

Edges and Switches, Tunnels and Bridges

D. Eppstein

M. van Kreveld

E. Mumford

B. Speckmann

Department of Information and Computing Sciences, Utrecht University

Technical Report UU-CS-2007-042

www.cs.uu.nl

ISSN: 0924-3275

Edges and Switches, Tunnels and Bridges

D. Eppstein¹ M. van Kreveld² E. Mumford³ B. Speckmann³

¹Department of Computer Science, University of California, Irvine, USA.

`eppstein@ics.uci.edu`

²Department of Information and Computing Sciences, Utrecht University, The Netherlands.

`marc@cs.uu.nl`

³ Department of Mathematics and Computer Science, TU Eindhoven, The Netherlands.

`e.mumford@tue.nl` and `speckman@win.tue.nl`

Abstract

Edge casing is a well-known method to improve the readability of drawings of non-planar graphs. A cased drawing orders the edges of each edge crossing and interrupts the lower edge in an appropriate neighborhood of the crossing. Certain orders will lead to a more readable drawing than others. We formulate several optimization criteria that try to capture the concept of a “good” cased drawing. Further, we address the algorithmic question of how to turn a given drawing into an optimal cased drawing. For many of the resulting optimization problems, we either find polynomial time algorithms or NP-hardness results.

1 Introduction

Drawings of non-planar graphs necessarily contain edge crossings. The vertices of a drawing are commonly marked with a disk, but it can still be difficult to detect a vertex within a dense cluster of edge crossings. *Edge casing* is a well-known method—used, for example, in electrical drawings and, more generally, in information visualization—to alleviate this problem and to improve the readability of a drawing. A *cased drawing* orders the edges of each crossing and interrupts the lower edge in an appropriate neighborhood of the crossing. One can also envision that every edge is encased in a strip of the background color and that the casing of the upper edge covers the lower edge at the crossing. See Fig. 1 for an example.

If there are no application specific restrictions that dictate the order of the edges at each crossing, then we can in principle choose freely how to arrange them. Certain orders will lead to a more readable drawing than others. In this paper we formulate several optimization criteria that try to capture the concept of a “good” cased drawing. Further, we address the algorithmic question of how to turn a given drawing into an optimal cased drawing.

Definitions. Let G be a graph with n vertices and m edges and let D be a drawing of G with k crossings. We want to turn D into a cased drawing where the width of the casing is given in the variable *casingwidth*. To avoid that the casing of an edge covers a vertex we assume that no vertex v of D lies on (or very close to) an edge e of D unless v is an endpoint of e . Further, no more than two edges of D cross in one point and any two crossings are far enough apart so that the casings of the edges involved do not interfere. With these assumptions we can consider crossings independently. Without these restrictions the problem changes significantly—optimization problems that are solvable in polynomial time can become NP-hard. In Section 5 we discuss in detail the effects of removing the restrictions on the input.

We define the *edge crossing graph* G_{DC} for D as follows. G_{DC} contains a vertex for every edge of D and an edge for any two edges of D that cross. Let C be a crossing between two edges e_1 and e_2 . In

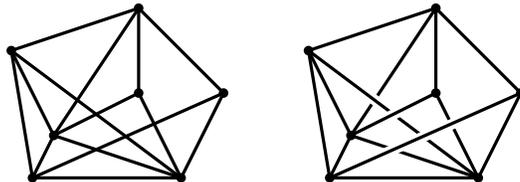


Figure 1: Normal and cased drawing of a graph.

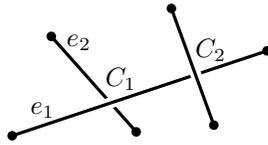


Figure 2: Tunnels and bridges.

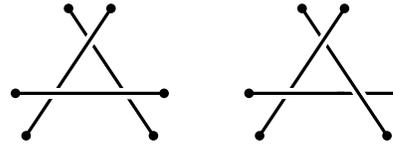


Figure 3: Stacking and weaving.

a cased drawing either e_1 is drawn on top of e_2 or vice versa. If e_1 is drawn on top of e_2 then we say that C is a *bridge* for e_1 and a *tunnel* for e_2 . In Fig. 2, C_1 is a bridge for e_1 and a tunnel for e_2 . The *length* of a tunnel is $\text{casewidth} / \sin \alpha$, where $\alpha \leq \pi/2$ is the angle of the edges at the crossing. A pair of consecutive crossings C_1 and C_2 along an edge e is called a *switch* if C_1 is a bridge for e and C_2 is a tunnel for e , or vice versa. In Fig. 2, (C_1, C_2) is a switch.

Stacking and weaving. When we turn a given drawing into a cased drawing, we need to define a drawing order for every edge crossing. We can choose to either establish a global top-to-bottom order on the edges, or to treat each edge crossing individually. We call the first option the *stacking model* and the second one the *weaving model*, since cyclic overlap of three or more edges can occur (see Fig. 3).

Quality of a drawing. Globally speaking, two factors may influence the readability of a cased drawing in a negative way. Firstly, if there are many switches along an edge then it might become difficult to follow that edge. Drawings that have many switches can appear somewhat chaotic. Secondly, if an edge is frequently below other edges, then it might become hardly visible. These two considerations lead to the following optimization problems for a drawing D .

MINTOTALSWITCHES Minimize the total number of switches.

MINMAXSWITCHES Minimize the maximum number of switches for any edge.

MINMAXTUNNELS Minimize the maximum number of tunnels for any edge.

MINMAXTUNNELLENGTH Minimize the maximum total length of tunnels for any edge.

MAXMINTUNNELDISTANCE Maximize the minimum distance between any two consecutive tunnels.

Cased drawings are also used in certain types of mazes where paths may cross (see, for example, Fig. 4). For this application the paths between junctions should be difficult to follow and the whole picture should look somewhat chaotic. This motivates the problem of maximizing the number of switches of a drawing D .

MAXTOTALSWITCHES Maximize the total number of switches.

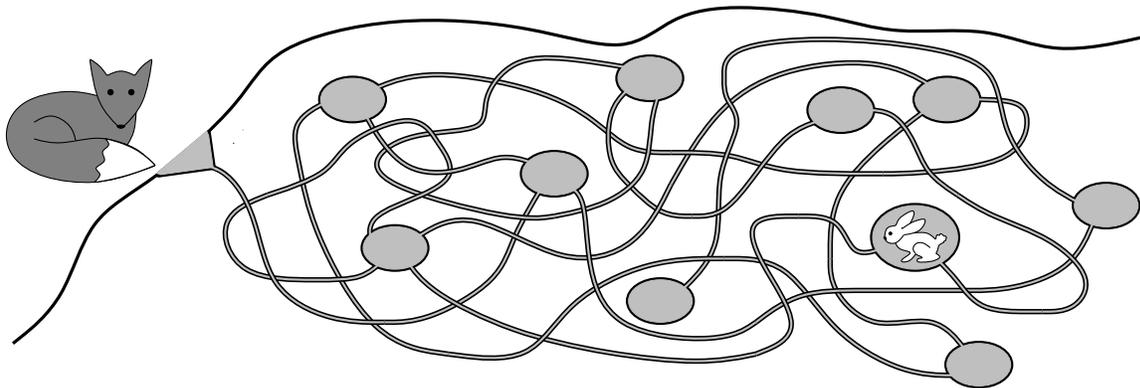


Figure 4: "Help the fox eat the rabbit."

