On Compatible Triangulations with a Minimum Number of Steiner Points

Anna Lubiw* Debajyoti Mondal*

Abstract

Two polygons are *compatible* if they have the same clockwise cyclic ordering of vertices. The definition extends to polygonal regions (polygons with holes) and to triangulations—for every face, the clockwise cyclic order of vertices on the boundary must be the same. It is known that every pair of compatible n-vertex polygonal regions can be extended to compatible triangulations by adding $O(n^2)$ Steiner points. Furthermore, $\Omega(n^2)$ Steiner points are sometimes necessary, even for a pair of polygons. Compatible triangulations provide piecewise linear homeomorphisms and are also a crucial first step in morphing planar graph drawings, aka "2D shape animation." An intriguing open question, first posed by Aronov, Seidel, and Souvaine in 1993, is to decide if two compatible polygons have compatible triangulations with at most k Steiner points. In this paper we prove the problem to be NP-hard for polygons with holes. The question remains open for simple polygons.

1 Introduction

For many computational geometry problems involving a polygon or polygonal region, the standard first step is to triangulate the region. However, for some problems, such as morphing of polygons, or finding a homeomorphism between polygons, the input consists of two polygons with a correspondence between them, and the desirable first step is to triangulate them in a consistent way. Unlike for a single polygon, it may be necessary to add new vertices, called "Steiner points". Our paper is about this harder problem, which was called "joint triangulation" by Saalfeld [12] and "compatible triangulation" by Aronov, Seidel, and Souvaine [3].

Research on finding compatible triangulations is motivated by applications in morphing [2] and 2D shape animation [5, 15], and in computing piecewise linear homeomorphisms of polygons.

Throughout, we deal with vertex-labelled straight-line planar drawings. The most general input we consider is a polygon with holes (a polygonal region), where we allow a hole to degenerate to a single point. Two polygons are *compatible* if they have the same clockwise cyclic ordering of vertices. Two polygonal regions P_1 and P_2 are *compatible* if their outer polygons are compatible, and their holes are compatible, i.e. each hole (considered as a polygon) in P_1 corresponds to a compatible hole in P_2. Note that the labelling provides the correspondence.

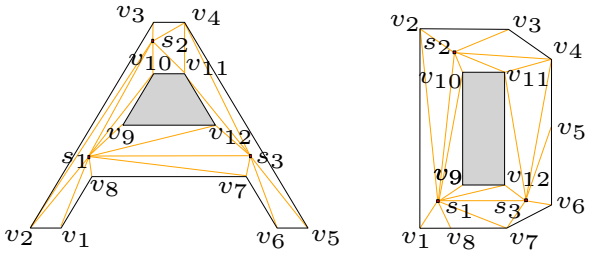

Figure 1: Two compatible polygons, each with one hole (shaded gray), and compatible triangulations of them with 3 Steiner points.

A *triangulation* $T(P)$ *of a polygonal region* P is a subdivision of its interior region into triangular faces. The vertices of $T(P) \setminus P$ are called *Steiner points* of $T(P)$. A pair of triangulations $T(P_1)$ and $T(P_2)$ of compatible polygonal regions P_1 and P_2, respectively, are *compatible* if their faces are compatible, i.e. every face of $T(P_1)$ (considered as a polygon) corresponds to a compatible face of $T(P_2)$. Again, the labelling provides the correspondence. Figure 1 illustrates a pair of compatible polygonal regions and their compatible triangulations.

Two special cases of compatible triangulations were studied independently. Saalfeld in 1987 [12] considered the case of two rectangles each with n points inside them (where the correspondence between the points is given) and showed that compatible triangulations always exist. Aronov et al., in 1993 [3] considered the case of simple compatible polygons. They showed that there exist compatible triangulations with $O(n^2)$ Steiner points. (A similar construction was given by Thomassen in 1983 [16, Theorem 4.1].) Furthermore, Aronov et al. gave an $O(n^2)$-time algorithm to compute such compatible triangulations, and they gave examples where $\Omega(n^2)$ Steiner points are necessary. They posed as an open problem to decide if two polygons have a compatible triangulation with k Steiner points, and observed that the case $k = 0$ can be decided in polynomial time via dynamic programming.

Our Result. We show that it is NP-hard to decide if

*Cheriton School of Computer Science, University of Waterloo, {alubiw,dmondal}@uwaterloo.ca

two compatible polygonal regions have compatible triangulations with at most k Steiner points, where $k \in \mathbb{N}$ is given as part of the input.

Further Background. There are a number of further results for the case of two simple polygons. Kranakis and Urrutia [9] gave an $O(n + r^2)$-time algorithm to find compatible triangulations of simple compatible polygons with $O(n + r^2)$ Steiner points, where r is the number of reflex vertices. Gupta and Wenger [8] gave a polynomial-time algorithm that provides an $O(\log n)$ approximation to the minimum number of Steiner points. A number of heuristic algorithms have been proposed—see e.g., [5, 15].

There is also a line of research on the case of polygons with point holes (Saalfeld's problem). Souvaine and Wenger [14] gave an $O(n^2)$-time algorithm to compute compatible triangulations with $O(n^2)$ Steiner points, and asked if there is a polynomial-time algorithm to construct compatible triangulations with the minimum number of Steiner points. Pach et al. [10] proved that $\Omega(n^2)$ Steiner points are sometimes necessary.

For the case of general polygonal regions—which encompasses both the above special cases—Babikov et al. [4] gave an $O(n^2)$-time algorithm to compute compatible triangulations with $O(n^2)$ Steiner points.

One approach to computing compatible triangulations is to first compute a triangulation for one of the polygonal regions, and then draw its underlying graph into the other polygonal region using polylines for drawing edges. The edge bends give rise to the Steiner points. This idea relates to the problem of drawing a planar graph on a given set of points, where the correspondence between vertices and the points is given. Pach and Wenger [11] gave an $O(n^2)$-time algorithm to compute such an embedding with $O(n^2)$ bends in total, and this was extended to deal with a bounding polygonal region in [6].

The version of the compatible triangulation problem where the correspondence between the two polygonal regions is *not* given is also well-studied and very relevant in practice, e.g. see [5]. In this setting, Aichholzer et al. [1] made the fascinating conjecture that for any two point sets each with n points, of which h lie on the convex hull, there is a mapping between them that permits compatible triangulations with no Steiner points.

2 Preliminaries

Let P be a polygon, possibly with holes. Two points a, b in P are *visible* if the line segment between them lies inside P; they are *1-bend visible* if there is a point c inside P that is visible to both a and b.

A *dent* on the boundary of P consists of three consecutive vertices u, d, v of P such that d is convex and u, v are reflex vertices, e.g., see the polygon P_1 in Figure 2.

