



An Approximation Algorithm for Fully Planar Edge-Disjoint Paths

Chien-Chung Huang, Mathieu Mari, Claire Mathieu, Kevin Schewior, Jens Vygen

► To cite this version:

Chien-Chung Huang, Mathieu Mari, Claire Mathieu, Kevin Schewior, Jens Vygen. An Approximation Algorithm for Fully Planar Edge-Disjoint Paths. SIAM Journal on Discrete Mathematics, 2021, 35 (2), pp.752-769. 10.1137/20M1319401 . hal-03099892

HAL Id: hal-03099892

<https://hal.science/hal-03099892>

Submitted on 6 Jan 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

An Approximation Algorithm for Fully Planar Edge-Disjoint Paths

Chien-Chung Huang*, Mathieu Mari†, Claire Mathieu‡, Kevin Schewior§, and Jens Vygen¶

January 7, 2020

Abstract

We devise a constant-factor approximation algorithm for the maximization version of the edge-disjoint paths problem if the supply graph together with the demand edges form a planar graph. By planar duality this is equivalent to packing cuts in a planar graph such that each cut contains exactly one demand edge. We also show that the natural linear programming relaxations have constant integrality gap, yielding an approximate max-multiflow min-multicut theorem.

1 Introduction

The edge-disjoint paths problem (EDP) is a fundamental problem in combinatorial optimization, consisting of connecting as many demand pairs as possible in a graph via edge-disjoint paths. There is a large body of literature studying this problem in various settings. A primary goal has been to find conditions under which there is a solution satisfying all demands, e.g. when the *cut condition* is sufficient; see Frank’s survey [14] or part VII of Schrijver’s book [31]. Unfortunately many cases of EDP are NP-hard, so it is natural to look for approximation algorithms. However there is no general theory and constant-factor approximations can only be expected in special cases.

One of the landmark results in this area is due to Seymour. Let $G + H$ denote the union (of the edge sets) of the supply graph G and the demand graph H . Seymour [34] proved that if $G + H$ is *planar and Eulerian*, the cut condition guarantees a solution to connect all demand pairs; such a solution can be found in polynomial time. Seymour’s pioneering result has motivated a sequence of follow-up works investigating EDP when $G + H$ is planar. We refer the readers to Frank’s surveys [14, 15] and Schrijver’s book [31, in particular Chapter 74.2] for an overview of these results. For example, one can decide in polynomial time whether

*CNRS, École Normale Supérieure, Université PSL, Paris, chien-chung.huang@ens.fr

†École Normale Supérieure, Université PSL, Paris, mathieu.mari@ens.fr

‡CNRS, IRIF, Paris-Diderot, claire.mathieu@irif.fr

§Universität zu Köln, kschewior@gmail.com. Supported by a DAAD PRIME grant.

¶Research Institute for Discrete Mathematics, Hausdorff Center for Mathematics, University of Bonn, vygen@or.uni-bonn.de

all demand pairs can be connected when $G + H$ is planar and the demand pairs lie on a bounded number of faces of G [26].

Unfortunately the *general* case that $G + H$ is planar is one of these cases in which EDP is NP-hard, as Middendorf and Pfeiffer [26] proved. Very little seems to be known about approximation in that setting. Korach and Penn [24] showed that, given the cut condition, one can satisfy all demands except one for each face of G , and such a solution can be found in polynomial time. However, this does not imply an approximation ratio in general.

1.1 Our Results

Before we present our main results, let us define the edge-disjoint paths problem formally:

Definition 1. *The EDGEDISJOINTPATHS problem (EDP) takes as input a supply graph $G = (V, E)$ and a demand graph $H = (V, D)$ and asks for a maximum-cardinality set of pairwise edge-disjoint cycles such that each cycle consists of an edge $\{u, v\}$ in D and a path between u and v in E .*

The conditions that the set \mathcal{C} of cycles in $G + H$ must satisfy can equivalently be written as: (1) $\forall C \in \mathcal{C}, |C \cap D| = 1$, (2) $\forall e \in D \cup E, |\{C \mid C \in \mathcal{C}, C \ni e\}| \leq 1$. We consider the case when $G + H$ is planar and present the first constant-factor approximation algorithm:

