Edge-weighted contact representations of planar graphs

Martin Nöllenburg [1]   *Roman Prutkin* [1]   *Ignaz Rutter* [1]

[1]Institute of Theoretical Informatics,
Karlsruhe Institute of Technology, Germany

Abstract

We study contact representations of edge-weighted planar graphs, where vertices are represented as interior-disjoint rectangles or rectilinear polygons and edges are represented as contacts of vertex boundaries whose contact lengths represent the edge weights.

For the case of rectangles, we show that, for any given edge-weighted planar graph whose outer face is a quadrangle, that is internally triangulated and that has no separating triangles, we can construct in linear time an *edge-proportional rectangular dual* (contact lengths are equal to the given edge weights and the union of all rectangles is again a rectangle) or report failure if none exists. In the case of arbitrarily many outer vertices, we show that deciding whether a square layout exists is NP-complete. If the orientation of each contact is specified by a so-called *regular edge labeling* and edge weights are lower bounds on the contact lengths, a corresponding rectangular dual that minimizes the area and perimeter of the enclosing rectangle can be found in linear time. On the other hand, without a given regular edge labeling, the same problem is NP-complete, as is the question whether a rectangular dual exists given lower *and* upper bounds on the contact lengths.

For the case of rectilinear polygons, we give a complete characterization of the polygon complexity required for representing connected internally triangulated graphs: For outerplanar graphs and graphs with a single inner vertex polygon, complexity 8 is sufficient and necessary, and for graphs with two adjacent or multiple non-adjacent internal vertices the required polygon complexity is unbounded.

Submitted: December 2012	Reviewed: April 2013	Revised: April 2013	Accepted: June 2013	Final: July 2013
		Published: July 2013		
	Article type: Regular paper	Communicated by: W. Didimo and M. Patrignani		

1   Introduction

Representing graphs by intersections or contacts of geometric objects has a long history in graph theory and graph drawing, which is covered in monographs and surveys [13,21]. For example, Koebe's circle packing theorem from 1936 establishes that every planar graph has a contact representation by touching disks (and vice versa) [16]; more recently it was shown that every planar graph is the intersection graph of line segments [6].

In this paper we are interested in a special class of contact representations for plane graphs, namely hole-free side-contact representations using rectangles and rectilinear polygons. In a *rectilinear representation* of a plane graph $G = (V, E)$ every vertex $v \in V$ is represented as a simple rectilinear polygon $P(v)$ and there is an edge $e = uv \in E$ if and only if $P(u)$ and $P(v)$ have a non-trivial common *boundary* or *contact path* $s(e)$ (i.e., length $|s(e)| > 0$). It is further required that the union $\bigcup_{v \in V} P(v)$ forms a simple rectilinear polygon itself, i.e., there are no holes in the representation. A standard assumption, which we will make throughout this paper, is that G is an internally triangulated plane graph. This excludes the degenerated case of four polygons that meet in a single point. A *rectangular dual* [17] of a graph G is a dissection of a rectangle into rectangles, which represents G as a contact graph; rectangular duals are thus an interesting special case of rectilinear representations, where all polygons and their union are rectangles. Rectangular duals and rectilinear representations with low-complexity polygons have practical applications, e.g., in VLSI design, cartography, or floor planning and surveillance in buildings [22]. In these applications, the area of vertex polygons and/or the boundary length of adjacent polygons often play an important role, e.g., in building surveillance polygon area is linked to the number of persons in a room and boundary length represents the number of transitions from one room to the next. This and similar examples immediately raise the question of representing a given weighted graph so that the weights control areas and lengths in a contact representation.

Previously, rectilinear representations and rectangular duals have been studied only for unweighted graphs [17, 19] and vertex-weighted graphs [2, 3, 10], where the polygon areas must be proportional to the vertex weights. This paper covers the remaining open aspect of representing edge-weighted graphs as touching rectilinear polygons. A natural way of encoding edge weights in a rectilinear representation is to require that the contact lengths of all adjacent vertex polygons are proportional to the given edge weights. So we define an *edge-proportional rectilinear representation* (EPRR) of an edge-weighted graph $(G, \omega : E \to \mathbb{R}^+)$ as a rectilinear representation in which additionally the contact length $|s(e)| = \omega(e)$ for every edge $e \in E$.

Related work.   It is known that unweighted graphs always have a rectilinear representation using rectangles, L-shaped and T-shaped polygons, i.e., at most 8-gons, and that there are some graphs for which complexity 8 is necessary [19, 26]. The class of unweighted graphs that have a rectangular dual is characterized as all plane triangulations without separating triangles [17, 18]. Orientation-constrained rectangular duals have also been considered [10].

For vertex-weighted graphs the goal is to find area-proportional rectilinear representations, in which the area of a polygon $P(v)$ is proportional to the weight of vertex v. In

a series of papers the polygon complexity that is sufficient to realize any weighted graph was decreased from 40 corners [7] over 34 corners [14], 12 corners [4], 10 corners [1] down to 8 corners [2], which is best possible due to the earlier lower bound of 8 [26]. Weighted rectangular duals have also been studied before, e.g., van Kreveld and Speckmann [25] presented several algorithms to compute rectangular duals with low area error. Eppstein et al. [10] gave a necessary and sufficient condition for rectangular duals to be area-universal, i.e., rectangular duals that can realize arbitrary vertex weights without changing their combinatorial structure. They also showed that, for a given combinatorial structure of the dual and given vertex weights, it can be efficiently tested whether these weights can be represented as the perimeters of the vertex rectangles rather than their areas. Biedl and Genc [3] showed that testing whether a rectangular representation with prescribed areas exists is NP-hard if the complexity of the outer face is unbounded.

Drawing planar graphs with edge weights as standard node-link diagram, where edge lengths are proportional to the edge weights, is an NP-hard problem [9] but can be decided in linear time for planar 3-connected triangulations [5].

Contribution. In Section 2 we consider rectangular duals. We present a linear-time algorithm that decides whether a given graph G has an *edge-proportional rectangular dual* (EPRD) with four outer rectangles and constructs it in the positive case. However, if G has arbitrarily many outer vertices, it can have exponentially many EPRDs, and it is NP-complete to decide whether there exists one such dual that forms a square. Moreover, if the combinatorial structure, i.e., the orientation of each edge of the dual, is specified, we use existing tools to find a rectangular dual where $|s(e)| \geq \omega(e)$ for all $e \in E$ and the size, i.e., the area or perimeter, of the outer rectangle is minimum. On the other hand, without a fixed combinatorial structure, we prove NP-completeness of the problem to find a representation where the lengths of the contact segments are lower and upper bounded. We also show that finding optimal duals for given lower bounds on $|s(e)|$ under various criteria over all combinatorial structures is NP-complete.

In Section 3, we consider edge-proportional rectilinear representations and show that for representing outerplanar graphs polygon complexity 8 is sometimes necessary and always sufficient. The class of outerpillars (outerplanar graphs whose weak dual is a caterpillar, i.e., a tree for which a path remains after removing all leaves) always has edge-proportional rectilinear representations of complexity 6, but already outerlobsters (outerplanar graphs whose weak dual is a lobster, i.e., a tree for which a caterpillar remains after removing all leaves) require complexity 8. If, on the other hand, the graph has two adjacent or multiple non-adjacent internal vertices, polygons of unbounded complexity are sometimes necessary. This completely characterizes the complexity of edge-proportional rectilinear representations for internally triangulated graphs.

2   Rectangular duals with contact length specifications

In this section we consider rectangular duals of edge-weighted planar graphs. In Section 2.1 we study duals with fixed contact lengths (EPRDs). In Section 2.2 we investigate the problem of finding optimal duals, i.e., duals with minimum width, height and area, for lower-bounded contact lengths and fixed combinatorial structure. In

Section 2.3 we show that for specified lower and upper bounds on contact lengths and variable combinatorial structure it is NP-complete to decide the existence of a suitable rectangular dual. Finally, in Section 2.4, we show that finding optimal duals with lower-bounded contact lengths over all possible combinatorial structures is NP-complete.

2.1   Rectangular duals with fixed contact lengths

In the first part of this section, we present a linear-time decision and construction algorithm for edge-proportional rectangular duals (EPRDs) with four outer rectangles. In the second part, we show that if arbitrarily many outer rectangles are allowed, the number of EPRDs might be exponential, and finding an EPRD that forms a square is NP-complete.

2.1.1   Four outer rectangles

He [12] proved that a planar graph G has a rectangular dual with four rectangles on the boundary if and only if (1) every interior face of G is a triangle and the outer face is a quadrangle, and (2) G has no separating triangles. We call a graph satisfying these conditions *properly triangular planar (PTP)*. Moreover, we denote the four vertices on the boundary of the outer face by v_N, v_W, v_S and v_E in counterclockwise order.

   A rectangular dual R of a PTP graph $G = (V, E)$ defines an orientation and a partition of the internal edges of G into two sets T_1 and T_2. The set T_1 contains the edges e for which $s(e)$ is horizonal, the remaining edges are in T_2. The orientation is obtained by orienting $uv \in T_1$ from u to v if $R(u)$ is below $R(v)$, similarly $uv \in T_2$ is oriented from u to v if $R(u)$ is to the left of $R(v)$. For a vertex v and one of the partitions T_i, $i = 1, 2$, we denote by $T_i^{\leftarrow}(v)$ and $T_i^{\rightarrow}(v)$ the incoming and outgoing edges of v that are contained in T_i, respectively. Note that for a specified botommost rectangle the orientation can be uniquely derived from the partition [11]. The orientation and partition then satisfies the following properties.

1. For each vertex v, a counterclockwise enumeration of its incident edges starting with the rightmost edge in $T_1^{\rightarrow}(v)$ encounters first the edges in $T_1^{\rightarrow}(v)$, then in $T_2^{\leftarrow}(v)$, then in $T_1^{\leftarrow}(v)$ and finally in $T_2^{\rightarrow}(v)$, and

2. all interior edges incident to v_N, v_W, v_S and v_E are in $T_1^{\leftarrow}(v_N)$, $T_2^{\rightarrow}(v_W)$, $T_1^{\rightarrow}(v_S)$ and $T_2^{\leftarrow}(v_E)$, respectively.

   We call any partition and orientation of the edges satisfying these properties a *regular edge labeling (REL)*. In his work, He [12] showed that every PTP graph admits a REL, and that a corresponding rectangular dual can be constructed from a REL in linear time.

   It is not hard to see that a REL derived from an edge-proportional rectangular dual has additional properties, following from the fact that for each rectangle the total length of the contacts on the left and right side as well as on the top and bottom side are the same, respectively.

$$\sum_{e \in T_1^{\leftarrow}(v)} \omega(e) = \sum_{e \in T_1^{\rightarrow}(v)} \omega(e), \quad \sum_{e \in T_2^{\leftarrow}(v)} \omega(e) = \sum_{e \in T_2^{\rightarrow}(v)} \omega(e). \tag{1}$$

Figure 1: Prior to the insertion of the next inner rectangle there is always a U-shape for which there exists a vertex whose corresponding rectangle needs to be inserted at the lower left corner of the U-shape in a unique way.

We call any REL satisfying this condition an *edge proportional REL* (EPREL). In the following we show that a weighted PTP graph $G = (V, E)$ has a unique EPREL, if one exists. Moreover, we show how to test the existence of such an EPREL in linear time and how to construct a corresponding edge-proportional rectangular dual.