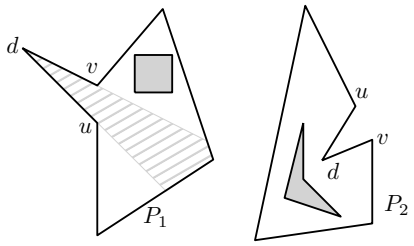

Figure 2: Illustration for Lemma 1. The visibility region of d is shown in gray stripes.

We refer to d as the *peak* of the dent. The *visibility region* of d consists of all the points inside P that are visible to P. An *inward dent* on the boundary of P consists of three consecutive vertices u, d, v of P such that d is reflex and u, v are convex vertices. The following simple lemma about dents in compatible triangulations of polygons will be a key ingredient of our NP-hardness proof.

Lemma 1 *Let P_1 and P_2 be a pair of compatible polygons. Assume that P_1 contains a dent u, d, v, and let Ψ be the visibility region of d in P_1. If u, v are not visible in P_2, then in any compatible triangulations d must be adjacent either to a Steiner point or a vertex (except u and v) inside Ψ.*

Proof. Any triangulation of P_1 (even with Steiner points) must use the edge (u, v) or an edge incident to d. In compatible triangulations of P_1 and P_2 the edge (u, v) is ruled out, and therefore d must be adjacent to a Steiner point or a vertex in $\Psi \setminus \{u, v\}$. □

3 NP-Hardness

In this section we prove that given a pair of compatible polygonal regions P_1, P_2, and $k \in \mathbb{N}$, it is NP-hard to decide if there are compatible triangulations of P_1 and P_2 with at most k Steiner points.

We reduce from the monotone rectilinear planar 3-SAT problem (MRP-3SAT), which is NP-complete [7]. The input of an MRP-3SAT instance I is a collection C of clauses over a set U of Boolean variables such that each clause contains at most three literals, and is either *positive* (consists of only positive literals), or *negative* (consists of only negative literals). Moreover, the corresponding *SAT-graph* G_I (the bipartite graph with vertex set $C \cup U$ and edge set $\{(c, x) \in C \times U : x \text{ appears in } c\}$) admits a planar drawing Γ satisfying the following properties:

- Each vertex in G_I is drawn as an axis-aligned rectangle in Γ.
- All the rectangles representing variables lie along a horizontal line ℓ.

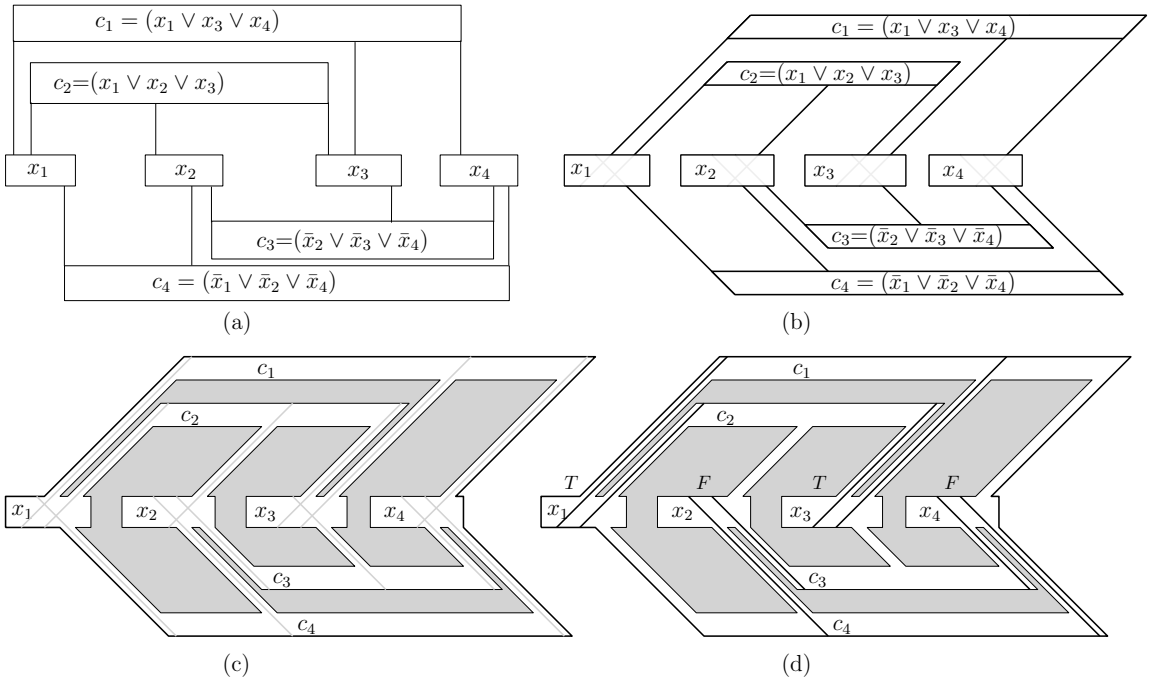

Figure 3: (a) An instance I of MRP-3SAT, and the corresponding drawing Γ. (b) Γ'. (c)–(d) Illustration for the hardness reduction.

- The rectangles representing positive (respectively, negative) clauses lie above (respectively, below) ℓ.
- Each edge (c, x) of G_I is drawn as a vertical line segment that connects the rectangles corresponding to c and x, e.g., see Figure 3(a).

The MRP-3SAT problem asks whether there is a truth assignment for U satisfying all clauses in C.

Given an instance $I = (U, C)$ of MRP-3SAT, we construct a pair of compatible polygonal regions P_1 and P_2 such that they admit compatible triangulations with at most $5|C|$ Steiner points, if and only if I is satisfiable.

Idea of the reduction: We first ensure that every clause in I has exactly three literals, by duplicating literals if necessary. Let the resulting instance be I'. It is straightforward to observe that I' is also an instance of MRP-3SAT, and I' is satisfiable if and only if I is satisfiable. Let Γ be the drawing corresponding to $G_{I'}$.

We modify the drawing Γ such that the edges and vertices corresponding to the positive (resp., negative) clauses become parallelograms, slanted 45° (resp., −45°) to the right, e.g., see Figure 3(b). For each clause $c \in C$, let $R(c)$ denote the parallelogram corresponding to c. We call $R(c)$ the "clause region". For each variable $u \in U$, let $B(u)$ denote the rectangle corresponding to u. We call $B(u)$ the "variable region". We call the edges of $G_{I'}$ *connectors* and we call the connectors that are incident to the top (resp., bottom) side of $B(u)$ *top* (resp., *bottom*) *connectors* of $B(u)$. We ensure that the extension of every top connector intersects the extensions of all the bottom connectors inside $B(u)$. Let the resulting drawing be Γ'. We construct P_1 and P_2 by modifying two distinct copies of Γ'.

