The hardness of routing two pairs on one face

Guyslain Naves

The hardness of routing two pairs on one face

Guyslain Naves

Received: date / Accepted: date

Abstract

We prove the NP-completeness of the integer multiflow problem in planar graphs, with the following restrictions: there are only two demand edges, both lying on the infinite face of the routing graph. This was one of the open challenges concerning disjoint paths, explicitly asked by Müller [5]. It also strengthens Schwärzler's recent proof of one of the open problems of Schrijver's book [9], about the complexity of the edge-disjoint paths problem with terminals on the outer boundary of a planar graph. We also give a directed acyclic reduction. This proves that the arc-disjoint paths problem is NP-complete in directed acyclic graphs, even with only two demand arcs.

1 Introduction

The multiflow problem has been studied in combinatorial optimization for many years, both because of its theoretical interests and applications. Basically, we want to find integer flows between pairs of terminals, respecting capacity constraints.

The general problem is NP-complete, with different types of constraints, see e.g. the survey of Frank [2]. The most general way to define constraints is to put capacities on the edges of the graph, in the same way as for the classical flow problem. When these capacities are 1 everywhere, this defines the edge-disjoint (or arc-disjoint) paths problem. Robertson and Seymour [8] proved that the multiflow problem is polynomial for undirected graphs, assuming that the total demand is fixed.

Special interest has been shown for solving the problem in planar graphs (directed or not). Kramer and Van Leeuwen [3] have shown that the undirected planar multiflow problem is NP-complete in the general case. Nevertheless, a good characterization theorem has been proved by Okamura and Seymour [6] for the edge-disjoint paths problem in planar Eulerian graphs, under the assumption that all terminals are on the boundary of a unique face of the graph. Despite sharpenings (see [1], [7]), the non-Eulerian case remained open until 2007, when Schwärzler [10] proved the NP-completeness of the edge-disjoint paths problem in planar graphs with all terminals on the boundary of the same face of the graph.

Between Robertson and Seymour's result, and Schwärzler's result, one could ask if there is a polynomial-time algorithm for the edge-disjoint paths problem in planar graphs, when the number of different pairs of terminals is fixed, but each may be repeated many times. Especially when there are only two parallel classes of demands, with terminals lying on a single face of the graph. Actually, Schwärzler's proof can easily be modified to prove that the problem is still NP-complete with three pairs of terminals. In this paper we give a new reduction, proving that the problem is NP-complete with only two pairs of terminals. This solves a question of Müller [5].

We will also give a directed version of our proof, showing that the arc-disjoint paths problem is NP-complete, even if G is planar with two opposite parallel classes of demand edges st and ts where vertices s and t belong to the boundary of the same face of G. Both results strengthen [5]. Finally, we prove the NP-completeness when G is a planar acyclic digraph and H consists of two pairs of terminals lying on the outer face of G.

2 Definitions

Let $G = (V,E)$ be an undirected graph, and let $c : E \to \mathbb{N}$ be a *capacity function* on the edges of G. Let $H = (T,D)$ be an undirected graph with $T \subseteq V$, and $r : D \to \mathbb{N}$ a *demand function* (or *request*). The *multiflow problem* is to find a multiset $\mathscr{C}$ of cycles of $G+H$ satisfying the following conditions :

Figure 1: *The two paths on the left do not cross, those on the right cross each other.*

(i) Each cycle of $\mathscr{C}$ contains exactly one edge of H.

(ii) For each edge of G, the number of cycles in $\mathscr{C}$ using it is less than its capacity.

(iii) For each edge of H, the number of cycles in $\mathscr{C}$ using it is exactly its request.

H is usually called the *demand graph*, T is the set of *terminals*. By cycle, we mean a closed sequence of distinct edges that are consecutive in the graph, or equivalently a connected Eulerian subgraph. The problem can easily be defined in digraphs, by replacing every occurrence of "cycle" by "directed cycle". Thus, an instance of the multiflow problem consists of a quadruple (G,H,r,c). In the following, c will always be supposed to be equal to 1. We will note $\mathscr{P}$ the set of paths obtained from $\mathscr{C}$ by ignoring the demand edges, and we will mainly speak about these paths instead of the cycles.
A *path* is formally a sequence of distinct edges that are consecutive. Thus, cycles are closed paths. Two paths are *edge-disjoint* if their edge-sets are disjoint. We define *directed paths* and *arc-disjoint paths* analogously. When c is constantly equal to 1, the multiflow problem is known as the *edge-disjoint paths problem* (respectively the *arc-disjoint paths problem*).
The graphs considered in this paper are always without loops, but parallel edges are allowed. Actually, whenever an edge $e \in E(G)$ has a capacity greater than 1, we replace it by $c(e)$ parallel edges. Let $U \in V$ be a subset of a vertex set of the graph. We note $\delta(U)$ the set of edges having exactly one extremity in U. Every set of edges that can be written as $\delta(U)$ for some U is called a *cut* of the graph. In directed graphs, $\delta^-(U)$ is the set of arcs entering U, $\delta^+(U)$ is the set of arcs leaving U. When $\delta(U) = \delta^+(U)$, we say that U is a *directed cut*. We define $d(U) := |\delta(U)|$ and similarly $d^+(U)$ and $d^-(U)$.
Let C be a cut of $G+H$. C is a *tight cut* if $c(C \cap E) - r(C \cap D) = 0$. If this difference is negative, the multiflow problem is not feasible. When H is reduced to a single edge, the famous *max-flow-min-cut* result states that a multiflow exists if and only if this difference is never negative (see Menger's theorem [4]). When C is a tight cut, each edge of C is entirely used in any solution of the multiflow problem : there are as many paths through each edge as its capacity. In directed graphs, tight cuts are the cuts $\delta(U)$ of $G+H$ with $c(\delta_G^+(U)) - r(\delta_H^-(U)) = 0$.
A *planar graph* is a graph that has an embedding in the plane without intersection of the edges (or arcs). Let P_1 and P_2 be two edge-disjoint paths in G. P_1 and P_2 *cross* at vertex $v \in V(G)$ if there are four edges $e_1, \ldots, e_4$ incident to v, appearing in this order around v, such that e_1 and e_3 are consecutive in P_1 and e_2, e_4 are consecutive in P_2. However, two paths may have a common vertex without crossing, see Figure 1. A *crossing* is a triple (P,Q,u) such that P and Q crosses at vertex u.

3 Outline of the proof

We will prove the NP-completeness of the edge-disjoint paths problem with only two pairs of terminals, by reduction from 3-SAT. Before giving the full proof, which is quite technical, we explain the main ideas of the reduction.
Consider an instance of 3-SAT, consisting of clauses over a set of variables. We build a graph in the form of a grid, with as many columns as there are clauses, and as many rows as twice the number of variables. In this grid, the intersection of each row with each column is one of two basic graphs. These two graphs, called XCH and LIC should have the following properties. They have two vertices of degree one in their left and right sides, and four — actually two pairs — at their top and their bottom. We suppose that there will always be two paths either from the two left top vertices or from the two right ones, going through them from top to bottom (*vertical paths*). We will also always route one or two paths from left to right (*horizontal paths*).

If there are two horizontal paths, then the two vertical paths can only be routed diagonally. The behaviours of XCH and LIC are different if there is only one horizontal paths. In that case, in XCH, the vertical paths are still forced to go diagonally through the graph, but in LIC, they can be routed vertically. Figures 2 and 3 illustrate the possible routing through the two graphs (and what we mean by *diagonally* and *vertically*).