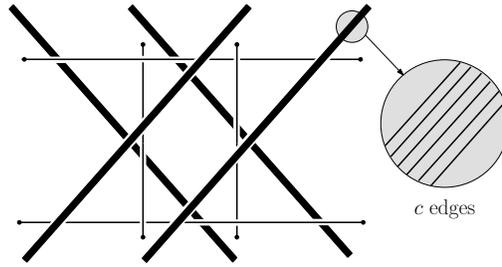


Figure 5: Optimal drawing in the weaving model for `MINTOTALSWITCHES`.

Fig. 5 illustrates that the weaving model is stronger than the stacking model for `MINTOTALSWITCHES`—no cased drawing of this graph in the stacking model can reach the optimum of four switches. For, the thickly drawn bundles of $c > 4$ parallel edges must be cased as shown (or its mirror image) else there would be at least c switches in a bundle, the four vertical and horizontal segments must cross the bundles consistently with the casing of the bundles, and this already leads to the four switches that occur as drawn near the midpoint of each vertical or horizontal segment. Thus, any deviation from the drawing in the casing of the four crossings between vertical and horizontal segments would create additional switches. However, the drawing shown is not a stacked drawing.

Related work. If we consider only simple arrangements of line segments in the plane as our initial drawing, then there is a third model to consider, an intermediate between stacking and weaving: drawings which are plane projections of line segments in three dimensions. We call this model the *realizable model*. Clearly every cased drawing in the stacking model is also a drawing in the realizable model, but not every cased drawing in the weaving model can be realized (see [10]). The optimal drawing in Fig. 5 can be realized, hence the realizable model is stronger than the stacking model.

Fig. 6 shows that the weaving model is stronger than the realizable model for `MINTOTALSWITCHES`—no cased drawing of this graph in the realizable model can reach the optimum number of 12 switches. The reasoning is quite similar to the one employed for the construction in Fig. 5. The thickly drawn bundles of $c > 12$ parallel edges must be cased as shown (or its mirror image) else there would be at least c switches in each bundle. The 8 vertical and horizontal “normal” edges must cross the bundles consistently with the casing of the bundles, and this already leads to the 12 switches as drawn in this figure. Thus, any deviation from the drawing in the casing of the 16 crossings between the normal edges would create additional switches. However, the normal edges as drawn here constitute a perfect 4×4 weaving which can not be realized (see [10]).

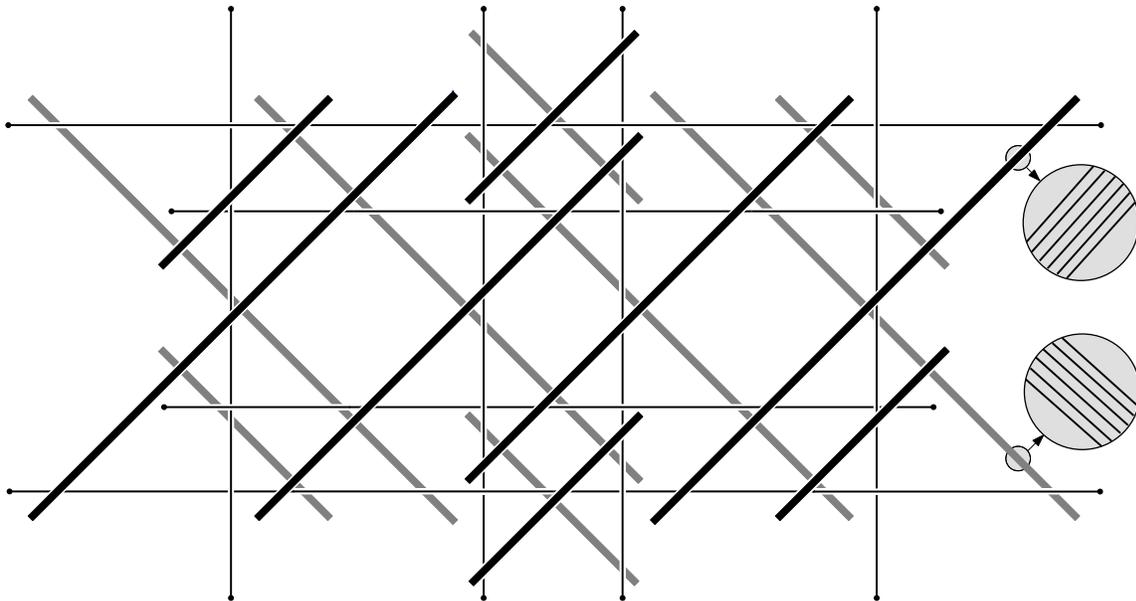


Figure 6: Optimal drawing in the weaving model for `MINTOTALSWITCHES`.

Model	Stacking	Weaving
MINTOTALSWITCHES	<i>open</i>	$O((q+1)k + q^{5/2} \log^{3/2} k)$
MAXTOTALSWITCHES	<i>open</i>	$O((t+1)k + t^{5/2} \log^{3/2} k)$
MINMAXSWITCHES	<i>open</i>	<i>open</i>
MINMAXTUNNELS	$O(m \log m + k)$ <i>exp.</i>	$O(m^4)$
MINMAXTUNNELLENGTH	$O(m \log m + k)$ <i>exp.</i>	NP-hard
MAXMINTUNNELDISTANCE	$O(m \log m + k \log m)$ <i>exp.</i>	$O((m+K) \log m)$ <i>exp.</i>

Table 1: Table of results: n is the number of vertices, m is the number of edges, $k = O(m^2)$ is the number of crossings, $K = O(m^3)$ is the total number of pairs of crossings on the same edge, $q = O(k)$ is the number of *odd face polygons*, and $t = O(k)$ is the number of *s-odd face polygons*.

Results. For many of the problems described above, we either find polynomial time algorithms or NP-hardness results in both the stacking and weaving models. We summarize our results in Table 1. All our bounds assume that $m = \Omega(n)$ and that $k = \Omega(m)$. Furthermore, we assume that the input drawing is given as an arrangement, in particular, all crossings are given as vertices of this arrangement. Our results are derived for straight line input drawings, but several of them also generalize to curved drawings. Section 2 presents the results concerning the optimization problems that seek to minimize the number of switches, Section 3 gives our results on maximizing the number of switches, and Section 4 discusses our solutions to the optimization problems that concern the tunnels. In Section 5 we show that MINTOTALSWITCHES becomes NP-hard in both the weaving and the stacking model if we allow more than three edges to cross in one point. We conclude with some open problems.