Theorem 1. *There is a polynomial-time 32-approximation algorithm for EDGEDISJOINTPATHS if $G + H$ is planar.*

We prove this theorem in Sections 3, 4, and 5. Our proof is based on rounding a fractional solution to a certain LP relaxation to obtain a subset of demand edges for which the cut condition is satisfied, i.e., no cut has more of these demand edges than supply edges: in the planar dual, this translates into what we call the NONNEGATIVECYCLES problem. The NONNEGATIVECYCLES problem is NP-hard even when $G + H$ is planar (Section 7), but we give an approximation algorithm. Once this rounding is done and we have an approximate solution of NONNEGATIVECYCLES, we can obtain edge-disjoint paths for at least half of these demand edges by building on ideas of Korach and Penn [24]. We use the four color theorem in several places.

Section 6 analyzes the integrality gaps of various LPs. By comparing the output of our algorithm to the value of the natural LP relaxation, we get an upper bound of 32 on its integrality gap. A special case, when every demand edge has an infinite (or large enough) number of parallel copies, has been called *maximum integer multiflow*. In this case the dual LP is a relaxation of the multicut problem, which asks for a smallest set of supply edges whose deletion destroys all paths for all demand edges. We show that this dual LP has an integrality gap of at most 2 if $G + H$ is planar. This yields:

Theorem 2. *If $G + H$ is planar, the minimum cardinality of a multicut is at most 64 times the maximum value of an integer multiflow.*

This is one of the few cases with a constant upper bound on the ratio of the smallest multicut and the largest *integer* multiflow (see, e.g., [36]). Another such case is when G arises from a tree by duplicating edges; then this ratio is 2 [18]; in general the ratio can be as large as $\Theta(|D|)$, even when G is planar.

1.2 Further Related Work

Hardness. The decision version of EDP is NP-complete [22], even in many special cases [28]. In terms of approximation, EDP is APX-hard [1]. Assuming that $\text{NP} \not\subseteq \text{DTIME}(n^{O(\log n)})$, where $n = |V|$, there is no $2^{o(\sqrt{\log n})}$ -approximation for EDP, even when G is planar and subcubic and each demand edge has one of its endpoints on the outer face of G [9]. Assuming that for some $\delta > 0$, not all problems in NP can be solved in randomized time 2^{n^δ} , there is no $n^{O(1/(\log \log n)^2)}$ -approximation even when G is a wall graph [10].

The difficulty is further illustrated by the integrality gap of a natural LP relaxation: even when G is planar and subcubic, the integrality gap is already in the order of $\Theta(\sqrt{n})$ [18].

Positive results. EDP can be solved in polynomial time when the number of demand edges is bounded by a constant [30]. The best known approximation guarantee is $O(\sqrt{n})$ [4], even when G is planar.

Nonetheless, there are some special cases for which significantly better approximation ratios are known: for instance, when G is planar and Eulerian, or planar and 4-edge-connected, there is an $O(\log n)$ -approximation [23]. When G is a grid, there is an $O(\log^2 n)$ -approximation [2]. When G is a wall graph, there is a $\tilde{O}(n^{1/4})$ -approximation [7]. When in addition all demand edges have one endpoint at the boundary of the wall, there is an $2^{O(\sqrt{\log n} \cdot \log \log n)}$ -approximation [10]. Yet none of those results gets to the range of a constant-factor approximation.

Variants. One way to relax EDP is to allow for congestion c , that is, in the solution, up to c paths may share the same edge. It is known that EDP becomes significantly easier with congestion: with congestion 2, there is a polylogarithmic approximation [6], and when in addition G is planar, there is a constant-factor approximation [33].

A closely-related but more difficult problem is the Node-Disjoint Paths (NDP) problem, in which two paths may not even share a common *node*. EDP can be reduced to NDP by taking the line graph of the supply graph, but this reduction does not preserve planarity. If $G + H$ is planar, however, Middendorf and Pfeiffer showed how to reduce EDP to NDP while preserving planarity. Hence the $\tilde{O}(n^{9/19})$ -approximation algorithm for NDP with planar supply graphs [8] implies the same approximation ratio for EDP with $G + H$ planar.