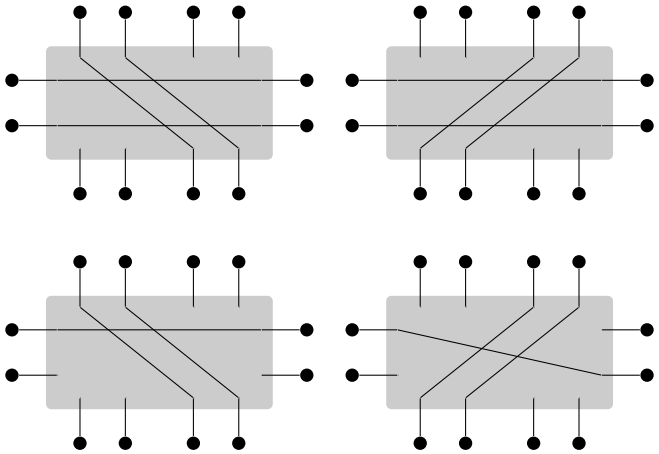

Figure 2: *Possible routings through XCH and LIC.*

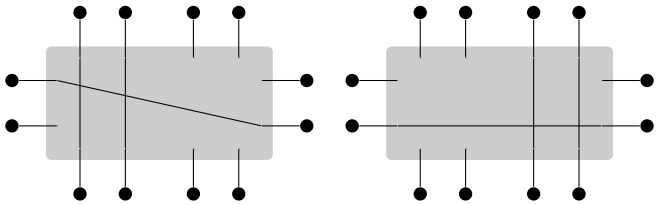

Figure 3: *Additional routings for LIC. When there is only one horizontal path, the vertical paths are not forced to be diagonal.*

The typical behaviour for these gadgets is to change the vertical paths from one side to the other (Lemmas 4 and 6). We say that they *shift* the paths. The only special case when the two vertical paths can stay on the same side is the following : there is a LIC and a single horizontal path. In that special case, the gadget *keeps* the paths (Lemma 7). The reduction is basically the following : there are two consecutive rows for each variable, and we route three paths along these two rows. Thus, one path will follow one row, and the two other paths will follow the other row, *deciding a variable assignment*. We also route two paths in each column, and we ask that *these two paths are kept an odd number of times*. Because of the properties of LIC and XCH, the two vertical paths can be kept in a particular gadget if and only if this gadget is a LIC and there is only one horizontal path in the corresponding row. We place LIC on those particular intersections for which the literal associated with the row appears in the clause encoded by the column. Thus, we ensure that whenever vertical paths are kept, the corresponding clause is made valid by the chosen assignment. Then we must guarantee that the two paths associated with each column are kept at least once (actually an odd number of times). This is done by asking these two paths to be routed from the two left pair of uppermost vertices of the column, to the two rightmost lower vertices. Because the number of rows is even (twice the number of variables), paths must indeed be kept an odd number of times.
Finally, we can add two common terminals for the vertical paths, and two others for the horizontal paths, achieving the desired restriction to two edges of demand. Unfortunately, there are several difficulties in implementing this reduction. Mainly, XCH and LIC does not exist, with the prescribed properties. Actually, the main difficulty is to enforce that the horizontal paths stay on their respective rows. We will prove that with our gadgets, *the horizontal paths cannot go through more than three rows in each column* (Lemma 8). Our solution is then to create a *buffer consisting of a large number of rows*, between the rows encoding the variables. Similarly, we must prove that *the vertical paths stay in their columns* (Lemmas 10 and 12). This

will be done, by using the property that our graph is *nearly Eulerian* (almost each vertex has even degree). This fact will help us to prove that the edges not in a solution induce cycles and a small set of paths, called *no-paths*. One of the main lemmas states that the extremities of each no-path are determined (Lemma 11). More exactly, there will be exactly one no-path per variable. Considering horizontal paths and no-paths together, many cuts are tight for them. Consequently, vertical paths cannot use the edges contained in these cuts : each vertical path intersects only one column. Then, local properties of XCH and LIC allow us to conclude the proof.

4 Preliminaries

4.1 Uncrossing the paths

We can suppose without loss of generality that each pair of paths induces at most one crossing :

Lemma 1. *Let (G,H) be an instance of the edge-disjoint paths problem with G planar. There is a solution for (G,H) if and only if there is a solution such that each pair of paths crosses at most once, and two paths with the same extremities do not cross.*

Proof. Let $\mathscr{P}$ be a solution minimizing the number of crossings. We can assume that every path is simple. Suppose that two paths P_1 and P_2 induce more than one crossing. Let u and v be the first and second crossings between P_1 and P_2, starting from one extremity of P_1.
We decompose P_i in three paths $Q_i \cup R_i \cup S_i$ where R_i has extremities u and v. Then replacing P_1 and P_2 by $P_1' = Q_1 \cup R_2 \cup S_1$ and $P_2' = Q_2 \cup R_1 \cup S_2$, we show that the number of crossings is reduced, contradicting the minimality of $\mathscr{P}$.
In every vertex except u and v, the paths are not locally modified, thus the number of crossings is not changed. The number of crossings between P_1' and P_2' is reduced by at least 2. Then, the neighbourhood of vertex u (and symmetrically v) is divided into four parts. A path going through two consecutive parts crosses one of P_1 and P_2, and one of P_1' and P_2'. A path going through two opposite parts crosses both P_1 and P_2. This proves the number of crossing is decreased.
Finally, if P_1 and P_2 have a common extremity u, and crosses at v, the same transformation applies again, decreasing the number of crossing by at least one. □

Solutions will always be supposed *uncrossed* (each pair of paths with different extremities induces at most one crossing) and simple. The following is an easy consequence of uncrossing :

Lemma 2. *Let G be a planar graph, a, b, c and d four vertices on the boundary of the infinite face of G, occuring in this order. Let $\mathscr{P}$ be an uncrossed set of (a,c)-paths and (b,d)-paths mutually edge-disjoint. Then, all (a,c)-paths cross the (b,d)-paths in the same order.*

Proof. There is a crossing neither between the (b,d)-paths, nor between the (a,c)-paths, so we can choose P among the (b,d)-paths such that all the other (b,d)-paths are on the same side. We prove the lemma by induction on the number of (b,d)-paths. P can be closed to a cycle by adding a curve on the infinite face of G. Then by Jordan's theorem, P separates the (b,d)-paths from a (say). All (a,c)-paths, starting from a must cross P before the other (b,d)-paths. Using induction on the (b,d)-paths minus P concludes the proof. □

4.2 Forbidding crossing

It happens to be useful to have in the proof an Eulerian graph. This is not possible since the edge-disjoint paths problem with terminals on the boundary of the infinite face is polynomially solvable as soon as all the inner vertices have even degrees [6], [1]. We introduce the following restriction : every vertex will be of degree four, but in some of them, paths will not be allowed to cross each other. This effect can easily be achieved by replacing those special vertices by a cycle of length 4, as described in Figure 4

Formally, we define the following problem, and show that it is equivalent to the original one :

Problem 1 (Extended Planar Edge-Disjoint Paths)**.**
INPUT *: a planar graph G, a demand graph H with $V(H) \subseteq V(G)$, and $U \subseteq V(G)$.*
OUTPUT *: Is there a solution to the edge-disjoint paths problem (G,H) such that for every $u \in U$, there is no crossing at u ?*

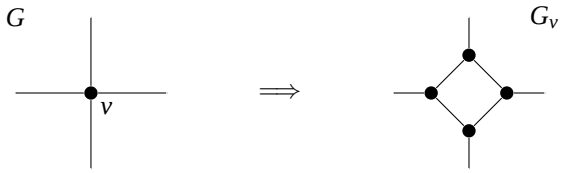

Figure 4: *How to forbid the crossing of paths at vertex v.*

Let G_v be the graph obtained from G by replacing a vertex v of degree 4 by a cycle of length 4, according to Figure 4. The equivalence with the planar edge-disjoint paths problem is a consequence of the following obvious lemma :

Lemma 3. *Let (G,H,U) be an instance of the extended planar edge-disjoint path problem. Let v be a vertex of G of degree 4. Then there is a solution to (G_v,H,U) if and only if there is a solution to $(G,H,U\cup\{v\})$.*

In the following, we will always study instances of the extended problem. Vertices of U are called *non-crossing vertices*. In the figures, we will represent vertices not in U (*crossing vertices*) by bold points. Note that there is no restriction over no-paths, they can cross other no-paths or paths at every vertex.

5 Implementing XCH and LIC

5.1 Basic graphs

We give the graphs encoding the gadgets XCH and LIC, and detail their respective properties. Let XCH be the graph depicted in Figure 5. The crossing vertices are a, b, c and d. We note $S=\{s_1,s_2,s_3,s_4\}$, $S'=\{s'_1,s'_2,s'_3,s'_4\}$, $T=\{t_1,t_2\}$ and $T'=\{t'_1,t'_2\}$.