2 Minimizing switches

In this section we discuss results related to the MINTOTALSWITCHES and MINMAXSWITCHES problems. We first discuss some non-algorithmic results giving simple bounds on the number of switches needed, and recognition algorithms for graphs needing no switches. As we know little about these problems for the stacking model, all results stated in this section will be for the weaving model.

Lemma 1 *Given a drawing D of a graph we can turn D into a cased drawing without any switches if and only if the edge crossing graph G_{DC} is bipartite.*

Corollary 2 *Given a drawing D of a graph we can decide in $O((n+m) \log(n+m))$ time if D can be turned into a cased drawing without any switches.*

Proof. We apply the bipartiteness algorithm of [3]. Note that this does not construct the arrangement, so there is no term with k in the runtime. \square

Define a *vertex-free cycle* in a drawing of a graph G to be a face f formed by the arrangement of the edges in the drawing, such that there are no vertices of G on the boundary of f . An *odd vertex-free cycle* is a vertex-free cycle composed of an odd number of segments of the arrangement.

Lemma 3 *Let f be an odd vertex-free cycle in a drawing D . Then in any casing of D , there must be a switch on one of the segments of f .*

Proof. Unless there is a switch, the segments must alternate between those that cross above the previous segment, and those that cross below the previous segment. However, this alternation cannot continue all the way around an odd cycle, for it would end up in an inconsistent state from how it started. \square

Lemma 4 *Given a drawing D of a graph the minimum number of switches of any cased drawing obtained from D is at least half of the number of odd vertex-free cycles in D .*

Proof. Let o be the number of odd vertex-free cycles in D . By Lemma 3, each odd vertex-free cycle must have a switch on one of its segments. Choose one such switch for each cycle; then each segment belongs to at most two vertex-free cycles, so these choices group the odd cycles into pairs of cycles sharing a common switch, together with possibly some unpaired cycles. The number of pairs and unpaired cycles must be at least $o/2$, so the number of switches must also be this large. \square

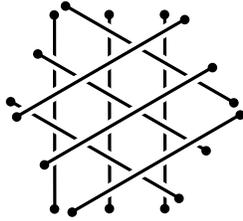


Figure 7: $O(n)$ edges and $\Omega(n^2)$ triangles.

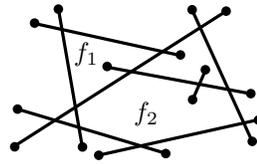


Figure 8: A degree-one graph, f_1 is an odd polygon and f_2 is an even polygon.

Lemma 5 For any n large enough, a drawing of a graph G with n vertices and $O(n)$ edges exists for which any crossing choice gives rise to $\Omega(n^2)$ switches.

Proof. A construction with three sets of parallel lines, each of linear size, gives $\Omega(n^2)$ vertex-free triangles, and each triangle gives at least one switch (see Fig. 7). \square

Lemma 6 For any n large enough, a drawing of a graph G with n vertices and $O(n^2)$ edges exists for which any crossing choice gives rise to $\Omega(n^4)$ switches.

Proof. We build our graph as follows: make a very elongated rectangle, place $n/6$ vertices equally spaced on each short edge, and draw the complete bipartite graph. This graph has $(n/6)^2$ edges. One can prove that there is a strip parallel to the short side of the rectangle, such that the parts of the edges inside the strip behave in the same way as parallel ones do with respect to creating triangles when overlapped the way it is described in the previous lemma. This gives us the desired graph with $\Omega(n^4)$ triangles, and hence with $\Omega(n^4)$ switches. \square

We define a *degree-one graph* to be a graph in which every vertex is incident to exactly one edge; that is, it must consist of a collection of disconnected edges (see Fig. 8).

Lemma 7 Let D be a drawing of a graph G . Then there exists a drawing D' of a degree-one graph G' , such that the edges of D correspond one-for-one with the edges of D' , casings of D correspond one-for-one to casings of D' , and switches of D correspond one-for-one with switches of D' .

Proof. Form G' by placing a small circle around each vertex of G . Given an edge $e = (u, v)$ in G , let u_e be the point where e crosses the circle around u and similarly let v_e be the point where e crosses the circle around v . Form D' and G' by replacing each edge $e = (u, v)$ in G by the corresponding edge (u_e, v_e) , drawn as the subset of edge e connecting those points.

As these replacements do not occur between any two crossings along any edge, they do not affect the switches on the edge. Both drawings have the same set of crossings, and any switch in a casing of one drawing gives rise to a switch in the corresponding casing of the other drawing. \square

In a drawing of a degree-one graph, define a *polygon* to be a sequence of segments of the arrangement formed by the drawing edges that forms the boundary of a simple polygon in the plane. Define a *face polygon* to be a polygon that forms the boundary of the closure of a face of the arrangement; note that there may be edges drawn in the interior of this polygon, as long as they do not separate it into multiple components.

Lemma 8 In a drawing of a degree-one graph, there can be no vertex on any segment of a polygon.

Proof. We have already required that no vertex can lie on an edge unless it is the endpoint of an edge. And, if a segment contains the endpoint of an edge, it cannot continue past the endpoint to form the boundary of a polygon. \square

Note, however, that a polygon can contain vertices in its interior. Define the *complexity* of a polygon to be the number of segments forming it, plus the number of graph vertices interior to the polygon. We say that a polygon is *odd* if its complexity is an odd number, and *even* if its complexity is an even number (see Fig. 8).

Lemma 9 *Let p be a polygon in a drawing of a degree-one graph. Then, modulo two, the complexity of p is equal to the sum of the complexities of the face polygons of faces within p .*

Proof. Each segment of p contributes one to the complexity of p and one to the complexity of some face polygon. Each vertex within p contributes one to the complexity of p and one to the complexity of the face that contains it. Each segment within the interior of p either separates two faces, and contributes two to the total complexity of faces within p , or does not separate any face and contributes nothing to the complexity. Thus in each case the contribution to p and to the sum of its faces is the same modulo two. \square

Lemma 10 *Let p be an odd polygon in a drawing of a degree-one graph. Then there exists an odd face polygon in the same drawing.*

Proof. By Lemma 9, the complexity of p has the same parity as the sum of the complexities of its faces. Therefore, if p is odd, it has an odd number of odd faces, and in particular there must be a nonzero number of odd faces. \square

Lemma 11 *Let D be a drawing of a degree-one graph. Then D has a casing with no switches if and only if it has no odd face polygon.*

