). The complement of a simple mosaic drawing with respect to its bounding rectangle is also a set of simply-connected tiles, which we call *channels*. Figure 1 shows two different representations of an outerplanar graph. In both cases the vertices have complexity 6 (rectangles and L-shapes), but the complexity of the channels differs substantially. We would like to argue that a higher complexity of the channels increases the visual complexity of the drawing and hence is likely to be an impediment to the user. However, more complex channels allow us to significantly reduce the area of a drawing. In this paper we hence investigate the trade-offs between these two competing quality measures.

Mosaic drawings. Following Cano *et al.* [3] we define mosaic drawings for *plane triangulated graphs*. Mosaic drawings are drawn on a *tiling* \mathcal{T} of the plane. For a given set $S \subseteq \mathcal{T}$ of tiles we define the *tile dual* of S to be the graph which has a vertex for each tile in S and an edge connecting two vertices if and only if the two corresponding tiles share a side. A *configuration* is a set of tiles with a connected tile dual. A configuration C is *simple* if the tiles in C are simply connected. The *boundary* of a simple configuration C is a simple rectilinear polygon and the *complexity* of a configuration is the number of vertices of its boundary.

Two configurations C_1 and C_2 are *adjacent* if and only if they contain tiles $t_1 \in C_1$ and $t_2 \in C_2$ such that t_1 and t_2 share a side. A *mosaic drawing* $D_{\mathcal{T}}(G)$ of a plane triangulated graph $G = (V, E)$ on \mathcal{T} represents every vertex $v \in V$ by a simple configuration $C(v)$ of tiles from \mathcal{T} . Two configuration $C(u)$ and $C(v)$ representing vertices u and v are adjacent if and only if $(u, v) \in E$.

Simple mosaic drawings.

A mosaic drawing is *simple* if (i) the union of its configurations is simply connected (there are no holes in the interior of the mosaic drawing) and (ii) two adjacent configurations share exactly one contiguous piece of boundary (two configurations do not touch two or more times).

and (iii) whenever four tiles meet in a point at least two adjacent tiles belong to the same configuration (note that this might be the outer configuration).

(see Fig. 2).

In this paper we focus on regular tilings, and specifically, the square tiling. In the following we hence denote mosaic drawings simply by $D(G)$.

We construct mosaic drawings within a bounding rectangle R . The area of a drawing $D(G)$ is the number of tiles inside R . The complement of a drawing $D(G)$ with respect to R is a set of simple configurations, which we call *channels*.

The boundary of a simple mosaic drawing $D(G)$ is a simple rectilinear polygon as well. The boundary of $D(G)$ naturally divides into maximal straight segments, which we refer to as the *boundary segments* of $D(G)$. Mosaic drawings respect the embedding of G . In particular, this implies that the order of the vertices of G around the outer face is the same as the order of the configurations around the boundary of $D(G)$.

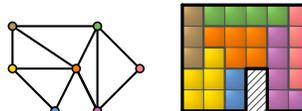


Fig. 2: Simple mosaic drawing $D(G)$ with one straight channel (shaded). The boundary of $D(G)$ is marked in black, thick gray lines indicate configuration boundaries.

Results and organization. We first consider mosaic drawing without channels, that is, mosaic drawings which are proper contact representations of a given triangulated outerplane graph G with n vertices. Alam *et al.* [2] prove that there are k -outerplanar graphs such that any mosaic drawing of these graphs requires $\Omega(kn)$ area. In Section 3 we show that there are outerplane (that is, 1-outerplane) graphs such that any mosaic drawing of these graphs *without channels* requires $\Omega(n^2)$ area. More specifically, there exist outerplane graphs with k ears, such that any mosaic drawing of these graphs has either $\Omega(n^2/k^2)$ area

or total channel complexity $\Omega(k)$. This bound is tight, since the algorithm by Chiang *et al.* [4] can be used to construct mosaic drawings without channels in $O(n^2)$ area. In Section 4 we consider *straight channels*. We prove that every outerplane graph with k ears requires at least $\Omega(kn)$ area if we are allowed to use only straight channels. Also this bound is tight for sufficiently small k : we describe an algorithm which draws outerplane graphs with k ears in $O(kn)$ area. Finally, in Section 5 we show that L-shaped channels are strictly more powerful than straight channels, but may still require $\Omega(n^{7/6})$ area.