Another more general version of EDP is its directed version (one can reduce EDP to its directed version by replacing each edge by a simple gadget). The directed version seems to be strictly harder than the undirected EDP; for example it is even NP-hard for two demand edges [13].

2 Preliminaries

We recall here some notions and well-known properties related to cuts, cycles, planar duality, and T -joins. Many readers may want to skip this section.

2.1 Cycles and Cuts

Throughout this work we work with a (multi)graph $(V, D \dot{\cup} E)$ whose edge set is partitioned into two sets, D and E . For example, we write $G + H = (V, D \dot{\cup} E)$, taking the union of the edge set of the supply graph $G = (V, E)$ and of the demand graph $H = (V, D)$. Each of these graphs can have parallel edges, and further parallel edges can arise by taking the (disjoint) union. We assume throughout, without loss of generality, that $G + H$ is connected.

A *cycle* (or circuit) in $(V, D \dot{\cup} E)$ is a set $C = \{e_1, \dots, e_k\}$ of edges for which there is a set $V(C) = \{v_1, \dots, v_k\}$ of k distinct vertices such that e_i has endpoints v_i and v_{i+1} for all $i = 1, \dots, k$ (where $v_{k+1} := v_1$). If the graph has loops, every loop would also constitute a cycle.

A D -cycle is a cycle that contains exactly one edge of D . Then EDGEDISJOINTPATHS asks for a maximum number of pairwise disjoint D -cycles.

We say an instance of EDP has a *complete solution* if there are $|D|$ pairwise disjoint D -cycles. For the existence of a complete solution, a necessary condition is the *cut condition*:

$$|C \cap D| \leq |C \cap E| \text{ for every cut } C. \quad (1)$$

It is necessary because every cycle intersects every cut in an even number of edges. In fact it is equivalent to requiring (1) only for *simple* cuts, i.e., edge sets $\delta(U)$ where both U and $V \setminus U$ induce connected subgraphs. A cut is simple if and only if it is a minimal cut.

The cut condition is very useful although in general it is NP-hard to check (this is an unpublished observation of Sebő). The reason is that in many special cases, e.g. when $G + H$ is planar and Eulerian [34], it is sufficient and then can also be checked in polynomial time. However, the cut condition is not sufficient in general, even when $G + H$ is planar. For instance, if $G + H$ is K_4 , the complete graph on four vertices, and D is a perfect matching, then the cut condition is satisfied but there is no complete solution.

2.2 Using Planarity

We exploit planarity in two different ways. First, we need an algorithmic version of the four color theorem. Robertson, Sanders, Seymour, and Thomas [29] improved on the original proof by Appel, Haken, and Koch and showed:

Theorem 3. *Given a planar graph, there exists an $O(n^2)$ -time algorithm to color its vertices with four colors such that no two adjacent vertices have the same color.*

Second, we use planar duality. It is well known that for any connected planar graph G there is a connected planar graph G^* (a planar dual) such that there is a one-to-one-correspondence between the edges of G and the edges of G^* that maps the cycles of G to the simple cuts of G^* . Such a graph can be computed in polynomial time (see, e.g., [27]). The planar dual is not always unique, but we can take an arbitrary one. We will not distinguish between edges and edge sets and their images in the planar dual.

Let us call a simple cut that contains exactly one edge of D a D -cut. Then an equivalent formulation of EDP, if $G + H$ is planar, asks for finding a maximum number of pairwise

disjoint D -cuts in the planar dual graph $(G + H)^*$. The cut condition translates to the *cycle condition* in the planar dual graph:

$$|C \cap D| \leq |C \cap E| \text{ for every cycle } C. \quad (2)$$

Deleting an edge in a planar graph is equivalent to contracting the corresponding edge in the planar dual. These operations preserve planarity.