Proof. As we have seen, D has a casing with no switches if and only if the edge crossing graph is bipartite. This graph is bipartite if and only if it has no odd cycles, and an odd cycle in the edge crossing graph corresponds to an odd polygon in D . For, if C is an odd cycle in the edge crossing graph, it must lie on a polygon p of D . Each crossing in C contributes one to the complexity of this polygon. Each edge of D that crosses p without belonging to C either crosses it an even number of times (contributing that number of additional segments to the complexity of p) and has both endpoints inside p or both outside p , or it crosses an odd number of times and has one endpoint inside p ; thus, it contributes an even amount to the complexity of p . Thus, p must be an odd polygon. By Lemma 10, there is an odd face polygon in D . Conversely, any odd face polygon in D can be shown to form an odd cycle in the edge crossing graph. \square

Theorem 12 *MINTOTALSWITCHES in the weaving model can be solved in time $O((q+1)k+q^{5/2} \log^{3/2} k)$, where k denotes the number of crossings in the input drawing and q denotes the number of its odd face polygons.*

Proof. Let D be the drawing which we wish to case for the minimum number of switches. By Lemma 7, we may assume without loss of generality that each vertex of D has degree one.

We apply a solution technique related to the Chinese Postman problem, and also to the problem of via minimization in VLSI design [2]: form an auxiliary graph G° , and include in G° a single vertex for each odd face polygon in D . Also include in G° an edge connecting each pair of vertices, and label this edge by the number of segments of the drawing that are crossed in a path connecting the corresponding two faces in D that crosses as few segments as possible. We claim that the minimum weight of a perfect matching in G° equals the minimum total number of switches in any casing of D .

In one direction, we can case D with a number of switches equal to or better than the weight of the matching, as follows: for each edge of the matching, insert a small break into each of the segments in the path corresponding to the edge. The resulting broken arrangement has no odd face cycles, for the breaks connect pairs of odd face cycles in D to form larger even cycles. Therefore, by Lemma 11, we can case the drawing with the breaks, without any switches. Forming a drawing of D by reconnecting all the break points adds at most one switch per break point, so the total number of switches equals at most the weight of the perfect matching.

In the other direction, suppose that we have a casing of D with a minimum number of switches; we must show that there exists an equally good matching in G° . To show this, consider the drawing formed by inserting a small break in each segment of D having a switch. This eliminates all switches in the drawing, so by Lemma 11, the modified drawing has no odd face polygons. Consider any face polygon in the modified drawing; by Lemma 10 it must include an even number of odd faces in the original drawing. Thus, the odd faces of D are connected in groups of evenly many faces in the modified drawing, and within each such group we can connect the odd faces in pairs by paths of breaks in the drawing, giving a matching in G° with total weight at most equal to the number of switches in D .

The number of vertices of the graph G^o is $O(q)$, where q is the number of odd face polygons in D . We can construct G^o in time $O(k + qk)$ where k is the number of crossings in D by using breadth-first search in the arrangement dual to D to find the distances from each vertex to all other vertices. A minimum weight perfect matching in a complete weighted graph with integer weights bounded by k can be found in time $O(q^{5/2} \log^{3/2} k)$ using the algorithm of Gabow and Tarjan [5]. Therefore the time for this algorithm is $O((q + 1)k + q^{5/2} \log^{3/2} k)$. \square

For MINMAXSWITCHES, we have the following weaker result.

Theorem 13 *If the edge crossing graph G_{DC} is planar, then we can assure at most four switches along any edge in the weaving model.*

Proof. Find an edge e that intersects at most five others, take e out, and recursively treat the rest. Put e back in such a way that it does not increase the switch count of any of the edges it intersects. Edge e itself will have at most four switches. \square

3 Maximizing switches

In this section we describe a solution for the MAXTOTALSWITCHES problem in the weaving model.

We define a *twin* along an edge e in a cased drawing to be a pair of consecutive crossings along e such that the crossings are either both bridges or both tunnels for e . Hence every pair of consecutive crossings along an edge e in a cased drawing forms either a twin or a switch for e . Let n_p be the number of all pairs of consecutive crossings along all the edges of a drawing D , and let n_t and n_s be the number of twins and switches, respectively, in a given casing of D . Clearly $n_p = n_s + n_t$. Thus the problem of finding the maximal number of switches for a cased drawing is equivalent to the problem of finding the minimal number of twins. We call a cased drawing without twins a *twin-free casing* or a *perfect weaving*. The cased drawing obtained from D_c by turning bridges into tunnels and vice versa is the *inversion* of D_c and denoted by $D_{\bar{c}}$. If D_c is a perfect weaving then so is $D_{\bar{c}}$.

Thanks to Lemma 7 we can again limit ourselves to considering degree-one graphs only. Let p be a vertex of a polygon (which is defined as in the previous section). We say that p is a *side vertex* of the polygon if the edges of the polygon form a 180° angle at p , and p is called a *corner vertex* otherwise (see Fig. 9: side vertices are depicted as white squares, corner vertices as black squares). We define the *s-complexity* of a polygon as the number of its side vertices. A polygon of odd s-complexity is called *s-odd*.

Lemma 14 *Let f be an s-odd polygon of a degree-one graph drawing D , then for any casing of D one of the edges forming f has a twin.*

Proof. Let e_1, \dots, e_k be the edges of D forming f in clockwise order. Then the first crossing of e_i along f in clockwise order is $c_{i,1} = e_i \cap e_{i-1}$, for $1 < i \leq k$, $c_{1,1} = e_1 \cap e_k$, and the last crossing of e_i along f in clockwise order is $c_{i,2} = e_i \cap e_{i-1}$ for $1 \leq i < k$, $c_{k,2} = e_k \cap e_1$.

Assume that a cased drawing D_c of D exists where none of e_i , $1 \leq i \leq k$, has a twin. Then if an edge e_i has an odd number of side vertices between $c_{i,1}$ and $c_{i,2}$, then $c_{i,1}$ and $c_{i,2}$ are cased the same way, otherwise they are cased differently. By symmetry we may assume that $c_{1,1}$ is a bridge. Since in total f has an odd number of side vertices, $c_{k,2}$ must be a bridge as well. But this is impossible because $c_{1,1}$ and $c_{k,2}$ are the same crossing. Hence for any cased drawing of D there is an edge in f that has a twin. \square

Lemma 15 *Let f be a polygon in a drawing of a degree-one graph. Then the parity of the number of side vertices of f is equal to the parity of the number of graph vertices inside f .*

Proof. If an edge e of D contains part of the boundary of f we call it a *boundary edge* for f , otherwise we say that e is an *external edge* for f . Let e be an external edge for f . Each crossing of e with the boundary of f is a side vertex of f —see Fig. 9. If e intersects the boundary of the polygon an odd number of times, then it has only one vertex inside it. If it intersects the boundary of f an even number of times, then it either has both vertices inside f or both outside f . Hence the parity of the number $n_{s,1}$ of times that edges external for f cross f 's boundary is equal to the parity of the number $n_{v,1}$ of the end vertices these edges contribute to the interior of f .