In the remainder of this section we first review mosaic drawings and introduce definitions and notation. We then briefly discuss mosaic drawings with general channels of arbitrary complexity. In particular, when using general channels, it follows directly from Alam *et al.* [2] that each triangulated outerplane graph has a simple mosaic drawing in $O(n)$ area, which is trivially tight.

General channels. If we allow general channels of arbitrary complexity then we can construct a compact mosaic drawings of a triangulated outerplanar graph G directly from an orthogonal grid drawing of the weak dual G^* of G . Specifically, one can represent each triangle in G , which corresponds to a vertex in the orthogonal grid embedding of G^* , using four tiles in a 2×2 layout (see Fig. 3). One can then connect these tiles into a mosaic drawing of at most 3 times the width and height of the initial orthogonal grid drawing. Dolev *et al.* [6] describe an algorithm which constructs orthogonal grid drawings of k -outerplanar graphs of maximum degree 4 in $O(kn)$ area. Since G^* has degree at most 3, we can directly use their result to construct simple mosaic drawings with arbitrary channels in $O(n)$ area.

Alam *et al.* [2] also use the algorithm by Dolev *et al.* [6] to construct mosaic drawings for k -outerplanar graphs in $O(kn)$ area. However, they do not consider the dual G^* but work directly with G . As a consequence they first need to convert G into a graph of maximum degree four, which can significantly increase the number of vertices (and hence the complexity of the drawing), while maintaining the general area bound.

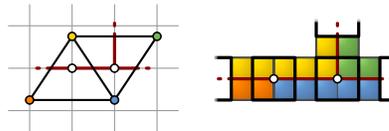


Fig. 3: A mosaic drawing of G based on an orthogonal grid embedding of G^* using 2×2 tiles per dual vertex and 1-wide connection pieces.

2 Preliminaries

Outer-path. Let G be a triangulated outerplane graph and let G^* be its weak dual. If G^* is a path then we call G an *outerpath*. The *length* of an outerpath is the number of vertices of G^* . A so-called *ear* of G is a triangle which is dual to a vertex of degree 1 in G^* . Each ear has exactly one vertex of degree 2, its so-called *tip*. Every outerpath has exactly two vertices of degree 2, the tips of its two ears. These two tips naturally divide the vertices of G into two consecutive sequences of vertices, the *upper* and the *lower sequence* of the outerpath G . In

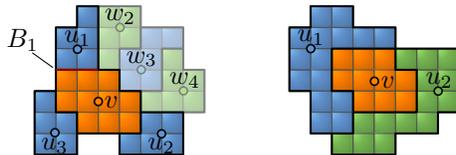


Fig. 4: Illustrations for Lemma 1.

our lower bound constructions we use a particular type of outerpath, namely a so-called *outerzigzag*, which is also known as a triangle strip. Its vertices have degree at most 4. The number of vertices in the upper and the lower sequence of an outerzigzag differ by at most one.

Next we show that we can restrict ourselves to natural mosaic drawings of maximal outerplanar graphs. A mosaic drawing is *natural* if it follows the unique outerplane embedding of a maximal outerplanar graph, where each configuration is adjacent to the outer face in the same order as in the outerplane embedding (or its reverse).

Lemma 1. *In any proper mosaic drawing $D(G)$ of a maximal outerplanar graph G , every configuration must be adjacent to the outer face.*

Proof. For the sake of contradiction, assume that a configuration $C(v)$ is not adjacent to the outer face. We perform a case analysis on the degree of v in G . Since G is maximally outerplanar, the degree of v is at least two. If v has at least 3 neighbors u_1, u_2 , and u_3 , then the boundary of $C(v)$ contains three parts: the shared boundaries $B_i = C(v) \cap C(u_i)$ for $i = 1, \dots, 3$. Now consider the part of the boundary of $C(v)$ between B_1 and B_2 (see Fig. 4 left). Since $D(G)$ has no point contacts and $C(v)$ is not adjacent to the outer face, there must be a sequence of vertices $u_1 = w_1, w_2, \dots, w_k = u_2$ such that $C(w_i)$ is adjacent to $C(w_{i+1})$ for $1 \leq i < k$. The same holds for the parts of the boundary of $C(v)$ between B_2 and B_3 , and B_3 and B_1 . This implies that G contains a K_4 $\{v, u_1, u_2, u_3\}$ as a minor, which contradicts the outerplanarity of G .

