

Morphing Rectangular Duals

Steven Chaplick¹, Philipp Kindermann², Jonathan Klawitter^{3,4},
Ignaz Rutter⁵, and Alexander Wolff³

¹ Maastricht University, The Netherlands

² Universität Trier, Germany

³ Universität Würzburg, Germany

⁴ University of Auckland, New Zealand

⁵ Universität Passau, Germany

Abstract. A rectangular dual of a plane graph G is a contact representations of G by interior-disjoint axis-aligned rectangles such that (i) no four rectangles share a point and (ii) the union of all rectangles is a rectangle. A rectangular dual gives rise to a regular edge labeling (REL), which captures the orientations of the rectangle contacts.

We study the problem of morphing between two rectangular duals of the same plane graph. If we require that, at any time throughout the morph, there is a rectangular dual, then a morph exists only if the two rectangular duals realize the same REL. Therefore, we allow intermediate contact representations of non-rectangular polygons of constant complexity. Given an n -vertex plane graph, we show how to compute in $\mathcal{O}(n^3)$ time a piecewise linear morph that consists of $\mathcal{O}(n^2)$ linear morphing steps.

Keywords: morphing · rectangular dual · regular edge labeling · lattice

1 Introduction

A *morph* between two representations (e.g., drawings) of the same graph G is a continuous transformation from one representation to the other. Preferably, a morph should preserve the user’s “mental map”, which means that, throughout the transformation, as little as necessary is changed to go from the source to target representation and that their properties are maintained [32]. For example, during a morph between two planar drawings, each intermediate drawing should also be planar. A *linear morph* moves each point along a straight-line segment at constant speed, where different points may have different speeds or may remain stationary. Note that a linear morph is fully defined by the source and target representation. A *piecewise linear morph* consists of a sequence of linear morphs, each of which is called a *step*.

Morphs are well studied for planar drawings. For example, it is known that piecewise linear planar morphs always exist between planar straight-line drawings [11] and that, for an n -vertex planar graph, $\mathcal{O}(n)$ steps suffice [1], which is worst-case optimal. Further research on morphs includes, among others, the study of morphs of convex drawings [3, 28], of orthogonal drawings [6, 35], on different surfaces [12, 26], and in higher dimensions [4].




Fig. 1: A rectangular dual \mathcal{R} for the graph G ; the REL \mathcal{L} induced by \mathcal{R} .

Less attention has been given to morphs of alternative representations of graphs such as intersection and contact representations. A *geometric intersection representation* of a graph G is a mapping \mathcal{R} that assigns to each vertex w of G a geometric object $\mathcal{R}(w)$ such that two vertices u and v are adjacent in G if and only if $\mathcal{R}(u)$ and $\mathcal{R}(v)$ intersect. In a *contact representation* we further require that, for any two vertices u and v , the objects $\mathcal{R}(u)$ and $\mathcal{R}(v)$ have disjoint interiors. Classic examples are interval graphs [9], where the objects are intervals of \mathbb{R} , or coin graphs [27], where the objects are interior-disjoint disks in the plane. Recently, Angelini et al. [2] studied morphs of right-triangle contact representations of planar graphs. They showed that one can test efficiently whether a morph exists (in which case a quadratic number of steps suffice). In this paper, we investigate morphs between contact representations of rectangles.

Rectangular duals. A *rectangular dual* of a graph G is a contact representation \mathcal{R} of G by axis-aligned rectangles such that (i) no four rectangles share a point and (ii) the union of all rectangles is a rectangle; see Fig. 1. Note that G may admit a rectangular dual only if it is planar and internally triangulated. Furthermore, a rectangular dual can always be augmented with four additional rectangles (one on each side) so that only these four rectangles touch the outer face of the representation. It is customary that the four corresponding vertices on the outer face of G are denoted by v_S , v_W , v_N , and v_E , and to require that $\mathcal{R}(v_S)$ is bottommost, $\mathcal{R}(v_W)$ is leftmost, $\mathcal{R}(v_N)$ is topmost, and $\mathcal{R}(v_E)$ is rightmost; see Fig. 1. The corresponding vertices are *outer*; the remaining ones are *inner*. Similarly, the four edges between the outer vertices are *outer*; the others are *inner*. A plane internally-triangulated graph has a representation with only four rectangles touching the outer face if and only if its outer face is a 4-cycle and it has no *separating triangle*, that is, a triangle whose removal disconnects the graph [28]. Such a graph is called a *properly-triangulated planar (PTP) graph*. For such a graph, a rectangular dual can be computed in linear time [25].

Historically, rectangular duals have been studied due to their applications in architecture [34], VLSI floor-planning [29, 37], and cartography [21]. Morphs between rectangular duals are of interest, e.g., due to their relation to rectangular cartograms. Rectangular cartograms were introduced in 1934 [33] and combine statistical and geographical information in thematic maps, where geographic




Fig. 2: Edge order at the four outer vertices and at an inner vertex in a REL.

regions are represented as rectangles and scaled in proportion to some statistic. There has been a lot of work on efficiently computing rectangular cartograms [10, 23, 36], see also the recent survey [31]. A morph between rectangular cartograms can visualize different data sets. Florisson et al. [18] implemented a method to construct rectangular cartograms by first extending the given map with “sea tiles” to obtain a rectangular dual, and then using a heuristic that moves maximal line segments until the area of the rectangles gets closer to the given data. They also used their heuristic to morph between two rectangular cartograms, but did not discuss when exactly this works and with what time complexity.

Regular edge labelings. A combinatorial view of a rectangular dual of a graph G can be described by a coloring and orientation of the edges of G [25]. This is similar to how so-called Schnyder woods describe contact representations of planar graphs by triangles [14]. More precisely, a rectangular dual \mathcal{R} gives rise to a 2-coloring and an orientation of the inner edges of G as follows. We color an edge $\{u, v\}$ blue if the contact segment between $\mathcal{R}(u)$ and $\mathcal{R}(v)$ is a horizontal line segment, and we color it red otherwise. We orient a blue (red) edge $\{u, v\}$ as uv if $\mathcal{R}(u)$ lies below (resp. left of) $\mathcal{R}(v)$; see Fig. 1. The resulting coloring and orientation has the following properties (Fig. 2):

- (1) For each outer vertex v_S , v_W , v_N , and v_E , the incident inner edges are blue outgoing, red outgoing, blue incoming, and red incoming, respectively.
- (2) For each inner vertex, the incident edges form four clockwise (*cw*) ordered non-empty blocks: blue incoming, red incoming, blue outgoing, red outgoing.

A coloring and orientation with these properties is called a *regular edge labeling (REL)* or *transversal structure*. We let $\mathcal{L} = (L_1, L_2)$ denote a REL, where L_1 is the set of blue edges and L_2 is the set of red edges. Let $L_1(G)$ and $L_2(G)$ denote the two subgraphs of G induced by L_1 and L_2 , respectively. Note that both $L_1(G)$ and $L_2(G)$ are st-graphs, that is, directed acyclic graphs with exactly one source and exactly one sink. Kant and He [25] introduced RELs as intermediate objects when constructing a rectangular dual of a PTP graph. It is well known that every PTP graph admits a REL and thus a rectangular dual [22, 25]. A rectangular dual \mathcal{R} *realizes* a REL \mathcal{L} if the REL induced by \mathcal{R} is \mathcal{L} .

We define the *interior* of a cycle to be the set of vertices and edges enclosed by, but not on the cycle. A 4-cycle is *separating* if there are other vertices both in its interior and in its exterior. A separating 4-cycle is *nontrivial* if its interior contains more than one vertex; otherwise it is *trivial*. We call non-separating 4-cycles also *empty 4-cycles*. (An empty 4-cycle contains exactly one edge.)




Fig. 3: Clockwise and counterclockwise rotations between RELs that recolor and reorient the edges inside an alternating 4-cycle $\langle a, b, c, d \rangle$ that is (a) empty or (b) separating.

If the edges of a cycle C alternate between red and blue, we say that C is *alternating*. We can move between different RELs of a PTP graph G by swapping the colors and reorienting the edges inside an alternating 4-cycle, see Fig. 3. This operation, which we call a *rotation* and which we define formally in Appendix A, connects all RELs of G . In fact, the RELs of G form a distributive lattice [19,20]. A 4-cycle C of G is called *rotatable* if it is alternating for at least one REL of G .