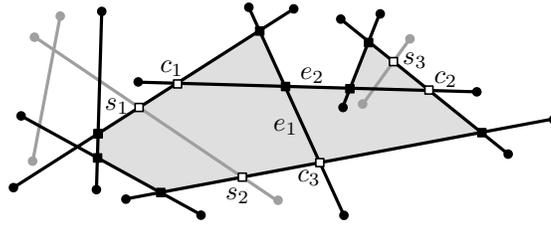


Figure 9: Gray edges do not contribute to the boundary of the gray polygon f ; s_1 , s_2 and s_3 are the side vertices formed by external edges that cross the boundary of f , and c_1 , c_2 and c_3 are the side vertices formed by pairs of boundary edges for f . c_1 and c_2 are charged to e_2 and c_3 is charged to e_1 .

Now let e be a boundary edge for f . We say that a side vertex v is charged to e if $v \in e$ but the segments of e incident to v are not part of f 's boundary, see Fig. 9 for an example. This way each side vertex formed by a pair of boundary edges is only charged to one of the edges.

If both ends of e lie outside f then an even number of side vertices of f is charged to e . The same happens if both ends of e lie inside f . If f contains only one end of e then e is charged by an odd number of side vertices of f . Hence the number $n_{s,2}$ of side vertices of f formed by pairs of boundary edges and the number $n_{v,2}$ of vertices contributed by boundary edges to the interior of f have the same parity. The total number of side vertices of f is $n_{s,1} + n_{s,2}$, and the total number of vertices inside f is $n_{v,1} + n_{v,2}$. These numbers have the same parity. \square

Lemma 16 *If a drawing has an s -odd polygon, then it also has an s -odd face polygon.*

Proof. Let f be an s -odd polygon in a drawing D , and assume it has k vertices inside it. Then k is the sum of the numbers of the vertices inside all face polygons inside f . Since k is odd, there is a face polygon of odd s -complexity inside f . \square

Consider an arrangement of closed or infinite curves in the plane such that every vertex of the arrangement has even degree. We can naturally extend the notion of a cased drawings to such an arrangement.

Lemma 17 *Let A be an arrangement of closed or infinite curves in the plane where every vertex has even degree. Then A has a perfect weaving.*

Proof. It is well known (see [6]) that the faces of the arrangement A can be two-colored such that no two edge-adjacent faces have the same color. Consider such a two-coloring of the faces into red and blue faces. At any crossing c of two curves, choose the order such that the curve that has red to its left when going away from c is above the curve that has red to its right. Then any edge bounding any face of the arrangement has a tunnel and a bridge. \square

Lemma 18 *A drawing D of a degree-one graph where all vertices lie on the outer face has a perfect weaving.*

Proof. We can see such a graph as an arrangement of lines in the plane (by extending the edges beyond their vertices). Then D has a perfect weaving by the previous lemma. \square

Lemma 19 *A drawing D of a degree-one graph has a twin-free cased drawing if and only if it has no s -odd polygons.*

Proof. Let D be a degree-one graph drawing without s -odd face polygons. For each internal face of D we pairwise connect the vertices inside it. The edges that have the vertices on the outer face we extend infinitely beyond their outer face vertices. The obtained arrangement of curves in the plane has a perfect weaving by Lemma 17. It is easy to see that removing the added connections leaves us with a perfect weaving of D . \square

Given a drawing, the casing with maximal number of switches can be found in almost exactly the same way as the drawing with minimal number of switches.

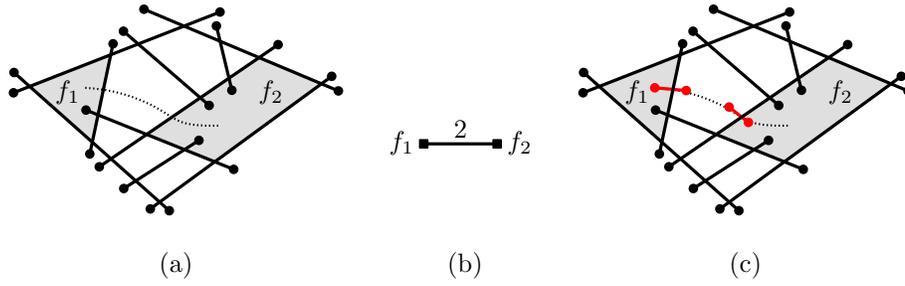


Figure 10: (a) A degree-one graph drawing D with two s -odd faces f_1 and f_2 . The dotted line represents the path connecting them that crosses minimum number of edges possible. (b) The auxiliary graph G^s for D . (c) D with (gray) helper edges added along the path connecting f_1 and f_2 .

Theorem 20 MAXTOTALSWITCHES in the weaving model can be solved in time $O((t+1)k+t^{5/2} \log^{3/2} k)$, where k denotes the number of crossings in the input drawing and t denotes the number of its s -odd face polygons.

Proof. We apply the same technique used in the proof of Theorem 12. We construct an auxiliary graph G^s that contains a vertex for each s -odd face polygon. Then we connect each pair of vertices with an edge and label them in exactly the same way as the edges of graph G^o . Then the minimum total number of twins in any casing of D is equal to the minimum weight perfect matching of G^s .

First we show that we can always have a casing with at most as many twins as the weight of the matching. For each edge of the matching, consider its path in the arrangement and the edges that it crosses. For each crossed edge e of the arrangement we add a short edge that crosses only edge e , see Fig. 10. We refer to these little edges and their end vertices as *helper* edges and vertices.

The resulting graph has no s -odd faces and can be cased as a perfect weaving. We remove the helper edges and vertices, getting at most one twin per removed edge. The resulting casing has as many twins as there were helper edges, hence the number of twins in the resulting graph is at most the weight of the matching.

Now assume we have a casing with minimum number of twins. We place a helper edge between the edges forming the twin along an edge e such that the helper edge crosses only e and case the crossing along e opposite with respect to the twins along e . The obtained drawing is cased and has no more twins. Hence it has no s -odd face polygons by Lemma 19.

We added an even number of helper vertices in total: an odd number of helper vertices to each s -odd face and an even number of helper vertices to each s -even face. Now for each face that has more than one helper vertex we pair the helper vertices arbitrarily and connect each pair with a so-called *connector* edge. As a result we have a set of paths of helper vertices, where helper and connector edges alternate in each path. Each s -odd face of our arrangement contains exactly one end point of such a path. Each path is non-trivial (it contains at least two vertices). Hence the set of paths forms a perfect matching with weight equal to the number of twins in the casing. \square

4 Minimizing tunnels

In this section we present three algorithms that solve MINMAXTUNNELS , $\text{MINMAXTUNNELLENGTH}$, and $\text{MAXMINTUNNELDISTANCE}$ in the stacking model. We also present algorithms for MINMAXTUNNELS and $\text{MAXMINTUNNELDISTANCE}$ in the weaving model. $\text{MINMAXTUNNELLENGTH}$ is NP-hard in the weaving model.