We prove that in any compatible triangulations with $5|C|$ Steiner points, for each clause c, there is a triangulation edge e_c that lies along one of the connectors incident to the clause region. If c is positive (resp., negative) then we can set the variable corresponding to e_c to true (resp., false) and this will satisfy the clause. We get a valid truth-value assignment because a variable region cannot contain extensions of both top and bottom connectors. Figures 3(c)–(d) illustrate a satisfying truth assignment for I. On the other hand, given a satisfying truth assignment, we show how to find compatible triangulations for P_1 and P_2 using $5|C|$ Steiner points.

3.1 Construction of Polygonal Region P_1

We modify a copy Γ'_1 of Γ' to construct P_1. First we create a *channel* of small non-zero width around each connector so that we have a polygon with holes. We denote the copies of $R(c)$ and $B(u)$ in P_1 by $R_1(c)$ and $B_1(u)$. We create nine dents with peaks $u, v, w, q, q', r, r', s, s'$ in the boundary of $R_1(c)$, as shown in Figures 4(a)–(b). The visibility region of each dent is illustrated using gray straight lines.

As illustrated in Figure 4(a), we place a hole h in the leftmost channel of $R_1(c)$, not intersecting the visibil-

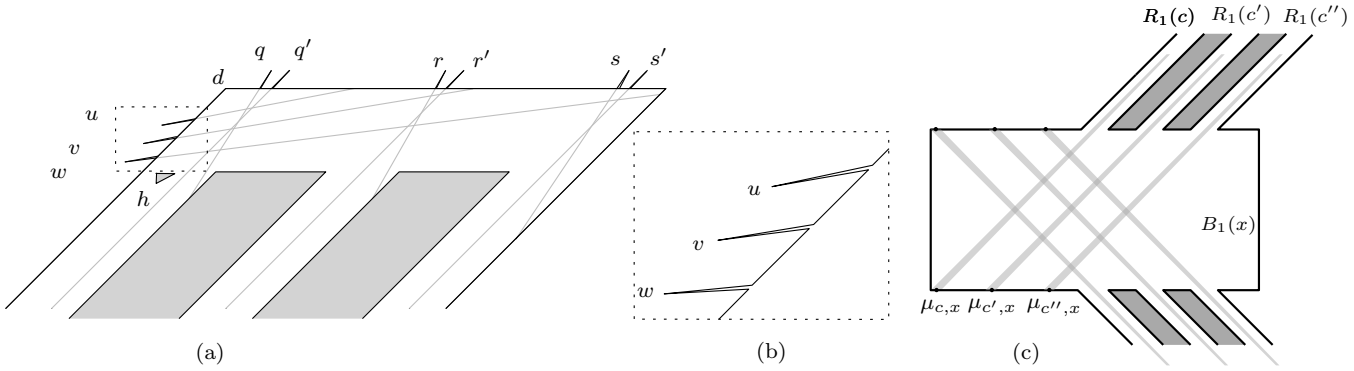

Figure 4: (a) A clause region in P_1. (b) A close-up of the dents corresponding to u, v, w. (c) Illustration for $B_1(x)$.

ity regions of the peaks $u, v, w, q, q', r, r', s, s'$. We refer the reader to Appendix A for the formal details of the construction and the precise placement of h.

We now modify the rectangles that correspond to the variables. Let x be a literal and let $B_1(x)$ be the corresponding rectangle in Γ'_1. See Figure 4(c). For every positive (resp., negative) clause c containing x, one or more[1] visibility regions corresponding to the peaks of $R_1(c)$ enter $B_1(x)$. We ensure that the visibility regions entering from the top (resp., bottom) of $B_1(x)$ are disjoint and only intersect the bottom (resp., top) side of $B_1(x)$. For each clause c containing x, we construct a vertex $\mu_{c,x}$ on the side of $B_1(x)$ such that $\mu_{c,x}$ is visible to the corresponding peak of $R_1(c)$. We refer to these newly constructed points as the *μ-points* of $B_1(x)$.

3.2 Construction of Polygonal Region P_2

We modify a copy Γ'_2 of Γ' to construct P_2. As in the construction of P_1, we create a channel of small non-zero width around each connector so that we have a polygon with holes. We denote the copies of $R(c)$ and $B(u)$ in P_2 by $R_2(c)$ and $B_2(u)$. We create four inward dents on the boundary of $R_2(c)$, and place the points $u, v, w, d, q, q', r, r', s, s'$, as shown in Figure 5(a). Finally, we place the hole h ensuring that no peak in $\{u, v, w\}$ is 1-bend visible to $\{q', r', s'\}$, e.g., see Figure 5(b). We refer the reader to Appendix A for the formal details of the construction.

We now modify the rectangles that correspond to the literals. Let x be a literal and let $B_2(x)$ be the corresponding rectangle in Γ'_2. The modification for $B_2(x)$ is analogous to that of $B_1(x)$. Specifically, for every visibility region (of some peak $p \in \{q', r', s'\}$) that intersects the box $B_1(x)$ in Γ'_1, we construct a point μ on the boundary of box $B_2(x)$ such that μ and p are visible in P_2. Figure 5(b) illustrates such visibilities with dashed lines.

[1]Recall that c may contain duplicates of a literal.

3.3 Properties of Compatible Drawings

In this section we prove some key properties of compatible triangulations $T(P_1)$ and $T(P_2)$ of P_1 and P_2, respectively. For clause c, let $\overline{R}_1(c)$ be the clause region $R_1(c)$ plus its three attached channels.

Lemma 2 *If c is a clause such that no peak q', r', s' is adjacent in $T(P_1)$ to a point outside $\overline{R}_1(c)$, then there are at least 6 Steiner points in $\overline{R}_1(c)$.*

Proof. Consider the 9 points $\{u, v, w, q, q', r, r', s, s'\}$. In P_1 each point in this set is the peak of a dent, so by Lemma 1, each of these 9 points must be adjacent in $T(P_1)$ to a vertex or a Steiner point. The only vertices visible to any of the 9 peaks are the μ-points visible to q', r', s', but they lie outside $\overline{R}_1(c)$. We assumed there is no edge from q', r', s' to a point outside $\overline{R}_1(c)$. The other 6 peaks are not visible to any point outside $\overline{R}_1(c)$. Thus each of the 9 peaks must be adjacent to a Steiner point in $\overline{R}_1(c)$. No point in $\overline{R}_1(c)$ is visible to more than two peaks. Thus we need at least $\lceil \frac{9}{2} \rceil = 5$ Steiner points. The only way that 5 Steiner points suffice is to use 4 Steiner points that are each adjacent to two peaks. Pairs of peaks that are visible to a common point in both P_1 and P_2 are indicated by edges in the graph H shown in Figure 6(a). We require a matching of size 4 in H. Observe that H is bipartite so the maximum size of a matching is equal to the minimum size of a vertex cover. The set $\{q, r, s\}$ is a vertex cover of size 3. Thus there is no matching of size 4, and the Lemma follows. □

Lemma 3 *For any clause c, there are at least 5 Steiner points in $\overline{R}_1(c)$.*

Proof. Consider the triangulation of P_1. The case where no peak q', r', s' has an incident edge to a point outside $\overline{R}_1(c)$ is covered by Lemma 2. It remains to consider the cases when there is such an edge.