2.3 T -Joins

Condition (2) is well known in the context of T -joins. For an edge set F , let $\text{odd}(F)$ denote the set of vertices whose degree is odd in (V, F) . For a subset $T \subseteq V$, a T -join is a set F of edges with $\text{odd}(F) = T$. If F is a minimum-cardinality $\text{odd}(F)$ -join, then F is called a *join*. Guan [20] observed that F is a join if and only if

$$|C \cap F| \leq |C \setminus F| \text{ for every cycle } C. \quad (3)$$

For any connected graph, any T with $|T|$ even, and any real edge weights, a minimum-weight T -join can be computed in polynomial time [12].

If we have a set of pairwise disjoint D -cuts and $D' \subseteq D$ is the set of demand edges that belong to one of these cuts, then we must have

$$|C \cap D'| \leq |C \cap E| \text{ for every cycle } C. \quad (4)$$

This is equivalent to saying that D' is a join after contracting the edges in $D \setminus D'$. Although this condition is again not sufficient, we will first find a set $D' \subseteq D$ satisfying (4) (in the planar dual graph).

3 Roadmap

3.1 Satisfying the Cycle Condition

The preceding considerations motivate the following problem in the planar dual graph.

Definition 2. The `NONNEGATIVECYCLES` problem (NNC) takes as input a supply graph $G = (V, E)$ and a demand graph $H = (V, D)$ and asks for a maximum-cardinality set $D' \subseteq D$ of demand edges such that

$$|C \cap D'| \leq |C \cap E| \text{ for every cycle } C \text{ in } G + H.$$

In other words, if edges in E have weight 1, edges in D' have weight -1 , and edges in $D \setminus D'$ have weight 0, then every cycle must have nonnegative total weight. This problem is NP-hard even when $G + H$ is planar; see Section 7. We devise a 16-approximation algorithm if $G + H$ is planar. It begins by solving the following linear programming relaxation, which

we call the *nonnegative cycle LP*:

$$\begin{aligned}
& \max \sum_{e \in D} x_e \\
& \sum_{e \in C \cap D} x_e \leq |C \cap E| \quad \forall \text{ cycles } C \\
& 0 \leq x_e \leq 1 \quad \forall e \in D
\end{aligned} \tag{5}$$

Lemma 1. *The nonnegative cycle LP (5) can be solved in polynomial time.*

Proof. By the equivalence of optimization and separation [19], it is sufficient to solve the separation problem. To this end, given $x \in \mathbb{R}^D$, define weights $w(e) := -x_e$ for $e \in D$ and $w(e) := 1$ for $e \in E$. Then the separation problem reduces to finding a negative-weight cycle or deciding that there is none. This can be done by computing a minimum-weight \emptyset -join J . If $w(J) \geq 0$, then there is no negative-weight cycle. Otherwise J can be decomposed into cycles, and at least one of them will have negative weight. \square

The second step of our algorithm will round the LP solution. In Section 4 we will prove:

Theorem 4. *Let x be a feasible solution to the nonnegative cycle LP (5). Then there is a polynomial-time algorithm to construct an integral feasible solution that has value at least $\frac{1}{16} \sum_{e \in D} x_e$.*

3.2 Packing T -Cuts

Let x be an integral feasible solution to the nonnegative cycle LP (in the planar dual), and let $J \subseteq D$ contain those demand edges e for which $x_e = 1$. If we delete the other demand edges (or contract them in the planar dual), the resulting instance now satisfies the cut condition (the cycle condition in the planar dual). Let $G' = (V, J \dot{\cup} E)^*$ be the graph that arises from the planar dual by contracting the edges in $D \setminus J$. Then J is a join in G' . Our goal is now to find as many pairwise disjoint J -cuts in G' as possible because after uncontracting, these will correspond to pairwise disjoint D -cuts in $(G + H)^*$ and hence to pairwise disjoint D -cycles in $G + H$.