If v has two neighbors u_1 and u_2 , then the boundary of $C(v)$ can be partitioned into two parts $B_1 = C(v) \cap C(u_1)$ and $B_2 = C(v) \cap C(u_2)$. This directly implies that $D(G)$ has a point contact or that $C(u_1) \cap C(u_2)$ is not contiguous (see Fig. 4 right). This is not allowed in a proper mosaic drawing. \square

Lemma 2. *Every proper mosaic drawing $D(G)$ of a maximal outerplanar graph G is natural.*

Proof. Lemma 1 shows that every configuration $C(v)$ must be adjacent to the outer face. Furthermore, the shared boundary between a configuration $C(v)$ and the outer face must be contiguous, since otherwise $C(v)$ forms a cut of $D(G)$, implying that v is a cut vertex of G . Now let v_1, \dots, v_n be the order of vertices implied by the unique outerplane embedding of G . Furthermore, let B_i be the shared boundary between $C(v_i)$ and the outer face. Note that, by definition, a set of consecutive vertices $\{v_i, v_{i+1}\}$ cannot be a cut of G . If B_i and B_{i+1} are

not consecutive along the boundary of $D(G)$, then $C(v_i) \cup C(v_{i+1})$ forms a cut of $D(G)$. Thus, $D(G)$ must be natural. \square

3 Mosaic drawings without channels

In this section we consider mosaic drawings without channels of an triangulated outerplane graph G with n vertices. As was shown in [4], we can always obtain such drawings with $O(n^2)$ area. Here we show that this bound is tight: there exist outerplane graphs that require $\Omega(n^2)$ area. In fact we show something stronger: to obtain sub-quadratic area, the number of boundary segments on the boundary of $D(G)$ must be proportional to the number of ears of G .

Lemma 3. *Any mosaic drawing $D(G)$ of an outerzigzag G of size n , where all vertex configurations are adjacent to the same or two consecutive boundary segments, requires $\Omega(n)$ width and $\Omega(n)$ height.*

Proof. Recall that mosaic drawings preserve the outerplane embedding of G and that hence the order of configurations along the boundary of $D(G)$ corresponds to the order of the vertices along the two sequences of G . The configurations $C(u_i)$ and $C(v_i)$ which correspond to a diagonal (u_i, v_i) necessarily form a cut in $D(G)$. These pairs of vertex configurations have to be nested $\Omega(n)$ times requiring $\Omega(n)$ width and $\Omega(n)$ height to draw (see Fig. 5). \square

We define a k -comb ($k \geq 3$) as an outerplane graph with n vertices and k ears, where each outerpath incident to an ear (a *leg*) is an outerzigzag with $\Theta(n/k)$ vertices, and all legs are on the same side with respect to the remaining part of the graph (see Fig. 6).

Lemma 4. *Any mosaic drawing of a k -comb G with n vertices requires $2k$ boundary segments or $\Omega(n/k)$ width and $\Omega(n/k)$ height.*

Proof. If the vertex configurations of one of the legs are adjacent to the same or two consecutive boundary segments of $D(G)$, then $D(G)$ requires $\Omega(n/k)$ height and $\Omega(n/k)$ width by Lemma 3. Otherwise, the vertex configurations of a single leg must be adjacent to at least 3 consecutive (the embedding is fixed)

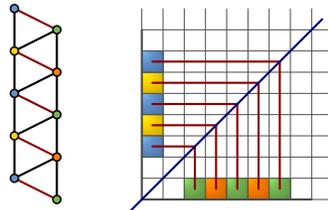


Fig. 5: $D(G)$ requires $\Omega(n)$ width and $\Omega(n)$ height since $\Omega(n)$ configurations have to cross the blue cut.

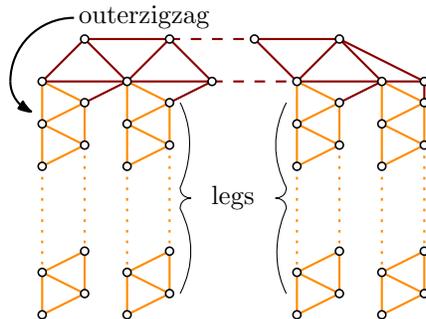


Fig. 6: A k -comb with k legs (orange) containing $\Theta(n/k)$ vertices each.