Lemma 1 *For an inner vertex v, any one of the sets $T_1^{\leftarrow}(v), T_1^{\rightarrow}(v), T_2^{\leftarrow}(v)$ or $T_2^{\rightarrow}(v)$ of an edge-proportional REL completely fixes the orientation and the partition of the edges incident to v. A corresponding orientation and partition can be found in $O(\deg(v))$ time if it exists.*

Proof: Assume $T_1^{\leftarrow}(v)$ is known, the other cases are symmetric. Let $\omega_1 = \sum_{e \in T_1^{\leftarrow}(v)} \omega(e)$ and let further $\omega_2 = (\sum_{uv \in E} \omega(uv) - 2\omega_1)/2$. It follows from condition (1) that necessarily $\sum_{e \in T_2^{\leftarrow}(v)} = \sum_{e \in T_2^{\rightarrow}(v)} = \omega_2$. Due to the requirement of the REL for the ordering of the edges around v, there is at most one way to orient and partition the edges incident to v such that condition (1) holds. It can be found in $O(\deg(v))$ time by a simple counterclockwise traversal of the edges incident to v, starting from the last edge in the known set $T_1^{\leftarrow}(v)$. $\square$

Observe that if the partition and orientation of the edges incident to a vertex v is determined, the shape of the rectangle representing v is completely fixed. Moreover, the conditions on the edges incident to v_N, v_W, v_S and v_E completely specify a rectangle R_I into which the remaining rectangles have to be inserted. We construct an ordering of the internal vertices $v_1, \ldots, v_{n-4}$ such that we can iteratively apply Lemma 1 to determine uniquely the shape of their rectangle as well as the position where they have to be inserted in R_I. Since we are completely guided by necessary conditions, this either results in a correct edge-proportional rectangular dual, or the procedure fails at some point, in which case an edge-proportional rectangular dual does not exist.

We maintain the following invariants in each step i.

1. The position and dimension of $R(v_1), \ldots, R(v_i)$ are uniquely determined.

2. All contacts between already inserted rectangles or the boundary polygon R_I have correct lengths.

3. The upper boundary of the polygon $\bigcup_{j=1}^{i} R(v_j) \cup R(v_S) \cup R(v_W) \cup R(v_E)$ is an x-monotone chain.

Note that initially $i = 0$ and all properties hold. By the third property there exists a *U-shape* on the upper boundary whose bottom side is horizontal, i.e., there are two vertical

segments adjacent to and above the bottom side. Let u be the lowest rectangle bounding this U-shape from the left and let $u_1, \ldots, u_k$ denote the rectangles bounding the U-shape from below; see Fig. 1. The corner at $R(u)$ and $R(u_1)$ implies that if G admits an edge-proportional rectangular dual, then there exists a unique vertex v that is not yet inserted, and that is incident to both u and u_1. We choose this vertex as the next vertex v_i. Its adjacencies to the vertices $u_1, \ldots, u_j$ for some $j \leq k$ completely determine its contacts from below, and hence $T_1^{\leftarrow}(v)$. By Lemma 1 its shape is completely determined. Moreover, the position is fixed as well due to the corner between $R(u)$ and $R(u_1)$. This implies Invariant 1. Invariant 2 is either satisfied or an edge-proportional rectangular dual does not exist since we only followed necessary conditions. Finally, Invariant 3 holds due to the choice of the U-shape. The whole algorithm can be implemented to run in linear time.

Theorem 1 *For an edge-weighted PTP graph G there exists at most one edge-proportional rectangular dual. It can be computed in linear time if it exists.*

Proof: The correctness is already shown, it remains to deal with the running time. The main issues are to a) quickly find a suitable U-shape b) find a corresponding vertex v that can be inserted and c) check whether the insertion produces only correct adjacencies.

For b) and c), observe that given the adjacent vertices u and u_1 of a U-shape as in Fig. 1, the vertex v must belong to the unique triangular face of G containing the edge uu_1 and a not-yet inserted vertex. Therefore, given u and u_1, v can be found in $O(1)$ time. Determining the shape of $R(v)$ with Lemma 1 takes $O(\deg(v))$ time. If the test in Lemma 1 fails, the algorithm terminates and reports that no edge-proportional rectangular dual of G exists. One can test in time proportional to the number of contacts of $R(v)$ to previously inserted rectangles, whether they all correspond to edges of G. This takes $O(\deg(v))$ time if the test is successful and at most $O(|V|)$ time if the test fails, but then the algorithm stops.

For a), we store the rectangles that have been inserted, but are not yet covered, in a doubly linked list, sorted from left to right. Each concave bend ($\llcorner$ or $\lrcorner$) in the upper contour of the inserted rectangles is a candidate for a left or right boundary of a U-shape as shown in Fig. 1. These candidates are stored as tuples $\llcorner(u,v)$ or $\lrcorner(u,v)$ for a pair of rectangles $R(u)$, $R(v)$ with a collinear vertical boundary, such that $R(u)$ is to the left of $R(v)$. Each candidate stores the pointers to both of its rectangles.

We store the candidates in a doubly linked list L, sorted by the x-coordinates of the common vertical boundaries of the associated rectangles. Each U-shape appearing during the insertion is characterized by a consecutive $\llcorner$ - $\lrcorner$ pair of candidates. After the insertion of a new rectangle, at most two new candidates arise, at most two disappear, and L can be updated in $O(1)$ time. The arising $\llcorner$ - $\lrcorner$ pairs in L that form a U-shape can be found in $O(1)$ time and are pushed onto a stack S, which is used to pick the next U-shape. Since this stack does not become empty before the last step, a U-shape suitable for insertion can be found in $O(1)$ time in each step. It follows that inserting a new rectangle $R(v)$ takes time $O(\deg(v))$, and hence the total running time is $O(|E|) = O(|V|)$. □
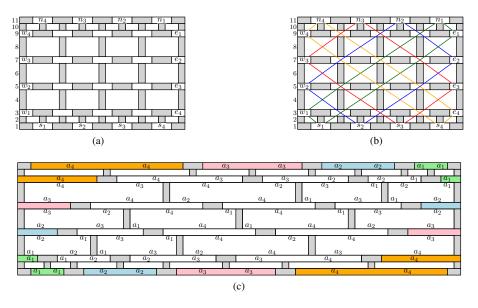
Figure 2: (a) graph G_4 defined by a rectangular dual. (b) Cycles E_1 (green), E_2 (blue), E_3 (red), E_4 (orange). (c) An EPRD of G_4 for the PARTITION instance $a_1 = 2$, $a_2 = 5$, $a_3 = 7$, $a_4 = 11$. Contacts are labeled by the assigned lengths.

2.1.2 Many outer rectangles

In the previous section, each of the four outer rectangles formed an entire side of the rectangle R_I into which all inner rectangles had to be inserted. Due to this fact, the orientation of each segment forming the boundary of R_I was fixed. For example, the left boundary of R_I was formed by the segments corresponding to the inner edges incident to v_W etc. This is no longer the case if we allow more then four outer rectangles. The left boundary of R_I can now be formed by segments which belong to several outer rectangles. Since the lengths of the edges on the boundary of the dual are unspecified, neither are the shapes of the outer rectangles. Thus the orientations of the boundary segments of R_I are no longer unique, which introduces a degree of freedom that makes the problem hard. In fact, now there can exist an exponential number of possible realizations, and optimizing over all of them is NP-complete, even if we fix the four corners of the dual:

Theorem 2 *Given an inner-triangulated edge-weighted plane graph $G = (V, E, \omega)$ without separating triangles and four or more outer vertices including v_{ll}, v_{lr}, v_{ur} and v_{ul}, it is NP-complete to decide whether there exists an edge-proportional rectangular dual $\mathcal{R}$ such that $\mathcal{R}$ forms a square whose lower left, lower right, upper left and upper right corner rectangles are $R(v_{ll})$, $R(v_{lr})$, $R(v_{ul})$ and $R(v_{ur})$, respectively.*

The proof is a reduction from the NP-complete problem PARTITION: given positive integers a_1, ..., a_m, decide whether there exists a subset $P \subseteq \{1, \ldots, m\}$, such that $\sum_{i \in P} a_i = \frac{1}{2} \sum_{i=1}^{m} a_i =: \sigma$.
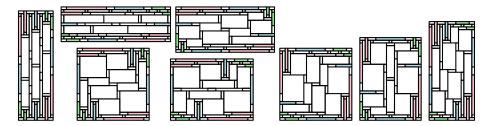
Figure 3: Layouts of G_3 corresponding to all possible partitions of $\{2,5,7\}$. The two square layouts correspond to the solutions $P = \{7\}$ and $P = \{2,5\}$.

In our reduction proof, for a given PARTITION instance, we shall define an inner-triangulated edge-weighted plane graph $G_m = (V, E, \omega)$ and prove that it has a square rectangular dual with specified contact lengths if and only if the PARTITION instance has a solution. We first show how to build G_m for an arbitrary m. For now, we ignore edge weights. It is easier to describe G_m by a corresponding rectangular dual $\mathcal{R}_m$; see Figure 2a. It is formed by $2m+3$ rows of rectangles. Every odd row starts and ends with a 2×1 rectangle and has $2m-1$ rectangles of dimension 3×1 in between. Rows 2 and $2m+2$ contain $4m+1$ rectangles with dimensions 1×1 and 2×1 placed in alternating order from left to right starting and ending with a 1×1 square. The remaining even rows contain $2m+1$ rectangles with dimensions 1×3 and 5×3 placed in alternating order from left to right starting and ending with a 1×3 rectangle. Obviously, G_m is inner triangulated and has no separating triangles. For ease of notation, we shall refer to vertex $v \in V$ by integer coordinates (x, y) if the corresponding rectangle $R(v)$ is the xth rectangle in row y in $\mathcal{R}_m$. The vertices $(1,1)$, $(2m+1, 1)$, $(2m+1, 2m+3)$, $(1, 2m+3)$ are selected as respective corners.

We define some special outer vertices: for $j = 1, \ldots, m$, let s_j denote the vertex $(2j, 1)$, n_j the vertex $(2(m-j)+2, 2m+3)$, w_j the vertex $(1, 2j+1)$ and e_j the vertex $(2m+1, 2(m-j)+3)$; see Figure 2a. We also define a subset $F \subseteq V$ (gray rectangles in Figure 2a): in each even row and also in the first and last row of $\mathcal{R}_m$, each rectangle in an odd position corresponds to a vertex in F and each rectangle in an even position to a vertex in $V \setminus F$. In all the remaining rows, each rectangle in an even position corresponds to a vertex in F and each rectangle in an odd position to a vertex in $V \setminus F$.

We now describe the edge weight assignment ω. For each vertex s_j, e_j, n_j and w_j, $j = 1, \ldots, m$, we set the weights of the five incident edges to 1, a_j, 1, a_j, 1 (in circular order, such that both outer edges have weight 1). From the specified contact lengths it follows that each such outer vertex has exactly two possible realizations in an EPRD: rectangles $R(s_j)$ and $R(n_j)$ are either $1 \times (1+a_j)$ or $(2a_j+1) \times 1$ and rectangles $R(w_j)$ and $R(e_j)$ are either $1 \times (2a_j+1)$ or $(1+a_j) \times 1$. We call the second and fourth edge of such a vertex *flippable*, and the corresponding contacts can be either both horizontal or vertical. We shall use these choices to encode to which partition the element a_j is assigned to. Figure 3 shows layouts of G_3 corresponding to each possible partition of $\{2,5,7\}$.

For each $j = 1, \ldots, m$, we define a set of edges $E_j \in E$ that forms a cycle in G_m. To define E_j, we start at w_j and follow its lower flippable edge. Then, at each inner vertex we follow the edge opposite to the previous one taken (note that all inner vertices of G_m have even degree). As we proceed, the first coordinate of visited vertices increases by 1 and the second decreases by 1 until we reach the third row at the vertex $(2j-1, 3)$. We proceed to the vertex $(4j-2, 2)$ and then reach s_j via its left flippable edge. We continue the cycle E_j in a similar manner: we leave s_j via its right flippable edge and follow opposite edges (after we pass the third row, both coordinates increase by 1) until we reach e_j via its lower flippable edge. We proceed analogously via n_j and finally return to w_j via its upper flippable edge. Figure 2b shows all such cycles for $m = 4$. Using some simple calculations on the coordinates of the vertices lying on the cycles E_j, we can prove the following facts: (1) each E_j is a simple cycle; (2) each edge uv with $u, v \in V \setminus F$ belongs to exactly one E_j, and each E_j contains only such edges between vertices in $V \setminus F$; (3) each vertex of degree 8 in G_m lies on exactly two cycles E_i and E_j; all remaining vertices in $V \setminus F$ lie on exactly one cycle.