Important related work. Other combinatorial structures of graph representations also form lattices; see the work by Felsner and colleagues [16,17]. In the context of morphs, Barrera-Cruz et al. [5] exploited the lattice structure of Schnyder woods of a plane triangulation to obtain piecewise linear morphs between planar straight-line drawings. While their morphs require $\mathcal{O}(n^2)$ steps (compared to the optimum of $\mathcal{O}(n)$), they have the advantage that they are “visually pleasing” and that they maintain a quadratic-size drawing area between any two steps. To this end, Barrera-Cruz et al. showed that there is a path in the lattice of length $\mathcal{O}(n^2)$ between any two Schnyder woods. We show an analogous result for RELs. In order to morph between right-triangle contact representations, Angelini et al. [2] leveraged the lattice structure of Schnyder woods. They showed that if no separating triangle has to be flipped (a flip is a step in the lattice) between the source and the target Schnyder wood, then a morph with $\mathcal{O}(n^2)$ steps exists (else, no morph exists that uses right-triangle contact representations throughout).

Contribution. We consider piecewise linear morphs between two rectangular duals \mathcal{R} and \mathcal{R}' of the same PTP graph G . If \mathcal{R} and \mathcal{R}' realize the same REL, then a single step suffices, but if \mathcal{R} and \mathcal{R}' realize distinct RELs of G , then no rectangular-dual preserving morph exists (Sect. 2.1). Therefore, we propose a new type of morph where intermediate drawings are contact representations of G using convex polygons with up to five corners (Sect. 2.2). We show how to construct such a *relaxed* morph as a sequence of $\mathcal{O}(n^2)$ steps that implement moves in the lattice of RELs of G (Sect. 2.3). To this end, we make use of the following two results on paths in this lattice (which we prove in Appendix A).

Proposition 1 (\star). *Given an n -vertex PTP graph G , the lattice of RELs of G has diameter $\mathcal{O}(n^2)$.*

Proposition 2 (★). *Let G be an n -vertex PTP graph with RELs \mathcal{L} and \mathcal{L}' . In the lattice of RELs of G , a shortest \mathcal{L} - \mathcal{L}' path can be computed in $\mathcal{O}(n^3)$ time.*

We ensure that between any two morphing steps, our drawings remain on a quadratic-size section of the integer grid – like those of Barrera-Cruz et al. [5]. In order to evaluate the intermediate representations when our drawings are not on the integer grid, we use the measure *feature resolution* [24], that is, the ratio of the length of the longest segment over the shortest distance between two vertices or between a vertex and a non-incident segment. We show that the feature resolution in *any* intermediate drawing is bounded by $\mathcal{O}(n)$.

Finally, we investigate executing rotations in parallel; see Sect. 3. As a result, we can morph between any pair $(\mathcal{R}, \mathcal{R}')$ of rectangular duals of the given graph using $\mathcal{O}(1)$ times the minimum number of steps needed to get from \mathcal{R} to \mathcal{R}' ; however, our polygons have up to eight corners.

We prove every statement marked with a (clickable) “★” in the appendix.

2 Morphing between Rectangular Duals

This section concerns morphs between two rectangular duals \mathcal{R} and \mathcal{R}' of a PTP graph G that realize the same REL, adjacent RELs, and finally any two RELs.

2.1 Morphing Between Rectangular Duals Realizing the Same REL

Theorem 3. *For a PTP graph G with rectangular duals \mathcal{R} and \mathcal{R}' , (i) if \mathcal{R} and \mathcal{R}' realize the same REL, then there is a linear morph between them; (ii) otherwise, there is no morph between them (not even a non-linear one).*

Proof. Biedl et al. [6] studied morphs of orthogonal drawings. They showed that a single (planarity-preserving) linear morph suffices if all faces are rectangular and all edges are parallel in the two drawings, that is, any edge is either vertical in both drawings or horizontal in both drawings. We can apply this result to two rectangular duals \mathcal{R} and \mathcal{R}' precisely when they realize the same REL. A linear morph between them changes the x-coordinates of vertical line segments and the y-coordinates of horizontal line segments but does not change their relative order.

Now assume that \mathcal{R} and \mathcal{R}' realize different RELs \mathcal{L} and \mathcal{L}' of G , respectively. Then w.l.o.g. some contact segment s changes from being horizontal in \mathcal{R} to being vertical in \mathcal{R}' . Since s must always be horizontal or vertical, it has to collapse to a point and then extend to a segment again. When s collapses, the intermediate representation is not a rectangular dual of G since four rectangles meet at a single point. Even worse, if a separating alternating cycle is rotated, then its interior contracts to a point, vanishes, and reappears rotated by 90° . \square




Fig. 4: A linear morph rotating the separating 4-cycle $\langle a, b, c, d \rangle$ emulates the rotation in the corresponding REL; see Fig. 3. The interior of the 4-cycle turns by 90° without changing its shape, while the outer contact segments move horizontally and vertically.

2.2 Morphing Between Rectangular Duals with Adjacent RELs

Let \mathcal{R} and \mathcal{R}' now realize different RELs \mathcal{L} and \mathcal{L}' of G , respectively. By Theorem 3, any continuous transformation between \mathcal{R} and \mathcal{R}' requires intermediate representations that are not rectangular duals of G , i.e., a morph in the traditional sense is not possible. We relax the conditions on a morph such that, in an intermediate contact representation of G , vertices can be represented by convex polygons of constant complexity – in this section, by 5-gons. However, we still require that these polygons form a tiling of the bounding rectangle of the representation. We call a transformation with this property a *relaxed morph*. (When we talk about linear morphing steps, we omit the adjective “relaxed”.) The following statement describes relaxed morphs when \mathcal{L} and \mathcal{L}' are adjacent, that is, $(\mathcal{L}', \mathcal{L})$ is an edge in the lattice of RELs of G .

Proposition 4. *Let \mathcal{R} and \mathcal{R}' be two rectangular duals of an n -vertex PTP graph G realizing two adjacent RELs \mathcal{L} and \mathcal{L}' of G , respectively. Then, we can compute in $\mathcal{O}(n)$ time a 3-step relaxed morph between \mathcal{R} and \mathcal{R}' . If \mathcal{R} and \mathcal{R}' have an area of at most $n \times n$ and feature resolution in $\mathcal{O}(n)$, then so has each representation throughout the morph.*

We assume w.l.o.g. that \mathcal{L}' can be obtained from \mathcal{L} by a cw rotation of an alternating 4-cycle C . The idea is to rotate the interior of the 4-cycle while simultaneously moving the contact segments that form the edges of C ; see Figs. 4 and 5. To ensure that, except for the vertices of C all regions remain rectangles and that moving the contact segments of C does not change any adjacencies, the representation needs to satisfy certain requirements. Therefore, our relaxed morph from \mathcal{R} to \mathcal{R}' consists of three steps. First a *preparatory morph* from \mathcal{R} to a rectangular dual \mathcal{R}_1 with REL \mathcal{L} for which C satisfies conditions stated below; second a *main morph* which transforms \mathcal{R}_1 to a rectangular dual \mathcal{R}_2 whose REL is \mathcal{L}' , and third a *clean-up morph* that transforms \mathcal{R}_2 into \mathcal{R}' .

We first describe the main morph $\langle \mathcal{R}_1, \mathcal{R}_2 \rangle$ in detail, as this allows us to also infer the conditions under which it can be executed successfully. Then we describe the preparatory morph $\langle \mathcal{R}, \mathcal{R}_1 \rangle$, whose sole purpose is to ensure the conditions that are required for the main morph.

Main morph $\langle \mathcal{R}_1, \mathcal{R}_2 \rangle$ to rotate C . Let a, b, c , and d be the vertices of C in cw order where a is the vertex with an outgoing red and outgoing blue edge in C , i.e., it corresponds to the bottom-left rectangle of C .




Fig. 5: A linear morph that rotates the inner contact segment of an empty 4-cycle.

Assume for now that C is separating. We have the following requirements for R_1 , which become apparent shortly.

(P1) The rectangle I_C bounding the interior of C is a square.

Next, we consider the four maximal segments of \mathcal{R} that contain one of the four borders of I_C , which we call *border segments*. Let s be the upper border segment of I_C and suppose its right endpoint lies on the left side of a rectangle $\mathcal{R}(x)$. Let S be the part in the horizontal strip defined by I_C that starts at I_C and ends at $\mathcal{R}(x)$.

(P2) The only horizontal segments that intersect S are border segments of I_C ; see Fig. 6c.

We define (P2) for the other three border segments of I_C analogously. Next, assume that C is empty. Then the rectangle I_C degenerates to a segment s , and we assume w.l.o.g. that s is horizontal. Now I_C still has two vertical border segments, but the two horizontal border segments share the segment s . Let s have again its right endpoint on the left side of $\mathcal{R}(x)$. Let S be the rectangular area of height 1 directly below s that starts at b and ends at $\mathcal{R}(x)$. We have the following requirement if C is empty.

(P2') The only horizontal segment that intersects S is s ; see Fig. 7c.

There is no requirement for the left side of s and the left vertical border segment of I_C . The requirements for the right vertical border segment is (P2).