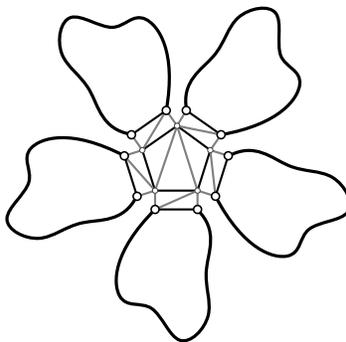


Fig. 7: By combining 5 copies of a graph we can assume that at least one copy is only adjacent to channels with tiles adjacent to a single side of R .

boundary segments of $D(G)$. Since two consecutive legs can share at most one of the boundary segments, we need at least 2 boundary segments for each leg, resulting in a total of $2k$ boundary segments. \square

Corollary 1. *Any mosaic drawing without channels of a 3-comb with n vertices requires $\Omega(n^2)$ area.*

4 Straight channels

In this section we consider mosaic drawings with straight channels. Lemma 4 directly implies that we need a total channel complexity of $\Omega(k)$ for drawings of outerplane graphs with k ears in subquadratic area, but multiple straight channels are generally more organized than fewer more complex channels. Here we show that outerplane graphs with k ears have mosaic drawings with straight channels of area $O(kn)$, and that this bound is tight for sufficiently small k .

To prove the lower bound, we consider *one-sided* mosaic drawings: every vertex configuration and channel must have a tile adjacent to the bottom of R . One-sided mosaic drawings only require proper channels on the bottom of R and

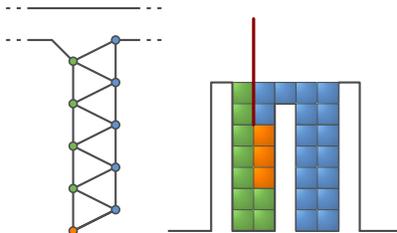


Fig. 8: The red line segment has to be crossed $\Omega(n)$ times.

can otherwise have an arbitrary boundary. Below we show that it is sufficient to consider one-sided mosaic drawings for our lower bound constructions.

Lemma 5. *Let G be an outerplane graph with n vertices and k ears such that a one-sided mosaic drawing $D_1(G)$ of G with certain channel restrictions requires $\Omega(f(n, k))$ area. Then there exists an outerplane graph G' with $O(n)$ vertices and $O(k)$ ears such that a mosaic drawing $D(G')$ of G' with the same channel restrictions also requires $\Omega(f(n, k))$ area.*

Proof. We construct G' by combining 5 copies of G , attaching them to a triangulated pentagon as illustrated in Figure 7. Since R has only 4 sides, there must be a copy of G such that all corresponding vertex configurations have a tile adjacent to the same side of R in $D(G')$. These vertex configurations thus form a one-sided mosaic drawing of G , possibly by rotating the drawing. \square

Lemma 6. *There exists an outerplane graph G with n vertices and k ears such that any one-sided mosaic drawing $D(G)$ with only straight channels requires $\Omega(\min(nk, n^2/k))$ area.*

Proof. Let G be the graph obtained by attaching another leg (outerzigzag) of size $n/2$ to a $(k - 1)$ -comb with $n/2$ vertices. Note that the other legs have size $\Theta(n/k)$. We will argue that any one-sided drawing $D(G)$ of G has width $\Omega(\min(k, n/k))$ and height $\Omega(n)$.

Consider the leg of size $n/2$. Let $v_1, \dots, v_{n/2}$ be the vertices of the leg in order along the outer face and assume w.l.o.g. that the corresponding configurations occur counter-clockwise along the boundary of $D(G)$. Since the drawing has only straight channels, the tiles of $C(v_1), \dots, C(v_{n/2})$ adjacent to the boundary of $D(G)$ must have non-decreasing x -coordinates. Let $v_{n/4}$ be the tip of the ear of this leg, and let t be the last tile of $C(v_{n/4})$ adjacent to the boundary (see Fig. 8). If we draw a line segment up from t as illustrated in Fig. 8, then this line segment must be crossed by all $\Omega(n)$ unions of the form $C(v_i) \cup C(v_{n/2-i+1})$. This directly implies that $D(G)$ must have height $\Omega(n)$.