We can now present the complete edge weight assignment ω. For an edge $e = uv \in E$ with u or v in F we set its weight $\omega(e)$ to either 1 or 3 according to the length of the corresponding contact in the dual $\mathcal{R}_m$ shown in Figure 2a. Each remaining edge $e \in E$ lies on exactly one cycle E_j. We set $\omega(e) = a_j$. This assignment has the following effect on possible realizations of inner rectangles.

Let $v \in V$ be an inner vertex of G_m. By construction, v has even degree, and each pair of opposing edges incident to v has the same weight. Then, every two contacts corresponding to such a pair of opposing edges must lie on opposite sides of the rectangle $R(v)$. Each inner vertex $v \in V \setminus F$ on a cycle E_j has one pair of opposing incident edges both lying on E_j. Thus, the corresponding contacts of length a_j must lie on opposite sides of the rectangle $R(v)$. This also holds for flippable edges incident to s_j, e_j, n_j and w_j. It follows that all contacts corresponding to edges on E_j must have the same orientation.

We claim that all rectangles corresponding to vertices in F (gray in Figure 2) have fixed dimensions and orientations. It is easy to see that rectangles $(1,4)$, $(1,6)$, $\ldots$, $(1,2m)$ on the left boundary as well as rectangles $(2m+1, 4)$, $(2m+1, 6)$, $\ldots$, $(2m+1, 2m)$ on the right boundary must be 1×3 rectangles. Both outer edges adjacent to each such vertex must correspond to horizontal contacts of length 1, and the remaining contact of length 3 must be vertical. Also, rectangles $(2,5)$, $(2,7)$, $\ldots$, $(2, 2m-1)$ as well as $(2m, 5)$, $(2m, 7)$, $\ldots$, $(2m, 2m-1)$ are 3×1 rectangles. By applying the above observation to all gray inner 1×3 or 3×1 rectangles of degree 4 and their neighbors of degree 8 in Figure 2 iteratively from left to right, we see that the shapes and orientations of all rectangles $R(v)$, $v \in F$, $\deg(v) = 4$, must be fixed. For example, we know that the contact between $(1,5)$ and $(2,5)$ must be vertical. By the above observation, so must be the contacts between $(2,5)$ and $(3,5)$, $(3,5)$ and $(4,5)$, etc.

Furthermore, the contact $(1,3)$-$(2,3)$ must be vertical. Thus, $(2,3)$ must be a 3×1 rectangle. This observation allows us to fix all remaining gray rectangles: since the contact $(2,3) - (2,2)$ is horizontal, so is $(1,1) - (2,2)$. Thus, $(1,1)$ is a 2×1 rectangle. By similar arguments, $(4,3)$, $(6,3)$, $\ldots$, $(2m,3)$, $(3,1)$, $(5,1)$, $\ldots$, $(2m-1,1)$, $(2, 2m+1)$, $(4, 2m+1)$, $\ldots$, $(2m, 2m+1)$, $(3, 2m+3)$, $(5, 2m+3)$, $\ldots$, $(2m-1, 2m+3)$ must be 3×1 rectangles, and the corner rectangles are 2×1. We can now prove Theorem 2.

Proof of Theorem 2: Let $a_1, \ldots, a_m$ be an instance of PARTITION. The graph G_m has $O(m^2)$ vertices and can be constructed in time polynomial in m as described above. Assume there exists an edge-proportional rectangular dual $\mathcal{R}$ of G_m that forms a square. We define

$P = \{j \mid 1 \le j \le m,$ contacts in $\mathcal{R}$ corresponding to edges on E_j are horizontal$\}.$

The width of $\mathcal{R}$ is $4 + 3(m-1) + m + 2\sum_{j \in P} a_j$ and the height is $4 + 3(m-1) + m + 2\sum_{j \in \{1,\ldots,m\} \setminus P} a_j$. Since $\mathcal{R}$ is a square, P solves the PARTITION instance.

Now let P be a solution of the PARTITION instance. We extend G_m to an edge-weighted PTP graph G'_m by adding four outer vertices v_W, v_S, v_E and v_N. We connect all vertices whose rectangle in $\mathcal{R}_m$ lies on the left boundary to v_W, those on the lower boundary to v_S etc. For such a new edge e incident to a vertex $v \in F$ we set $\omega(e)$ to 1, 2 or 3 according to the length of the corresponding boundary segment of the gray rectangle $R(v)$ in Figure 2. For a new edge e incident to a vertex $v \notin F$ we set $\omega(e)$ as follows: if v has been on the lower or upper boundary, i.e., $v \in \{n_j, s_j\}$ for some $j = 1, \ldots, m$, we set $\omega(e) = 2a_j + 1$ if $j \in P$ and $\omega(e) = 1$ if $j \notin P$. If v has been on the left or right boundary, i.e. $v \in \{w_j, e_j\}$ for some $j = 1, \ldots, m$, we set $\omega(e) = 2a_j + 1$ if $j \notin P$, and $\omega(e) = 1$ if $j \in P$. We now color the edges of the PTP graph G'_m: if an edge e is incident to a vertex $v \in F$, we color e red if the corresponding segment $s(e)$ lies on the upper or lower boundary of the gray rectangle $R(v)$ and blue otherwise. For an edge $e \in E_j$, we color e red if $j \in P$ and blue otherwise. All edges incident to v_S and v_N are colored red, and all edges incident to v_W and v_E blue. By considering all types of vertices in G'_m, it can be easily verified that this coloring partitions incident edges of each inner vertex into four non-empty contiguous subsets of alternating colors. This induces an *undirected* REL of G'_m. Fusy [11] showed that there is a bijection between directed and undirected RELs that preserves the coloring. Also, for each vertex, the two red subsets of edges have the same total weight, and so do the two blue subsets. Thus, there exists an edge-proportional REL of G'_m. As will be shown in the next subsection (see Corollary 2), it follows that G'_m has an edge-proportional rectangular dual $\mathcal{R}$, such that red edges of G_m correspond to horizontal contacts in $\mathcal{R}$ and blue to vertical. Since P solves the PARTITION instance, the inner rectangles in $\mathcal{R}$ must form a square.

This decision problem is contained in $\mathcal{NP}$: we can guess an orientation for each contact corresponding to an edge $e \in E$ of the given graph $G = (V, E)$. Then we can check whether it induces an edge-proportional rectangular dual of G that forms a square with given corners in linear time, e.g. using an insertion algorithm similar to the one in the previous subsection. $\square$

Theorem 2 implies that minimizing the aspect ratio and maximizing the area of an edge-proportional rectangular dual are difficult.

Corollary 1 *Maximizing the area and minimizing the aspect ratio of a rectangular dual with given contact lengths and unspecified contact orientations is NP-hard.*

2.2 Rectangular duals with minimum contact lengths

Next we consider a slightly relaxed version of the problem, where we assume that the input consists of a REL, which combinatorially describes the rectangular dual, and a

weight function specifying minimum contact lengths for all edges. The task is then to find a rectangular dual according to the given REL that minimizes the total size of the layout. Note that in this setting any instance is feasible since any given rectangular dual can be scaled to become a feasible solution.

Using the method of He [12] we can construct in linear time a rectangular dual R of the PTP graph G that realizes the given REL, but does not yet satisfy the edge-length constraints. We can either modify He's algorithm to directly compute a suitable rectangular dual in linear time, or take a slightly different perspective on the problem. The rectangular dual R of G can also be seen as an orthogonal representation with rectangular faces of the dual graph $G^\star$ of G, where every degree-3 vertex corresponds to a face of G and every orthogonally drawn edge corresponds to two adjacent faces. This allows us to use a modified version of a linear-time compaction algorithm for orthogonal drawings [8, Chapter 5.4] that respects the minimum contact length $\omega(e)$ for each $e \in E$ as the minimum length of the corresponding dual edge $e^\star$. The main idea of the approach is to define two independent planar edge-weighted st-graphs N_{hor} and N_{ver}, the first one using the edges in T_1, the other one the edges in T_2. Tamassia [24] described an algorithm to compute two weighted topological numberings on N_{hor} and N_{ver} from which the coordinates of all vertices of R (or $G^\star$) can be obtained. These numberings immediately minimize the total height, total width and area of R subject to the length constraints.

Theorem 3 *For a weighted PTP graph* (G, ω) *with a given REL, a corresponding rectangular dual with minimum width, height, and area of the inner rectangles can be computed in linear time such that each edge e is represented by a contact of length at least* $\omega(e)$.

In particular, if the given REL is edge-proportional, the algorithm computes an edge-proportional rectangular dual. Conversely, an edge-proportional rectangular dual directly induces an edge-proportional REL.

Corollary 2 *A weighted PTP graph admits an edge-proportional REL if and only if it admits an edge-proportional rectangular dual.*

2.3 Rectangular duals with minimum and maximum contact lengths and variable REL

Unlike in the case of precisely specified contact lengths (and no REL specification) or lower-bounded contact lengths with fixed REL covered in the previous sections, it becomes NP-hard to decide the existence of a rectangular dual if no REL is specified and we are given lower and upper bounds for the contact lengths that must be respected.

Theorem 4 *Given a PTP graph* $G = (V, E)$ *with two edge-weight functions* $\alpha, \beta : E \to \mathbb{R}^+$ *with* $\alpha(e) \leq \beta(e)$ *for all* $e \in E$, *it is NP-complete to decide if G has a rectangular dual* $\mathcal{R} = \{R(v) \mid v \in V\}$ *so that for every edge* $e \in E$ *the contact segment s(e) has length* $\alpha(e) \leq |s(e)| \leq \beta(e)$.

Note that if the REL is fixed, the same problem can be solved in polynomial time, e.g. by linear programming or by adding upper edge capacities to the MinCost Flow approach in [8, Chapter 5.4].

The proof is a gadget proof reducing from the NP-complete problem PLANAR 3SAT [20]. PLANAR 3SAT is the satisfiability problem for Boolean formulae ϕ in conjunctive normal form with at most three variables per clause, which are *planar* in the sense that the induced bipartite *variable-clause graph* H_ϕ consisting of a vertex for every variable, a vertex for every clause, and an edge for every occurrence of a variable in a clause, is planar. Such a graph H_ϕ can be drawn on a grid of polynomial size with all variable vertices placed on a horizontal line and the clause vertices connected in a comb-shaped manner from above or below that line [15]. In our reduction, we create an edge-weighted PTP graph G_ϕ for a PLANAR 3SAT formula ϕ, which has a rectangular dual (mimicking the above mentioned drawing of H_ϕ) if and only if ϕ is satisfiable. In the next three subsections we describe how to construct G_ϕ in detail: for each component type of the variable-clause graph H_ϕ like variables, pipes and clauses we present corresponding edge-weighted subgraphs of G_ϕ and their realizations as rectangular layouts with suitable contact lengths. We show that all realizations are unique up to truth values of the associated variables of ϕ. Finally, we use this fact to show that rectangular layouts of G_ϕ with suitable contact lengths correspond to satisfying assignments of ϕ, and vice versa.

2.3.1   Variables and pipes

The basic building block for the variable gadgets and their links to the clause gadgets is a 5-vertex graph (type-2 gadget) flanked by three auxiliary isomorphic 7-vertex graphs (type-1 gadgets), see Figure 4. The important property of this subgraph is that it has only two valid realizations as a rectangle contact graph, one of which encodes the value *true*, the other one the value *false*, and both of which have the same outer shape. We shall show that any other attempt to realize this subgraph violates either the edge length constraints or requires non-rectangular vertex regions. We now describe both gadgets in detail. We say that an edge e has weight x if $\alpha(e) = \beta(e) = x$ and weight $x : y$ if $\alpha(e) = x$ and $\beta(e) = y$.