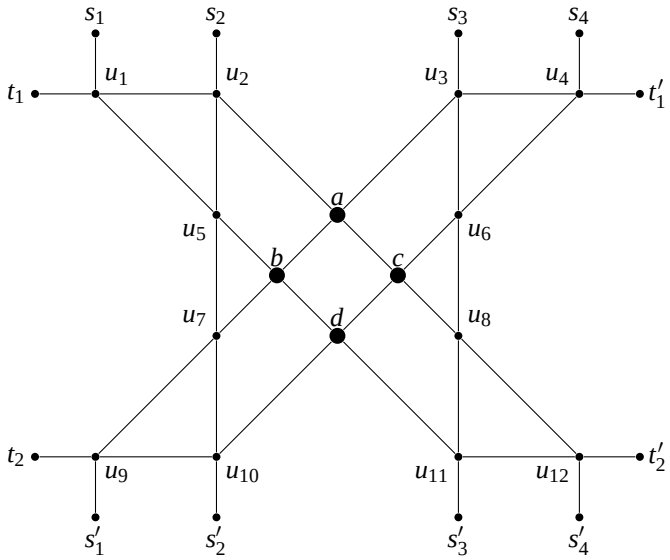

Figure 5: *The graph XCH. There are only four crossing vertices a, b, c and d. All the* inner *vertices have degree* 4.

Lemma 4. *Let $\mathscr{P}=\{S_1,S_2,T_1,T_2\}$ be an uncrossed set of edge-disjoint paths in XCH, satisfying :*

(i) S_1 and S_2 are (S,S')-paths,

(ii) T_1 and T_2 are (T,T')-paths.

Then S_1 and S_2 are either an (s_1,s'_3)-path and an (s_2,s'_4)-path, or an (s_3,s'_1)-path and an (s_4,s'_2)-path.

Proof. Let $\mathscr{P}$ be as described in the lemma. As the (S,S')-paths must cross the (T,T')-paths, there are at least 4 crossings in $\mathscr{P}$. We know that these crossings occur in vertices a, b, c and d, and the (T,T')-paths (resp. the (S,S')-path) do not cross each other.
Suppose ab is in an (S,S')-path, then ac and bd must belong to distinct (T,T')-paths, and cd is in the second (S,S')-path. Then au_2, bu_5, cu_8, du_{11} are in (T,T') paths, and au_3, bu_7, cu_6 and du_{10} are in (S,S')-paths. As there is no other crossing except in the four central vertices, the (S,S')-paths are connected to s'_1, s'_2, s_3 and s_4. The case when ab belongs to a (T,T')-path is similar and gives the other solution. □

These paths exist, as shown by Figure 6.

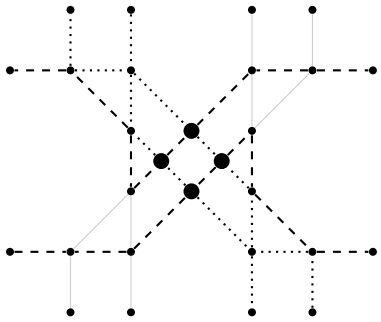

Figure 6: *Existence of the paths for Lemma 4.*

Lemma 5. *Let $\mathscr{P}$ be a set of three edge-disjoint paths in XCH that satisfies :*

(i) $\mathscr{P}$ contains exactly two $(\{s_1,s_2\},S'')$-paths, and S'' is either $\{s'_1,s'_2\}$ or $\{s'_3,s'_4\}$,

(ii) $\mathscr{P}$ contains exactly one (T,T')-path.

Then, $S'' = \{s'_3,s'_4\}$.

Here we find the first differences between the graph XCH and its ideal model given in Section 3. We need to suppose that the two vertical paths come from the same side (through s_1 and s_2 or s_3 and s_4), and leave also at the same side . We need to prove that this is actually the case, but we can already remark that it is true in the particular case when the gadgets above and below are XCH with two horizontal paths, by Lemma 4.

Proof. If not, then there is a set $\mathscr{P}$ of three edge-disjoint paths, a (t,t')-path Q, an (s_1,s'_1)-path P_1 and an (s_2,s'_2)-path P_2. As Q must cross the two other paths, all paths contain at least one of a,b,c and d. Then, Q uses one edge of u_2a, u_5b, u_7b, $u_{10}d$, and both P_1 and P_2 use two of these edges, contradicting the edge-disjointness of the paths. □

Let LIC be the graph depicted in Figure 7. We note again $S = \{s_1,s_2,s_3,s_4\}$, $S' = \{s'_1,s'_2,s'_3,s'_4\}$, $T = \{t_1,t_2\}$ and $T' = \{t'_1,t'_2\}$.

Lemma 6. *Let $\mathscr{P}$ be a set of edge-disjoint paths in LIC that satisfy :*

(i) $\mathscr{P}$ contains exactly 2 $(\{s_1,s_2\},S'')$-paths, where S'' is either $\{s'_1,s'_2\}$ or $\{s'_3,s'_4\}$,

(ii) $\mathscr{P}$ contains exactly 2 (T,T')-paths.

Then, $S'' = \{s'_3,s'_4\}$. Moreover, there cannot be another $(S\cup T,S'\cup T')$-path.

Proof. Suppose $S'' = \{s'_1,s'_2\}$. Let C be the cut $\{u_7,u_9,u_{10}\}$. Let Q be the (t_2,T')-path, P_1 the (s_1,s'_1)-path and P_2 the (s_2,s'_2)-path. All three different paths meet C, and $d(C) = 6$. Moreover, there is no crossing in C, thus considering $\delta(C)$, u_5u_7 is used by Q and bu_7, du_{10} by P_1 and P_2. Now, there are four distinct paths entering $C' = \{u_1,u_2,u_5\}$, but $d(C') = 6$, contradiction.
As the edges of $\delta(\{a,b,c,d\})$ are all used by $\mathscr{P}$, there is no other $(S\cup T,S'\cup T')$-path. □

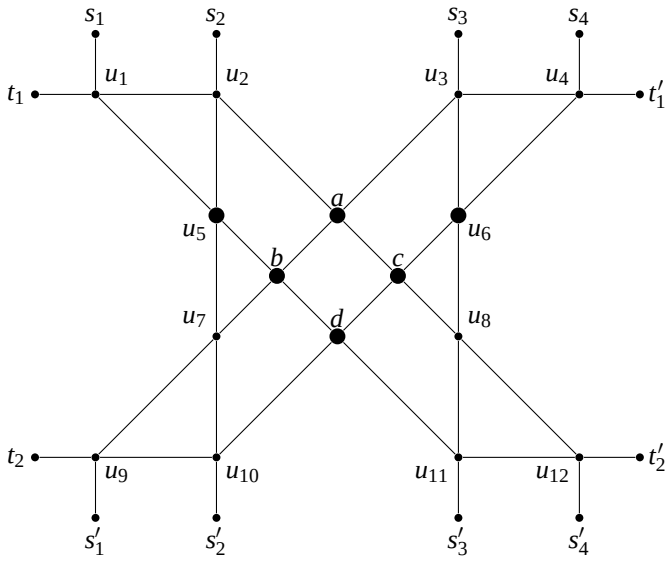

Figure 7: *Graph LIC. Except two additional crossing vertices, it is similar to XCH.*

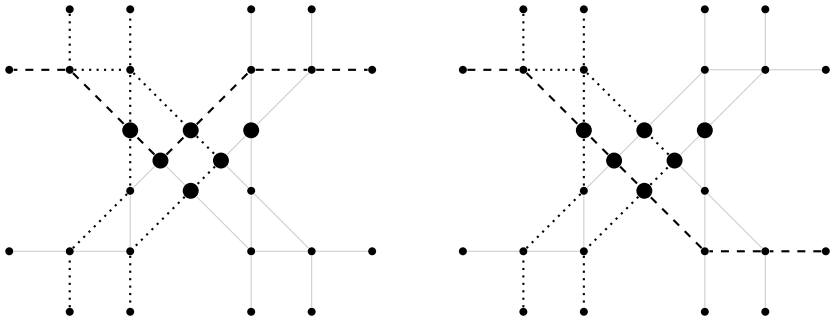

Figure 8: *Existence of the paths for Lemma 7.*

Lemma 7. *There exist a (T,T')-path P, an (s_1,s'_1)-path P_1 and an (s_2,s'_2)-path P_2, pairwise edge-disjoint, in LIC.*
There exist a (T,T')-path P, an (s_3,s'_3)-path P_1 and an (s_4,s'_4)-path P_2, pairwise edge-disjoint, in LIC.