Next we consider the other legs. Lemma 3 implies that we need $\Omega(k)$ boundary segments, or $D(G)$ has width at least $\Omega(n/k)$. Since every channel can add at most 4 boundary segments, we need at least $\Omega(k)$ channels in the first case, which require a tile each. As a result, the width of $D(G)$ is at least $\Omega(\min(k, n/k))$, implying a total area of at least $\Omega(\min(nk, n^2/k))$. \square

Corollary 2. *There exists an outerplane graph G with n vertices such that any mosaic drawing $D(G)$ with only straight channels requires $\Omega(n\sqrt{n})$ area.*

Drawing algorithm. We now present an algorithm for making a mosaic drawing $D(G)$ with straight channels for a maximal outerplane graph G . The algorithm is incremental, starting with a single edge of G and repeatedly adding a vertex attached to the endpoints of an existing edge. Note that every maximal outerplane graph can be constructed in this way. An edge in this construction is called *open* if we still need to add a vertex to it to obtain G . That is, an edge is open if and only if it is the initial edge or an internal edge of G . In every step we arbitrarily choose an open edge to extend with a vertex.

An open edge (a, b) is represented in $D(G)$ in two possible ways, as shown at the top of Fig. 9: one of the two vertex configurations is an L-shape and lies either on top of the other configuration (Case (I)) or touches the other configuration from the side (Case (II)). The horizontally mirrored case, where b is an L-shape instead of a , is also possible, but completely symmetric. Therefore we will consider only the cases shown in the figure. Furthermore note that the vertex configurations can still extend further as indicated by the open borders in the figure. In particular, both a and b can be L-shapes. However, such extensions are not relevant for the different cases.

We now show how to extend the drawing when a vertex c is added to an open edge (a, b) . There are four cases (see Fig. 9):

- (i): **(a, c) and (c, b) are not open.** In both Case (I) and (II) we simply fill up the remaining hole with c .
- (ii): **(a, c) is open and (c, b) is not open.** In Case (I) we place c below b and shift a down such that it forms a new Case (I) with c . In Case (II) we extend

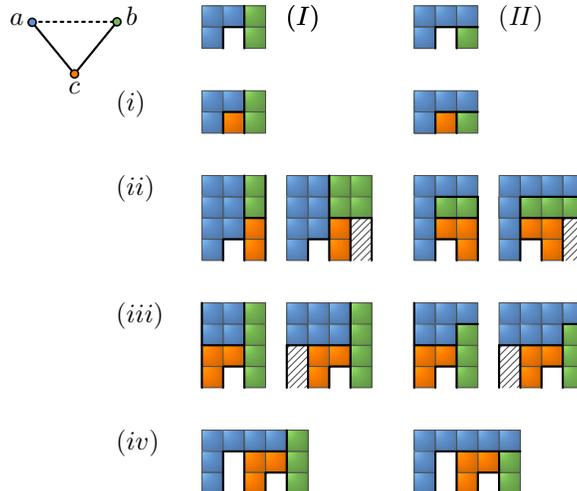


Fig. 9: The cases of the incremental drawing algorithm.

b into the hole and place c below it as an L-shape pointing to the left. The open edge (a, c) then forms the symmetric version of Case (I).

(iii): (a, c) is not open and (c, b) is open. In both Case (I) and (II) we extend a into the hole and place c below it as an L-shape pointing to the right. The open edge (c, b) then forms a new Case (I).

(iv): (a, c) and (c, b) are open. In both Case (I) and (II) we put c in the hole as an L-shape pointing to the right. Then the open edge (a, b) forms a new Case (II) and the open edge (c, b) forms a new Case (I).

In the above construction Case (iv) is a special case. We call a vertex c in Case (iv) a *splitter*. The configuration of a splitter can and sometimes must extend a channel downwards. It must do so before it is no longer part of an open edge. Therefore, some cases must be handled slightly differently when a or b is a splitter. This is shown on the right side of each case in Fig. 9 whenever this is relevant. Furthermore, splitters are part of two open edges simultaneously. Thus, in the independent construction of the two open edges, its configuration may get a different height. This can be fixed by extending the construction of one of the open edges vertically until the height of the configuration of the splitter matches. This way the configuration of the splitter will not have any additional complexity or incorrect adjacencies.