A type-1 gadget is a 6-cycle with one inner vertex, see Figure 5. Let a, b, c, d, e and f be the outer vertices ordered counterclockwise around the inner vertex g. Each outer vertex has degree 5 in G_ϕ. Let the edge weights α and β be set as in Figure 5. This edge weight assignment has the following effect on the dimensions of the rectangles: each of the rectangles $R(a)$, $R(b)$, $R(c)$, $R(d)$, $R(e)$ and $R(f)$ must have dimensions 1×2 or 2×1, and $R(g)$ must be a 1×3, 2×2 or 3×1 rectangle. This leads us to the only three possible realizations of a type-1 subgraph shown in Figure 5b, 5c and 5d (disregarding rotations). The same holds if we merge the two dashed edges with both weights 1 adjacent to the vertex b or e, respectively, to single edges with weight 2.

A type-2 gadget is a 4-cycle with one inner vertex, see Figure 6. Let h, i, j and k be the outer vertices ordered counterclockwise around the inner vertex ℓ. The rectangle $R(\ell)$ must be a 1×1 square. According to the edge weight assignment, the rectangles $R(h)$, $R(i)$, $R(j)$ and $R(k)$ must have perimeters between 13 and 15. If the contacts corresponding to the four edges hi, ij, jk and kh have alternating orientations
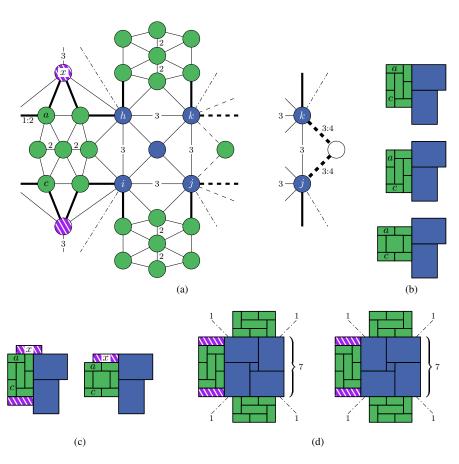
Figure 4: The standard building block formed by a type-2 gadget and three type-1 gadgets. (a) Edge-weighted contact graph. Thick edges have weights $\alpha(e) = 1$, $\beta(e) = 2$. (b) We make sure only the middle case is possible. Due to the purple dashed rectangles and the degrees of *a* and *c*, the layouts in (c) are not possible. The only two possible layouts are shown in (d). If the total length of the contacts corresponding to the dashed edges is at most 7, the contacts of length 1 corresponding to the dash-dotted edges force the lower and upper type-1 gadget to form 5×3 rectangles horizontally centered at the type-2 gadget. The left one encodes the value *true* and the right one *false*.
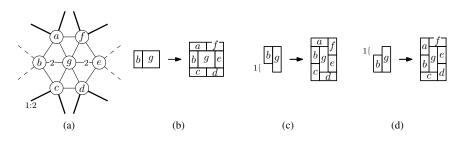
Figure 5: The type-1 gadget. The edge weights in (a) are $\alpha(e) = \beta(e) = d$ for an edge e labeled with a single value d (with $d = 1$ for unlabeled edges). For an edge e labeled with $x : y$ it is $\alpha(e) = x$ and $\beta(e) = y$. For the thick edges, it is $\alpha(e) = 1$ and $\beta(e) = 2$. The subgraph in (a) has exactly three possible realizations with respect to the specified edge weights: (b), (c) and (d).
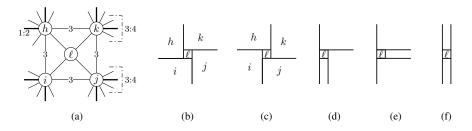

Figure 6: The type-2 gadget and its realizations.

like in Figure 6b or 6c, each of the four outer rectangles must have dimensions 3×4 or 4×3. Otherwise (see Figure 6d, 6e and 6f), one of these rectangles must have height or width at least 7. But then it has perimeter at least 16. Thus, the two realizations in Figure 6b and 6c are the only possible ones, and the complete type-2 gadget has dimensions 7×7.

To form a standard building block for our reduction proof, we shall connect a type-2 gadget to at least three type-1 gadgets, see Figure 4. Using additional vertices (like the purple vertex x in the figure) we can force the type-1 gadgets to be centered at the type-2 gadget they are connected to. In other words, only the middle case in Figure 4b should be possible. Due to the edge weights, the rectangle $R(x)$ must have dimensions 3×1. The vertices a, c, d, f in Figure 4a have degree 5 in G_ϕ. Therefore, the upper and lower boundaries of the corresponding green rectangles $R(a)$, $R(c)$, $R(d)$, $R(f)$ must be completely covered by the purple 3×1 rectangles, so the layouts in Figure 4c are not possible. Thus, the left type-1 gadget forms a 3×5 rectangle centered vertically at the type-2 gadget.

We shall use two options for the dashed edges on the right of the subgraph in Figure 4b. In the first option, there are six dashed edges adjacent to k and j with weights 1, 1:2, 1, 1, 1:2, 1 (from top to bottom), see Figure 4b, left. In the second
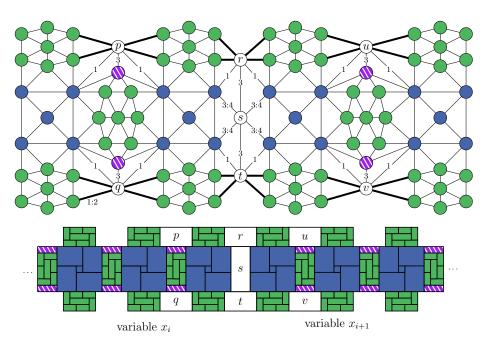
Figure 7: Two variable gadgets connected by buffer rectangles $R(r)$, $R(s)$, $R(t)$. The relative positions of the gadgets are fixed and form a single horizontal row.

option, the dashed edges adjacent to k and j respectively are merged to a single edge with weight $3 : 4$, see Figure 4b, right. In both cases, the total length of the contacts corresponding to the dashed edges is between 6 and 8. We shall make sure that this length is always at most 7. For example, this is the case when the dashed edges belong to another type-1 gadget to the right of the subgraph in Figure 4b. As we shall see later, this assumption also holds if we have other gadgets (an inverter or a replicator) on the right. Also, it always holds for the second option, because a single rectangle can only have contacts with total length at most 7 to a type-2 gadget. Since the dash-dotted edges in Figure 4b correspond to contacts of length 1, both the lower and the upper type-1 gadget must be 5×3 rectangles centered horizontally at the type-2 gadget. The only two possible layouts are shown in 4d. Let the left one encode the value *true* and the right one *false*.

Several copies of the building block can be attached to each other both vertically and horizontally so that the green 7-vertex subgraphs (type-1 gadgets) link two adjacent blue 5-vertex subgraphs (type-2 gadgets). This synchronizes the states of all blocks: either all linked blocks are in the *true* state or all are in the *false* state. This allows us to create horizontal variable gadgets with vertical branches leading towards the clause gadgets. We create 90° pipe bends in the same way. Two different variable gadgets are separated by three *buffer vertices* (or rectangles) that do not link the gadget states while fixing relative positions of the gadgets, see Figure 7.
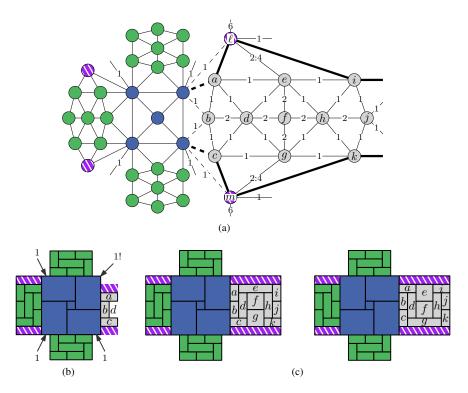
Figure 8: (a) The inverter gadget subgraph. Thick edges have weight $1 : 2$. (b) $R(a)$ and $R(c)$ must have different orientations. (c) The only two possible realizations.

2.3.2   Inverters and replicators

For the reduction, we shall need the ability to invert the state of the incoming truth value. We shall achieve this by defining an *inverter* gadget, see Figure 8. The rectangles $R(a)$, $R(b)$, $R(c)$, $R(i)$, $R(j)$, $R(k)$ must have dimensions 1×2 or 2×1, $R(\ell)$ and $R(m)$ must be 6×1 rectangles and $R(f)$ a 2×2 square. Both $R(d)$ and $R(h)$ are either 1×3 or 2×2. Furthermore, $R(b)$ must be oriented vertically. Due to the shape of $R(d)$, the rectangles $R(a)$ and $R(c)$ can not be vertical at the same time, so the total length of the contacts corresponding to dashed edges in Figure 8b is at most 7. If $R(d)$ is 2×2, both $R(a)$ and $R(c)$ must be oriented horizontally, see Figure 8b. This is not possible, since each non-covered boundary piece of the blue rectangles in Figure 8b must have length 1. Thus, either $R(a)$ or $R(c)$, but not both, are oriented horizontally, and $R(d)$ is a 1×3 rectangle.

Due to the specified edge weights, the rectangles $R(e)$ and $R(g)$ are either 2×2, 3×1 or 4×1. If $R(a)$ is horizontal, then $R(e)$ is 2×2, $R(h)$ is 1×3 and $R(g)$ is 4×1. The case in which $R(c)$ is horizontal is symmetric. Due to an argument similar to the one for the building block in Figure 4, the purple rectangles force the inverter to be centered vertically at the type-2 gadget. Thus, we have exactly one realization of the inverter for each truth value of the connected variable, see Figure 8c.
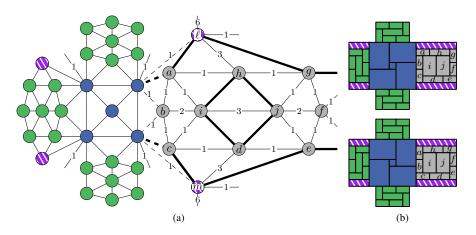
Figure 9: (a) The replicator gadget subgraph. Thick edges have weight 1 : 2. (b) The only two possible realizations.

To unify our construction, we define *replicator* gadgets. These gadgets do not invert the truth value of a variable, but they have the same dimensions (6×7) as the inverters. The corresponding edge-weighted contact subgraph is depicted in Figure 9. Due to the specified edge weights, $R(a)$, $R(b)$, $R(c)$, $R(e)$, $R(f)$ and $R(g)$ must be 1×2 or 2×1 rectangles, $R(i)$ and $R(j)$ must be 2×3 and $R(d)$ and $R(h)$ must be 3×1. Thus, we have exactly one realization for each truth value, see Figure 9b.

2.3.3 Clauses

It remains to describe the clause gadget, whose rectangular layout is shown in Figure 10. It takes three inputs, two from the left side and one from below or above depending on whether the clause gadget is placed above or below the variable row. Note that the input from below or above is duplicated. Each input port consists either of an inverter or a replicator gadget. The type of the port gadget depends not only on whether the literal in the clause is positive or negative, but also on the position of the port in the clause gadget: The top left and the bottom right ports use an inverter for a positive literal and a replicator for a negative one; the bottom left and top right ports use a replicator for a positive literal and an inverter for a negative one. This configuration has the following effect on the two core rectangles $R(l)$ and $R(r)$ of the clause gadget, whose contact length is bounded by $\alpha(lr) = 19$ and $\beta(lr) = 20$. Every false literal stretches its adjacent rectangle $R(l)$ or $R(r)$ vertically by a length of 1 (in fact by a length of 2 for $R(r)$ since the last literal is duplicated). If all literals are *true* then both $R(l)$ and $R(r)$ have height 19 and also $|s(lr)| = 19$. By inspecting all cases one can see that as long as one literal is *true* we have $19 \leq |s(lr)| \leq 20$, but as soon as all three literals are false the contact length becomes $|s(lr)| = 21$ violating the specified upper bound. This is exactly the behavior needed for the clause gadget.
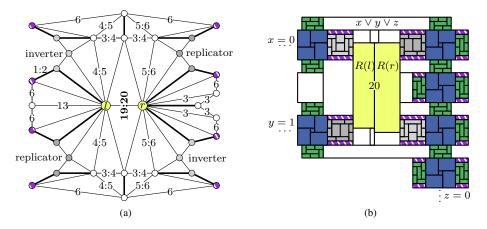
Figure 10: (a) Clause gadget subgraph for the clause $x \vee y \vee z$ and (b) the layout for the state $x = 0$, $y = 1$, $z = 1$. The contact length $|s(lr)|$ of the two yellow rectangles $R(l)$ and $R(r)$ is 20.