Our argument will be partly about the graph H (in Figure 6(a)) of pairs of peaks that are visible to a common point in both P_1 and P_2, and partly about the

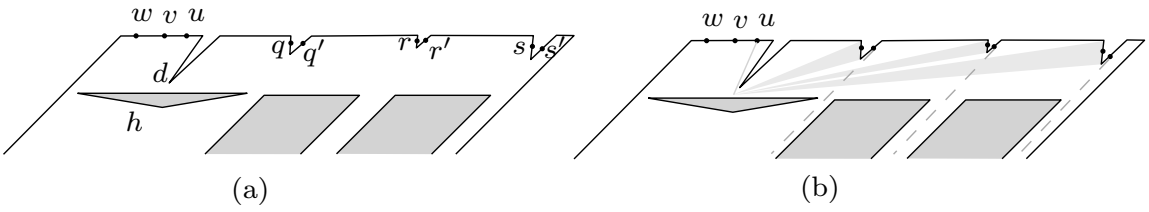

Figure 5: (a) A clause region in P_2. (b) The vertex u is not 1-bend visible to q', r', s'.

geometry of P_1. First we note that the argument used above in the proof of Lemma 2 can be strengthened to show that if we use just one edge from a peak to a point outside the clause region then we still need 5 Steiner points inside the region. In graph H, observe that if one of q', r', s' is removed, then we have 8 vertices, and a maximum matching of size 3, which means that we can use 3 edges (Steiner points) to cover 6 vertices, leaving 2 vertices that need one Steiner point each, for a total of 5 Steiner points. It remains to consider the cases where at least two of the points q', r', s' have an incident edge to a point outside the clause region. We deal with the case where q' has such an edge and the case where r' has such an edge but q' does not.

Suppose there is an edge e from q' to a point outside $\overline{R}_1(c)$. Observe that edge e cuts off the visibility regions of v and w. The effect on graph H is to remove the edges of H incident to v and w, e.g., see Figure 6(b). Thus we need one Steiner point for each of v and w, one Steiner point for r (irrespective of how r' is connected), one Steiner point for s and one more for u, a total of at least 5.

Next suppose there is no edge from q' to a point outside of $\overline{R}_1(c)$, but there is an edge e' from r' to a point outside of $\overline{R}_1(c)$. The edge e' cuts off the visibility region of w. The effect on graph H is to remove the edges (w, r) and (w, s), e.g., see Figure 6(c). We then need a Steiner point for s (irrespective of how s' is connected), and for the remaining 6 vertices $\{u, v, w, q, q', r\}$, we have a subgraph with a minimum vertex cover $\{q, r\}$ of size 2, thus a maximum matching of 2 edges (Steiner points) to cover 4 vertices, leaving 2 vertices that need one Steiner point each, for a total of 5 Steiner points. □

Lemma 4 *If $T(P_1)$ and $T(P_2)$ use $5|C|$ Steiner points each, then for any clause c, there is an edge in $T(P_1)$ from at least one of q', r', s' to a μ-point.*

Proof. By Lemma 3 every region $\overline{R}_1(c)$ has at least 5 Steiner points. Thus every such region must have exactly 5 Steiner points and there are no Steiner points in the variable regions. Suppose there is a clause c such that $T(P_1)$ has no edge from q', r' or s' to a μ-point. Then there is no edge from q', r' or s' to a point outside $\overline{R}_1(c)$. But then by Lemma 2 the clause region must have at least 6 Steiner points, a contradiction. □

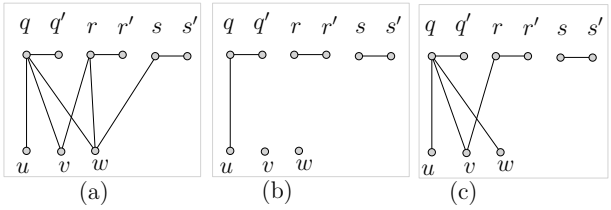

Figure 6: (a) Graph H of pairs of peaks that that are visible to a common Steiner point in both P_1 and P_2. (b)–(c) Illustration for Lemma 3.

3.4 Reduction

Theorem 5 *The following problem is NP-hard: Given a pair of compatible polygonal regions P_1, P_2, and $k \in \mathbb{N}$, decide if P_1 and P_2 have compatible triangulations with at most k Steiner points.*

Proof. Let $I = (U, C)$ be an instance of MRP-3SAT, and let P_1 and P_2 be the corresponding compatible polygons, as described in Sections 3.1–3.2. Appendix A presents further details on how to construct P_1 and P_2 using a polynomial number of bits, so this is a polynomial-time reduction. We now prove that P_1 and P_2 admit a pair of compatible triangulations, each with at most $5|C|$ Steiner points, if and only if I admits a satisfying truth assignment.

We first assume that P_1 and P_2 admit compatible triangulations with at most $5|C|$ Steiner points. By Lemma 4, for any clause c there is an edge in the triangulation of P_1 from at least one peak $z \in \{q', r', s'\}$ to a μ-point, say $\mu_{c,x}$. We use the edge $(z, \mu_{c,x})$ to assign a truth value to variable x. If c is a positive (resp., negative) clause, then we set x to true (resp., false). Clearly we have satisfied each clause. If there is a variable u whose truth value is not assigned yet, then setting the truth value of u arbitrarily would still keep the clauses satisfied. It remains to show that the truth-value assignment is consistent. Suppose there is a variable u such that some clause c forces u to be true, and some other clause c' forces u to be false. Without loss of generality we may assume that c is positive and c' is negative. Consequently, in each of $R_1(c)$ and $R_1(c')$, there exists a peak that is incident to some μ-point in $B_1(x)$. By construction of the μ-points in $B_1(x)$, the

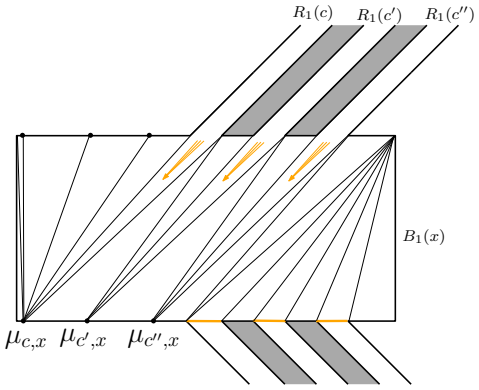

Figure 7: A triangulation for $B_1(x)$, where $x = true$.

two corresponding edges cross, a contradiction.

Assume now that I admits a satisfying truth assignment. We will find corresponding compatible triangulations of P_1 and P_2. For each variable x, if x is set to true, then we close the channels of the negative clauses and construct the compatible triangulations of the rectangles $B_1(x)$ and $B_2(x)$ using the μ-points on the bottom side of these rectangles, e.g., see Figure 7. The construction when x is set to false is symmetric.