Edmonds and Johnson [12] and Lovász [25] showed that a perfect half-integral packing of J -cuts exists and can be computed in polynomial time:

Theorem 5. *For every graph G and every join J in G there exist vertex sets $U_1, \dots, U_{2|J|}$ (not necessarily distinct) that form a laminar family, such that $|\delta(U_i) \cap J| = 1$ for all i and for every edge e of G there are at most two indices i with $e \in \delta(U_i)$. Such sets can be computed in polynomial time.*

This theorem has been called the T -cut packing theorem because if J is a T -join (for some vertex set T) then the cuts will be T -cuts, i.e., cuts $\delta(U)$ with $|U \cap T|$ odd. However, we prefer to continue speaking of J -cuts because this is what we need. See [15] and Theorem 7 for extensions of this result. However, it is not easy (and without the planarity assumption in general impossible) to deduce from this a large number of pairwise disjoint J -cuts. Nevertheless we will show in Section 5:

Theorem 6. *Let G be a planar graph and J a join in G . Then there is a family of at least $\frac{|J|}{2}$ pairwise disjoint J -cuts in G , and such a family can be computed in polynomial time.*

3.3 Overall Algorithm and Proof of Theorem 1

We have now all ingredients to prove our main theorem. The overall algorithm can be summarized as follows. The input consists of two graphs $G = (V, E)$ and $H = (V, D)$ such that $G + H$ is planar.

Step 0: Construct a planar dual graph $(G + H)^*$

Step 1: Solve the nonnegative cycle LP (5) in the planar dual graph $(G + H)^*$, to get a fractional solution x ;

Step 2: Use the algorithm of Theorem 4 (proved in Section 4) to deduce an integral feasible solution x' to NNC with $\sum_{e \in D} x'_e \geq \frac{1}{16} \sum_{e \in D} x_e$. Let $J := \{e \in D : x'_e = 1\}$.

Step 3: Contract the edges in $D \setminus J$ (in the planar dual graph) and use the algorithm of Theorem 6 (proved in Section 5) to compute at least $\frac{|J|}{2}$ pairwise disjoint J -cuts. By planar duality, these correspond to at least $\frac{|J|}{2}$ pairwise disjoint D -cycles in $G + H$. Output these.

We now prove our main result.

Proof. (of Theorem 1) A J -cut in $(V, J \dot{\cup} E)^*$ indeed corresponds to a J -cycle in $(V, J \dot{\cup} E)$ and hence in $G + H$. Therefore the output is a feasible solution. By Lemma 1, Theorem 4, and Theorem 6, the algorithm runs in polynomial time. The number of D -cycles that it computes is at least $\frac{1}{32}$ times the LP value.

On the other hand, consider an optimum solution, say consisting of OPT many D -cycles. Then there is a set $\bar{D} \subseteq D$ such that $|\bar{D}| = \text{OPT}$ and $(G, (V, \bar{D}))$ has a complete solution. This instance must therefore satisfy the cut condition, in other words we have $|C \cap \bar{D}| \leq |C \cap E|$ for every cut C in $G + H$. Therefore, in the planar dual, setting $\bar{x}_e := 1$ for $e \in \bar{D}$ and $\bar{x}_e := 0$ for $e \in D \setminus \bar{D}$ defines a feasible solution to the nonnegative cycle LP. Hence the LP value is at least OPT .

We conclude that the algorithm is a polynomial-time 32-approximation algorithm for EDP if $G + H$ is planar. \square

4 Rounding the Nonnegative Cycle LP (Proof of Theorem 4)

In this section we prove Theorem 4, whose statement we now recall.

Theorem 4. *Let x be a feasible solution to the nonnegative cycle LP (5). Then there is a polynomial-time algorithm to construct an integral feasible solution that has value at least $\frac{1}{16} \sum_{e \in D} x_e$.*

4.1 Algorithm

Starting from a feasible solution x of the nonnegative cycle LP (5), we first contract all cycles in (V, D) . Every edge e that vanishes in this contraction has $x_e = 0$, therefore this preprocessing will not change $\sum_{e \in D} x_e$. Now (V, D) is a forest. Then we execute the following two algorithms and pick the better of the two solutions S_1 and S_2 obtained from Algorithms 1 and 2 respectively.

The first algorithm is independent of the fractional solution x and simply picks a subset of edges incident to leaves. It is the better choice if (V, D) has many leaves (degree-1 vertices).

Algorithm 1: The leaf algorithm

Input: an NNC input $G + H$ such that (V, D) is a forest, and a feasible solution x to the nonnegative cycle LP (5).

Output: a feasible solution S_1 to NONNEGATIVECYCLES.