2.3.4   Putting all blocks together

All presented building blocks can be laid out on an orthogonal grid subdivided into 10×10 tiles (most of these tiles have a 7×7 block in one corner, see Figure 11, which shows the rectangular dual of the entire graph G_ϕ for an example formula ϕ). In order to create an actual PTP graph, the remaining gaps between the described gadgets must be filled by dummy rectangles, i.e., dummy vertices in G_ϕ. This is always possible and can be done systematically. All building blocks, except for the clause gadget, have the same top-level structure: they are all formed by 7×7 blocks, with 3×5 blocks connected to their left and right and 5×3 blocks connected to their top and bottom. As can be verified in Figure 11, the clause gadgets are compatible with this pattern.

The remaining holes are filled with dummy rectangles with dimensions 7×7, 5×3, 3×5, 3×1 and 1×7 in a manner demonstrated in Figure 11. In the resulting layouts, no four rectangles meet in a single point. Thus, the contact graph G_ϕ is inner triangulated. We add four outer vertices to make it PTP. The lower and upper bounds on the lengths of the contacts between dummy and non-dummy rectangles can be extracted from the presented gadget subgraphs (the corresponding edges are indicated as stubs in the drawings). For a contact $s(e)$ between two dummy rectangles, we set the weights $\alpha(e) = \beta(e)$ to the desired length of the contact according to Figure 11. Note that for our proof we do not have to show that the dummy rectangles have unique realizations: it suffices to show that they can be used to fill the holes in the fixed "skeleton" of the rectangular dual formed by the gadgets (colored rectangles in Figure 11).

We use the presented gadgets to prove Theorem 4.

Proof of Theorem 4: Let ϕ be a PLANAR 3SAT formula. As mentioned above, the variable-clause graph H_ϕ can be drawn on an orthogonal grid of size polynomial in the size of ϕ. Furthermore, the number of vertices in each gadget subgraph is $O(1)$ or polynomial for the variable gadgets. Thus, the graph G_ϕ can be constructed in
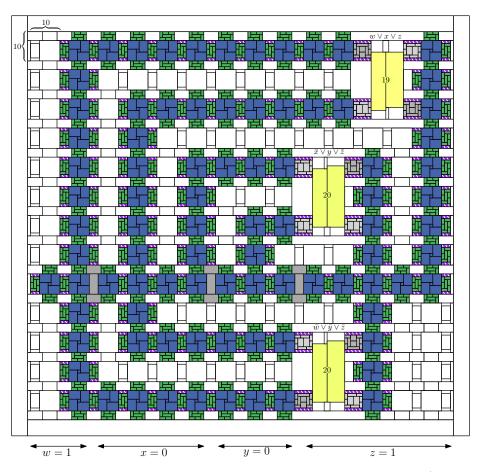
Figure 11: Rectangular dual of the graph G_ϕ for the planar 3Sat formula $\phi = (w \vee x \vee z) \wedge (\bar{x} \vee y \vee \bar{z}) \wedge (\bar{w} \vee \bar{y} \vee \bar{z})$ with the satisfying variable assignment $w = 1$, $x = 0$, $y = 0$, $z = 1$.

polynomial time. As we have shown above, the positions of type-1 and type-2 gadgets in a layout whose contact lengths respect α and β are fixed, and the respective choice of realization depends only on the truth value of the corresponding variable.

If ϕ is satisfiable, we choose realizations of the gadgets according to the truth value of the corresponding variable in a satisfying variable assignment. By the properties of the clause gadgets, all rectangles can be realized in a way that respects the prescribed bounds on the contact lengths. Thus, a suitable rectangular dual of G_ϕ exists.

Now let $\mathcal{R}$ be a rectangular dual of G_ϕ that respects the bounds on the contact lengths specified by α and β. Then, in each clause gadget, the contact between the two yellow rectangles $R(l)$ and $R(r)$ must have length at most 20, so at least one of the connected type-2 gadgets must have the realization corresponding to the truth value *true*. Since all type-1 and type-2 gadgets connected to the same variable gadget are
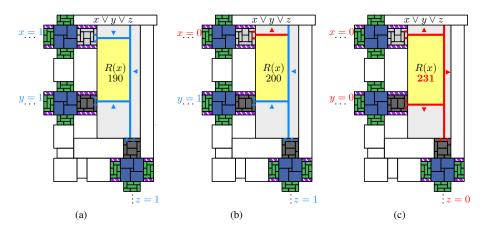
Figure 12: Clause gadget for the proof of Theorem 5 for the clause $x \vee y \vee z$. Each *false* literal stretches either width or height of the yellow rectangle $R(x)$ by 1. Thus, only in the case (c) for $x = y = z = 0$ its area exceeds $\gamma(x) = 220$.

synchronized among each other, we can extract a satisfying truth value assignment for ϕ directly from $\mathcal{R}$. This shows NP-hardness.

The decision problem is contained in $\mathcal{NP}$. If we guess a REL of G which fixes the orientation of each contact, we can test the existence of a rectangular dual which respects α and β in polynomial time, e.g., using linear programming. □

A variant of the problem with lower bounded contact lengths and upper bounded rectangle areas turns out to be NP-hard as well as the next theorem shows.

Theorem 5 *Given a PTP graph $G = (V, E)$ with an edge-weight function $\omega : E \to \mathbb{R}^+$ and a vertex-weight function $\gamma : V \to \mathbb{R}^+$, it is NP-hard to decide if G has a rectangular dual $\mathcal{R} = \{R(v) \mid v \in V\}$ so that*

(i) for every edge $e \in E$ the contact segment $s(e)$ has length $|s(e)| \geq \omega(e)$ and

(ii) for every vertex $v \in V$ the rectangle $R(v)$ has area $|R(v)| \leq \gamma(v)$.

Sketch of the proof: The proof is very similar to the proof of Theorem 4. See the diploma thesis of Roman Prutkin [23] for more details. Except for the clause gadget, we use the same gadgets. However, the uniqueness of the realization up to the truth value must now be forced by a suitable assignment of minimal contact lengths specified by ω and maximal rectangle areas specified by γ. These assignments can be defined according to the layouts in Figures 5 to 9. We omit detailed specifications of the weighted subgraphs at this point and present only the new clause gadget.

The clause gadget is shown in Figure 12. Note that the positions of type-1 and type-2 gadgets are fixed. The gadget takes three inputs: two from the left and one from below (if the clause lies above the variable row) or from above (if the clause lies below

the variable row). At each input port we again place either an inverter or a replicator gadget. The upper left and the upper input port (if the clause lies below the variable row) use an inverter for a positive literal and a replicator for a negative one. The lower left and the lower input port (if the clause lies above the variable row) use a replicator for a positive literal and an inverter for a negative one.

This definition has the following effect on the area of the yellow rectangle $R(x)$, which is upper-bounded by $\gamma(x) = 220$: The position of the left boundary of $R(x)$ is fixed, and the coordinates of the upper, lower and right boundaries are determined by the truth value of the incoming variable, such that each *false* literal increases either the width or height of $R(x)$ by 1. Thus, the area of $R(x)$ is $231 > \gamma(x)$ only in the case when all three literals are *false*, and in all other seven cases it is at most $220 = \gamma(x)$. Now it is straightforward to adapt the proof of Theorem 4.

It is not obvious whether the problem is contained in $\mathcal{NP}$. Unlike for the case in Theorem 4, it is unclear whether the decision problem in Theorem 5 can be solved in polynomial time for a fixed REL: a similar linear programming approach doesn't seem to work here, since the area constraints are non-linear and also not necessarily positive semidefinite. □

2.4   Minimizing layout size for specified minimum contact lengths

In Section 2.2 we have demonstrated that for a given PTP graph $G = (V, E)$, a fixed REL and lower bounds on the contact lengths specified by ω, a rectangular layout of minimum size can be found in linear time. In this section, we consider the problem of finding the smallest layout of G with respect to the lower bounds ω over all RELs of G. First, we shall consider area minimization and show the following theorem:

Theorem 6 *Given an edge-weighted PTP graph $G = (V, E, \omega)$ and a bound $A \in \mathbb{Q}^{+}$, it is NP-complete to decide whether a rectangular dual $\mathcal{R} = \{R(v) \mid v \in V\}$ exists such that the total area of $\mathcal{R}$ is at most A and for each edge $e \in E$ the corresponding contact segment $s(e)$ in $\mathcal{R}$ has length at least $\omega(e)$.*

The proof is a reduction from the NP-hard problem PARTITION. Let $a_1, \ldots, a_m$ be an instance of PARTITION and $\sigma := \frac{1}{2}\sum_{i=1}^{m} a_i$. We set $\varepsilon = \frac{1}{4(2m+5)}$ and $a_{m+1} = \varepsilon$. For our reduction proof, we define graphs G_j, $j = 1, \ldots, m$ recursively; see Figure 13. The Graph G_1 is a 4-cycle s_1, e_1, n_1, w_1 with the inner node s_0. The edges $w_1 s_1$ and $e_1 n_1$ have weight a_1; see Figure 13a. Figure 13d shows how to construct G_{j+1} from G_j. We insert G_j into a 4-cycle $s_{j+1}, e_{j+1}, n_{j+1}, w_{j+1}$ and connect s_{j+1} to s_j and e_j, then e_{j+1} to e_j and n_j, then n_{j+1} to n_j and w_j and, finally, w_{j+1} to w_j and s_j. For $i = 1, \ldots, j$, edges $w_i s_i$ and $e_i n_i$ have weight a_i, and all other edges in G_j have weight ε.

From now on we shall consider only duals of G_j in which $R(s_j)$ is the lower left, $R(w_j)$ the upper left, $R(n_j)$ the upper right and $R(e_j)$ the lower right rectangle. The graph G_1 has exactly two such duals; see Figure 13b and 13c. There are exactly two possible realizations of the four outer rectangles $R(s_j)$, $R(w_j)$, $R(n_j)$ and $R(e_j)$ of G_j; see Figure 13e and 13f. The gray rectangle in the figures is replaced by a suitable dual of G_{j-1}. Note that in Figure 13f this dual is rotated by $90°$ clockwise compared to the layout in Figure 13e.
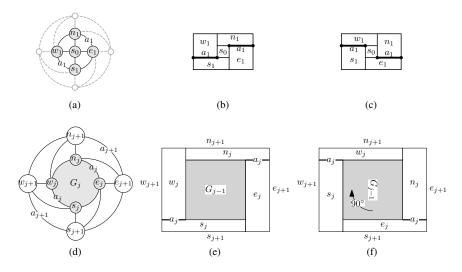
Figure 13: The recursively defined graphs G_j, $j = 1, \ldots, m$. Gray dashed edges do not belong to G_j. Black unlabeled edges have weight ε. (a), (b), (c): Graph G_1 (black edges) and both its realizations. (d) Recursive construction of G_{j+1} from G_j. (e), (f): Each dual of G_{j-1} can be extended to a dual of G_j in exactly two ways.