Since every clause contains at least one true literal, for every clause c, there exist one or more peaks in P_1 that are visible to their corresponding μ-points. We show that in each scenario, the corresponding clause gadgets can be triangulated in a compatible fashion. We include the details in Appendix B. □

4 Conclusion

We have proved that computing compatible triangulations with at most k Steiner points is NP-hard for polygons with holes. The following questions are open:

1. Is the problem in NP? Is it complete for existential theory of the reals [13]?

2. What is the complexity of the problem for a pair of simple polygons? For a pair of rectangles with points inside?

3. How hard is it to decide if two polygonal regions, or two rectangles with points inside, have compatible triangulations with no Steiner points? For simple polygons, this can be decided in polynomial-time [3].

References

[1] O. Aichholzer, F. Aurenhammer, F. Hurtado, and H. Krasser. Towards compatible triangulations. *Theoretical Computer Science*, 296(1):3–13, 2003.

[2] 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. *to appear in SIAM Journal on Computing*, 2017.

[3] B. Aronov, R. Seidel, and D. L. Souvaine. On compatible triangulations of simple polygons. *Comput. Geom.*, 3:27–35, 1993.

[4] M. Babikov, D. L. Souvaine, and R. Wenger. Constructing piecewise linear homeomorphisms of polygons with holes. In *Proceedings of the 9th Canadian Conference on Computational Geometry, Kingston, Ontario, Canada*, 1997.

[5] W. V. Baxter III, P. Barla, and K.-i. Anjyo. Compatible embedding for 2D shape animation. *IEEE Transactions on Visualization and Computer Graphics*, 15(5):867–879, 2009.

[6] T. M. Chan, F. Frati, C. Gutwenger, A. Lubiw, P. Mutzel, and M. Schaefer. Drawing partially embedded and simultaneously planar graphs. *Journal of Graph Algorithms and Applications*, 19(2):681–706, 2015.

[7] M. de Berg and A. Khosravi. Optimal binary space partitions in the plane. In *Proceedings International Computing and Combinatorics Conference (COCOON 2010)*, volume 6196 of *LNCS*, pages 216–225. Springer, 2010.

[8] H. Gupta and R. Wenger. Constructing pairwise disjoint paths with few links. *ACM Transactions on Algorithms*, 3(3):26, 2007.

[9] E. Kranakis and J. Urrutia. Isomorphic triangulations with small number of Steiner points. *International Journal of Computational Geometry & Applications*, 9(2):171–180, 1999.

[10] J. Pach, F. Shahrokhi, and M. Szegedy. Applications of the crossing number. *Algorithmica*, 16(1):111–117, 1996.

[11] J. Pach and R. Wenger. Embedding planar graphs at fixed vertex locations. *Graphs and Combinatorics*, 17(4):717–728, 2001.

[12] A. Saalfeld. Joint triangulations and triangulation maps. In *Proceedings of the Third Annual Symposium on Computational Geometry (SoCG)*, pages 195–204. ACM, 1987.

[13] M. Schaefer. Complexity of some geometric and topological problems. In *17th International Symposium on Graph Drawing, (GD 2009)*, volume 5849 of *LNCS*, pages 334–344. Springer, 2010.

[14] D. L. Souvaine and R. Wenger. Constructing piecewise linear homeomorphisms. Technical report, DIMACS, New Brunswick, New Jersey, 1994.

[15] V. Surazhsky and C. Gotsman. High quality compatible triangulations. *Engineering with Computers*, 20(2):147–156, 2004.

[16] C. Thomassen. Deformations of plane graphs. *Journal of Combinatorial Theory, Series B*, 34(3):244–257, 1983.

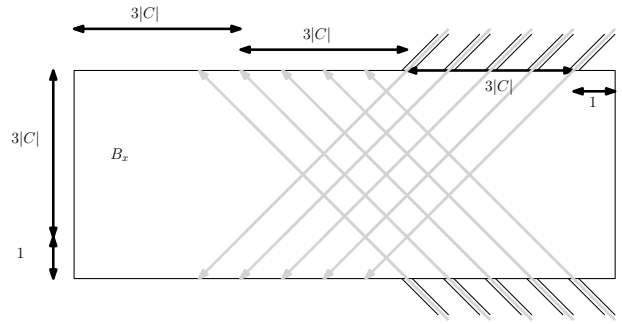

Figure 8: Transforming Γ into Γ'.

Appendix A

In this section we describe the construction details of P_1 and P_2. Furthermore, we show that the construction can be accomplished in polynomial time, in particular, with a polynomial number of bits for the coordinates of all points.

The construction has two stages. In the first stage we construct Γ' from Γ (see Figures 3(a)–(b)) and in the second stage we construct P_1 and P_2 from copies of Γ' by adding the appropriate dents (see Figures 4 and 5).

For the first stage we claim that Γ' can be constructed on a polynomial-sized grid (thus, with a logarithmic number of bits per coordinate). We make each variable rectangle $B(x)$ of width $9|C|+1$ and height $3|C|+1$, as illustrated in Figure 8. The length allocated for channel attachments is $3|C|$ to incorporate possible variable duplicates. The distance between successive μ-points on the bottom (top) of $B(x)$ is at least 1. We make each clause parallelogram R_c of height 1. The resulting drawing Γ' has height $O(|U|+|C|)$ and width $O(|U||C|)$.

We now turn to the second stage, the construction of P_1 and P_2. We refer the reader to Figure 11, which illustrates for each positive clause c, the correspondence between the vertices of $R_1(c)$ and $R_2(c)$. Observe that the left, top, and right sides of holes h_1 and h_2 have no dents added to them and remain straight line segments.

We construct the peaks $u, v, w, q, q', r, r', s, s'$ of $R_1(c)$ iteratively. Our plan is to first construct the visibility lines of the peaks (see Figure 9) and later enlarge these to visibility cones. Start by placing the points q', r', s' above the top boundary of $R_1(c)$ so that the lines from them to the corresponding μ-points are parallel to the channel sides and centered in the channels. Next, place points u, v, w to the left of $R_1(c)$ and in the top half of $R_1(c)$, and choose points u', v', w' on the boundary of $B(x)$ to be the endpoints of the visibility lines emanating from u, v, w respectively. Specifically, choose u' and v' by taking the midpoints of the tops of h_1 and h_2 and projecting upward at 45°.

We can explicitly compute the equations of the 6 visibility lines emanating from q', r', s', u, v, w and we can compute the intersection points formed by them. Let $p_{u,q'}$ be the intersection point of the visibility line of u and the visibility line of q', etc. We can next choose points a, b, c on the visibility lines of q', r', s', respectively, such that each point is in-between the appropriate p points. For example, point a is the midpoint of $p_{v,q'}$ and $p_{u,q'}$. From these, we can construct points q, r, s and their visibility lines through a, b, c to appropriate points on the channels. Observe that because a, b, c lie in the upper half of $B_1(x)$, points q, r, s lie in the extensions of the channels, and remain to the right of d, u', v', respectively.