- 1 Find a 4-coloring of G , consider a color class containing the largest number of leaves of (V, D) , and let I be the set of leaves of that color.
 - 2 Let S_1 be the set of edges in D that are incident to a leaf in I .
 - 3 **return** S_1
-

The second algorithm does a delicate rounding of the fractional LP solution; see Figure 1 for an example.

Algorithm 2: The internal algorithm

Input: an NNC input $G + H$ such that (V, D) is a forest, and a feasible solution x to the nonnegative cycle LP (5).

Output: a feasible solution S_2 to NONNEGATIVECYCLES.

- 1 For each connected component of the forest (V, D) : Root the tree arbitrarily. Let B denote the set of vertices with at least two children, and let L denote the set of leaves.
 - 2 For each $b' \in B \cup L$ and $b \in B \setminus \{b'\}$ such that b is an ancestor of b' and each inner vertex of the b - b' -path has exactly one child: subdivide the first edge on this path, thus inserting an artificial edge $a_{b,b'}$ incident to b , and set $x_{a_{b,b'}} = 0$. Let A denote the set of artificial edges.
 - 3 Define a budget function $y : D \cup A \rightarrow \mathbb{R}_{\geq 0}$. Initially $y_e = x_e$ for all $e \in D \cup A$.
 - 4 For each artificial edge $a_{b,b'} \in A$ follow the path from b down to b' . For each traversed edge $e \in D$, decrease y_e and increase $y_{a_{b,b'}}$ by the same amount. Do this until $y_{a_{b,b'}} = 1$ or we reach b' .
 - 5 Consider the edges $e = (u, v) \in D$ in an order of non-increasing distance from the root, and for each such edge e , if $y_e > 0$ then: Follow the path from u up to the root. For each traversed edge $e' \in D \cup A$, decrease $y_{e'}$ and increase y_e by the same amount. Do this until $y_e = 2$ or we reach the root.
 - 6 $S_2^0 \leftarrow \{e \in D : y_e > 0\}$.
 - 7 Contract each tree of (V, D) , find a 4-coloring of the resulting graph, inducing a 4-coloring of the trees, choose the color class maximizing the number of edges of S_2^0 in trees of that color, and let S_2 denote that subset of S_2^0 .
 - 8 **return** S_2
-




Figure 1: Illustrating Algorithm 2. (a): the original tree with an arbitrarily chosen root; vertices of B are red squares; next to each edge e the value x_e is shown. (b): the rooted tree after inserting the artificial edges (red, dotted) and the initial budgets y_e (after step 3). (c): each artificial edge has collected up to one unit from below (in step 4). (d): edges e_1, \dots, e_5 (shown in blue, bold) are collecting up to two units from above (in step 5) and are included into S_2^0 (in step 6). We see the final distribution if the edges are considered in this order. The green and yellow shades show from where the final budgets y_{e_1}, \dots, y_{e_5} come from.

4.2 Analysis

These algorithms are well-defined and polynomial time by Theorem 3. We will prove that S_1 and S_2 are feasible NONNEGATIVECYCLES solutions (Lemmas 2 and 4, respectively) and that $\max\{|S_1|, |S_2|\} \geq \frac{1}{16} \sum_{e \in D} x_e$ (Lemma 5), thus establishing Theorem 4.

Lemma 2. S_1 is a feasible solution for NONNEGATIVECYCLES.

Proof. Let C be a cycle in $G + H$. By definition of S_1 each edge of $C \cap S_1$ is incident to at least one vertex of I , and since vertices of I are leaves of (V, D) those vertices are all distinct, so $|C \cap S_1| \leq |V(C) \cap I|$. Each vertex of $V(C) \cap I$ is incident to at most one edge of $C \cap D$, hence to at least one edge of $C \cap E$. Since I is an independent set in G , those edges are all distinct, so $|V(C) \cap I| \leq |C \cap E|$. \square

To prove that S_2 is also a feasible solution for NONNEGATIVECYCLES, we first show the following key inequality:

Lemma 3. Let P be the set of edges of a path in (V, D) . Then

$$\sum_{e \in P} x_e \geq 2|S_2 \cap P| - 2. \quad (6)$$

Proof. We may assume that $S_2 \cap P \neq \emptyset$, for the assertion is trivial otherwise. Let v be the vertex in P that is closest to the root (picked in step 1 of Algorithm 2) of the connected component of (V, D) that contains P . Then P is the union of two paths P_1 and P_2 that both begin in v and go down the tree. One of these paths may be empty. Let $i \in \{1, 2\}$ with

$S_2 \cap P_i \neq \emptyset$, and let e_i^h be the edge of $S_2 \cap P_i$ that is closest to v (and to the root). After step 5 of Algorithm 2, every edge e of $(S_2 \cap P_i) \setminus \{e_i^h\}$ has $y_e = 2$, and these budgets come from edges of P_i . The latter is because

- (i) e_i^h is the only edge in $S_2 \cap P_i$ to which budget that is originally from closer to the root than v can be moved, and
- (ii) budget that was moved towards the root (to an artificial edge) in step 4 remains at this artificial edge or is moved in step 5 to an edge that is not closer to the root than the edge where the budget was initially.

Thus,

$$\sum_{e \in P_i} x_e \geq \sum_{e \in S_2 \cap P_i, e \neq e_i^h} y_e = 2(|S_2 \cap P_i| - 1).$$

This already implies (6) if $S_2 \cap P = S_2 \cap P_i$.

It remains to consider the case that neither $S_2 \cap P_1$ nor $S_2 \cap P_2$ is empty. Then $v \in B$, and after step 2 there is an artificial edge incident to v on the path from v to e_i^h (for each $i \in \{1, 2\}$). Let f_i be the last artificial edge on the path from v to e_i^h . Since $e_i^h \in S_2$, we had $y_{e_i^h} > 0$ after step 4 (i.e., even before we started augmenting $y_{e_i^h}$ in step 5). Therefore $y_{f_i} = 1$ after step 4 (and, by the same argument as before, this budget comes from edges of P_i), and $y_{e_i^h} \geq 1$ after step 5. Thus

$$\sum_{e \in P_i} x_e \geq 1 + \sum_{e \in S_2 \cap P_i, e \neq e_i^h} y_e = 1 + 2(|S_2 \cap P_i| - 1) = 2|S_2 \cap P_i| - 1.$$

Adding this for P_1 and P_2 yields (6). □

Lemma 4. S_2 is a feasible solution for NONNEGATIVECYCLES.

Proof. Let C be a cycle in $G+H$ that contains at least one demand edge and at least one supply edge (recall that (V, D) is a forest). The cycle C alternates between D -segments (maximal subpaths of C all whose edges belong to D) and E -segments. Suppose there is some cycle C with $|C \cap S_2| > |C \cap E|$. Among those cycles C , choose one such that C has as few D -segments as possible.

We claim that for every tree T of D , $C \cap T$ is connected. Assume not; then there is a path $P \subseteq D$ whose endpoints u and v belong to $V(C)$, but no edge of P belongs to C . The two endpoints u and v of P partition C into two u -to- v paths, P_1 and P_2 . One of the two cycles $C_1 = P_1 \cup P$ and $C_2 = P_2 \cup P$ must have $|C_i \cap S_2| > |C_i \cap E|$ and has strictly fewer D -segments than C , contradicting the choice of C . This proves the claim.

Thus each D -segment of C belongs to a different tree. By Lemma 3, every D -segment P satisfies (6). Moreover, since x is a feasible solution to the nonnegative cycle LP, $\sum_{e \in C \cap D} x_e \leq |C \cap E|$.

If $|C \cap E| = 1$, then (6) yields $2|C \cap S_2| \leq 2 + \sum_{e \in C \cap D} x_e \leq 2 + |C \cap E| = 3$, implying $|C \cap S_2| \leq 1 = |C \cap E|$. Otherwise, thanks to selecting an independent set in the contracted graph in step 7, if there is more than one D -segment, every E -segment has at least two edges.