We now describe the main morph for the case that C is a separating 4-cycle. In this case, the interior of C forms a square in \mathcal{R}_1 by (P1). Recall that the rotation of C from \mathcal{L} to \mathcal{L}' turns the interior of C by 90° . During the morph, we move each corner of I_C to the coordinates of the corner that follows in cw order around I_C in \mathcal{R}_1 ; see Fig. 4. All other points in I_C are expressed as a convex combination of the corners of I_C and then move according to the movement of the corners. Furthermore, we move all points on the left border segment of I_C that are outside the boundary of I_C horizontally to the x-coordinate of the right side of I_C . We move the points on other border segments of I_C analogously.

This describes a single linear morph that results in a rectangular dual R_2 that realizes the REL \mathcal{L}' . Since I_C starts out as a square by (P1), throughout the morph, I_C remains a square, and by similarity all rectangles inside I_C remain rectangles. The rectangles a , b , c , and d become convex 5-gons. Furthermore, since outside I_C the horizontal border segments of I_C move an area that contains

no other horizontal segments by (P2), no contact along a vertical segment arises or vanishes. Analogously, for the vertical border segments, no contact along a horizontal segment arises or vanishes. Hence, we maintain the same adjacencies.

Note that, if I_C would not be a square, then its corners would move at different speeds and I_C would deform to a rhombus where the inner angles are not 90° , and so would all the rectangles inside I_C .

Next, consider the case that C is an empty 4-cycle. Recall that in this case, the rectangle I_C degenerates to a segment s and we assume that s is horizontal. Note that \mathcal{R}' has a vertical contact between a and c , since we assume a cw rotation from \mathcal{L} to \mathcal{L}' . We then move the right endpoint of s vertically down by 1 and horizontally to the x-coordinate of the left endpoint of s ; see Fig. 5. We also move all points on the border segments that contain the right endpoint of s accordingly. The rectangles a and c become convex 5-gons. Furthermore, since outside I_C only the horizontal border segments of I_C lie inside the area of height 1 below s by (P2'), no contact along a vertical segment arises or vanishes. Analogously, due to condition (P2) for the vertical border segments, no contact along a horizontal segment arises or vanishes. Hence, maintaining the same adjacencies.

To show that the feature resolution remains in $\mathcal{O}(n)$, note that both \mathcal{R}_1 and \mathcal{R}_2 are drawn on a $n \times n$ grid. Furthermore, the rectangles inside I_C are scaled during the morph, but since I_C is a square, the whole area inside I_C is scaled by at most $\sqrt{2}$. The distances outside I_C cannot become smaller than 1.

Lemma 5. *Let \mathcal{R}_1 and \mathcal{R}_2 be two rectangular duals of an n -vertex PTP graph G realizing two adjacent RELs \mathcal{L} and \mathcal{L}' of G , respectively, such that \mathcal{R}_1 satisfies (P1) and (P2) (or (P1) and (P2')). Then, we can compute in $\mathcal{O}(n)$ time a relaxed morph between \mathcal{R}_1 and \mathcal{R}_2 . If \mathcal{R}_1 and \mathcal{R}_2 have an area of at most $n \times n$ and feature resolution in $\mathcal{O}(n)$, then so has each representation throughout the morph.*

Preparatory morph $\langle \mathcal{R}, \mathcal{R}_1 \rangle$. We consider again the case where C is separating first. To obtain \mathcal{R}_1 from \mathcal{R} , we extend G to an auxiliary graph \hat{G} that is almost a PTP graph but that contains empty chordless 4-cycles (which are represented by four rectangles touching in a single point). For \hat{G} , we compute an auxiliary REL $\hat{\mathcal{L}}$ where the empty chordless 4-cycles of \hat{G} are colored alternately. We then use the second step of the linear-time algorithm by Kant and He [25] to compute an (almost) rectangular dual $\hat{\mathcal{R}}_1$ of \hat{G} that realizes $\hat{\mathcal{L}}$. By reversing the changes applied to G to obtain \hat{G} , we derive \mathcal{R}_1 from $\hat{\mathcal{R}}_1$. We explain in Appendix B the algorithm by Kant and He and why it also works for \hat{G} .

We start with the changes to ensure (P2) for the upper border segment s of I_C ; it works analogously for the other border segments. Let s end to the right again at $\mathcal{R}(x)$. Let P_d be the leftmost path in $L_2(G)$ from d to x . Let $y_1(R)$ and $y_2(R)$ denote the lower and upper y-coordinate of a rectangle R , respectively. Note that (P2) holds if for each vertex v on P_d we have $y_1(\mathcal{R}_1(v)) < y_2(\mathcal{R}_1(a))$. Therefore, from G to \hat{G} , we duplicate P_d by splitting each vertex v on $P_d \setminus \{x\}$ into two vertices v_1 and v_2 ; see Fig. 6. We then connect v_1 and v_2 with a blue edge. Let y be the successor of d on P_d . We assign the edges cw between (and including) dy and ad to d_1 , and the edges cw between ad and dy to d_2 . If $x = y$,




Fig. 6: We compute \mathcal{R}_1 via an auxiliary rectangular dual $\hat{\mathcal{R}}_1$ and an auxiliary REL $\hat{\mathcal{L}}$.


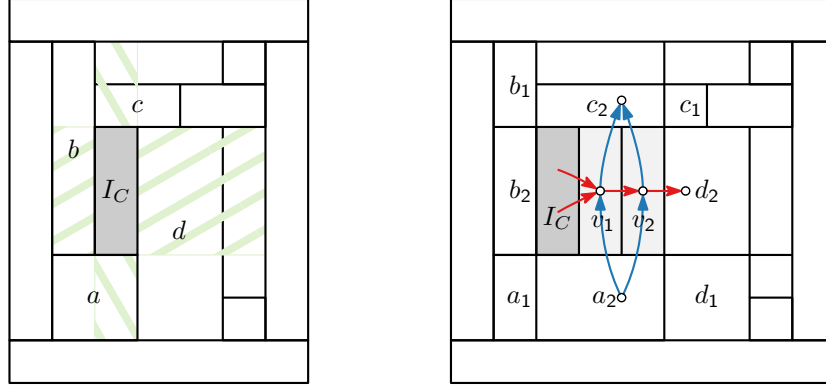


Fig. 7: Preparatory morph analogous Fig. 6 for the case when C is empty.

the edge dy is assigned to both d_1 and d_2 ; otherwise we replace dy with d_1y_1 and d_2y_2 . For all other vertices v on $P_d \setminus \{d, x\}$, let u be the predecessor and let w be the successor of v on P_d . We assign the edges cw between vw and uv to v_1 , and the edges cw between uv and vw to v_2 ; furthermore, we add the edges u_1v_1 , v_1w_1 , u_2v_2 , and v_2w_2 . As a result, there is a path from a to x in $\hat{\mathcal{R}}_1$ through the “upper” copies of the vertices in $P_d \setminus \{x\}$, and the bottom side of their corresponding rectangles are aligned. Hence, $y_1(\hat{\mathcal{R}}_1(v_1)) < y_1(\hat{\mathcal{R}}_1(v_2)) = y_2(\hat{\mathcal{R}}_1(a))$ for every $v \in P_d \setminus \{x\}$, and $y_1(\hat{\mathcal{R}}_1(x)) < y_2(\hat{\mathcal{R}}_1(a))$. We obtain for each v on $P_d \setminus \{x\}$ the rectangle $\mathcal{R}_1(v)$ by merging $\mathcal{R}_1(v_1)$ and $\mathcal{R}_1(v_2)$. This works analogously if C is empty; see Fig. 7.

Next, we describe how to ensure (P1) in \mathcal{R}_1 , i.e., that the interior I_C of C is a square. Let w_C and h_C be the minimum width and height, respectively, of a rectangular dual of I_C . These values can be computed in $\mathcal{O}(I_C)$ time. Note that because of (P2), the algorithm by Kant and He [25] will draw I_C with minimum width and height in \mathcal{R}_1 : the algorithm draws every horizontal line segment as low as possible, and because of (P2) there is no horizontal line segment to the right of I_C that forces the upper boundary segment of I_C to be higher; a symmetric argument applies to the left boundary segment of I_C . Hence, if $w_C = h_C$, then



(a) If I_C is not a square in \mathcal{R}_1 with (P2), ... (b) ... then we extend $\hat{\mathcal{R}}_1$ and $\hat{\mathcal{L}}$ with dummy vertices v_i inside $\langle a, b, c, d \rangle$.