It remains to enlarge the visibility lines to cones. We can do this by explicitly computing a tolerance τ such that if the width of every cone is at most τ inside P_1 then no visibility cone will contain points it should not, and no two visibility cones will intersect when they should not. Since we only need a lower bound on τ, this can be done with a polynomial number of bits. Finally, from τ we can explicitly choose the points where each visibility cone intersects the boundary of P_1 (see for example points v_a and v_b in Figure 10), and from these we can compute the dent vertices for each peak (see points v_1 and v_2 in the same figure).

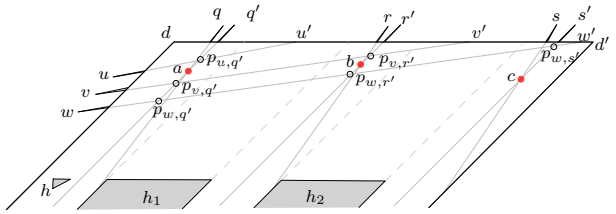

Figure 9: Details of the construction of $R_1(c)$.

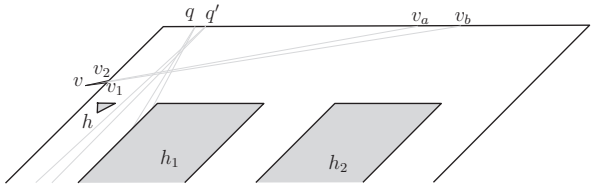

Figure 10: Details of visibility cone construction for v.

Finally, we place the hole h in the first channel of $R_1(c)$ such that the base of h is aligned with the top sides of h_1 and h_2. Furthermore, we ensure that h remains to the left of the visibility region of q'.

The construction of the channels and inward dents for $R_2(c)$ is simpler compared to $R_1(c)$. We choose the inward dent D with peak d such that the entire left side of each remaining inward dent is visible to d, as illustrated using dotted lines in Figure 11(b). Furthermore, we ensure that exactly one μ-point is visible to each of q', r', s'. Figure 11(b) illustrates these visibilities with dashed lines.

The placement of the triangular hole h is similar to that of $R_1(c)$. Here we ensure an additional constraint that the base of h must be large enough to block any 1-bend visibility between $\{u, v, w\}$ and $\{q', r', s'\}$.

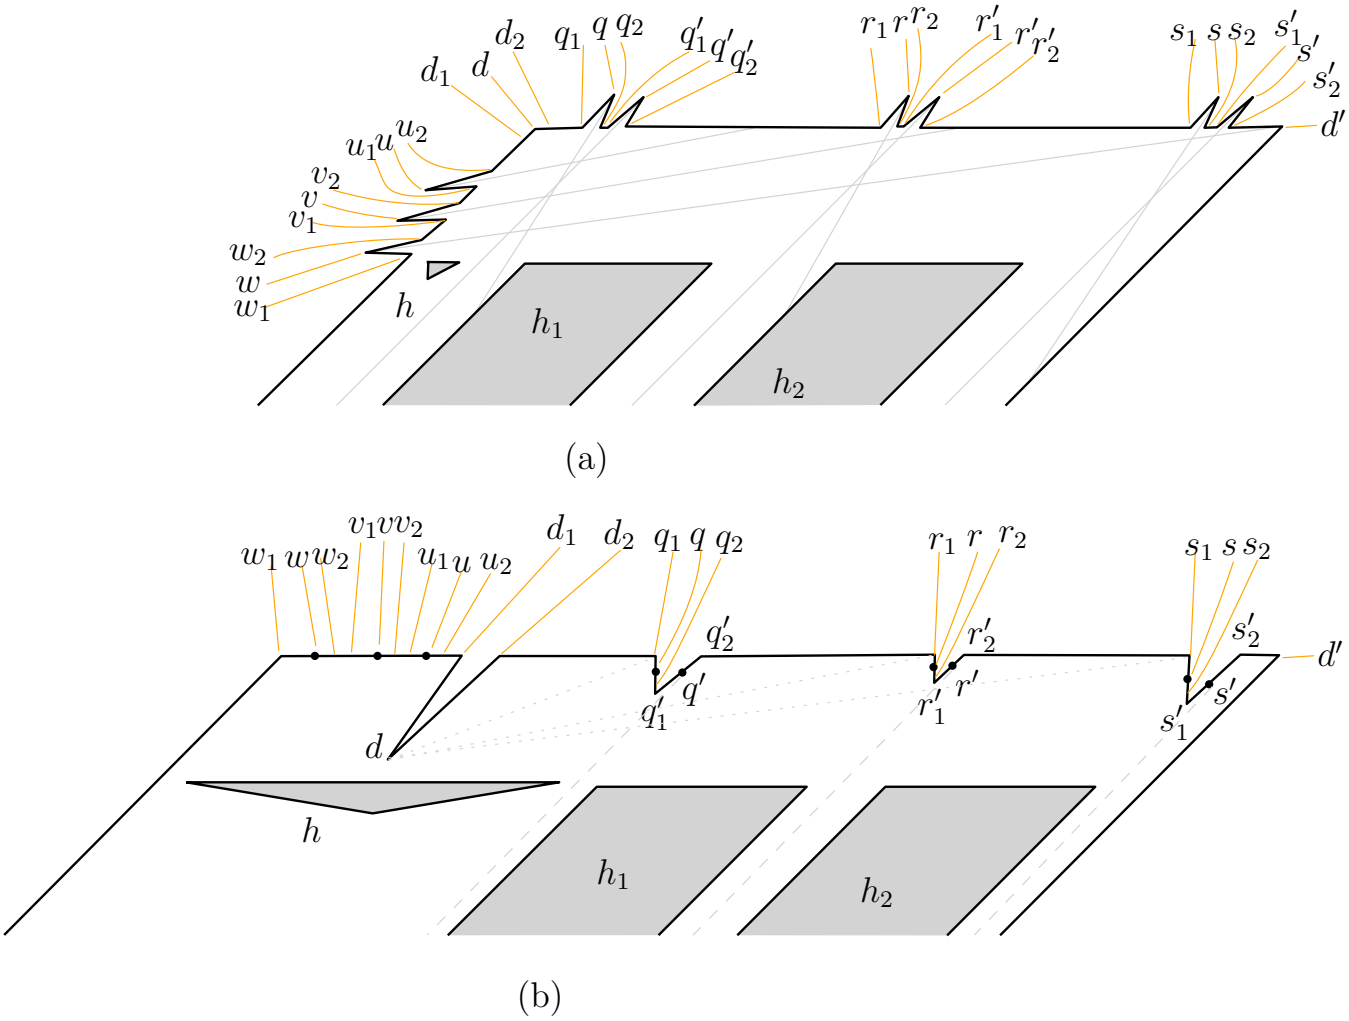

Figure 11: (a)–(b) Schematic representations for $R_1(c)$ and $R_2(c)$.