Let $\mathcal{R}_j$ be a rectangular dual of G_j. We define

$$H(\mathcal{R}_j) = \{i \in \mathbb{N} \mid 1 \leq i \leq j, \text{ contact } s(w_i s_i) \text{ is horizontal in } \mathcal{R}_j\},$$
$$V(\mathcal{R}_j) = \{i \in \mathbb{N} \mid 1 \leq i \leq j, \text{ contact } s(w_i s_i) \text{ is vertical in } \mathcal{R}_j\}.$$

Note that the contacts corresponding to edges $w_i s_i$, $s_i e_i$, $e_i n_i$, $n_i w_i$ must have alternating orientations. The main idea of our reduction proof is to encode a partition by orientations of the contacts with minimum lengths a_i, $i = 1, \ldots, m$. The following lemma shows that each partition can be encoded by a dual of G_m. We say that a rectangular dual $\mathcal{R}$ of G_j with the lower left rectangle $R(s_j)$, upper left rectangle $R(w_j)$, upper right rectangle $R(n_j)$ and lower right rectangle $R(e_j)$ *respects* a partition $\mathcal{P} = (P_H, P_V)$ of $\{1, \ldots, j\}$, $j \in P_H$, if $\mathcal{R}$ respects ω, $H(\mathcal{R}_j) = P_H$ and $V(\mathcal{R}_j) = P_V$.

Lemma 2 *Consider a graph G_j, $1 \leq j \leq m$, and a partition $\mathcal{P} = (P_H, P_V)$ of $\{1, \ldots, j\}$, $j \in P_H$. The following statement holds. (1) There exists a rectangular dual $\mathcal{R}$ that respects $\mathcal{P}$. (2) Each dual $\mathcal{R}$ of G_j that respects $\mathcal{P}$ has width and height at least*

$$w_{\mathcal{P}} := 2 \sum_{i \in P_H} a_i + (2|P_V| + 1)\varepsilon,$$
$$h_{\mathcal{P}} := 2 \sum_{i \in P_V} a_i + (2|P_H| + 1)\varepsilon$$

respectively. (3) There exists a dual $\mathcal{R}_j$ of G_j with width $w_{\mathcal{P}}$ and height $h_{\mathcal{P}}$ that respects $\mathcal{P}$.
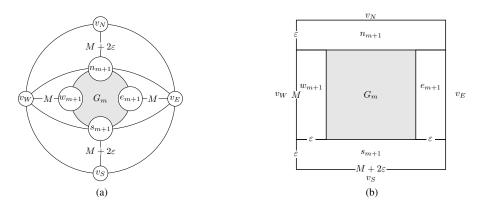
Figure 14: The PTP graph G used for the proof of Theorem 6 and its rectangular dual. Edges are labeled by their weights, unlabeled edges have weight ε.

Proof: We prove (1) first. The claim obviously holds for $j = 1$. Now let $j \geq 2$, and assume the claim holds for $j - 1$. We consider two different cases.

Case 1: $j - 1 \in P_H$. Then, we define $P'_H = P_H \setminus \{j\}$, $P'_V = P_V$ and apply the claim inductively to $\mathcal{P}' = (P'_H, P'_V)$ and G_{j-1}. Thus, there exists a rectangular dual $\mathcal{R}'$ of G_{j-1} that respects $\mathcal{P}'$. Therefore, $H(\mathcal{R}') = P'_H = P_H \setminus \{j\}$ and $V(\mathcal{R}') = P_V$. Then, we obtain $\mathcal{R}$ with $H(\mathcal{R}) = P_H$ and $V(\mathcal{R}) = P_V$ by inserting $\mathcal{R}'$ into the gray rectangle in Figure 13e.

Case 2: $j - 1 \in P_V$. We define $P'_H = P_V$, $P'_V = P_H \setminus \{j\}$ and apply the claim inductively to $\mathcal{P}' = (P'_H, P'_V)$ and G_{j-1}. Thus, there exists a rectangular dual $\mathcal{R}'$ of G_{j-1} that respects $\mathcal{P}'$. Therefore, $H(\mathcal{R}') = P'_H = P_V$ and $V(\mathcal{R}') = P_H \setminus \{j\}$. We obtain $\mathcal{R}$ with $H(\mathcal{R}) = P_H$ and $V(\mathcal{R}) = P_V$ by rotating $\mathcal{R}'$ by 90° clockwise and then inserting it into the gray rectangle in Figure 13f.

(2) The lower bounds on width and height can be easily verified. For each $i \in H(\mathcal{R})$, the contacts corresponding to edges $s_i w_i$ and $n_i e_i$ are horizontal, as well as the contacts corresponding to $s_i e_i$ and $n_i w_i$ for $i \in V(\mathcal{R})$. Their projections onto the x-axis do not overlap in the interior, hence the lower bound on the total width holds. A similar argument holds for the height.

(3) With the tools from Section 2.2 we can construct a dual $\mathcal{R}_j$ that respects $\mathcal{P}$ with width $w_\mathcal{P}$ and height $h_\mathcal{P}$.   □

The PTP graph $G = (V, E)$ used for the reduction is depicted in Figure 14. We define $M = 2\sigma + (2m+1)\varepsilon < 2\sigma + \frac{1}{4}$. The figure shows how the four outer vertices v_S, v_E, v_N, v_W are connected to the vertices s_{m+1}, e_{m+1}, n_{m+1} and w_{m+1}. The weights of the corresponding edges are chosen in a way such that in each rectangular dual $\mathcal{R}$ of G both the total width and height of the inner rectangles are at least $M + 2\varepsilon$. If they are both exactly $M + 2\varepsilon$, then $\mathcal{R}$ induces a rectangular dual $\mathcal{R}_m$ of G_m with height and width at most M; see the gray rectangle in Figure 14. We can now prove Theorem 6.

Proof of Theorem 6: We set $A = (M + 4\varepsilon)^2 = (2\sigma + (2m+5)\varepsilon)^2 = (2\sigma + \frac{1}{4})^2$. First, assume there exists a solution P of the PARTITION instance $a_1, \ldots, a_m$ and $\mathcal{P} =$

$(P, \{1, \ldots, m\} \setminus P)$. Without loss of generality, let $m \in P$. Then, according to Lemma 2, there exists a rectangular dual $\mathcal{R}_m$ of G_m with $H(\mathcal{R}_m) = P$ with width $w_P \leq 2\sigma + (2m + 1)\varepsilon = M$ and height $h_P \leq 2\sigma + (2m + 1)\varepsilon = M$. The dual $\mathcal{R}_m$ can then be extended to a rectangular dual $\mathcal{R}$ of G with height and with at most $M + 4\varepsilon$ and total area A.

Now assume the PARTITION instance has no solution. Since each rectangular dual $\mathcal{R}_m$ induces a partition of $\{1, \ldots, m\}$ by $H(\mathcal{R}_m)$, it has height or width at least $2\sigma + 1$. Therefore, each rectangular dual $\mathcal{R}$ of G either has width at least $2\sigma + 1$ and height at least M, or vice versa. Thus, its area is at least $M(2\sigma + 1) \geq 2\sigma(2\sigma + 1) > (2\sigma + \frac{1}{4})^2 = A$. The graph G can be constructed in time polynomial in m. This proves NP-hardness in Theorem 6.

The decision problem is contained in $\mathcal{NP}$. If we guess an orientation for each contact in a rectangular dual of G, we can construct a dual $\mathcal{R}$ that respects ω and has minimum width and height in linear time using tools from Section 2.2.   □

With similar gadgets we can show that minimizing the perimeter, $\max(\text{width}, \text{height})$ or the total edge length of a rectangular dual of a PTP graph G with respect to specified lower bounds on the contact lengths ω over all RELs of G is NP-complete:

Theorem 7 *Given an edge-weighted PTP graph $G = (V, E, \omega)$ and a bound $A \in \mathbb{Q}^+$, it is NP-complete to decide whether a rectangular dual $\mathcal{R} = \{R(v) \mid v \in V\}$ exists such that the perimeter (or $\max(\text{width}, \text{height})$, or the total edge length) of $\mathcal{R}$ is at most A and for each edge $e \in E$ the corresponding contact segment $s(e)$ in $\mathcal{R}$ has length at least $\omega(e)$.*

3   Length-universal rectilinear layouts

In this section we consider the number of bends required for constructing edge-proportional rectilinear representations (or EPRRs for short) of internally triangulated planar graphs $G = (V, E, \omega)$. In our proofs we assume that the graphs are biconnected because every internally triangulated graph can be made biconnected by adding vertices in the outer face. Since our representations preserve the outerplanar embedding, the removal of the corresponding polygons does not create holes.

The *complexity* of a rectilinear polygon p is its number of bends, and is denoted by $k(p)$. The complexity of an EPRR P is $k(P) := \max_{u \in V} k(P(u))$. The complexity of a graph G with weight function ω is $k(G, \omega) = \min_{P \in \mathcal{P}} k(P)$, where $\mathcal{P}$ denotes all EPRRs of (G, ω). For a graph class $\mathcal{G}$, the complexity is the maximum complexity for any graph from $\mathcal{G}$ with any weight function ω, i.e., $k(\mathcal{G}) = \max_{G \in \mathcal{G}} \max_{\omega : E \to \mathbb{R}^+} k(G, \omega)$. We are interested in determining $k(\mathcal{G})$ for different classes of graphs.

Theorem 8 *For any biconnected internally triangulated graph G with two adjacent internal vertices and any positive integer k_0, there exists a weight function ω such that $k(G, \omega) \geq k_0$.*

Let $\mathcal{I}$ denote the class of biconnected internally triangulated graphs containing no adjacent internal vertices; we have $k(\mathcal{I}) \geq k_0$ for any positive integer k_0.

Proof: For the first part of the theorem let G be a biconnected internally triangulated graph with two adjacent internal vertices u and v. We define ω such that $\omega(e) = 1$
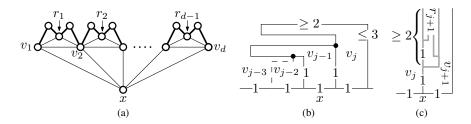
Figure 15: A graph that does not admit an edge-proportional rectilinear representation with 6-gons for $d = 55$; thin edges have weight 1 and thick edges have weight 4.

for all $e \neq uv$ and $\omega(uv) = k_0 \cdot (\deg(u) + \deg(v) - 2) =: M$. Now, in any contact representation the polygon $P(u) \cup P(v)$ has a boundary of length $\deg(u) + \deg(v) - 2$. On the other hand, this polygon necessarily contains the contact path of length M corresponding to the edge uv. This path has at least $M/(\deg(u) + \deg(v) - 2) = k_0$ bends as it would cross the boundary of $P(u) \cup P(v)$ otherwise.

For the second part, consider the graph K_4 with internal vertex x, outer vertices a, b, c and all edge weights set to 1 except for $\omega(ax) = 2$. Since $\omega(ax) \geq \omega(bx) + \omega(cx)$ the path $s(ax)$ must have a bend. Now consider a fan-graph on $k_0 + 2$ vertices with center vertex a, and insert into each fan triangle T a new internal vertex x_T connected with edges of weight 2 to a and weight 1 to the other two vertices of T. By the above observation the polygon $P(a)$ needs one bend per path $s(ax_T)$ for all k_0 fan triangles T. $\square$

This shows that to achieve positive results, we may allow only few isolated interior vertices. Thus we consider outerplanar graphs and graphs with one internal vertex.

3.1   Outerplanar graphs

In this subsection we study the complexity of edge-proportional rectilinear representations for internally triangulated outerplanar graphs.

Proposition 1 *For the class $\mathcal{O}$ of biconnected internally triangulated outerplanar graphs $k(\mathcal{O}) \geq 8$.*

Proof: Consider the family of graphs depicted in Figure 15a. We show that if all thin edges have weight 1 and the thick edges have weight 4, then for $d = 55$, the corresponding weighted graph (G, ω) does not admit a representation with complexity less than 8. Assume for contradiction that P is a representation with complexity at most 6. (Note that the complexity of a rectilinear polygon is always even.)

Claim: There exists a chain $Q = \{v_i, \dots, v_{i+8}\}$ such that the contacts between $P(x)$ and $P(v)$ with $v \in Q$ all lie on a common line.