Proof. See Figure 8. □

With these four lemmas, the ideal behaviour of XCH and LIC is not reached by the two gadgets: paths can still go from $S \cup S'$ to $T \cup T'$, and we did not prove that the vertical paths must use either the two rightmost edges or the two leftmost edges between two consecutives gadgets.

5.2 Aggregating gadgets

In order to build the graph for the reduction, we need to aggregate gadgets in the form of a grid, using XCH and LIC in the crossing of rows and columns. This is done by linking the edges incident to corresponding vertices of degree 1. Figure 9 shows how we build the grid, and gives the notation that we will use. Moreover, we define $X := \{x_i \ : \ i \in [\![1,4n]\!]\}$ and $X' := \{x'_i \ : \ i \in [\![1,4n]\!]\}$ where n is the number of columns, and similarly $Y := \{y_i \ : \ i \in [\![1,2p]\!]\}$ and $Y' := \{y'_i \ : \ i \in [\![1,2p]\!]\}$ where p is the number of rows. We call i^{th} *vertical cut* the set of edges $V_i := \{f_k^{i,j} \ : \ j \in [\![1,p]\!], k \in [\![1,2]\!]\}$, and j^{th} *horizontal cut* the set of edges $H_j := \{e_k^{i,j} \ : \ i \in [\![1,n]\!], k \in [\![1,4]\!]\}$. Vertex v in $M(i,j)$ will be denoted $v^{i,j}$.

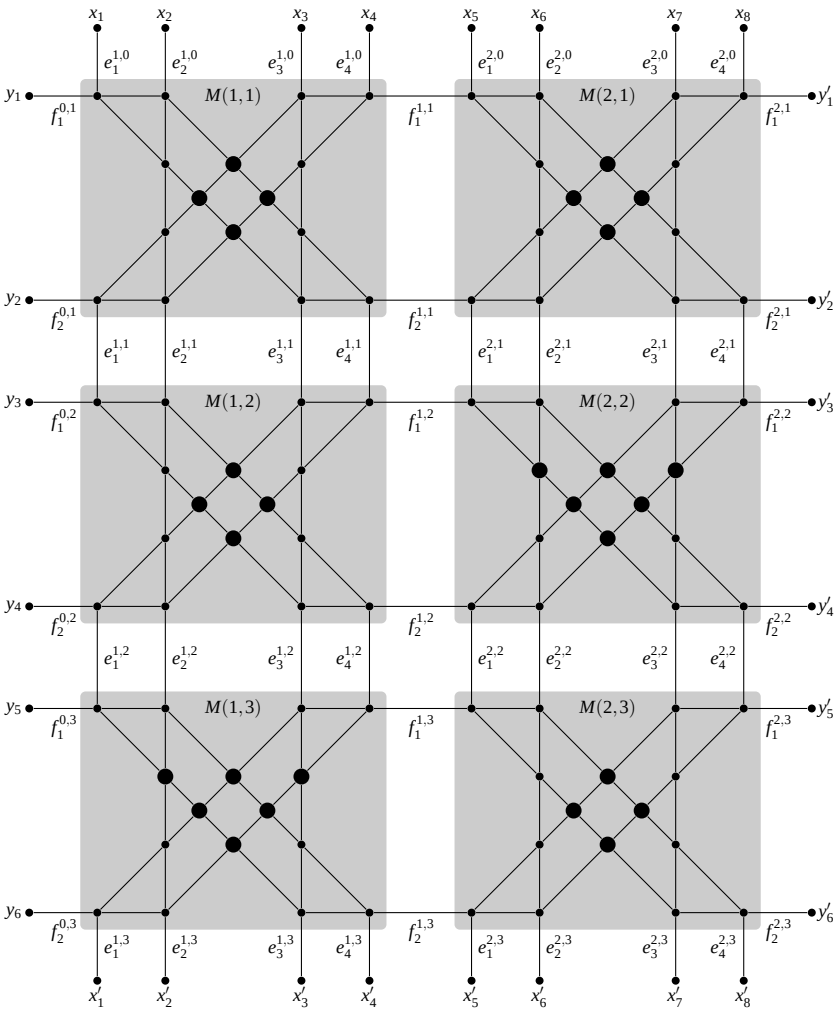

Figure 9: *A grid of dimension* 2×3.

The following lemma will help us to explain why the horizontal paths cannot go through more than two different gadgets in each column, justifying the buffers. More exactly, the no-paths can be vertically moved by at most 2 rows in each column, so the buffers must contain at least 4 times the number of columns plus 2.

Lemma 8. *Let G be a grid of dimension 1×3, built exclusively with LIC. Let $\mathscr{P}$ be a set of edge-disjoint paths, consisting of two (X,X')-paths A and B, a $(\{y_1,y_2\},y'_1)$-path C, a (y_3,y'_2)-path D, a (y_4,y'_3)-path E, a (y_5,y'_4)-path F, a (y_6,y'_5)-path H and an (X',y'_6)-path I. Then y_1 and y_2 are disconnected from X' in $G \setminus E(\mathscr{P})$ (or equivalently, there is no no-path from $\{y_1,y_2\}$ to X').*

Note that this lemma is still true if some of the LIC graphs are replaced by XCH. Figure 10 shows the extremities of each path.

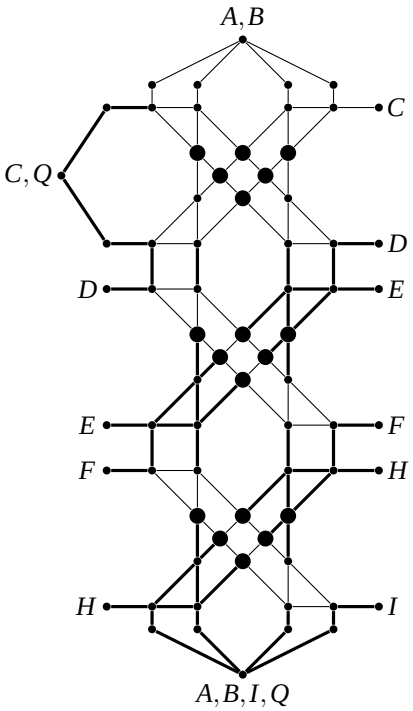

Figure 10: *There is no solution to this edge-disjoint paths problem, even if Q can cross other paths at every vertex. Edges contained in some tight cut are drawn thick.*

Proof. Suppose that there is a $(\{y_1,y_2\},X')$-no-path Q. Note that Q can cross other paths at non-crossing vertices. We can take $\mathscr{P}$ uncrossed (uncrossing the paths does not change the set of edges used by the solution), and without loss of generality, A is routed on the left of B.
Because of tight cuts, $M(1,1)$ contains C, $M(1,2)$ contains E, and $M(1,3)$ contains H. In $M(1,2)$, E goes through u_6, (u_3 and u_4 are non-crossing vertices, and at least one of the paths A, B or D goes through these vertices) and uses u_6c or u_6u_8. For the same reason, E goes through u_7 and uses u_7u_5 or u_7b. Consider in $M(1,2)$ the cuts $C_1 := \{u_5u_7, u_5b, ab, ac, u_6c, u_6u_8\}$ and $C_2 := \{u_5u_7, u_7b, bd, cd, cu_8, u_6u_8\}$. There are exactly four paths routed through these two cuts, and E must use three edges from at least one of them, because it can use neither $u_{10}d$ nor au_3. Then F cannot go through C_1, proving that it can cross A or B only at vertex d in $M(1,2)$.
Similarly, in $M(1,3)$, using path H and the same two cuts (we call them C'_1 and C'_2), F can only cross A or B at vertex a. Because of Lemma 2 and the tight cut between $M(1,2)$ and $M(1,3)$, F crosses A in a of $M(1,3)$ and then crosses B in d of $M(1,2)$. As F cannot use edges in both C_1 and C'_2, F must traverse d of $M(1,2)$ and a of $M(1,3)$ from left to right or from right to left, an odd number of times. This is indeed a contradiction. □