Appendix B

In this appendix we fill in the remaining details of the proof of Theorem 5. These details concern the second half of the proof, where we assume that I admits a satisfying truth assignment and we find corresponding compatible triangulations of P_1 and P_2. As explained in the main text, for each variable x, if x is set to true, then we close the channels of the negative clauses and construct the compatible triangulations of the rectangles $B_1(x)$ and $B_2(x)$ using the μ-points on the bottom side of these rectangles, e.g., see Figure 7. The construction when x is set to false is symmetric.

Since every clause contains at least one true literal, for every clause c, there exist one or more peaks in P_1 that are visible to their corresponding μ-points. It remains to show that in each scenario, the corresponding clause gadgets can be triangulated in a compatible fashion.

We only describe the case when c is a positive clause. The case when c is negative is symmetric.

Let x_q, x_r, x_s be the literals of c, and assume that their corresponding peaks q', r', s' appear in this order from left to right in $R_1(c)$.

Case 1 ($x_q = true$): Figures 12(a)–(b) illustrate the compatible triangulations of $R_1(c)$ and $R_2(c)$ for the case when $x_q = true$ and $x_r = x_s = false$. Figures 12(c)–(d) illustrate the compatible triangulations of $R_1(c)$ and $R_2(c)$ for the case when $x_q = true$ and $x_r = x_s = true$. The scenarios when $x_r = true$ and $x_s = false$, or vice versa, can be handled by switching between the local configurations corresponding to the true and false values.

Case 2 ($x_q = false$, $x_r = true$): Figures 13(a)–(b) illustrate the compatible triangulations of $R_1(c)$ and $R_2(c)$ for the case when $x_s = true$. The scenario when $x_s = false$ can be handled by switching between the local configurations corresponding to the true and false values.

Case 3 ($x_q = x_r = false$, $x_s = true$): The compatible triangulations of $R_1(c)$ and $R_2(c)$ for this case are illustrated in Figures 13(c)–(d).