This follows easily from the fact that there are 55 vertices on the path $v_1, \dots, v_d$, and to avoid all such chains Q, $P(x)$ would need to bend at least once every 9th contact. But then we get at least $\lceil 55/9 \rceil = 7$ bends on $P(x)$.
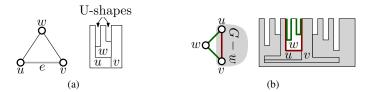
Figure 16: Construction of a rectilinear representation by 8-gons for an outerplanar graph. (a) A triangle with reference edge e, forming the base case. (b) Inserting $P(w)$ into the U-shape of uv (red), creating new U-shapes for uw and vw (green).

Claim: Let i and Q be chosen as in the previous claim, and without loss of generality assume that their contacts lie on a common horizontal line. Then for $j = i+3, \ldots, i+5$, we have that $P(v_j)$ has height more than 3.

Assume for a contradiction that the height of $P(v_j)$ is at most 3. Then, as $P(v_j)$ has perimeter at least $2 \cdot 4 + 5 = 13$, it must be realized as an L-shape with an overhang of width at least 2, say to the left. It follows that $P(v_{j-1})$ has height at most 1; see Fig. 15b. But then $P(v_{j-2})$ has perimeter at most 4 as it is enclosed in a 1×1-box, a contradiction. The case that the overhang is to the right is symmetric. This proves the second claim.

Now consider $P(v_{i+4})$. Either its left or right side does not have a bend, and hence is a vertical segment of length at least 3. Without loss of generality assume that it is the right side. We then consider v_{i+4} and v_{i+5}, and their common neighbor r_{i+4}. The situation is depicted in Fig. 15c. The path $s(r_{i+4}v_{i+5})$ has length 1, and thus bends at the reflex point of $P(v_{i+5})$. Since both $P(v_{i+4})$ and $P(v_{i+5})$ have height at least 3 and $P(r_{i+4})$ has perimeter 10, $P(r_{i+4})$ needs two bends in order to achieve the correct contact lengths with both of them; a contradiction to the assumption that P has complexity 6. □

On the other hand, we describe an algorithm that produces for any outerplanar graph G with weight function ω a representation with complexity 8.

Proposition 2 *For the class $\mathcal{O}$ of internally triangulated outerplanar graphs $k(\mathcal{O}) \leq 8$.*

Proof: First, we make a given $G \in \mathcal{O}$ biconnected while keeping it outerplanar by adding dummy vertices, see Figure 17. Then, we show that for any biconnected outerplanar graph $G = (V, E)$ with weight function ω and a reference edge $e \in E$ on the outer face, there exists an edge-proportional rectilinear representation P such that for each edge uv on the outer face with $uv \neq e$, there exists a U-shape whose left and right boundary are formed by the polygons $P(u)$ and $P(v)$, whose open side points to the top, and whose width is at most $\varepsilon/2$, where ε is the smallest weight of all edges.

For a triangle uvw with reference edge uv this is obviously possible; see Fig. 16a. We construct the drawing inductively. Let G be an arbitrary graph with reference edge e. Since G is outerplanar and has more than three vertices, it has a degree-2 vertex w that is not adjacent to e. By induction $G - w$ has a desired representation P with respect to the reference edge e. Let u and v denote the two neighbors of w, which are connected by an edge on the outer face of $G - w$. Note that the presence of edge uv implies that $G - w$ remains biconnected. By the properties of P, there is a U-shape for the edge uv. We
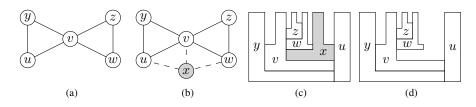
Figure 17: Edge-proportional rectilinear representation with 8-gons for outerplanar inner triangulated graphs with separating vertices. (a), (b): after inserting x and edges xw, xv, xu, vertex v is no longer separating. (c) We construct an EPRR with octagons; (d) then we delete $P(x)$.

then insert a new polygon $P(w)$ into this U-shape as illustrated in Fig. 16b. Obviously, this preserves all invariants. □

Propositions 1 and 2 imply that for the class $\mathcal{O}$ of outerplanar graphs, we have $k(\mathcal{O}) = 8$. We remark that our technique for representing outerplanar graphs with 8-gons extends to graphs with a single internal vertex, by wrapping the drawing around this central vertex while creating the U-shapes for the outer edges. Moreover, in the outerplanar case, it is simultaneously possible to achieve given areas for all vertices by suitably stretching the polygons to satisfy the area demands. Next, we consider special cases and show that outerpaths (outerplanar graphs whose weak dual is a path) require six bends, and that six bends suffice for outerpillars.

Proposition 3 *For the class $\mathcal{P}$ of internally triangulated outerpaths $k(\mathcal{P}) \geq 6$.*

Proof: We show that rectangles are not sufficient, even for outerpaths. Consider the fan-graph G on 17 vertices consisting of a path $v_1, \dots, v_{16}$, whose edge weights are $2, 1, 2, 2, 1, 2, \dots, 2, 1, 2$, and a center vertex u that is connected to all vertices of the path by edges of weight 1. Assume for a contradiction that G admits an edge-proportional rectilinear representation by rectangles. Consider four vertices $v_i, v_{i+1}, v_{i+2}, v_{i+3}$ with contact lengths 2-1-2. We claim that $P(u)$ has a bend for any such sequence.

Assume this is not the case, then all contact paths of v_i, v_{i+1}, v_{i+2} and v_{i+3} with $P(u)$ form adjacent, without loss of generality, horizontal segments. Each of these segments has length 1, and hence the widths of $P(v_i), P(v_{i+1}), P(v_{i+2})$ and $P(v_{i+3})$ are fixed to 1. The remaining degree of freedom is to choose their heights. However, $\omega(v_i v_{i+1}) = 2$ requires that $P(v_i)$ and $P(v_{i+1})$ have height at least 2, and the same argument applies for $P(v_{i+2})$ and $P(v_{i+3})$. However, $\omega(v_{i+1} v_{i+2}) = 1$ requires that the height of $P(v_{i+1})$ or $P(v_{i+2})$ is 1, a contradiction, and $P(u)$ must have a bend during this sequence. Since there are five such sequences for $i = 1, 4, 7, 10, 13$, it follows that $P(u)$ has at least six bends. □

Proposition 4 *For the class $\mathcal{P}'$ of internally triangulated outerpillars $k(\mathcal{P}') \leq 6$.*

Proof: Let G' be an outerpillar. As a first step, we pick the outerpath $G \subseteq G'$ that is obtained by removing all but the two outermost degree-2 vertices of G'; see Fig. 18.
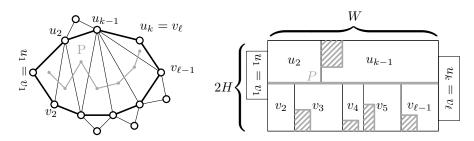
Figure 18: Construction of an edge-proportional rectilinear representation with 6-gons for outerpillars. Input graph with an outerpath subgraph, drawn with a thick boundary (left). An EPRR of the outerpath with U-shapes for the remaining vertices shown as tiled rectangles (right).

Let P be the path that is dual to G. The path P splits the boundary of the outer face of G into two paths $\pi_1 = u_1, \ldots, u_k$ and $\pi_2 = v_1, \ldots, v_\ell$ that are internally disjoint, and that share exactly their endpoints, i.e., $u_1 = v_1$ and $u_k = v_\ell$, both of which have degree 2 in G. Let H be larger than the maximum weight, and let W denote the total weight of all internal edges. Construct a $2H \times W$ box, split it horizontally into two boxes of size $H \times W$. We then split the upper box into $k-2$ rectangles $P(u_2), \ldots, P(u_{k-1})$ such that the width of $P(u_i)$ is the sum of the weights of all internal edges of G incident to u_i. We split the lower box into boxes for $v_2, \ldots, v_{\ell-1}$ analogously; see Fig. 18. Observe that this ensures correct contact lengths for all internal edges of G. Next, we place rectangles for u_1 and u_k as boxes to the left and right of the drawing such that they have the correct contact lengths. This ensures correct contact lengths except for internal vertices of π_1 and π_2, respectively, that are adjacent. They touch in a segment of length H, which is too long. To remedy this, we remove for any such pair $v_i v_{i+1}$ occurring in this order on π_1 or π_2 a corner of the rectangle of v_{i+1}. This corner is chosen such that its width is at most half the smallest contact length, and such that afterwards $|s(v_i v_{i+1})| = \omega(v_i v_{i+1})$. This finishes the construction for outerpaths. For the more general case of outerpillars observe that there is a small U-shaped gap between any adjacent pair of vertices on the outer face, and we can hence use the same approach as in the proof of Proposition 2 to attach further leaves to the central outerpath determined by the spine P. Note that the polygons can be stretched such that one contact has the correct length, so that only six bends are necessary.   □

This completely characterizes the complexity of length-universal layouts for internally triangulated graphs. As we have seen outerplanar graphs whose dual is a caterpillar require complexity 6. The dual of the example graph showing that the 8 bends are necessary is a lobster. Hence our results are best possible.

One disadvantage is that our drawings have an outer face of high complexity. However, we can show that one cannot do better. If one limits the complexity of the outer face to some fixed number K, then there exist outerplanar graphs that require complexity k_0 for any positive integer $k_0 > K$.

Proposition 5 *For every given k_0, $K \in \mathbb{N}$, k_0, $K \geq 4$, there exists an edge-weighted*
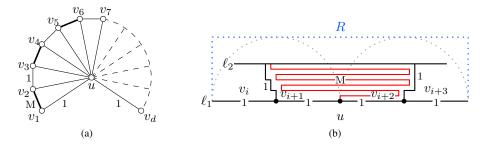
Figure 19: (a) Fan graph G_d used in the proof of Proposition 5. Thick edges have weight M, thin non-dashed edges have weight 1. (b)

outerpath $(G, \omega) \in \mathcal{O}$ such that each edge-proportional rectilinear representation P of G with complexity $k(P) \leq k_0$ has more than K bends on the outer face.

Proof: Consider the family of outerpaths G_d in Figure 19a, which are triangle fans around the vertex u containing $d - 1$ triangles. We set $M = 4k_0$ and select the value of d to be $5 \cdot (k_0 + K + 1) + 1$. Assume there exists an edge-proportional rectilinear representation P with polygon complexity $k(P) \leq k_0$ and at most K bends on the outer face. Then there must exist five consecutive neighbors $v_i, \ldots, v_{i+4}$ of u, $i \in \{0, \ldots, d-4\}$, such that: (1) the contacts between the polygons $P(u)$ and $P(v_i)$, ..., $P(u)$ and $P(v_{i+4})$ lie on a single line ℓ_1, and (2) the outer boundaries of the polygons $P(v_i), \ldots, P(v_{i+4})$ lie on a single line ℓ_2; see Figure 19b. Otherwise, any five consecutive neighbors v_j, ..., v_{j+4} would be responsible for at least one bend of the polygon $P(u)$ or the outer contour P_{out}. But then, the total number of bends of $P(u)$ and P_{out} would be at least $\frac{d}{5} - 1$. Thus, it would hold $k_0 + K \geq \frac{d}{5} - 1$ and $5 \cdot (k_0 + K + 1) \geq d$, a contradiction.

Without loss of generality, let $\omega(v_i v_{i+1}) = 1$ (otherwise, use $i+1$ instead of i). Then it is also $\omega(v_{i+2} v_{i+3}) = 1$. Thus, the line ℓ_1 must be parallel to the line ℓ_2, and the two lines have distance at most 1 to each other. Therefore, the contact $s(v_{i+1} v_{i+2})$ between the polygons $P(v_{i+1})$ and $P(v_{i+2})$ must completely fit into a rectangle R of width 4 and height 1 (dotted blue rectangle in Figure 19b). Thus, $P(v_{i+1})$ must have at least $k_0 + 1$ bends, a contradiction. □

3.2   Graphs with one internal vertex

Next, we show that polygon complexity 8 also suffices to represent every inner triangulated graph with one inner vertex as an edge-proportional rectilinear representation.