Lemma 7. *The above algorithm computes a simple mosaic drawing $D(G)$ that correctly represents G .*

Proof. By construction the resulting drawing $D(G)$ contains all adjacencies in G . Thus, we simply need to argue that $D(G)$ does not have any adjacencies not in G . We use the following invariant: for every open edge (a, b) the left side of the configuration of a is on the boundary of $D(G)$ or a is a splitter. The same holds for b and the right side. It is straightforward to argue that this invariant is maintained throughout all the cases. Then it is also easy to argue that no unwanted adjacencies are introduced when adding a vertex c . Therefore, $D(G)$ correctly represents G . \square

Theorem 1. *For every maximal outerplane graph G with n vertices and k ears there is a mosaic drawing $D(G)$ with the following properties:*

1. $D(G)$ has $O(k)$ straight channels.
2. $D(G)$ has $O(nk)$ area.
3. All vertex configurations of $D(G)$ are L-shaped or rectangular.

Proof. We show that the above algorithm computes a mosaic drawing with the desired properties. First note that only splitters introduce (straight) channels (one each). Since every open edge must correspond to a unique ear, and the number of open edges increases for every splitter (Case (iv)), there can be at most k channels. For the complexity of the vertex configurations, note that the complexity of the configurations never increases when adding a vertex. Therefore the complexity of a configuration can be at most its initial complexity: rectangular or L-shaped. To argue the area of $D(G)$, we show that $D(G)$ has $O(k)$

width and $O(n)$ height. The width only increases in Case (iv) (by 2 columns) or when a channel is introduced (by 1 column), which both involves a splitter. Every splitter is involved in exactly one Case (iv) and can introduce a channel only once. Therefore, the width of $D(G)$ is linear in the number of splitters, which is $O(k)$. To argue the height we use a compaction argument. Consider two rows of $D(G)$ such that no vertex configuration has a horizontal boundary on the horizontal line separating the two rows. Then we can replace these two rows by a single row without changing adjacencies and without increasing the complexities of the vertex configurations. Now, since every line separating two rows must contain a horizontal boundary of a vertex configuration, and there are at most $3n$ such horizontal boundaries (vertex configurations are at most L-shapes), the height of $D(G)$ is $O(n)$. \square

5 L-Shaped channels

In this section we prove a lower bound on the area of mosaic drawings with L-shaped channels. We first show that the lower bound construction for straight channels cannot directly be extended to L-shaped channels. For this we define *generalized k -combs*: k -combs where legs can have different sizes. Note that the example used for the lower bound in Lemma 6 is a generalized k -comb.

Lemma 8. *Every generalized k -comb G with n vertices allows a one-sided mosaic drawing $D(G)$ with L-shaped channels of $O(n)$ area.*

Proof. By using one L-shaped channel per leg which almost immediately bends to the left, a complete leg can be drawn with $O(1)$ height (see Fig. 10). We can then draw all legs next to each other to obtain a mosaic drawing of $O(n)$ width. Care has to be taken when drawing the first and last leg of the comb to ensure all vertex configurations are adjacent to a channel originating from the bottom. Finally, one big L-shaped channel spanning the entire drawing is used to border the “spine” of the generalized k -comb, resulting in a drawing of $O(n)$ area. \square

Before we specifically consider L-shaped channels, we give a more general result on mosaic drawings with channels of constant complexity. As implied by

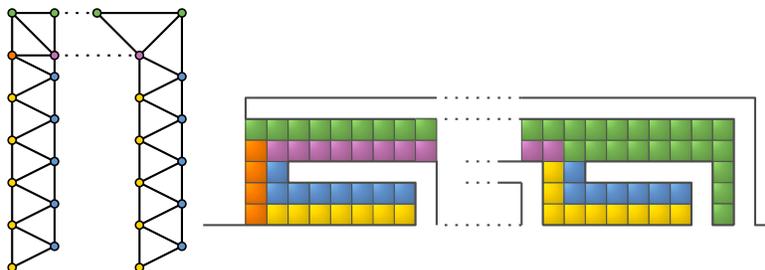


Fig. 10: Using L-shaped channels we can create a one-sided drawing of a generalized k -comb in linear space.