4.1 Stacking model

In the stacking model, some edge e has to be bottommost. This immediately gives the number of tunnels of e , the total length of tunnels of e , and the shortest distance between two tunnels of e . The idea of the algorithm is to determine for each edge what its value would be if it were bottommost, and then choose the edge that is best for the optimization to be bottommost (smallest value for MINMAXTUNNELS and $\text{MINMAXTUNNELLENGTH}$, and largest value for $\text{MAXMINTUNNELDISTANCE}$). The other $m - 1$ edges are stacked iteratively above this edge. It is easy to see that such an approach indeed maximizes the

minimum, or minimizes the maximum. We next give an efficient implementation of the approach. The idea is to maintain the values of all not yet selected edges under consecutive selections of bottommost edges instead of recomputing it.

We start by computing the arrangement of edges in $O(m \log m + k)$ expected time, for instance using Mulmuley's algorithm [9]. This allows us to determine the value for all edges in $O(k)$ additional time.

For MINMAXTUNNELS and MINMAXTUNNELLENGTH, we keep all edges in a Fibonacci heap on this value. One selection involves an EXTRACT-MIN, giving an edge e , and traversing e in the arrangement to find all edges it crosses. For these edges we update the value and perform a DECREASE-KEY operation on the Fibonacci heap. For MINMAXTUNNELS we decrease the value by one and for MINMAXTUNNELLENGTH we decrease by the length of the crossing, which is $\text{casingwidth} / \sin \alpha$, where α is the angle the crossing edges make. For MINMAXTUNNELS and MINMAXTUNNELLENGTH this is all that we need. We perform m EXTRACT-MIN and k DECREASE-KEY operations. The total traversal time along the edges throughout the whole algorithm is $O(k)$. Thus, the algorithm runs in $O(m \log m + k)$ expected time.

For MAXMINTUNNELDISTANCE we use a Fibonacci heap that allows EXTRACT-MAX and INCREASE-KEY. For the selected edge we again traverse the arrangement to update the values of the crossing edges. However, we cannot update the value of an edge in constant time for this optimization. We maintain a data structure for each edge that maintains the minimum tunnel distance in $O(\log m)$ time under updates. The structure is an augmented balanced binary search tree that stores the edge parts in between consecutive crossings in its leaves. Each leaf stores the distance between these crossings. Each internal node is augmented such that it stores the minimum distance for the subtree in a variable. The root stores the minimum distance of the edge if it were the bottommost one of the remaining edges. An update involves merging two adjacent leaves of the tree and computing the distance between two crossings. Augmentation allows us to have the new minimum in the root of the tree in $O(\log m)$ time per update. In total this takes $O(m \log m + k \log m)$ expected time.

Theorem 21 *Given a straight-line drawing of a graph with n vertices, $m = \Omega(n)$ edges, and k edge crossings, we can solve MINMAXTUNNELS and MINMAXTUNNELLENGTH in $O(m \log m + k)$ expected time and MAXMINTUNNELDISTANCE in $O(m \log m + k \log m)$ expected time in the stacking model.*

4.2 Weaving model

In the weaving model, the polynomial time algorithm for MINMAXTUNNELS comes from the fact that the problem of directing an undirected graph, and minimizing the maximum indegree, can be solved in time quadratic in the number of edges [11]. We apply this on the edge crossing graph of the drawing, and hence we get $O(m^4)$ time. For minimizing tunnel length per edge, we can show:

Theorem 22 *MINMAXTUNNELLENGTH is NP-hard in the weaving model.*

Proof. The reduction is from PLANAR 3-SAT, shown NP-hard by Lichtenstein [8]. The reduction is similar to the one for maximizing minimum visible perimeter length in sets of opaque disks of unit size [1]. Note that the proof implies that no PTAS exists. The reduction only uses edges that intersect two or three other edges, so restricting the number of intersections per edge to be constant leaves the problem NP-hard. Also, the number of orientations of edges is constant.

A cased drawing of a set of line segments has property (A) if every line segment has at most two tunnels at crossings with a perpendicular segment, or one tunnel at a crossing with a non-perpendicular segment. Our reduction is such that a PLANAR 3-SAT instance is satisfiable if and only if a set of line segments has a cased drawing with property (A).

We arrange a set of line segments of equal length, using only four orientations. The slopes are -4 , $-\frac{1}{4}$, $+\frac{1}{4}$, and $+4$. If two perpendicular line segments cross, then one has tunnel length equal to the width w of the casing at the crossing. If two other line segments cross, then one edge has tunnel length $w / \sin(\gamma) = 2,125 \cdot w$ at the crossing, where $\gamma = 2 \cdot \arctan(\frac{1}{4})$ is the (acute) angle between the line segments. Therefore, a cased drawing with property (A) has tunnel length at most $2,125 \cdot w$, whereas a cased drawing that does not satisfy property (A) has an edge that has tunnel length at least $3 \cdot w$. This shows the direct relation between property (A) and MINMAXTUNNELLENGTH, and provides the gap that shows that no PTAS exists.

A Boolean variable x_i is modeled by a cycle of crossing line segments as in Fig. 11. Along the cycle, crossings alternate between perpendicular and non-perpendicular, and hence it has even length. The

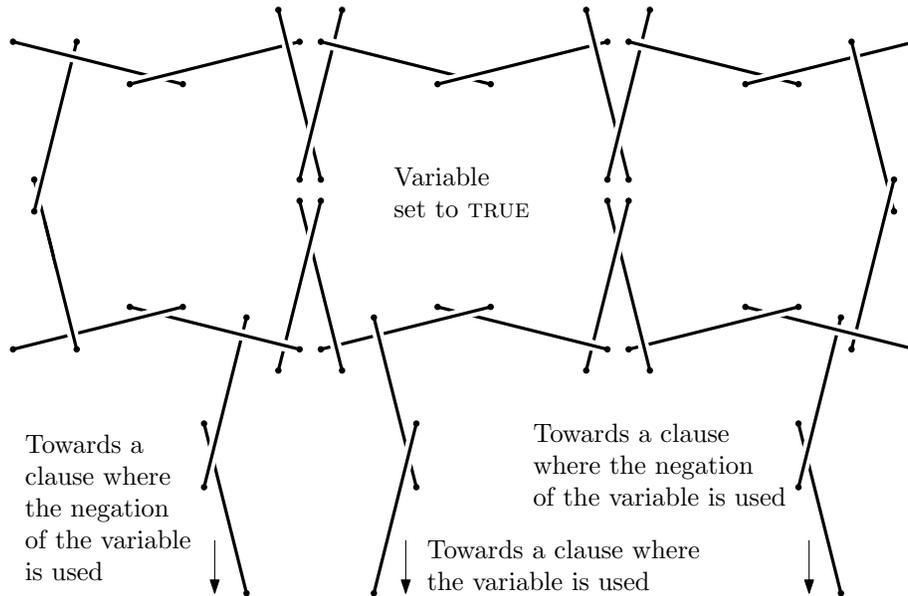


Figure 11: Boolean variable and the connection of channels.