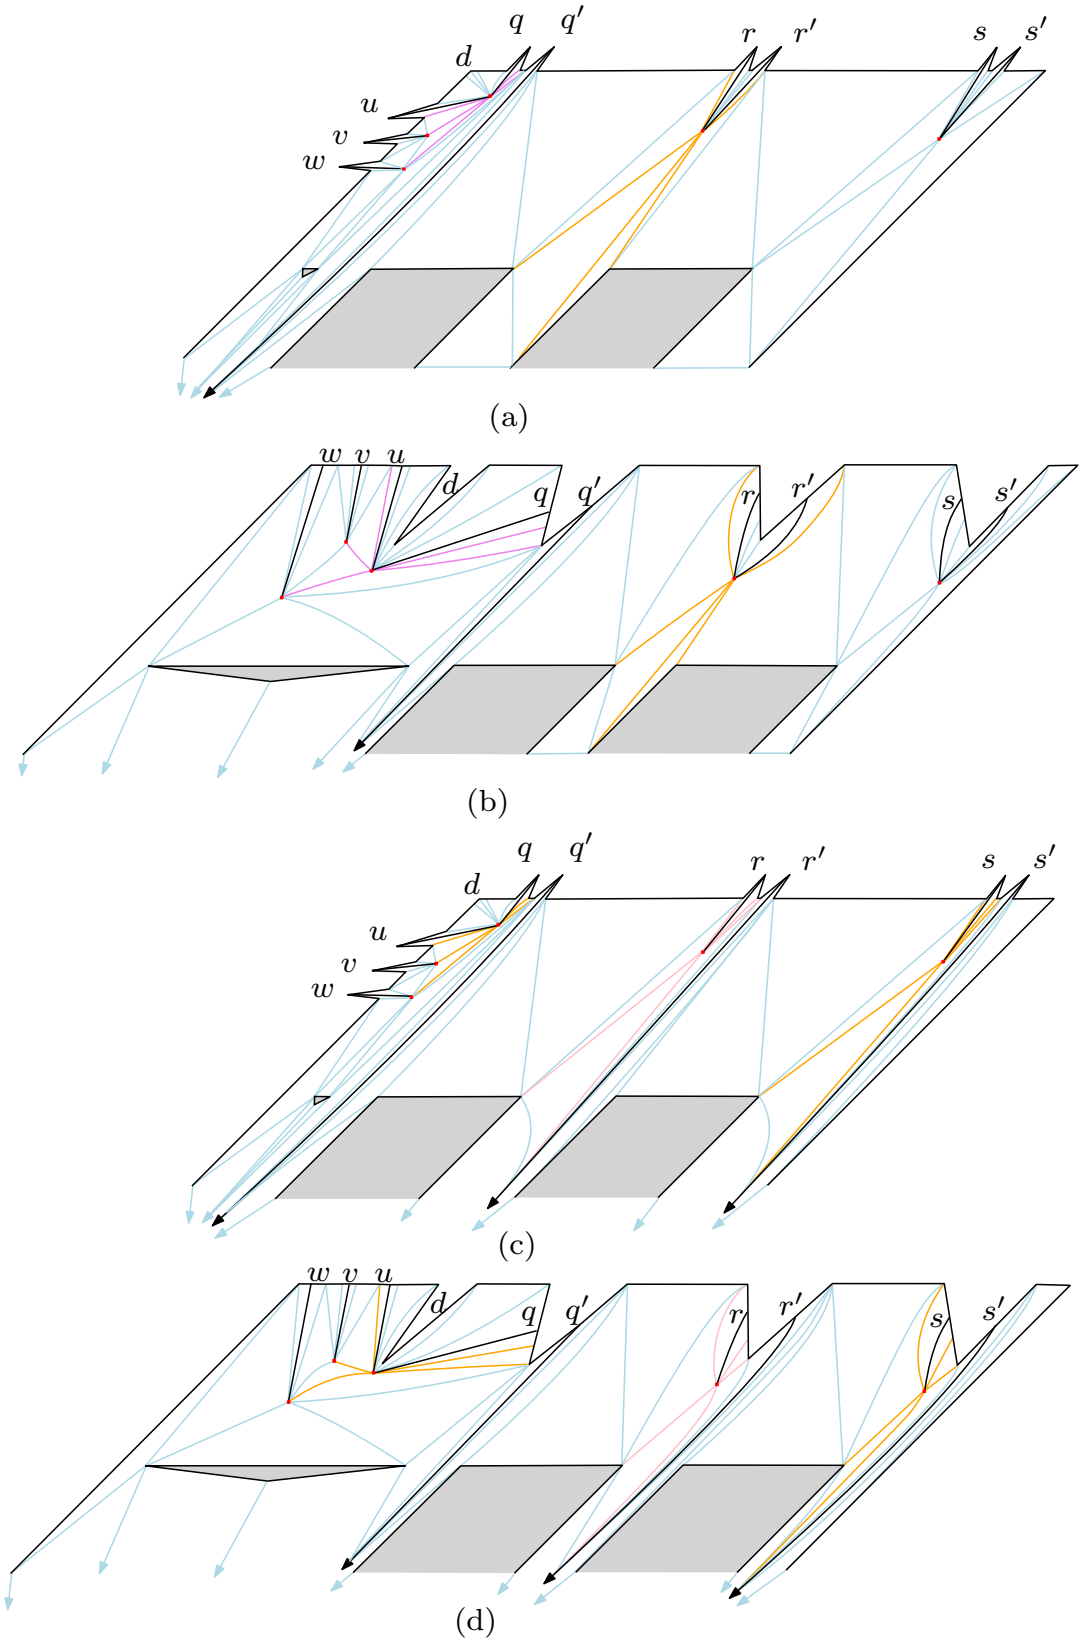

Figure 12: Illustration for Case 1. Colors are used to better illustrate the correspondence between the two drawings. Some edges are drawn curved but that is only to make the drawing more readable.

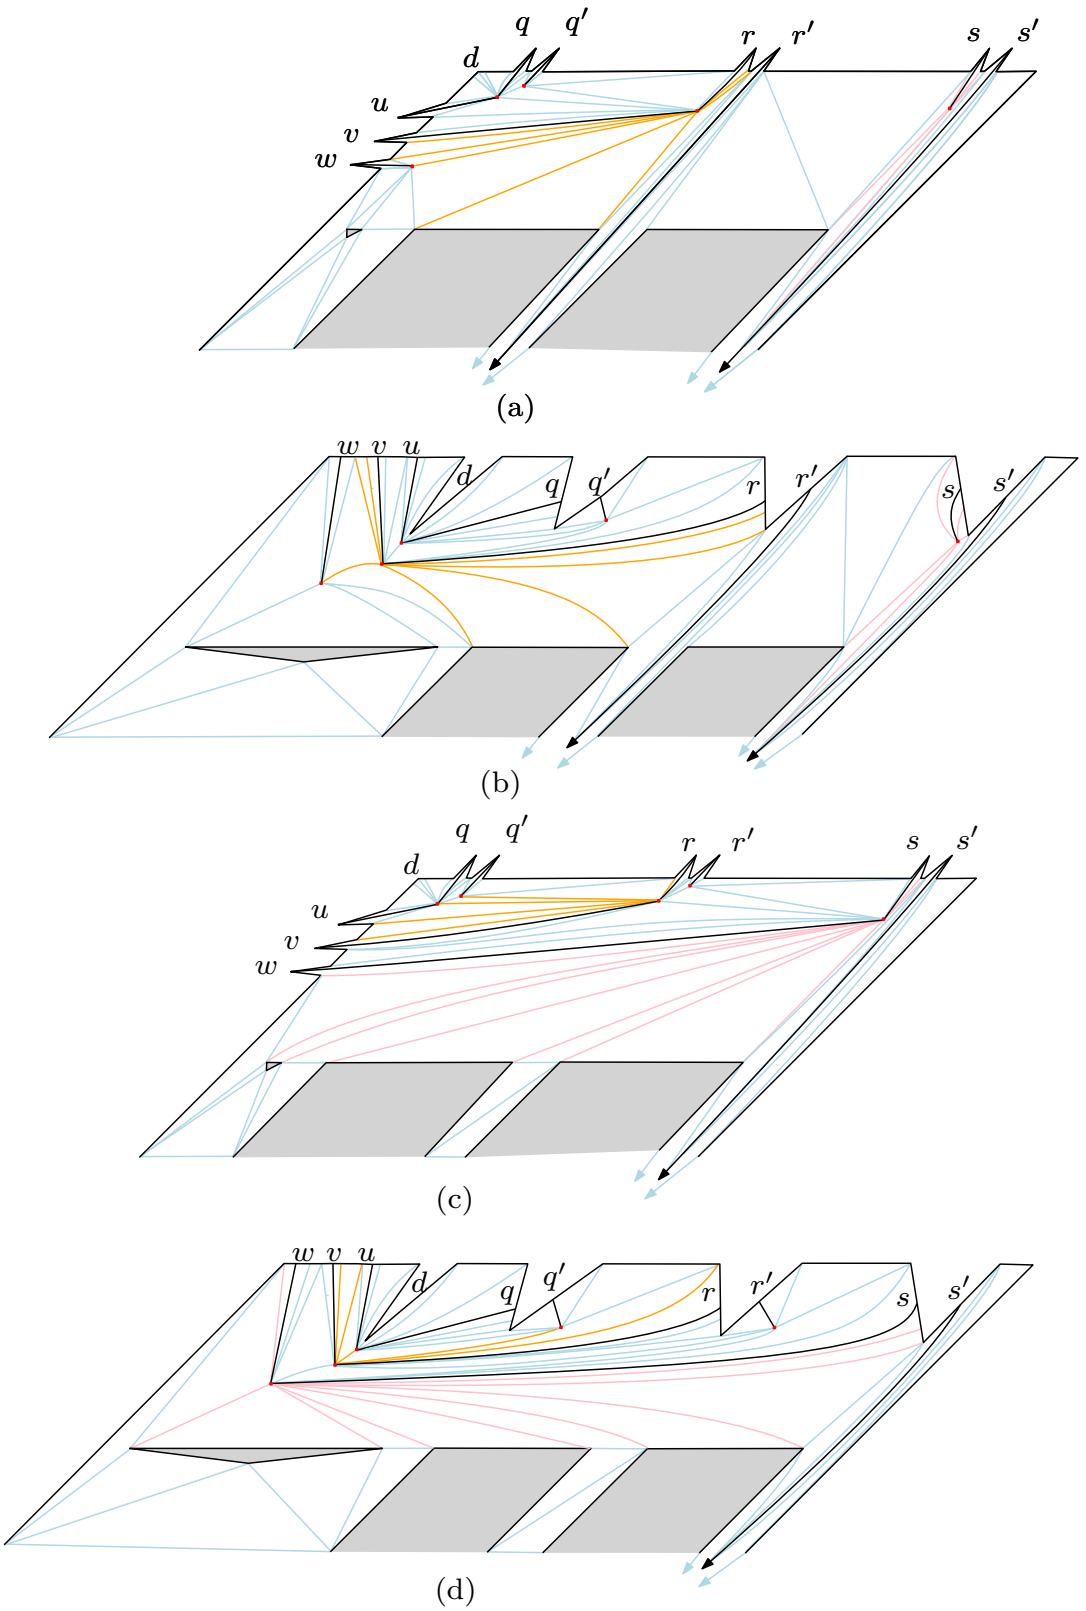

Figure 13: Illustration for (a)–(b) Case 2, and (c)–(d) Case 3. Some edges are drawn curved but that is only to make the drawing more readable.