Fig. 8: To ensure (P1), i.e., that I_C is a square in \mathcal{R}_1 , we extend $\hat{\mathcal{R}}_1$ and $\hat{\mathcal{L}}$ further.

no further changes to \hat{G} are required. Otherwise, if w.l.o.g. $w_C < h_C$, we add $h_C - w_C$ many buffer rectangles between I_C and d as follows; see Fig. 8. Let $\Delta = h_C - w_C$. From G to \hat{G} , we add vertices v_1, \dots, v_Δ with a red path through them and, for $i \in \{1, \dots, \Delta\}$, we add the blue edges av_i and $v_i c$. All incoming red edges of d in G from the interior of C become incoming red edges of v_1 , and we add a red edge (v_Δ, d_2) . In \hat{G} , the minimum width and height of I_C are now the same and I_C is drawn as a square in $\hat{\mathcal{R}}_1$. To obtain \mathcal{R}_1 from $\hat{\mathcal{R}}_1$, we remove the buffer rectangles and stretch all right-most rectangles of I_C to $y_1(\mathcal{R}_1(d))$.

Concerning the running time, note that we can both find and split the paths for (P2) and add the extra vertices for (P1) in $\mathcal{O}(n)$ time. Since \hat{G} and $\hat{\mathcal{L}}$ have a size in $\mathcal{O}(n)$, the algorithm by Kant and He [25] also runs in $\mathcal{O}(n)$ time.

Finally, we show that the area of \mathcal{R}_1 is bounded by $n \times n$. Observe that each triangle in G corresponds to a T-junction in \mathcal{R} and thus to an endpoint of a maximal line segment. There are $2n - 4$ triangles in G and thus $n - 2$ inner maximal line segments besides the four outer ones. The algorithm by Kant and He [25] ensures that each x- and y-coordinate inside a rectangular dual contains a horizontal or vertical line segment, respectively. Note that \mathcal{R}_1 contains exactly Δ more maximal line segments than \mathcal{R} . These were added inside C if in I_C the number of horizontal and vertical maximal line segments differed by at least Δ . Hence, \mathcal{R}_1 contains at most $n - 2$ vertical and at most $n - 2$ horizontal inner maximal line segments. Thus, the area of \mathcal{R}_1 is bounded by $n \times n$. Lastly, note that $\hat{\mathcal{R}}_1$ and \mathcal{R}_1 have the same size. Furthermore, we move points only away from each other, so the feature resolution remains in $\mathcal{O}(n)$.

Lemma 6. *Let \mathcal{R} be a rectangular dual of an n -vertex PTP graph G realizing a REL \mathcal{L} of G . Let C be an alternating separating 4-cycle in \mathcal{L} . Then, we can compute in $\mathcal{O}(n)$ time a rectangular dual \mathcal{R}_1 of G realizing \mathcal{L} that satisfies the*

requirements (P1) and (P2). If \mathcal{R} has an area of at most $n \times n$ and feature resolution $\mathcal{O}(n)$, then so has \mathcal{R}_1 and each representation throughout the morph.

To prove Lemma 6, we do not use zig-zag moves, which were introduced for morphing orthogonal drawings [6, 35], since then we would not be able to bound the area by $n \times n$ throughout the morph. In order to keep a bound of $\mathcal{O}(n) \times \mathcal{O}(n)$, it seems that we would need a re-compactification step after each zig-zag move. Therefore, we keep the modifications in our morph as local as possible.

Let us now consider the morph $\langle \mathcal{R}_1, \mathcal{R}_2 \rangle$ again. Since \mathcal{R}_1 now satisfies (P1) and (P2), only the inside I_C of C , the four rectangles of C , and the border segments of I_C move. The target positions of these can be computed in $\mathcal{O}(n)$ time. The linear morph is then defined fully by the start and target positions. Furthermore, \mathcal{R}_2 and all intermediate representations have the same area as \mathcal{R}_1 .

Proof (of Prop. 4). By Theorem 3 and Lemma 5, we can get from \mathcal{R} via \mathcal{R}_1 and \mathcal{R}_2 to \mathcal{R}' using three steps. The claims on the running time and the area follow from Theorem 3, Lemmas 5 and 6, and the observations above. \square

2.3 Morphing Between Rectangular Duals

Combining results from the previous sections, we can now prove our main result.

Theorem 7. *Let G be an n -vertex PTP graph with rectangular duals \mathcal{R} and \mathcal{R}' . We can find in $\mathcal{O}(n^3)$ time a relaxed morph between \mathcal{R} and \mathcal{R}' with $\mathcal{O}(n^2)$ steps that executes the minimum number of rotations. If \mathcal{R} and \mathcal{R}' have an area of at most $n \times n$ and feature resolution in $\mathcal{O}(n)$, then so does each representation throughout the morph.*

Proof. Let \mathcal{L} and \mathcal{L}' be the RELs realized by \mathcal{R} and \mathcal{R}' , respectively. By Prop. 2 a shortest path between \mathcal{L} and \mathcal{L}' in the lattice of RELs of G can be computed in $\mathcal{O}(n^3)$ time, and its length is $\mathcal{O}(n^2)$ by Prop. 1. For each rotation along this path, we construct a relaxed morph with a constant number of steps in $\mathcal{O}(n)$ time by Prop. 4. The area and feature resolution also follow from Prop. 4. \square

3 Morphing with Parallel Rotations

We now show how to reduce the number of morphing steps by executing rotations in parallel. We assume that all separating 4-cycles in our PTP graph G are trivial.

Consider two cw rotatable separating 4-cycles C and C' that share a maximal horizontal line segment s as border segment; see Fig. 9. If C contains the left endpoint of s , a rotation of C would move s downwards while a rotation of C' would move s upwards. Therefore, such a morph skews angles such that they are not multiples of 90° even at vertices that are not incident to the interior of C or C' . To avoid such morphs, we say that C and C' are *conflicting*. For a set of cw rotatable separating 4-cycles \mathcal{C} for \mathcal{R} , this gives rise to a *conflict graph* $K(\mathcal{C})$ with vertex set \mathcal{C} . Note that a separating 4-cycle can be in conflict with at most four other separating 4-cycles. Therefore, $K(\mathcal{C})$ has maximum degree four.




Fig. 9: Two conflicting separating 4-cycles that share the interior segment s .




Fig. 10: When we rotate four cw alternating 4-cycles that share a vertex v in a single relaxed morph from \mathcal{R} to \mathcal{R}' , then v is temporarily represented by a convex 8-gon.

Next, consider a separating 4-cycle C that shares a maximal horizontal line segment s with an empty 4-cycle C' ; see Fig. 10. In this case, we can rotate and translate the inner contact segment of C' downwards, which allows us to simultaneously rotate C and C' without creating unnecessary skewed angles. Also note that two cw rotatable empty 4-cycles may only overlap with one edge but may not contain an edge of the other. Hence, they are not conflicting.

To rotate a set \mathcal{C} of alternating 4-cycles using $\mathcal{O}(1)$ steps, we divide \mathcal{C} into color classes based on $K(\mathcal{C})$ and rotate one color class at a time.

Proposition 8 (\star). *Let \mathcal{R} be a rectangular dual of a PTP graph G with REL \mathcal{L} whose separating 4-cycles are all trivial. Let \mathcal{C} be a set of alternating 4-cycles of \mathcal{R} . Let \mathcal{L}' be the REL obtained from \mathcal{L} by executing all rotations in \mathcal{C} . There exists a relaxed morph with $\mathcal{O}(1)$ steps from \mathcal{R} to a rectangular dual \mathcal{R}' realizing \mathcal{L}' . The morph can be computed in linear time.*

Note that there exist rectangular duals with a linear number of alternating 4-cycles – extend Fig. 10 into a grid structure. Hence, parallelization can reduce the number of morphing steps by a linear factor. Even more, using Prop. 8, we obtain the following approximation result.

Theorem 9 (\star). *Let G be a PTP graph whose separating 4-cycles are all trivial. Let \mathcal{R} and \mathcal{R}' be two rectangular duals of G , and let OPT be the minimum number of steps in any relaxed morph between \mathcal{R} and \mathcal{R}' . Then we can construct in cubic time a relaxed morph consisting of $\mathcal{O}(\text{OPT})$ steps.*

4 Concluding Remarks

In the parallelization step, we considered only PTP graphs whose separating 4-cycles are trivial. It remains open how to parallelize rotations for RELs of PTP

graphs with nontrivial separating 4-cycles, in particular, to construct morphs that execute rotations of nested 4-cycles in parallel. It would also be interesting to guarantee area bounds for morphs with parallel rotations.

During our relaxed morphs, we allow rectangles to temporarily turn into convex 5-gons (with four edges axis-aligned). Alternatively, one could insist that the intermediate objects remain ortho-polygons. This would require up to six vertices per shape and would force not only the outer rectangles in Fig. 4 to change their shape, but also the rectangles in the interior. We find our approach more natural.