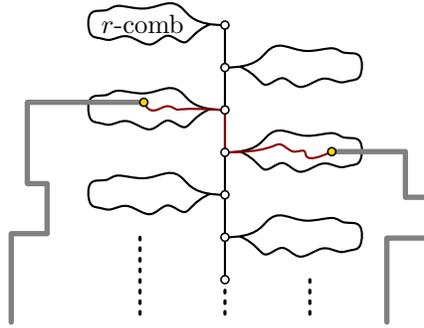


Fig. 11: The graph dual of a (k, r) -signpost, with r -combs as legs.

Lemma 5, it is sufficient to consider one-sided mosaic drawings. We use the following trivial observation.

Observation 2. *A channel of complexity c adds at most $4c$ boundary sequences to a drawing $D(G)$.*

To prove the following lower bounds, we need a special type of outerplane graph: a (k, r) -signpost. A (k, r) -signpost is like a k -comb, but with two important differences: (1) the legs are connected on alternating sides of the “spine”, and (2) every leg is an r -comb. The graph dual of a (k, r) -signpost is shown in Fig. 11.

Lemma 9. *Any one-sided mosaic drawing $D(G)$ with constant complexity channels of a (k, r) -signpost G with n vertices must have width and height $\Omega(\min(k, n/k))$, for some constant r .*

Proof. If the channels of $D(G)$ can have complexity c , then we require that $r > c/2$. Since r is constant, the legs of the r -combs in a (k, r) -signpost have size $\Omega(n/k)$. By Lemma 4 we know that each r -comb either requires width and height $\Omega(n/k)$, or is adjacent to $2r$ boundary segments of $D(G)$. In the latter case, since $2r > c$, every representation of an r -comb must contain a tile adjacent to the bottom of R . Therefore, the union of the configurations of two consecutive r -combs on opposite sides of the “spine” of the (k, r) -signpost separates $D(G)$ into two parts (see Fig. 11). Since the (k, r) -signpost contains $\Omega(k)$ of such pairs of consecutive r -combs, $D(G)$ must have width and height $\Omega(k)$. Combining both cases results in the claimed lower bound of $\Omega(\min(k, n/k))$ for the width and height of $D(G)$. \square

We can now prove lower bounds for mosaic drawings with L-shaped channels. In the following, let the x -coordinate of an L-shaped channel be the x -coordinate of its vertical leg.

Lemma 10. *Let $D(G)$ be a one-sided mosaic drawing of an outerzigzag G on n vertices with m L-shaped channels, width w , and height h . Then $dm + w + h = \Omega(n)$, where d is the difference between the x -coordinates of the rightmost and leftmost channels.*

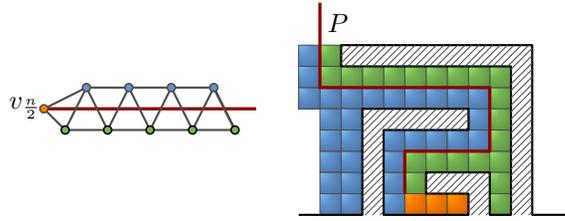


Fig. 12: Illustration for Lemma 10: the y -monotone rectilinear path P has length $\Omega(n)$.

Proof. Following the proof of Lemma 6, let v_1, \dots, v_n be the vertices of G in order along the outer face, and let $v_{n/2}$ be the tip of an ear of G . Starting from a tile of $C(v_{n/2})$ we draw a y -monotone rectilinear path P in $D(G)$ until it escapes $D(G)$ using as few horizontal segments as possible (see Fig. 12). Note that such a path exists and can have at most m horizontal segments. Furthermore, P must be crossed by all $\Omega(n)$ unions of the form $C(v_i) \cup C(v_{n-i+1})$, which implies that the length of P is $\Omega(n)$. The sum of the lengths of the vertical segments of P is at most h . The horizontal segments of P have length at most d , except maybe the highest segment, which can have length w . We directly obtain that $dm + w + h = \Omega(n)$. \square

Lemma 11. *Any one-sided mosaic drawing $D(G)$ with only L-shaped channels of a k -comb G with $k = n^{1/3}$ requires $\Omega(n^{2/3})$ width or height.*