variable satisfies property (A) iff the cycle has cyclic overlap, which can be clockwise or counterclockwise. One state is associated with $x_i = \text{TRUE}$, the other is associated with $x_i = \text{FALSE}$. In each state, the line segments of the cycle alternate in allowing an additional, perpendicular line segment to have a bridge over the line segment of the cycle. In the figure, where the cycle is in the TRUE-state, the line segments with slope $+\frac{1}{4}$ and $+4$ allow such an extra tunnel under a line segment that is not from the cycle. If the cycle is in the FALSE-state, the line segments with slope -4 and $-\frac{1}{4}$ allow the extra tunnel. We use the line segments of slope $-\frac{1}{4}$ to make connections and channels to clauses where \bar{x}_i occurs, and the line segments with slope $+\frac{1}{4}$ for clauses where x_i occurs. Note that the variable can be made larger easily to allow more connections, in case the variable occurs in many clauses.

Channels are formed by line segments that do not cross perpendicularly. So any line segment of the channel can have a tunnel at at most one of its two crossings, or else property (A) is violated. Note that a sequence of crossing line segments with slopes such as $-4, +4, +\frac{1}{4}, -\frac{1}{4}$ gives a turn in the channel. The exact position of the crossing is not essential and hence we can easily reach any part of the plane with a channel, and ending with a line segment of any orientation. A 3-SAT clause is formed by a single line

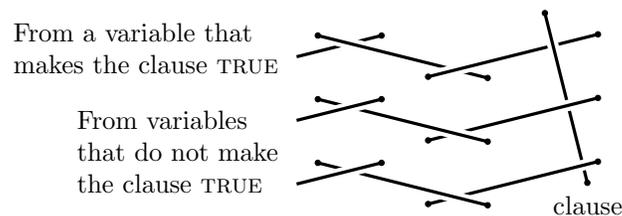


Figure 12: A clause construction.

segment that is crossed perpendicularly by three other line segments, see Fig. 12. Property (A) holds if the clause line segment has at most two tunnels. This corresponds directly to satisfiability of the clause.

With this reduction, testing if property (A) holds is equivalent to testing if the PLANAR 3-SAT instance is satisfiable, and NP-hardness follows. \square

In the remainder of this section we show how to solve MAXMINTUNNELDISTANCE. We observe that there are polynomially many possible values for the smallest tunnel distance, and perform a binary search on these, using 2-SAT instances as the decision tool.

We first compute the arrangement of the m edges to determine all crossings. Only distances between two—not necessarily consecutive—crossings along any edge can give the minimum tunnel distance. One

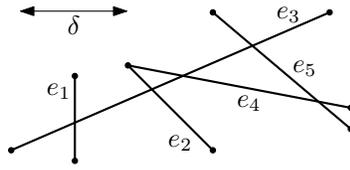


Figure 13: The 2-SAT formula $(\bar{x}_{13} \vee \bar{x}_{23}) \wedge (\bar{x}_{23} \vee x_{34}) \wedge (\bar{x}_{23} \vee x_{35}) \wedge (x_{34} \vee x_{35})$.

edge crosses at most $m - 1$ other edges, and hence the number of candidate distances, K , is $O(m^3)$. Obviously, K is also $O(k^2)$. From the arrangement of edges we can determine all of these distances in $O(m \log m + K)$ time. We sort them in $O(K \log K)$ time to set up a binary search. We will show that the decision step takes $O(m + K)$ time, and hence the whole algorithm takes $O(m \log m + K \log K) = O((m + K) \log m)$ time.

Let δ be a value and we wish to decide if we can set the crossings of edges such that all distances between two tunnels along any edge is at least δ . For every two edges e_i and e_j that cross and $i < j$, we have a Boolean variable x_{ij} . We associate x_{ij} with TRUE if e_i has a bridge at its crossing with e_j , and with FALSE otherwise. Now we traverse the arrangement of edges and construct a 2-SAT formula. Let e_i , e_j , and e_h be three edges such that the latter two cross e_i . If the distance between the crossings is less than δ , then e_i should not have the crossings with e_j and e_h as tunnels. Hence, we make a clause for the 2-SAT formula as follows (Fig. 13): if $i < j$ and $i < h$, then the clause is $(x_{ij} \vee x_{ih})$; the other three cases ($i > j$ and/or $i > h$) are similar. The conjunction of all clauses gives a 2-SAT formula that is satisfiable if and only if we can set the crossings such that the minimum tunnel distance is at least δ . We can construct the whole 2-SAT instance in $O(m + K)$ time since we have the arrangement, and satisfiability of 2-SAT can be determined in linear time [4].

Theorem 23 *Given a straight-line drawing of a graph with n vertices and $m = \Omega(n)$ edges, we can solve MAXMINTUNNELDISTANCE in $O((m + K) \log m)$ expected time in the weaving model, where $K = O(m^3)$ is the total number of pairs of crossings on the same edge.*

5 Removing the restrictions

The restrictions that we impose on the input drawing—no vertex lies on a or very close to an edge, no more than two edges cross in one point, and crossings are not too close—have a significant influence on the problems that we study. For example, if we allow edges to partly overlap with vertices, then it is natural to always draw the vertex on top with a casing around it. In this case, however, there might be no cased drawing in the stacking model, because three edges that necessarily give cyclic overlap can easily be constructed. Testing if cyclic overlap occurs, and hence, if a cased drawing in the stacking model exists can be done by topological sort.

If we remove the restriction that edge crossings are not too close, we may have three edges that intersect so close that it is not possible for the edges to have a tunnel at the one crossing and a bridge at the other. Consequently, three such edges must (locally) be stacked, where one edge has a single bridge over both other edges, and both other edges have a single tunnel. Allowing such “triple crossings” makes the problem MINTOTALSWITCHES NP-hard.

Theorem 24 *If triple crossings of edges are allowed, then MINTOTALSWITCHES is NP-hard in both the weaving and the stacking model.*

Proof. By reduction from PLANAR MAX-2SAT shown NP-hard by [7]. Assume an instance of such a problem which has n Boolean variables and m clauses. The objective is to satisfy as many of the 2SAT clauses as possible. We reduce the instance to an instance of cased drawing such that the number of switches is the same as the number of unsatisfied clauses in the PLANAR MAX-2SAT instance.