References

1. S. Alamdari, P. Angelini, F. Barrera-Cruz, T. M. Chan, G. Da Lozzo, G. Di Battista, F. Frati, P. Haxell, A. Lubiw, M. Patrignani, V. Roselli, S. Singla, and B. T. Wilkinson. How to morph planar graph drawings. *SIAM J. Comput.*, 46(2):824–852, 2017. doi:10.1137/16M1069171.
2. P. Angelini, S. Chaplick, S. Cornelsen, G. Da Lozzo, and V. Roselli. Morphing contact representations of graphs. In G. Barequet and Y. Wang, editors, *Symp. Comput. Geom. (SoCG)*, volume 129 of *LIPIcs*, pages 10:1–10:16. Schloss Dagstuhl – LZI, 2019. doi:10.4230/LIPIcs.SoCG.2019.10.
3. P. Angelini, G. Da Lozzo, F. Frati, A. Lubiw, M. Patrignani, and V. Roselli. Optimal morphs of convex drawings. In L. Arge and J. Pach, editors, *Symp. Comput. Geom. (SoCG)*, volume 34 of *LIPIcs*, pages 126–140. Schloss Dagstuhl – LZI, 2015. doi:10.4230/LIPIcs.SoCG.2015.126.
4. E. Arseneva, P. Bose, P. Cano, A. D’Angelo, V. Dujmović, F. Frati, S. Langerman, and A. Tappini. Pole dancing: 3D morphs for tree drawings. *J. Graph Alg. Appl.*, 23(3):579–602, 2019. doi:10.7155/jgaa.00503.
5. F. Barrera-Cruz, P. Haxell, and A. Lubiw. Morphing Schnyder drawings of planar triangulations. *Discrete Comput. Geom.*, 61(1):161–184, 2019. doi:10.1007/s00454-018-0018-9.
6. T. Biedl, A. Lubiw, M. Petrick, and M. Spriggs. Morphing orthogonal planar graph drawings. *ACM Trans. Algorithms*, 9(4), 2013. doi:10.1145/2500118.
7. G. Birkhoff. Rings of sets. *Duke Math. J.*, 3(3):443–454, 1937. doi:10.1215/S0012-7094-37-00334-X.
8. G. Birkhoff and S. A. Kiss. A ternary operation in distributive lattices. *Bull. Amer. Math. Soc.*, 53(8):749–752, 1947.
9. K. S. Booth and G. S. Lueker. Testing for the consecutive ones property, interval graphs, and graph planarity using PQ-tree algorithms. *J. Comput. Syst. Sci.*, 13(3):335–379, 1976. doi:10.1016/S0022-0000(76)80045-1.
10. K. Buchin, B. Speckmann, and S. Verdonschot. Evolution strategies for optimizing rectangular cartograms. In N. Xiao, M. Kwan, M. F. Goodchild, and S. Shekhar, editors, *GIScience*, volume 7478 of *LNCS*, pages 29–42. Springer, 2012. doi:10.1007/978-3-642-33024-7_3.
11. S. S. Cairns. Deformations of plane rectilinear complexes. *Amer. Math. Monthly*, 51(5):247–252, 1944. doi:10.1080/00029890.1944.11999082.
12. E. W. Chambers, J. Erickson, P. Lin, and S. Parsa. How to morph graphs on the torus. In D. Marx, editor, *Symp. Discrete Algorithms (SODA)*, pages 2759–2778, 2021. doi:10.1137/1.9781611976465.164.

13. P. Dasgupta and S. Sur-Kolay. Slicible rectangular graphs and their optimal floorplans. *ACM Trans. Design Automation Elect. Syst.*, 6(4):447–470, 2001. doi:10.1145/502175.502176.
14. H. de Fraysseix, P. O. de Mendez, and P. Rosenstiehl. On triangle contact graphs. *Combin. Prob. Comput.*, 3(2):233–246, 1994. doi:10.1017/S0963548300001139.
15. 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.
16. S. Felsner. Rectangle and square representations of planar graphs. In J. Pach, editor, *Thirty Essays on Geometric Graph Theory*, pages 213–248. Springer, 2013. doi:10.1007/978-1-4614-0110-0_12.
17. S. Felsner and F. Zickfeld. On the number of α -orientations. In A. Brandstädt, D. Kratsch, and H. Müller, editors, *Graph-Theoretic Concepts Comput. Sci. (WG)*, volume 4769 of *LNCS*, pages 190–201. Springer, 2007. doi:10.1007/978-3-540-74839-7_19.
18. S. Florisson, M. J. van Kreveld, and B. Speckmann. Rectangular cartograms: construction & animation. In J. S. B. Mitchell and G. Rote, editors, *ACM Symp. Comput. Geom. (SoCG)*, pages 372–373, 2005. doi:10.1145/1064092.1064152.
19. É. Fusy. Transversal structures on triangulations, with application to straight-line drawing. In P. Healy and N. S. Nikolov, editors, *Graph Drawing (GD)*, volume 3843 of *LNCS*, pages 177–188. Springer, 2006. doi:10.1007/11618058_17.
20. 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.
21. K. R. Gabriel and R. R. Sokal. A new statistical approach to geographic variation analysis. *Syst. Biol.*, 18(3):259–278, 1969. doi:10.2307/2412323.
22. X. He. On finding the rectangular duals of planar triangular graphs. *SIAM J. Comput.*, 22(6):1218–1226, 1993. doi:10.1137/0222072.
23. R. Heilmann, D. A. Keim, C. Panse, and M. Sips. Recmap: Rectangular map approximations. In M. O. Ward and T. Munzner, editors, *IEEE Symp. Inform. Vis. (InfoVis)*, pages 33–40, 2004. doi:10.1109/INFVIS.2004.57.
24. M. Hoffmann, M. J. van Kreveld, V. Kusters, and G. Rote. Quality ratios of measures for graph drawing styles. In *Canadian Conf. Comput. Geom. (CCCG)*, 2014. URL: <http://www.cccg.ca/proceedings/2014/papers/paper05.pdf>.
25. G. Kant and X. He. Regular edge labeling of 4-connected plane graphs and its applications in graph drawing problems. *Theoret. Comput. Sci.*, 172(1):175–193, 1997. doi:10.1016/S0304-3975(95)00257-X.
26. S. G. Kobourov and M. Landis. Morphing planar graphs in spherical space. *J. Graph Alg. Appl.*, 12(1):113–127, 2008. doi:10.7155/jgaa.00162.
27. P. Koebe. Kontaktprobleme der konformen Abbildung. *Berichte über die Verhandlungen der Sächsischen Akademie der Wiss. zu Leipzig. Math.-Phys. Klasse* 88, pages 141–164, 1936.
28. K. Koźmiński and E. Kinnen. Rectangular duals of planar graphs. *Netw.*, 15(2):145–157, 1985. doi:10.1002/net.3230150202.
29. S. M. Leinwand and Y. Lai. An algorithm for building rectangular floor-plans. In *21st Design Automation Conference*, pages 663–664, 1984. doi:10.1109/DAC.1984.1585874.
30. E. Mumford. *Drawing graphs for cartographic applications*. PhD thesis, Faculty of Mathematics and Computer Science, TU Eindhoven, 2008. URL: <https://research.tue.nl/en/publications/drawing-graphs-for-cartographic-applications>.

31. S. Nusrat and S. G. Kobourov. The state of the art in cartograms. *Comput. Graph. Forum*, 35(3):619–642, 2016. doi:10.1111/cgf.12932.
32. H. C. Purchase, E. Hoggan, and C. Görg. How important is the “mental map”? – An empirical investigation of a dynamic graph layout algorithm. In M. Kaufmann and D. Wagner, editors, *Graph Drawing (GD)*, volume 4372 of *LNCS*, pages 184–195. Springer, 2007. doi:10.1007/978-3-540-70904-6_19.
33. E. Raisz. The rectangular statistical cartogram. *Geogr. Review*, 24(2):292–296, 1934. URL: <http://www.jstor.org/stable/208794>, doi:10.2307/208794.
34. P. Steadman. Graph theoretic representation of architectural arrangement. *Archit. Research Teach.*, pages 161–172, 1973.
35. A. van Goethem, B. Speckmann, and K. Verbeek. Optimal morphs of planar orthogonal drawings II. In D. Archambault and C. D. Tóth, editors, *Graph Drawing and Network Visualization (GD)*, pages 33–45. Springer, 2019. doi:10.1007/978-3-030-35802-0_3.
36. M. J. 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.
37. G. K. H. Yeap and M. Sarrafzadeh. Sliceable floorplanning by graph dualization. *SIAM J. Discrete Math.*, 8(2):258–280, 1995. doi:10.1137/S0895480191266700.

Appendix

A The Lattice of RELs

In this section, we review, for a given PTP graph, the structure of its RELs. In particular, we detail how rotations allow us to switch from one REL to another. We also remind the reader of the lattice of RELs induced by such rotations. Moreover, we prove that the lattice has a height of $\mathcal{O}(n^2)$ (Prop. 1) and that a shortest path can be computed in $\mathcal{O}(n^3)$ time (Prop. 2).