Proof. Let w and h be the width and height of $D(G)$, respectively. For the i^{th} leg of G , let L_i be the set of channels adjacent to the configurations of leg i . Note that the set L_i must be consecutive along the boundary of $D(G)$. Furthermore, let $m_i = |L_i|$ and let d_i be the difference between the x -coordinates of the rightmost and leftmost channels in L_i . By Lemma 10 we have that $d_i m_i + w + h = \Omega(n^{2/3})$ for each leg. Furthermore, $d_i \geq m_i$ and $\sum_i d_i \leq w$, since $D(G)$ is one-sided. If $d_i m_i = \Omega(n^{2/3})$ for all i , then $d_i = \Omega(n^{1/3})$ for all i . As a result, $w \geq \sum_i d_i = n^{1/3} \Omega(n^{1/3}) = \Omega(n^{2/3})$. Otherwise we obtain that $w + h = \Omega(n^{2/3})$, which directly implies the claimed result. \square

Lemma 12. *There exists an outerplane graph G with n vertices such that any mosaic drawing $D(G)$ with only L-shaped channels requires $\Omega(n^{7/6})$ area.*

Proof. We construct G by connecting a (k, r) -signpost with $k = \sqrt{n}$ with a k' -comb with $k' = n^{1/3}$, each part having $n/2$ vertices. Lemma 9 implies that the width and height of $D(G)$ are both $\Omega(\sqrt{n})$. Furthermore, Lemma 11 implies that the width or height of $D(G)$ is $\Omega(n^{2/3})$. As a result, $D(G)$ has area $\Omega(n^{7/6})$. Although we have only argued this for a one-sided mosaic drawing now, this result also holds for a general mosaic drawing with L-shaped channels due to Lemma 5. \square

6 Conclusions and open problems

We investigated the trade-offs between two complexity measures for mosaic drawings of outerplanar graphs: channel complexity and area. Both measures clearly contribute to the quality of mosaic drawings. However, in the future we would like to also consider the complexity of the shapes representing vertices. Specifically, it would be interesting to determine under which conditions it is possible to represent vertices with a constant number of tiles or with configurations of constant boundary complexity.

Acknowledgments. Q.B. and B.S. are supported by the Netherlands Organisation for Scientific Research (NWO) under grant number 639.023.208. K.V. is supported by the Netherlands Organisation for Scientific Research (NWO) under grant number 639.021.541.

References

1. M. J. Alam, T. Biedl, S. Felsner, M. Kaufmann, S. G. Kobourov, and T. Ueckerdt. Computing cartograms with optimal complexity. *Discrete and Computational Geometry*, 50(3):784–810, 2013.
2. M. J. Alam, T. Bläsius, I. Rutter, T. Ueckerdt, and A. Wolff. Pixel and voxel representations of graphs. In *Proc. 23rd International Symposium on Graph Drawing*, LNCS 9411, pages 472–486, 2015.
3. R. G. Cano, K. Buchin, T. Castermans, A. Pieterse, W. Sonke, and B. Speckmann. Mosaic drawings and cartograms. *Computer Graphics Forum*, 34(3):361–370, 2015.
4. Y.-T. Chiang, C.-C. Lin, and H.-I. Lu. Orderly spanning trees with applications. *SIAM Journal on Computing*, 34(4):924–945, 2005.
5. M. de Berg, E. Mumford, and B. Speckmann. On rectilinear duals for vertex-weighted plane graphs. *Discrete Mathematics*, 309(7):1794–1812, 2009.
6. D. Dolev, T. Leighton, and H. Trickey. Planar embedding of planar graphs. In *Advances in Computing Research*, volume 2, pages 147–161, 1984.
7. J. J. Fowler. Strongly-connected outerplanar graphs with proper touching triangle representations. In *Proc. 21st International Symposium on Graph Drawing*, LNCS 8242, pages 155–160, 2013.
8. P. Koebe. Kontaktprobleme der konformen Abbildung. *Berichte über die Verhandlungen der Sächsischen Akademie der Wissenschaften zu Leipzig. Math.-Phys. Klasse*, 88:141–164, 1936.
9. K. Koźmiński and E. Kinnen. Rectangular duals of planar graphs. *Networks*, 5(2):145–157, 1985.
10. H. C. Purchase. Metrics for graph drawing aesthetics. *Journal of Visual Languages & Computing*, 13(5):501–516, 2002.
11. E. Raisz. The rectangular statistical cartogram. *Geographical Review*, 24(2):292–296, 1934.
12. P. Ungar. On diagrams representing maps. *Journal of London Mathematical Society*, s1–28(3):336–342, 1953.