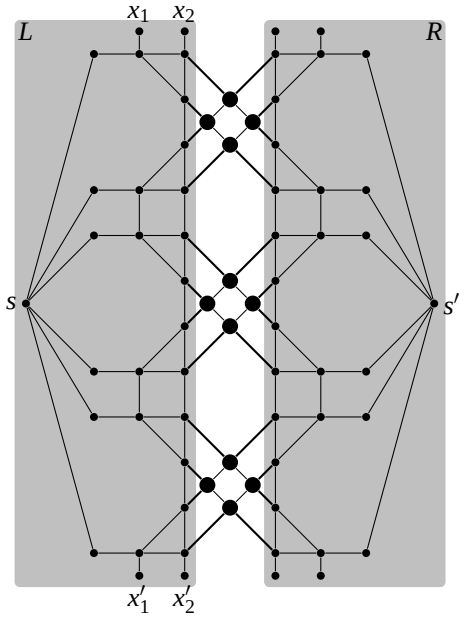

Figure 11: *There is no solution to this edge-disjoint paths problem, with* 5 *paths between* s *and* s'*, see Lemma 9.*

We need a last "local" lemma, proving that without the presence of a LIC, the vertical paths of a column are shifted. This fills the holes of Lemma 5. As the no-paths only use three rows of each column (this is not fully proved yet), we will only need to study the case with a grid of 3 rows.

Lemma 9. *Let G be a grid of XCH of dimension 1×3. Then there is no set of edge-disjoint paths consisting of five (Y,Y')-paths, one (x_1,x'_1)-path and one (x_2,x'_2)-path.*

Proof. Suppose that these paths exist. We distinguish two special cuts $\delta(L)$ and $\delta(R)$. There are exactly 12 vertices for crossings, and 10 are needed. The (Y,Y')-paths use 5 edges of $\delta(L)$ and 5 of $\delta(R)$. The (x_1,x'_1)-path A and the (x_2,x'_2)-path B both use an even number of edges in these two cuts. Moreover, they can do at most 2 crossings in each of the three groups of four crossing vertices (corresponding to the crossing vertices of an XCH graph), thus they go through each of these groups. Then each uses at least 6 edges in the two cuts, and because of parity, A uses 4 edges of $\delta(L)$ and 2 of $\delta(R)$, and B uses 2 edges of $\delta(L)$ and 4 edges of $\delta(R)$. Because $d(L) = d(R) = 12$, there cannot be more. Thus, each (Y,Y')-path uses exactly one edge of $\delta(R)$, and one of $\delta(L)$. Then, in the central XCH graph, there are exactly one edge of $\delta(L)$ and one edge of $\delta(R)$ used by A, and two edges of $\delta(R)$ used by B. At least two edges of $\delta(L)$ must be used by the (Y,Y')-paths, and at most one of $\delta(R)$. But this leads to a contradiction, as each (Y,Y')-path cannot use more than one edge in any of the two cuts. □

6 Reduction

Let φ be a Boolean formula in conjunctive normal form, every clause is composed of 3 literals, with $n \geq 3$ clauses over a set of $p' \geq 3$ variables. We encode the formula in a graph G_φ, a grid of XCH and LIC with n columns. We need two rows for each variable, plus between every two of these rows, and after the last one, a buffer consisting of $q = 4(p'+3)n+2$ rows. Then, the grid has exactly $p = 2p'(q+1)$ rows.
The i^{th} variable X_i corresponds to rows $1+2(i-1)(1+q)$ and $q+2+2(i-1)(1+q)$. Intuitively, one no-path will be routed through one of these two rows (but actually, we can only enforce that it will stay near one of them). All the other rows are parts of buffers, and then contain only XCH. Row $1+2(i-1)(1+q)$ corresponds to the assignment of *true* to X_i, so there is a LIC in each column corresponding to the clause where X_i appears positively. Similarly, row $q+2+2(i-1)(1+q)$ corresponds to the assignment of *false*

to X_i, so there is a LIC in each column corresponding to the clause where X_i appears negatively. All other subgraphs are XCH. In summary, $M(i,j)$ is a LIC if and only if :

- either $j = 1+2(i-1)(1+q)$ and X_i appears positively in the j^{th} clause,
- or $j = q+2+2(i-1)(1+q)$ and X_i appears negatively in the j^{th} clause.

We add two terminals for the vertical paths, x and x'. As the number of rows is even, and we want the vertical paths to be shifted an odd number of times, we add the following edges : xx_{4k+1}, xx_{4k+2}, $x'_{4k+3}x'$ and $x'_{4k+4}x'$ for each $k \in [\![0, n-1]\!]$. Thus we enforce the parity of the number of shifts. Moreover, to reduce the number of odd vertices, we add the edges $x_{4k+3}x_{4k+4}$ and $x'_{4k+1}x'_{4k+2}$ for each $k \in [\![0, n-1]\!]$.
For each variable X_i, $i \in [\![1, p']\!]$, we add two new vertices w_i and w'_i, and the edges w_iy_j and $w'_iy'_j$ for all $j \in [\![4(i-1)(q+1)+1, 4(i-1)(q+1)+2q+4]\!]$ (that is w_i is connected to the vertices of Y between the two rows encoding X_i). These new vertices will be the extremities of the no-paths. Then we add two terminals for the horizontal paths y and y'. y is connected to w_i by $2q+3$ parallel edges, and to every vertex of Y that still have a degree one by a single edge. Symmetrically, y' is connected to w'_i by $2q+3$ parallel edges, and to every vertex of Y' not adjacent to some w'_i. Thus y and y' have degrees $2p-p'$.
Finally we demand to find $2n$ (x,x')-paths and $2p-p'$ (y,y')-paths, defining the demand graph H_φ. The vertices of odd degree in $G_\varphi + H_\varphi$ are the vertices of $W := \{w_i \,:\, i \in [\![1, p']\!]\}$ and $W' := \{w'_i \,:\, i \in [\![1, p']\!]\}$. The following cuts are tight : $\delta(x)$, $\delta(x')$, $\delta(y)$ and $\delta(y')$. Figure 12 gives an overview of the reduction.
We define precisely the *vertical paths* to be the (x,x')-path of a solution, and the *horizontal paths* are the (y,y')-paths. We say that two paths are *parallel* if they are both horizontal, or both vertical, otherwise they are *orthogonal*.

6.1 No-paths

Let $\mathscr{P}$ be a solution to the edge-disjoint paths problem (G_φ, H_φ). By considering $\mathscr{P}$ as a set of cycles of $G_\varphi + H_\varphi$ containing exactly one edge from $E(H_\varphi)$, the vertices of G_φ have the same degrees in $G_\varphi + H_\varphi$ and $G_\varphi + H_\varphi \setminus E(\mathscr{P})$. So the edges of the complementary graph of the solution $\mathscr{P}$ is a $W \cup W'$-join $Q \subset E(G_\varphi)$. The following lemma proves that the no-paths cannot traverse the buffers, implying that the vertical paths stay in their respective columns. This last fact will allow us to apply the previous lemmas.