Let G be a PTP graph, let \mathcal{L} be a REL of G , and let C be a 4-cycle in G and \mathcal{L} . Recall that C is called *alternating* if the edges of C alternate between red and blue in \mathcal{L} . By swapping the colors of the edges inside C and by then fixing the orientations of the recolored edges (uniquely), we obtain a different REL \mathcal{L}' of G ; see Fig. 3. Considering rectangular duals that realize \mathcal{L} and \mathcal{L}' , respectively, note that the operation “rotates” the interior of the cycle (when C is separating) – or the interior contact segment – by 90° . Hence, we call such an operation either a *clockwise (cw)* or a *counterclockwise (ccw) rotation*. A 4-cycle C of G is called *rotatable* if it is alternating for at least one REL of G .

The lattice. Recall that a *lattice* is a poset in which each pair (a, b) has a unique smallest upper bound – the *join* $a \vee b$ of a and b – and a unique largest lower bound – the *meet* $a \wedge b$ of a and b . A lattice is *distributive* if the join and meet operation are distributive with respect to each other. Fusy [19, 20] showed that the set of RELs of G forms a distributive lattice, where $\mathcal{L} \leq \mathcal{L}'$ if and only if \mathcal{L}' can be obtained from \mathcal{L} by ccw rotations. Let $\mathcal{H}(G)$ denote the lattice of RELs of G . The minimum element \mathcal{L}_{\min} of $\mathcal{H}(G)$ is the unique REL that admits no cw rotation; the maximum element \mathcal{L}_{\max} is the unique REL that admits no ccw rotation.

Shortest paths. A path in the REL lattice $\mathcal{H}(G)$ is *monotone* if it uses only cw rotations or only ccw rotations. By Birkhoff’s theorem [7], the length of two monotone paths between two elements in a distributive lattice is equal. Furthermore, any three elements a, b, c of a distributive lattice have a unique median $(a \vee b) \wedge (a \vee c) \wedge (b \vee c) = (a \wedge b) \vee (a \wedge c) \vee (b \wedge c)$ that lies on a shortest path between any two of them [8]. Applying this to a , b , and $a \vee b$, yields that $a \vee b$ lies on a shortest path from a to b . Hence, for any pair of RELs, there exists a shortest path that contains their meet and one that contains their join.

Upper bound on path lengths. Recall that a 4-cycle is separating if there are other vertices both in its interior and its exterior. A separating 4-cycle is *nontrivial* if its interior contains more than one vertex; otherwise it is *trivial*. Note that the four vertices of a separating 4-cycle C can be seen as the outer cycle of a PTP subgraph and thus, by the coloring rules of RELs, each vertex of C has edges in only one color to vertices in the interior of C . We call non-separating 4-cycles also *empty 4-cycles*.



(a) A PTP graph with a $K_{2,5}$ subgraph. (b) Two properly intersecting 4-cycles.

Fig. 11: (a) Only the outermost 4-cycle of this $K_{2,5}$ subgraph of a PTP graph can be rotatable; (b) hence, two 4-cycles C and C' with properly intersecting interiors cannot be rotatable.

Consider an n -vertex PTP graph G where all separating 4-cycles are trivial. Then any path in $\mathcal{H}(G)$ has length at most $\mathcal{O}(n^2)$ [15]. To extend this result to general PTP graphs, we use a recursive decomposition of G following Eppstein et al. [15] (as well as Dasgupta and Sur-Kolay [13] and Mumford [30]). For a nontrivial separating 4-cycle C of a PTP graph G , we define two *separation components* of G with respect to C : The *inner separation component* G_C^{in} is the maximal subgraph of G with C as outer face. The *outer separation component* G_C^{out} is the minor of G obtained by replacing the interior of C with a single vertex v_C . Note that both G_C^{in} and G_C^{out} are PTP graphs. A *minimal separation component* of G is a separation component of G that cannot be split any further. Partitioning G into minimal separation components takes linear time [15].

Eppstein et al. [15] pointed out that there can be a quadratic number of nontrivial separating 4-cycles in G , but that these can be represented in linear space by finding all maximal complete bipartite subgraphs $K_{2,i}$ of G . Such a representation can be found in linear time [15]. Note that, among the 4-cycles of a subgraph $K_{2,i}$ (with u and v forming the small partition), only the outermost 4-cycle C is rotatable. This can be seen as C is the outer 4-cycle of a smaller PTP graph and thus any of the smaller cycles has a monochromatic path from u to v ; see Figs. 11a and 11b.

Observation 10 *Let G be a PTP graph with two rotatable 4-cycles C and C' . Then the interiors of C and C' are either disjoint or one lies inside the other.*

Proof. We prove that it is not possible that the interiors of C and C' intersect properly. Assume otherwise and observe that then C and C' must intersect in two non-adjacent vertices u and v such that each contains one vertex of the other; see Fig. 11b. However, then u and v form one partition of a subgraph $K_{2,i}$. Since neither C nor C' is the outermost 4-cycle of this $K_{2,i}$ neither can be rotatable. \square

Note that two rotatable 4-cycles C and C' may *overlap* in the sense that one of them, say, C' is empty and its interior contains exactly one edge of C ; see Fig. 12. Indeed, if C rotates multiple times on a path from \mathcal{L}_{\min} to \mathcal{L}_{\max} , then each edge of C is interior to an empty rotatable 4-cycle.




Fig. 12: Two rotatable 4-cycles C and C' may overlap if, say, C' is empty and its interior contains exactly one edge of C .

Observation 10 implies that a rotation of a 4-cycle C' for a REL \mathcal{L} of G can be seen as a rotation on \mathcal{L} restricted to G_C^{in} if C' lies inside of C , or restricted to G_C^{out} if their interiors are disjoint or if C lies inside of C' . In the last case, to combine the RELs of G_C^{out} and G_C^{in} into a REL of G , we have to rotate the REL of G_C^{in} once. This yields the following.

Lemma 11. *Let G be a PTP-graph with a nontrivial rotatable separating 4-cycle C . A rotation sequence σ on a REL \mathcal{L} of G can be partitioned into a rotation sequence that acts on \mathcal{L} restricted to G_C^{out} and a rotation sequence that acts on \mathcal{L} restricted to G_C^{in} .*

Proposition 1 (\star). *Given an n -vertex PTP graph G , the lattice of RELs of G has diameter $\mathcal{O}(n^2)$.*

Proof. For a PTP graph that contains only empty and trivial separating 4-cycles or, equivalently, for a minimal separation component, this is known [15]. We partition G into two separation components and apply an inductive argument based on Lemma 11. Since the bound on the diameter is on the number of inner vertices (that is, we do not count the outer 4-cycle), the sum of the lengths of the sequences in the two separation components is still in $\mathcal{O}(n^2)$. \square

Construction of a shortest path. Let G be an n -vertex PTP graph whose separating 4-cycles are trivial. Let \mathcal{R} and \mathcal{R}' be two rectangular duals of G that realize the RELs \mathcal{L} and \mathcal{L}' , respectively. For a rotatable 4-cycle of G , we defined the rotation count $f_C(\mathcal{L})$ as the number of rotations of C in a monotone sequence from the minimum REL \mathcal{L}_{\min} to \mathcal{L} . Let $\mathcal{P}(G)$ be a poset where each vertex is a tuple (C, i) consisting of a rotatable 4-cycle C and a rotation count i . Furthermore, $(C, i) \leq (C', j)$ when for each REL \mathcal{L} with $f_C(\mathcal{L}) \leq i$ it holds that $f_{C'}(\mathcal{L}) \leq j$; that is, it is not possible to increase the rotation count of C' from j to $j + 1$ prior to increasing the rotation count of C from i to $i + 1$. Eppstein et al. [15] showed that each REL \mathcal{L} of G can be represented by the partition of $\mathcal{P}(G)$ into a downward-closed set $L(\mathcal{L})$ and an upward-closed set $U(\mathcal{L})$ such that $(C, i) \in L(\mathcal{L})$ when $i < f_C(\mathcal{L})$ and $(C, i) \in U(\mathcal{L})$ otherwise. In fact, they showed that $\mathcal{P}(G)$ is order-isomorphic to the poset defined by Birkhoff's representation theorem [7]. Hence, the meet of two RELs \mathcal{L} and \mathcal{L}' in $\mathcal{H}(G)$ (the lattice of RELs of G) is represented by $L(\mathcal{L}) \cap L(\mathcal{L}')$. By Obs. 10 and Prop. 1, we can extend this result again to any PTP graph G by considering the minimum

separation components of G separately. The rotation count of each rotatable four-cycle is then defined only with respect to the minimum separation component of G it lies in.