A variable is an even cycle of bundles of edges. Each bundle consists of $m + 1$ parallel edges that are either horizontal or vertical. A horizontal bundle intersects two vertical bundles, and each vertical bundle intersects two horizontal bundles, see Fig. 14. Since the variables part of the construction is bipartite, it need not contain any switches. Either all horizontal edges are on top of the vertical ones, or vice versa. This corresponds to the TRUE and FALSE assignment of the variable, respectively.

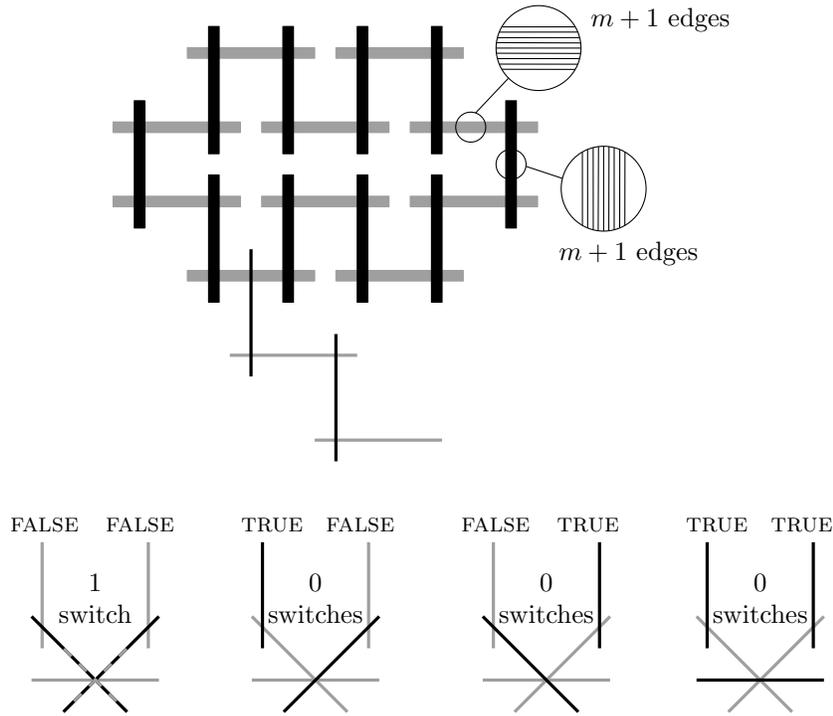


Figure 14: A variable, a channel, and four cases of a clause construction.

We tap off a channel from a variable construction with a single vertical edge as in Fig. 14. The edge crosses all edges of a horizontal bundle, and to avoid switches, it must be below or above all of them (depending on whether the vertical edges in the variable are below or above). The channel itself is a sequence of single edges that leads from the variable construction to a clause construction. A channel has no switches if the edges alternate in having bridges at both crossings and having tunnels at both crossings. The parity of the number of edges in a channel determines if the channel arrives at a clause with an edge that already has a tunnel, or with an edge that already has a bridge. By using horizontal, vertical, and diagonal edges, we can always get the parity as desired and end with an edge in any orientation. So far, no switches are needed yet in a `MINTOTALSWITCHES` cased drawing.

A clause consists of a single edge that has a triple crossing with the two channels that come from variables that occur in the clause. The clause edge itself cannot have a switch, because it only has the triple intersection. To avoid switches on the last edges of the channels, we can put the clause edge on top if both channel edges already have a tunnel, or we can put the clause edge at the bottom otherwise. If both channel edges already have a bridge, then there is no way to avoid a switch because one of the them will have to go below the other. There will be exactly one switch; in all other cases there are no switches at the clause. So we let the parity of any chain be such that the last channel edge be such that it has a bridge if the variable has a state that gives `FALSE` in the clause.

The maximum number of switches needed is obviously at most m , the number of clauses of the `PLANAR MAX-2SAT` instance. We used bundles of $m+1$ parallel edges in the variables to make sure that a variable cannot have a mixed state without giving at least $m+1$ switches already. If the bundles were just single edges, then with only two switches, one part of the channels can use the `TRUE` state and the other part the `FALSE` state of the variable, possibly satisfying many more clauses. The use of bundles prevents this possibility.

Note that the minimum number of switches will be achieved with a stacking drawing. Hence, the problem is NP-hard in both models. \square

6 Conclusions and Open Problems

We presented polynomial time algorithms or NP-hardness results for a number of optimization problems that are motivated by cased drawings. Naturally, we would like to establish the difficulty of the `MINMAXSWITCHES` problem. We would also like to implement our algorithms to visually evaluate the quality of the resulting drawings.

References

- [1] S. Cabello, H. Haverkort, M. van Kreveld, and B. Speckmann. Algorithmic aspects of proportional symbol maps. In *14th European Symposium on Algorithms*, number 4168 in LNCS, pages 720–731, 2006.
- [2] R.-W. Chen, Y. Kajitani, and S.-P. Chan. A graph-theoretic via minimization algorithm for two-layer printed circuit boards. *IEEE Transactions on Circuits and Systems*, 30(5):284–299, 1983.
- [3] D. Eppstein. Testing bipartiteness of geometric intersection graphs. In *Proc. 15th ACM-SIAM Symposium on Discrete Algorithms*, pages 853–861, 2004.
- [4] S. Even, A. Itai, and A. Shamir. On the complexity of timetable and multicommodity flow problems. *SIAM Journal on Computing*, 5(4):691–703, 1976.
- [5] H. N. Gabow and R. E. Tarjan. Faster scaling algorithms for general graph-matching problems. *Journal of the ACM*, 38(4):815–853, 1991.
- [6] B. Grünbaum. Two-coloring the faces of arrangements. *Periodica Mathematica Hungarica*, 11:181–185, 1980.
- [7] L. J. Guibas, J. E. Hershberger, J. S. B. Mitchell, and J. S. Snoeyink. Approximating polygons and subdivisions with minimum link paths. *International Journal of Computational Geometry and Applications*, 3(4):383–415, 1993.
- [8] D. Lichtenstein. Planar formulae and their uses. *SIAM Journal on Computing*, 11(2):329–343, 1982.
- [9] K. Mulmuley. *Computational Geometry: An Introduction through Randomized Algorithms*. Prentice Hall, 1994.
- [10] J. Pach, R. Pollack, and E. Welzl. Weaving patterns of lines and line segments in space. *Algorithmica*, 9(6):561–571, 1993.
- [11] V. Venkateswaran. Minimizing maximum indegree. *Discrete Applied Mathematics*, 143:374–378, 2004.