Proposition 6 *For the class $\mathcal{I}_1$ of plane inner triangulated graphs with exactly one inner vertex $k(\mathcal{I}_1) \leq 8$.*

Proof: Consider a graph $G = (V, E, \omega) \in \mathcal{I}_1$. We make G biconnected by adding new vertices as described earlier. Let u be the only inner vertex in G, let d denote its degree, and let $v_1, \ldots, v_d$ be its neighbours in counterclockwise order. We set $\sigma = \frac{1}{2} \sum_{i=1}^{d} \omega(uv_i)$,
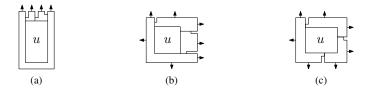
Figure 20: Proof of Proposition 6 for $d = 3$. For $d > 3$ we acquire similar octagonal constructions by subdividing polygons.

$\ell = \max_{i=1,\dots,d} \omega(uv_i)$, $j = \operatorname{argmax}_{i=1,\dots,d} \omega(uv_i)$ and ε to be the minimum edge weight in G. We can always draw $P(u)$ as a rectangle with correct contact lengths and wrap the polygons $P(v_i)$, $i = 1, \dots, d$ around it such that the contacts between $P(u)$ and $P(v_i)$ have no bends except for at most three polygons $P(v_i)$ whose contacts to $P(u)$ have at most three bends in total:

Case 1: if $\ell > \sigma$, we draw $P(v_j)$ as a U-shaped polygon containing the left, lower and right boundary of the $(2\sigma - \ell) \times (\ell - \sigma)$ rectangle $P(u)$ and draw the rest of the polygons $P(v_i)$, $i \neq j$, as L- or $\perp$-shapes on the upper boundary of $P(u)$; see Figure 20a.

Case 2: If $\ell = \sigma$, we draw $P(v_j)$ as an L-shaped polygon containing two sides of $P(u)$ and wrap the other polygons $P(v_i)$, $i \neq j$, around the remaining free corner of $P(u)$; see Figure 20b.

Case 3: If $\ell < \sigma$, we can wrap the polygons $P(v_i)$ around $P(u)$ such that for at most three polygons $P(v_i)$ the contact between $P(v_i)$ and $P(u)$ has exactly one bend each and for the rest of the polygons it has no bends; see Figure 20c.

In all three cases we can draw the polygons $P(v_i)$ around $P(u)$ with at most eight bends each such that for each consecutive pair $P(v_i)$ and $P(v_{i+1})$ there is a U-shape of width at most $\varepsilon/2$ pointing outwards. Furthermore, the depth of such a U-shape can be increased arbitrarily; see arrows in Figure 20.

Consider a pair v_i, v_{i+1} of consecutive neighbors of u, $i = 1, \dots, m$. If u is the only common neighbor of v_i and v_{i+1}, the U-shape between the polygons $P(v_i)$ and $P(v_{i+1})$ remains empty. Otherwise, denote the second common neighbor by u_i. The removal of v_i and v_{i+1} disconnects G. Let V_i' be the vertex set of the component containing u_i and let G_i be the subgraph of G induced by $V_i' \cup \{v_i, v_{i+1}\} \subseteq V$. Then, $G_i \in \mathcal{O}$. We choose $v_i v_{i+1}$ as the reference edge and use the construction from the proof of Proposition 2. Thus, $G_i - \{v_i, v_{i+1}\}$ can be drawn inside the corresponding $\varepsilon/2$ wide U-shape with $\perp$-shaped polygons. This completes the construction of an edge-proportional rectilinear representation of G with complexity 8. $\square$

4   Conclusions

In this work, we have introduced the new notion of edge-proportional contact representations for edge-weighted planar graphs by encoding the edge weights as contact lengths. We have presented a constructive linear-time decision algorithm for the existence of edge-proportional rectangular duals (EPRDs) with four outer rectangles; see

Section 2.1.1. If arbitrarily many outer rectangles are allowed, it is NP-complete to decide whether an EPRD that forms a square exists; see Section 2.1.2. We have studied the problem of finding small rectangular duals for lower-bounded contact lengths. If the combinatorial structure (REL) of the dual is fixed, optimal duals can be constructed efficiently with existing tools; see Section 2.2. In Section 2.4, we have proved NP-completeness for the case of variable RELs. Also, deciding the existence of a dual under specified lower and upper bounds on the contact lengths is NP-complete, and a similar decision problem is NP-hard for lower-bounded contact lengths and upper-bounded rectangle areas; see Section 2.3. Furthermore, in Section 3 we have considered edge-proportional rectilinear representations (EPRRs) of internally triangulated plane graphs and have given tight bounds on the necessary and sufficient polygon complexity.

Open problems. We have shown that edge-proportional rectilinear representations that have bounded polygon complexity for an arbitrary edge-weight function exist only for a rather restricted class of graphs. To extend the proposed approach to a wider class of graphs, we have to allow contact lengths that deviate from the specified edge weights. One option is to interpret the weights as lower bounds on the edge lengths (thus one could highlight a subsegment of optimal length for each edge) and construct a dual that minimizes the error. As we have shown, this problem is NP-complete for rectangular duals. It can be solved exactly using integer linear programming, but so far no efficient approximation algorithms are known. One could also study this problem for rectilinear polygons of complexity 6 and above. Other possible trade-offs between the error and representability of wider classes of graphs are of interest, too.

We have shown that the number of edge-proportional rectangular duals of a graph is sometimes exponential, and deciding whether there exists an EPRD that forms a square is NP-complete. Is it also hard to decide whether an EPRD exists at all? Another open question is whether, for a PTP graph with a fixed REL, we can efficiently decide the existence of a rectangular dual with lower-bounded contact lengths and upper-bounded rectangle areas. This would answer the question whether the same decision problem for variable RELs is contained in $\mathcal{NP}$.

A very interesting challenge is to investigate the complexity of finding layouts with minimum size for the case that all contact lengths have lower bounds 1. We have shown NP-completeness for arbitrary edge weights only. In fact, the restriction to unit weights is equivalent to a long-standing open question posed by He [12]: Given a PTP graph G, find a rectangular dual of G that can be drawn with minimum area (or minimum perimeter, or minimum width) such that each contact has length at least 1.

References

[1] M. J. Alam, T. Biedl, S. Felsner, A. Gerasch, M. Kaufmann, and S. G. Kobourov. Linear-time algorithms for proportional contact graph representations. In *Proc. 22nd International Symposium on Algorithms and Computation (ISAAC'11)*, volume 7074 of *Lecture Notes Comput. Sci.*, pages 281–291. Springer-Verlag, 2011. `doi:10.1007/978-3-642-25591-5_30`.

[2] M. J. Alam, T. Biedl, S. Felsner, M. Kaufmann, S. G. Kobourov, and T. Ueckerdt. Computing cartograms with optimal complexity. In *Proc. 28th ACM Symposium on Computational Geometry (SoCG'12)*, pages 21–30. ACM, 2012. `doi:10.1145/2261250.2261254`.

[3] T. Biedl and B. Genc. Complexity of octagonal and rectangular cartograms. In *Proc. 17th Canadian Conference on Computational Geometry (CCCG'05)*, pages 117–120, 2005.

[4] T. Biedl and L. E. R. Velázquez. Orthogonal cartograms with few corners per face. In *Proc. 12th International Symposium on Algorithms and Data Structures (WADS'11)*, volume 6844 of *Lecture Notes Comput. Sci.*, pages 98–109. Springer-Verlag, 2011. `doi:10.1007/978-3-642-22300-6_9`.

[5] S. Cabello, E. D. Demaine, and G. Rote. Planar embeddings of graphs with specified edge lengths. *J. Graph Algorithms Appl.*, 11(1):259–276, 2007. `doi:10.7155/jgaa.00145`.

[6] J. Chalopin and D. Gonçalves. Every planar graph is the intersection graph of segments in the plane. In *Proc. 41st Annual ACM Symposium on Theory of Computing (STOC'09)*, pages 631–638. ACM, 2009. `doi:http://doi.acm.org/10.1145/1536414.1536500`.

[7] M. de Berg, E. Mumford, and B. Speckmann. On rectilinear duals for vertex-weighted plane graphs. *Discrete Mathematics*, 309(7):1794–1812, 2009. `doi:10.1016/j.disc.2007.12.087`.

[8] G. Di Battista, P. Eades, R. Tamassia, and I. G. Tollis. *Graph Drawing: Algorithms for the Visualization of Graphs*. Prentice Hall, 1999.

[9] P. Eades and N. C. Wormald. Fixed edge-length graph drawing is NP-hard. *Discrete Applied Mathematics*, 28(2):111–134, 1990. `doi:10.1016/0166-218X(90)90110-X`.

[10] D. Eppstein, E. Mumford, B. Speckmann, and K. Verbeek. Area-universal and constrained rectangular layouts. *SIAM J. Comput.*, 41(3):537–564, 2012. `doi:10.1137/110834032`.

[11] É. Fusy. Transversal structures on triangulations: A combinatorial study and straight-line drawings. *Discrete Math.*, 309(7):1870 – 1894, 2009. `doi:10.1016/j.disc.2007.12.093`.

[12] X. He. On finding the rectangular duals of planar triangular graphs. *SIAM J. Comput.*, 22(6):1218–1226, 1993. `doi:10.1137/0222072`.

[13] P. Hliněný and J. Kratochvíl. Representing graphs by disks and balls (a survey of recognition-complexity results). *Discrete Mathematics*, 229(1–3):101–124, 2001. `doi:10.1016/S0012-365X(00)00204-1`.

[14] Kawaguchi and H. Nagamochi. Orthogonal drawings for plane graphs with specified face areas. In *Proc. Theory and Application of Models of Computation (TAMC'07)*, volume 4484 of *Lecture Notes Comput. Sci.*, pages 584–594. Springer-Verlag, 2007. `doi:10.1007/978-3-540-72504-6_53`.

[15] D. E. Knuth and A. Raghunathan. The problem of compatible representatives. *SIAM J. Discrete Math.*, 5(3):422–427, 1992. `doi:10.1137/0405033`.

[16] P. Koebe. Kontaktprobleme der konformen Abbildung. *Ber. Sächs. Akad. Wiss. Leipzig, Math.-Phys. Klasse*, 88:141–164, 1936.

[17] K. A. Kozminski and E. Kinnen. Rectangular dualization and rectangular dissections. *IEEE Trans. Circuits and Systems*, 35(11):1401–1416, 1988. `doi:10.1109/31.14464`.

[18] S. M. Leinwand and Y.-T. Lai. An algorithm for building rectangular floor-plans. In *Proc. 21st Design Automation Conference*, pages 663–664, 1984.

[19] C.-C. Liao, H.-I. Lu, and H.-C. Yen. Compact floor-planning via orderly spanning trees. *J. Algorithms*, 48(2):441–451, 2003. `doi:10.1016/S0196-6774(03)00057-9`.

[20] D. Lichtenstein. Planar formulae and their uses. *SIAM J. Comput.*, 11(2):329–343, 1982. `doi:10.1137/0211025`.

[21] T. A. McKee and F. R. McMorris. *Topics in Intersection Graph Theory*. SIAM, 1999.

[22] T. Nishizeki and M. S. Rahman. Rectangular drawing algorithms. In R. Tamassia, editor, *Handbook of Graph Drawing and Visualization*, chapter 13. CRC Press, 2012. To appear.

[23] R. Prutkin. Kontaktrepräsentationen von kantengewichteten planaren Graphen. Diploma thesis, Karlsruhe Institute of Technology, 2012.

[24] R. Tamassia. Drawing algorithms for planar st-graphs. *Australasian J. Combinatorics*, 2:217–235, 1990.

[25] M. van Kreveld and B. Speckmann. On rectangular cartograms. *Comput. Geom. Theory Appl.*, 37(3):175–187, 2007. `doi:10.1016/j.comgeo.2006.06.002`.

[26] K.-H. Yeap and M. Sarrafzadeh. Floor-planning by graph dualization: 2-concave rectilinear modules. *SIAM J. Comput.*, 22(3):500–526, 1993. `doi:10.1137/0222035`.