Let C be a rotatable 4-cycle of G . We define the *rotation count* $f_C(\mathcal{L})$ as the number of rotations of C in a monotone sequence from the minimum REL \mathcal{L}_{\min} to \mathcal{L} . Let $\bar{\mathcal{L}} = \mathcal{L} \wedge \mathcal{L}'$, the meet of \mathcal{L} and \mathcal{L}' . By a result of Eppstein et al. [15], for every rotatable 4-cycle C of G , it holds that $f_C(\bar{\mathcal{L}}) = \min\{f_C(\mathcal{L}), f_C(\mathcal{L}')\}$. Details are given in Appendix A.

Proposition 2 (\star). *Let G be an n -vertex PTP graph with RELs \mathcal{L} and \mathcal{L}' . In the lattice of RELs of G , a shortest \mathcal{L} - \mathcal{L}' path can be computed in $\mathcal{O}(n^3)$ time.*

Proof. Let $\bar{\mathcal{L}} = \mathcal{L} \wedge \mathcal{L}'$ be the meet of \mathcal{L} and \mathcal{L}' . As noted in Appendix A, there exists a shortest path from \mathcal{L} to \mathcal{L}' via $\bar{\mathcal{L}}$. Therefore, we first compute $\bar{\mathcal{L}}$ and then construct monotone paths from \mathcal{L} and \mathcal{L}' to $\bar{\mathcal{L}}$, respectively. To this end, we first find all cw rotatable 4-cycles in \mathcal{L} and then execute rotations until we reach \mathcal{L}_{\min} . With each step, we count the number of times the involved cycles have been rotated so far and check whether this enables cw rotations of other 4-cycles. When we reach \mathcal{L}_{\min} , we have the rotation counts for all (involved) rotatable 4-cycles for \mathcal{L} (and \mathcal{L}'). For each rotatable 4-cycle C of G , we compute $f_C(\bar{\mathcal{L}})$ as $\min\{f_C(\mathcal{L}), f_C(\mathcal{L}')\}$. One such rotation step takes $\mathcal{O}(n)$ time and by Prop. 1 the number of steps is at most $\mathcal{O}(n^2)$.

Finally, to compute two shortest paths from \mathcal{L} and \mathcal{L}' to $\bar{\mathcal{L}}$, we greedily rotate cw rotatable 4-cycles whose rotation count is great than $f_C(\bar{\mathcal{L}})$. This takes $\mathcal{O}(n^2)$ time. \square

In Sect. 2.3, we construct a relaxed morph between two rectangular duals \mathcal{R} and \mathcal{R}' along a shortest paths between them. This morph requires at most a quadratic number of steps, since a shortest path between \mathcal{R} and \mathcal{R}' has by Prop. 2 at most quadratic length. Note that for this asymptotic bound on the number of steps, it would suffice to morph along, for example, a path from \mathcal{R} via a rectangular dual that realizes \mathcal{L}_{\min} to \mathcal{R}' . By Obs. 10, this path also has a length in $\mathcal{O}(n^2)$. However, with Theorem 9, we show that by executing rotations in parallel and morphing along a shortest path, we can construct a morph that uses only $\mathcal{O}(1)$ times the minimum number of steps needed to get from \mathcal{R} to \mathcal{R}' .

B Using Kant and He on Auxiliary Graph \hat{G}

In this section we explain the algorithm by Kant and He [25] and explain why we can use it on the auxiliary graph \hat{G} with auxiliary REL $\hat{\mathcal{L}}$ in the proof of Lemma 6.

Kant and He [25] introduced RELs and described two linear-time algorithms that compute a REL for a given PTP graph; one algorithm is based on edge contractions, the other is based on canonical orderings. They then use the algorithm by He [22] to construct in linear time a rectangular dual that realizes this REL and where the coordinates are all integers.

The algorithm by He works as follows; see Fig. 13. Given a PTP graph G with REL (L_1, L_2) . Consider the weak dual $L_1^*(G)$ of $L_1(G)$. In $L_1^*(G)$, there is a vertex f for every interior face f in the plane embedding of $L_1(G)$, and there are two vertices \bar{f}_L and \bar{f}_R for the outer face \bar{f} in the plane embedding of $L_1(G)$. There is an edge (f, f') if f is the interior face to the left and f' is the interior face to the right of some edge in $L_1(G)$; an edge (\bar{f}_L, f') if \bar{f} is the face to the left and f' is the interior face to the right of some edge in $L_1(G)$; and an edge (f, \bar{f}_R) if f is the interior face to the left and \bar{f} is the face to the right of some edge in $L_1(G)$. Then $L_1^*(G)$ is a planar st-graph with source \bar{f}_L and sink \bar{f}_R . Compute a topological numbering d of $L_1^*(G)$. For any vertex v of G , let $\text{left}(v)$ be the face to the left of v in $L_1(G)$, and let $\text{right}(v)$ be the face to the right of v in $L_1(G)$; these are the two faces between an incoming and an outgoing edge of v . Then the algorithm sets the x-coordinate of the left side of $\mathcal{R}(v)$ to $d(\text{left}(v))$ and the x-coordinate of the right side of $\mathcal{R}(v)$ to $d(\text{right}(v))$. The y-coordinates for the bottom and top side of $\mathcal{R}(v)$ are calculated analogously from the weak dual $L_2^*(G)$ of $L_2(G)$.

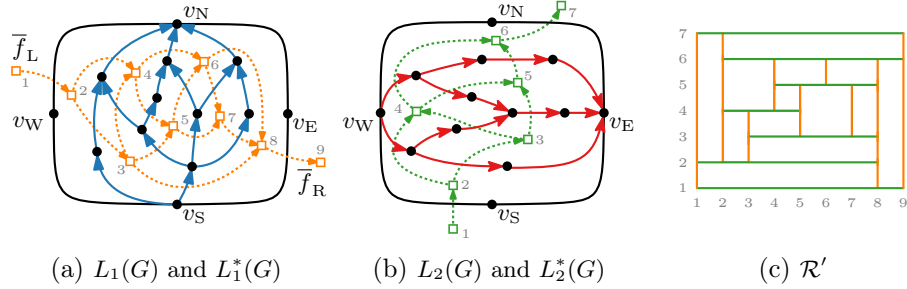


Fig. 13: Illustration of the algorithm by He [22]. (a+b) Topological numberings (gray) for the weak duals $L_1^*(G)$ and $L_2^*(G)$ (dotted) of $L_1(G)$ and $L_2(G)$; (c) resulting rectangular dual \mathcal{R} .

Note that the maximal vertical segments in the resulting rectangular dual \mathcal{R} are in bijection with the vertices of $L_1^*(G)$ and the maximal horizontal segments are in bijection with the vertices of $L_2^*(G)$; see Fig. 13c.

He [22] did not show this explicitly, but for a given REL, his linear-time algorithm yields a rectangular dual of minimum width and height, and thus of minimum area and perimeter.

In the proof of Lemma 6 we extended a PTP graph G to an auxiliary graph \hat{G} where some faces are chordless 4-cycles. We also described how the REL \mathcal{L} of G can be extended to an auxiliary REL $\hat{\mathcal{L}}$ of \hat{G} that adheres to the REL coloring rules. Also, each chordless 4-cycle C bounding a face of \hat{G} is colored alternately. In a rectangular dual $\hat{\mathcal{R}}_1$ of \hat{G} , the four rectangles corresponding to the vertices of C meet in a single point; there are thus not only T-junctions, but also crossings. The vertices of $\hat{L}_1^*(G)$ and $\hat{L}_2^*(G)$ thus still correspond to

maximal vertical and horizontal line segments in $\hat{\mathcal{R}}_1$, respectively. See also again Figs. 6 and 7. Hence, applying the algorithm by Kant and He [25] to \hat{G} and $\hat{\mathcal{L}}$, we obtain the (almost) rectangular dual $\hat{\mathcal{R}}_1$ described above.

C Proofs Omitted in Section 3

Proposition 8 (\star). *Let \mathcal{R} be a rectangular dual of a PTP graph G with REL \mathcal{L} whose separating 4-cycles are all trivial. Let \mathcal{C} be a set of alternating 4-cycles of \mathcal{R} . Let \mathcal{L}' be the REL obtained from \mathcal{L} by executing all rotations in \mathcal{C} . There exists a relaxed morph with $\mathcal{O}(1)$ steps from \mathcal{R} to a rectangular dual \mathcal{R}' realizing \mathcal{L}' . The morph can be computed in linear time.*

Proof. We assume that \mathcal{C} contains only cw rotatable cycles; otherwise we apply the same algorithm to the ccw rotatable cycles afterwards. Our algorithm to construct a morph from \mathcal{R} to some \mathcal{R}' consists of the following steps. (i) We construct the conflict graph $K(\mathcal{C})$. This can be done in linear time by traversing the four border segments of each cycle in \mathcal{C} . (ii) We greedily compute a 5-coloring of $K(\mathcal{C})$, again in linear time, and pick a set \mathcal{C}_i , $i \in \{1, \dots, 5\}$, corresponding to one color. We add all empty 4-cycles of \mathcal{C} to \mathcal{C}_1 . (iii) We compute and execute a preparatory linear morph such that, (iv) all rotations in \mathcal{C}_i can safely be executed in parallel. Hence, we need a constant number of linear morphs to arrive at a rectangular dual \mathcal{R}' realizing \mathcal{L}' .