Lemma 10. *Let G be a grid of dimension $n \times p$. Let $\mathscr{P}$ be an uncrossed set of (X,X')-paths and (Y,Y')-paths pairwise edge-disjoint. Suppose there exists $i \in [\![2, p-2]\!]$ such that for all $j \in [\![1, n]\!]$, $M(i,j)$ and $M(i+1,j)$ are XCH and there are exactly four crossings of paths of $\mathscr{P}$ in $M(i,j)$ and in $M(i+1,j)$. Then, there is no no-path between a vertex of row $i-1$ and a vertex of row $i+2$.*

Proof. Let $i \in [\![2, p-2]\!]$ such that each crossing vertex of rows i and $i+1$ is used for a crossing, and suppose there is a no-path Q between rows $i-1$ and $i+2$. Then Q must pass through u_5u_7 or u_6u_8 in some $M(j,i)$, say u_6u_8 by symmetry, and then passes through $e_3^{j,i}$, $e_4^{j,i}$, $e_1^{j+1,i}$ or $e_2^{j+1,i}$ (because all the a, b, c and d vertices are used by paths).
Let P_1, P_2, P_3, P_4 be the paths going through edges $c^{j,i}u_8^{j,i}$, $d^{j,i}u_{11}^{j,i}$, $a^{j,i+1}u_3^{j,i+1}$, $c^{j,i+1}u_6^{j,i+1}$ respectively. We consider the cut $C = \delta(U)$, with $U = \{u_8^{j,i}, u_{11}^{j,i}, u_{12}^{j,i}, u_3^{j,i+1}, u_4^{j,i+1}, u_6^{j,i+1}\}$, intersected by these four paths and Q. Because $|C| = 8$, at most four different paths and no-paths intersect C. As the solution is uncrossed, P_1 and P_2 are different and parallel, and P_3 and P_4 also. There is no crossing vertex among U, thus $P_2 = P_3$. By Lemma 1, P_1, P_2 and P_4 are crossed by their orthogonal paths in the same order, then $P_1 = P_4$. We remark that $j \neq n$, otherwise the routing is not possible.
At least one of P_1, P_2 and Q must go through $U' = \{u_7^{j+1,i}, u_9^{j+1,i}, u_{10}^{j+1,i}, u_1^{j+1,i+1}, u_2^{j+1,i+1}, u_5^{j+1,i+1}\}$. By a similar argument, there are paths P'_1 using $b^{j+1,i}u_7^{j+1,i}$ and $u_5^{j+1,i+1}b^{j+1,i+1}$ and P'_2 using $d^{j+1,i}u_{10}^{j+1,i}$ and $u_2^{j+1,i+1}a^{j+1,i+1}$. Exactly one of the four considered paths must use two of the edges $u_6^{j,i}u_8^{j,i}$, $u_5^{j+1,i}u_7^{j+1,i}$, $u_6^{j,i+1}u_8^{j,i+1}$ and $u_5^{j+1,i+1}u_7^{j+1,i+1}$, say P_1 (the other cases are isomorphic). Then, we can consider the cycle consisting of the subpath of P_1 between $c^{j,i}$ and $c^{j,i+1}$, the edges $c^{j,i+1}a^{j,i+1}$ and $c^{j,i}d^{j,i}$ and the subpath of P_2 between $d^{j,i}$ and $a^{j,i+1}$. Because of Lemma 1, no path can enter inside this cycle, but it contains at least one crossing vertex in rows i and $i+1$, leading to a contradiction.

□

Lemma 11. *The complementary Q of $\mathscr{P}$ can be decomposed into cycles and one (w_i, w'_i)-path for each $i \in [\![1, p']\!]$.*

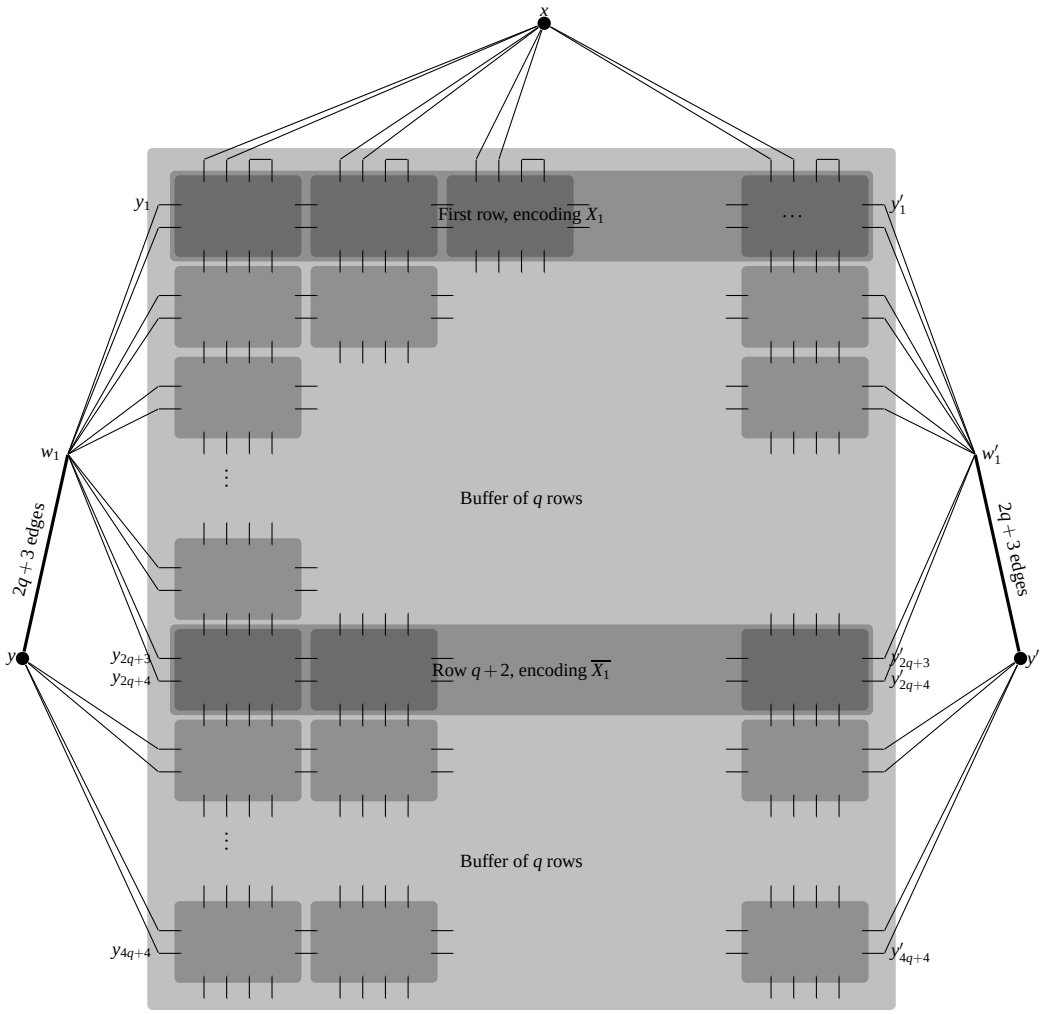

Figure 12: *The first $2(q+1)$ rows of the graph obtained by reduction. These are the rows corresponding to the encoding of the first variable.*

Proof. There are exactly np gadgets in the graph, among which exactly $3n$ are LIC. The number of crossing vertices is exactly $4np+6n$. Moreover, there are $2n$ vertical paths, crossing each of the $2p-p'$ horizontal paths. At most $2n(p'+3)$ crossing vertices are not used to cross paths. As the number of rows in a buffer is $q=4(p'+3)n+2$, there are at least two consecutive rows where all the crossing vertices are used to cross paths. Then we can apply Lemma 10 : there is no no-path going from the top to the bottom of a buffer.
Because of the parity of vertices, $G \setminus \mathscr{P}$ is a $W \cup W'$-join, and can be decomposed in cycles and some paths with extremities in $W \cup W'$, but there can only be (w_i, w'_i)-paths, as all other possible paths would have to go through a buffer. □

This leads to a key consequence :

Lemma 12. *Every vertical path is contained in one column.*

Proof. $2p-p'$ horizontal paths and p' no-paths are routed through each vertical cut, which contains exactly $2p$ edges. Then vertical paths cannot use edges of vertical cuts. □

6.2 Proof of the reduction

Theorem 1. *The planar edge-disjoint paths problem is strongly NP-complete, even if the demand graph has only two edges, with terminals lying on the boundary of the infinite face of the input graph.*

Proof. We use the graph of polynomial size built in Section 6.
Suppose that the formula is satisfiable, and consider an assignment satisfying φ. We route two horizontal paths through each row with the following exceptions, where we route only one horizontal path :

- in row $1+2(k-1)(1+q)$ if the value of variable k is *true*,
- in row $q+2+2(k-1)(1+q)$ if the value of variable k is *false*

Then, for each column, we switch the two vertical paths, using Figure 6, except in the row corresponding to the first variable satisfying the clause associated with the column, where we keep the vertical paths. This is possible by Lemma 7, as there is a LIC at this intersection. The vertical paths of each column are kept exactly once. Then these paths are valid, whence the edge-disjoint paths problem has a solution.
Suppose now that there is a solution to the edge-disjoint paths problem. By Lemma 12, the vertical paths do not intersect vertical cuts, and each no-path or horizontal path intersects each vertical path only once. We show that each no-path can use at most 3 distinct rows in each column (and thus are separated by at least $q-2n-1$ rows), iteratively on the columns of the grid. This is done by applying Lemma 8 (because there is no other no-path in the two nearest rows). Then, for each $k \in [\![1, p']\!]$, the (w_k, w'_k)-path cannot intersect both the rows between $2(q+1)(k-1)-2$ and $2(q+1)(k-1)+4$, and the rows between $2(q+1)(k-1)+q-1$ and $2(q+1)(k-1)+q+5$, as $q > 2n+5$. If it intersects the first group, we set the variable V_k to *true*, otherwise we set it to *false*.
We consider an arbitrary column. In each row except for p' distinct groups of three consecutive rows, there are two horizontal paths traversing the gadget from left to right, as in the hypothesis of Lemmas 4 and 6. Moreover, by construction, there are XCH above and below every LIC, so we can effectively apply Lemmas 4 and 6 to all the gadgets, except a block of five consecutive gadgets for each variable, which contains the intersection of the corresponding no-path with the given column.
Now, for each of these blocks, if there are only XCH, the first and fifth gadgets follow Lemma 4, and by Lemma 9, the vertical paths cannot be shifted here. Otherwise, there is a LIC and a no-path in one of the three nearest rows, thus the variable assignment validates the corresponding clause (the fact that the no-path is not forced to pass through the LIC has no consequence, as long as it must be close enough). Then whenever the vertical paths are not shifted, the clause associated with the column is satisfied. As each column sees its vertical paths kept at least once, each clause is satisfied, thus the assignment is feasible for φ. □

7 Directed case

Using a folkloric reduction, the following result is an obvious consequence of Theorem 1.

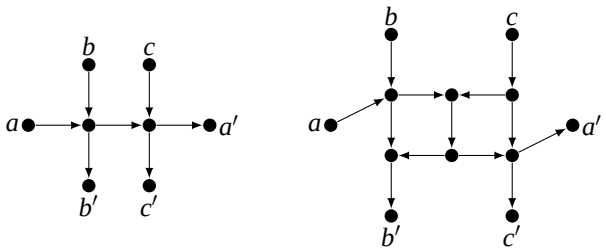

Figure 13: *In NO, there is no path from c to b', whereas it is possible in YES, as long as no other path goes through the graph.*

Corollary 1. *The arc-disjoint paths problem is strongly NP-complete, even if G is planar, the demand graph has only two arcs and the terminals lie on the boundary of the infinite face of G.*

We use again the grids of subgraphs, but with different subgraphs, to prove the following theorem :

Theorem 2. *The planar arc-disjoint paths problem is NP-complete, even if G is acyclic and H consists of two sets of parallel edges.*

We will reduce from SATISFIABILITY. Let $C_1 \wedge \ldots \wedge C_n$ be a formula with n clauses, over the set of variables $\{X_1, \ldots, X_p\}$. Let G_1 be a grid with n columns and $2p$ rows, where each point $G_1(i,j), i \in [\![1,2p]\!], j \in [\![1,n]\!]$ of the grid is a special subgraph, defined as follows (see figure 13):

- $G_1(2i-1,j)$ is the graph YES if X_i appears positively in C_j,
- $G_1(2i,j)$ is the graph YES if X_i appears negatively in C_j,
- $G_1(i,j)$ is NO in all other cases.

Claim 1. *The formula is satisfiable if and only if there is a set $\mathscr{P}$ of arc-disjoint paths in G_1 such that :*

(i) for each $j \in [\![1,n]\!]$, there is a path P_j in $\mathscr{P}$ from $c \in G_1(1,j)$ to $b' \in G_1(2p,j)$,

(ii) for each $i \in [\![1,p]\!]$, there is a path Q_i in $\mathscr{P}$ either from $a \in G_1(2i-1,1)$ to $a' \in G_1(2i-1,n)$ or from $a \in G_1(2i,1)$ to $a' \in G_1(2i,n)$.

Proof. Suppose that $\mathscr{P}$ exists. For all $i \in [\![1,p]\!]$, if Q_i has extremities $a \in G_1(2i-1,1)$ and $a' \in G_1(2i-1,n)$, then assign value *false* to X_i, otherwise assign value *true*. Horizontal cuts and vertical cuts are directed, thus every path Q_i, $(i \in [\![1,p]\!])$ is contained in a single row and every path $P_j, j \in [\![1,n]\!]$, is contained in a single column. For each path $P_j, j \in [\![1,n]\!]$, let i be the index of the first row where P_j goes through the left edge between $G_1(i,j)$ and $G_1(i+1,j)$. Then P_j is the only path that goes through $G_1(i,j)$, and $G_1(i,j)$ is a YES graph. If i is even, it means that $X_{\frac{i}{2}}$ appears negatively in C_j and this variable has value *false*, thus C_j is satisfied. Otherwise $X_{\frac{i+1}{2}}$ appears positively in C_j and the value of this variable is *true*, thus C_j is also satisfied. Then the formula is satisfied. The converse is obvious. □

We just have to enforce paths to be as required in the previous claim. Condition (i) is easy to satisfy. To check condition (ii) we need some gadgets, depicted in Figure 14.

Claim 2. *Let P_1 be a path between a and one of a_i, $i \in \{1,2\}$, and P_2 be a path between b and b_j, $j \in \{1,2\}$, in IF, LL or TT. If P_1 and P_2 are arc-disjoint, then $i = j$.* □

Claim 3. *There are not two arc-disjoint paths in VV, one from b_2 to b and the other from a_1 to a.* □

Claim 2 and Claim 3 can be readily checked. We will also need the graph ON introduced in Figure 15. We now describe the full graph for the reduction. G is build from G_1 in the model of Figure 16. G is built from a grid with $2p+n$ columns and $2p$ rows. The subgrid defined by columns $p+1$ to $p+n$ and rows 1 to p is G_1. Note that two rows in G_1 correspond to one row in G. Squares $G(i,i)$, $G(p+1-i, n+p+i)$,

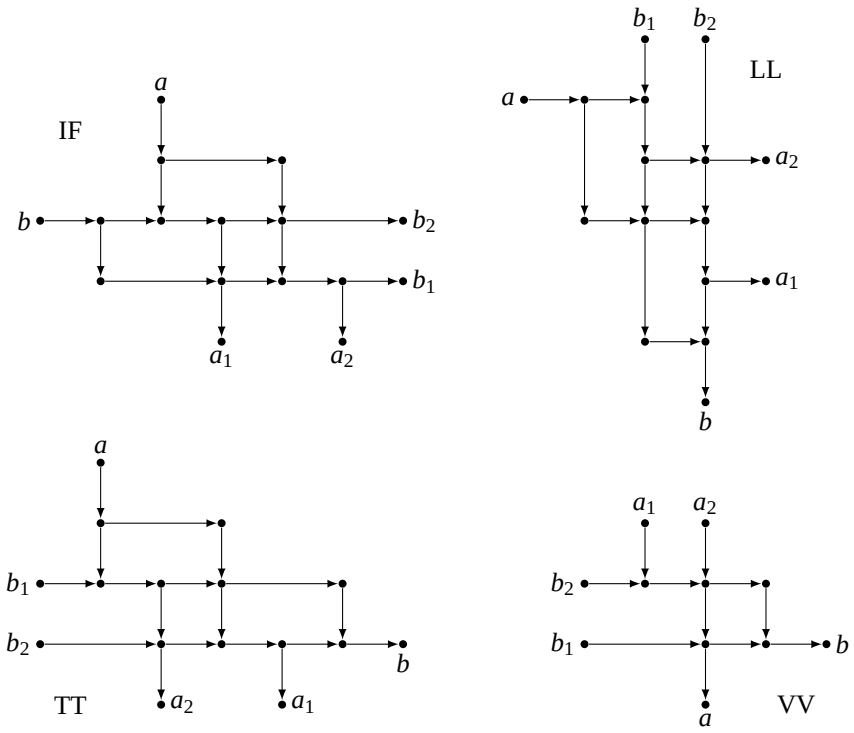

Figure 14: *Four special subgraphs, from top to bottom and left to right : IF, LL, TT, VV. First three have the property that if there are two arc-disjoint paths, one between a and a_i and the other between b and b_j, then $i = j$.*