The preparatory linear morph for \mathcal{C}_i works similar to the preparatory morph in the serial case. Using an auxiliary PTP graph \hat{G} and an auxiliary REL $\hat{\mathcal{L}}$ we find a target rectangular dual \mathcal{R}_1 that also realizes \mathcal{L} . We can then use a linear morph from \mathcal{R} to \mathcal{R}_1 according to Sect. 2.1. For each 4-cycle in \mathcal{C}_i that does not share a segment with another 4-cycle in \mathcal{C}_i , we extend G and \mathcal{L} towards \hat{G} and $\hat{\mathcal{L}}$, respectively, as before. Since these cycles are non-conflicting, this can be done nearly independently for each 4-cycle; however a rectangle might now be split both horizontally and vertically (meaning, its vertex gets duplicated twice).

For a separating 4-cycle C and empty 4-cycles C_1, \dots, C_k that share a maximal line segment s as border segment, we assume w.l.o.g. that s is the upper border segment of C . Further, let C_1, \dots, C_k appear along s from left to right and let $C = \langle a, b, cd \rangle$ as in the proof of Lemma 6. This case is illustrated in Fig. 14a. The rotation of C moves s downwards. Note that the rotation of C breaks s and moves the part right of the interior I_C of C down. Thus, each 4-cycle C_i , $i \in \{1, \dots, k\}$ has to follow this downward motion. Furthermore, C_i move the part of s right of it also one down and hence C_i moves down the height of I_C plus i ; see Fig. 14c. Therefore, to create enough vertical space below s that contains not other horizontal segment, we ensure that in \mathcal{R}_1 each rectangle $\mathcal{R}_1(v)$ with v in $P_d \setminus \{x\}$ satisfies $y_1(\mathcal{R}_1(v)) < y_2(\mathcal{R}_1(a)) - k$; see Fig. 14b. Recall that the preparatory step for C in the serial case duplicates all vertices on P_d . We do the same here with a minor change, namely, we split d into the vertices d_1, \dots, d_{k+2} while all other vertices in $P_d \setminus \{d, x\}$ are only duplicated. Let y be the successor of d on P_d . We assign the edges clockwise between (and including) dy and ad to d_1 , and the edges clockwise between ad and dy to d_{k+2} . For

$i \in \{1, \dots, k+1\}$, we connect d_i and d_{i+1} with a blue edge and add a red edge ad_i . Furthermore, we add the red edge d_1y_1 and, for $i \in \{2, \dots, k+2\}$, the red edge d_iy_i . For the vertices $v \in P_d \setminus \{d\}$, we add edges as before. As a result we get that $\hat{\mathcal{R}}_1(d_1), \dots, \hat{\mathcal{R}}_1(d_k)$ stack on top of each other and, for all v in $P_d \setminus \{x\}$, we have $y_1(\hat{\mathcal{R}}_1(v)) < y_2(\hat{\mathcal{R}}_1(a)) - k$. In other words, there is enough vertical space below s in $\hat{\mathcal{R}}_1$ and thus also in \mathcal{R}_1 .

The case where s is the interior segment shared only by empty 4-cycles is handled analogously to the previous case and Lemma 6.

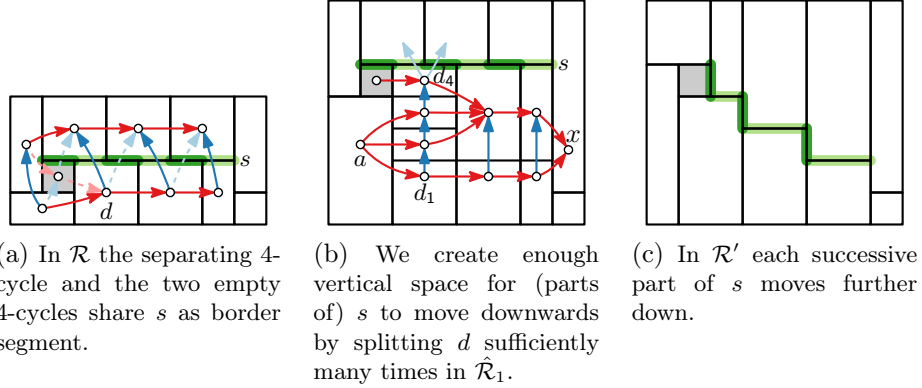


Fig. 14: To enable a parallel rotation of all three alternating 4-cycles in the rectangular dual \mathcal{R} that share the border segment s , we first morph to a rectangular dual \mathcal{R}_1 derived from the auxiliary rectangular dual $\hat{\mathcal{R}}_1$. In doing so, we ensure that there is enough space below s such that its parts can move down to their positions in \mathcal{R}' .

The auxiliary graph \hat{G} and the auxiliary REL $\hat{\mathcal{L}}$ can be constructed locally around each involved 4-cycle. Furthermore, since each vertex is either only duplicated a constant number of times or one vertex linear many times in the number of involved empty 4-cycles (such as d above), the sizes of \hat{G} and $\hat{\mathcal{L}}$ are linear in the size of G . Hence, the algorithm by Kant and He [25] computes $\hat{\mathcal{R}}_1$ in $\mathcal{O}(n)$ time and we can derive \mathcal{R}_1 in $\mathcal{O}(n)$ time. Finally, to find \mathcal{R}' , we compute the positions of all rectangles after the rotations. Because of the preparatory step, this can be again done locally on \mathcal{R}_1 in overall $\mathcal{O}(n)$ time. \square

Theorem 9 (\star). *Let G be a PTP graph whose separating 4-cycles are all trivial. Let \mathcal{R} and \mathcal{R}' be two rectangular duals of G , and let OPT be the minimum number of steps in any relaxed morph between \mathcal{R} and \mathcal{R}' . Then we can construct in cubic time a relaxed morph consisting of $\mathcal{O}(\text{OPT})$ steps.*

Proof. Let \mathcal{R} and \mathcal{R}' be realized by RELs \mathcal{L} and \mathcal{L}' , respectively. Let $\bar{\mathcal{L}} = \mathcal{L} \wedge \mathcal{L}'$, the meet of \mathcal{L} and \mathcal{L}' . We first consider the path from \mathcal{L} to $\bar{\mathcal{L}}$. Let X be the set of cw rotations between \mathcal{L} and $\bar{\mathcal{L}}$. Not all cw rotations in X can be executed

in parallel because other cw rotations in X might have to be executed first. Let k be the minimum number of linear morphs of a relaxed morph M from \mathcal{L} to $\bar{\mathcal{L}}$. This morph has to execute all rotations in X , so it partitions X into sets X_1, \dots, X_k where X_i contains all cw rotations in X executed by M in step i . Our algorithm repeatedly applies Prop. 8 to execute all possible cw rotations in X within $\mathcal{O}(1)$ steps. Hence, in the first application of Prop. 8, our algorithm executes all rotations in X_1 (and possibly more). Subsequently, in the second application of Prop. 8, our algorithm executes all rotations in X_2 (except those it had already executed in the first application). Let x be a rotation that is executed in the j -th application of Prop. 8. Then we have $x \in X_i$ with $i \geq j$. Thus, our algorithm requires at most k applications of Prop. 8 to execute all rotations in $X = X_1 \cup \dots \cup X_k$, each of which requires $\mathcal{O}(1)$ steps. In total, our algorithm uses $\mathcal{O}(k)$ steps to get from \mathcal{L} to $\bar{\mathcal{L}}$.

Symmetrically, let X' be the set of ccw rotations between $\bar{\mathcal{L}}$ and \mathcal{L}' . If ℓ is the minimum number of linear morphs of a relaxed morph from $\bar{\mathcal{L}}$ to \mathcal{L}' , then our algorithm requires $\mathcal{O}(\ell)$ steps between $\bar{\mathcal{L}}$ and \mathcal{L}' , so in total $\mathcal{O}(k + \ell)$ steps from \mathcal{L} to \mathcal{L}' .

Let M^* be a relaxed morph between \mathcal{R} and \mathcal{R}' that uses the minimum number of linear morphs. Because of the lattice structure, M^* has to execute all cw rotations in X and all ccw rotations in X' . Hence, M^* uses at least $\max\{k, \ell\} \geq (k + \ell)/2$ linear morphs.

The running time follows from Theorem 7 and Prop. 8. \square