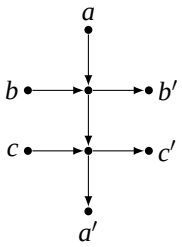

Figure 15: *Graph ON, it has the same property as graph NO.*

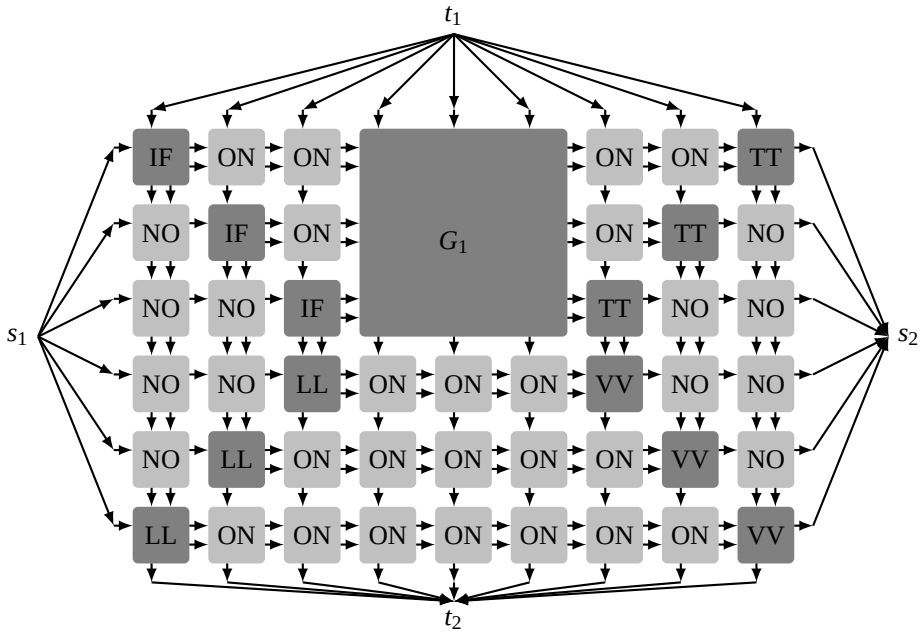

Figure 16: *The graph for the reduction with a formula containing three clauses over three variables.*

$G(2p+1-i,i)$ and $G(p+i,n+p+i)$, for all $i \in [\![1,p]\!]$, are special graph IF, TT, LL and VV respectively. Others are either NO or ON, according to the figure. Rows, columns, vertical cuts and horizontal cuts are defined in the same way as in Section 6. We add four terminals, one for each side of the grid (see the figure).

We add a demand of $2p$ from s_1 to s_2, and $2p+n$ from t_1 to t_2.

Claim 4. *G is acyclic.*

Proof. Observe that in the grid, all arcs are from left to right or from top to bottom, and the special graphs are all acyclic. □

Claim 5. *There is exactly one path going from top to bottom in each column of G, this path never leaves the column. There is exactly one path going from left to right in each row of G, this path never leaves the row.*

Proof. Because vertical and horizontal cuts are directed, and $\{s_1\}$, $\{s_2\}$, $\{t_1\}$, $\{t_2\}$ are tight cuts. □

Claim 6. *If a horizontal path leaves G_1 by the lower edge of its row in the $(n+p)^{th}$ vertical cut, then the same path enters G_1 by the lower edge of its row in the p^{th} vertical cut.*

Proof. In each square ON or NO, there is exactly one path from b to b' or from c to c', and one path from a to a', because of properties of these gadgets. Consequently, a vertical path leaves an IF gadget by the left if and only if it enters the LL gadget of the same column by the left. Similarly for paths between LL and VV, for paths between TT and VV, for paths between IF and G_1 and for paths between G_1 and TT. If a path leaves G_1 by the lower edges of its row, say row $i \in [\![1,p]\!]$, then it enters $G(i,2p+n+1-i)$ by vertex b_2. The vertical path of column $2p+n+1-i$ leaves $G(i,2p+n+1-i)$ by vertex a_2 by Claim 2, and then enters $G(2p+1-i,2p+n+1-i)$ by vertex a_1. Thus the horizontal path in row $2p+1-i$ goes in $G(2p+1-i,2p+n+1-i)$ using vertex b_1 by Claim 3, and leaves $G(2p+1-i,i)$ by vertex a_1. Using Claim 2, the vertical path in column i goes in $G(2p+1-i,i)$ by vertex b_1, thus goes out $G(i,i)$ by vertex a_1. By Claim 2 again, the horizontal path of row i leaves $G(i,i)$ by b_1, and then enters G_1 by the lower edge, proving the claim. □

Proof. (of Theorem 2) Claim 6 proves that the path in G_1 satisfies the condition (ii) of Claim 1: if there is a solution to the arc disjoint path problem, the formula is satisfiable. The converse is also true, it is sufficient

to extend the solution for G_1 naturally.
As the construction is obviously polynomial, we found a Karp reduction between the two problems. The arc-disjoint paths problem being in NP, Theorem 2 is proved. Note that we can identify t_1 with s_2, and t_2 with s_1, thus proving that it is still true with $G+H$ planar and only two terminals. □

Corollary 2. *The arc-disjoint paths problem in planar graphs is NP-complete, even if the demand graph has only two arcs, with one of request* 1 *(one flow plus one path).*

Proof. We modify the preceding reduction. We remove t_1 and t_2, and add arcs from the bottom of a column, to the top of the next column to the left. This preserves planarity. Then we add a demand arc from the bottom of the leftmost column to the top of the rightmost column, with demand 1. We keep the arc s_2s_1. The new demand must be routed through the new arcs because of the vertical tight cuts. Thus, this transformation preserves the property of the original reduction. □

The problem where the total amount of demands is fixed remains open in digraphs, and in particular, the special case where we want to find a cycle in a planar digraph, that goes through two specified vertices. This last problem is mentioned in [9].

References

[1] Frank, A.: Edge-disjoint paths in planar graphs. Journal of Combinatorial Theory B **39** , 164-178 (1985)

[2] Frank, A.: Packing paths, circuits and cuts - a survey. In: Korte, B., Lovász, L., Prömel, H.J., Schrijver, A. (eds.) Paths, Flows, and VLSI-Layout, 47-100. Springer, Berlin (1990)

[3] Kramer, M.R, Van Leeuwen, J.: The Complexity of Wire-Routing and Finding the Minimum Area Layouts for Arbitrary VLSI Circuits. In: F.P. Preparata : Advances in Computing Research **2** : VLSI Theory, 129-146. JAI press, London (1984)

[4] Menger, K.: Zur allgemeinen Kurventheorie. Fundamenta Mathematicae **10**, 96-115 (1927).

[5] Müller, D.: On the complexity of the planar directed edge-disjoint paths problem. Mathematical Programming, vol. **105**, No.2-3, 275-288 (2006)

[6] Okamura, H., Seymour, P.D.: Multicommodity flows in planar graphs. Journal of Combinatorial Theory B **31**, 75-81 (1981)

[7] Okamura, H.: Multicommodity flows in graphs. Discrete Applied Mathematics **6**, 55-62 (1983)

[8] Robertson, N., Seymour, P.D.: Graph Minors XIII. The Disjoint Paths Problem. Journal of Combinatorial Theory B **63**, 65-110 (1995)

[9] Schrijver, A.: Combinatorial Optimization : Polyhedra and Efficiency. Springer (2003)

[10] Schwärzler, W.: On the complexity of the planar edge-disjoint paths problem with terminals on the outer boundary. Combinatorica **29**, 121-126 (2009).