Summing (6) over all D -segments P yields $2|C \cap S_2| \leq |C \cap E| + \sum_{e \in C \cap D} x_e \leq |C \cap E| + |C \cap E|$. Therefore, $|C \cap S_2| \leq |C \cap E|$ as required. \square

Lemma 5.

$$\max\{|S_1|, |S_2|\} \geq \frac{1}{16} \sum_{e \in D} x_e.$$

Proof. The construction of Algorithm 2 maintains the invariant $\sum_{e \in D} x_e = \sum_{e \in D \cup A} y_e$. Since edges of S_2^0 have $y_e \leq 2$, artificial edges have $y_e \leq 1$, and other edges have $y_e = 0$, we can write $\sum_{e \in D \cup A} y_e \leq 2|S_2^0| + |A|$. Observe that the number of artificial edges is at most $2|L| - 2$. Finally, by construction $|L| \leq 4|S_1|$ and $|S_2^0| \leq 4|S_2|$.

Combining, we conclude

$$\sum_{e \in D} x_e \leq 2|S_2^0| + 2|L| \leq 8(|S_2| + |S_1|) \leq 16 \max\{|S_1|, |S_2|\}.$$

\square

5 Packing Cuts (Proof of Theorem 6)

In this section we prove Theorem 6, whose statement we now recall.

Theorem 6. *Let G be a planar graph and J a join in G . Then there is a family of at least $\frac{|J|}{2}$ pairwise disjoint J -cuts in G , and such a family can be computed in polynomial time.*

Recall that a J -cut is a simple cut that contains exactly one edge from J . An example by Korach and Penn [24] shows that the factor 2 in Theorem 6 is best possible. Without planarity we cannot hope for any constant factor, e.g. if G is a complete graph and J a perfect matching in G , then there are no two disjoint J -cuts.

Let J be a join in a planar graph G . Middendorf and Pfeiffer [26] showed that it is NP-complete to decide whether there are $|J|$ pairwise disjoint J -cuts. Hence also finding the maximum number of pairwise disjoint J -cuts is NP-hard. However, we give an approximation algorithm in the following. Our proof is inspired by the paper of Korach and Penn [24].

Recall that a family of subsets of V is *laminar* if for any two of those sets either they are disjoint or one is a subset of the other. It is *cross-free* if for any two of those sets, either they are disjoint, or one is a subset of the other, or their union is V . We will show the following:

Lemma 6. *Let $G = (V, E)$ be a connected planar graph. Let \mathcal{L} be a laminar family of (not necessarily distinct) vertex sets of G such that $\delta(U)$ is a simple cut for all $U \in \mathcal{L}$ and every edge is contained in at most two of these cuts. Then there exists a polynomial-time algorithm to compute a sub-family $\mathcal{L}' \subseteq \mathcal{L}$ such that $|\mathcal{L}'| \geq \frac{1}{4}|\mathcal{L}|$ and the cuts $\delta(U)$ for $U \in \mathcal{L}'$ are pairwise disjoint.*

Again, the example by Korach and Penn [24] shows that the factor 4 in Lemma 6 is best possible. Without planarity we cannot hope for any constant factor, e.g. if G is a complete graph and \mathcal{L} contains all singletons.

Before we prove Lemma 6, let us first see how it implies Theorem 6.

Proof. (Theorem 6) Let J be a join in G . By Theorem 5, we can compute vertex sets $U_1, \dots, U_{2|J|}$ (not necessarily distinct) that form a laminar family \mathcal{L} , such that $|\delta(U_i) \cap J| = 1$ for all i and every edge is contained in at most two of the cuts $\delta(U_i)$, $i = 1, \dots, 2|J|$. Before applying Lemma 6, we want to make sure that these cuts are all simple cuts, and this will be achieved in the following by a sequence of transformations.

For $i = 1, \dots, 2|J|$ let $U'_i \subseteq U_i$ be the vertex set of the connected component of $G[U_i]$ for which $|\delta(U'_i) \cap J| = 1$. Then $\delta(U'_i) \subseteq \delta(U_i)$. We claim that the sets $U'_1, \dots, U'_{2|J|}$ form a laminar family \mathcal{L}' . Let $1 \leq i < j \leq 2|J|$. If $U_i \cap U_j = \emptyset$, then $U'_i \cap U'_j = \emptyset$. If $U_i \subseteq U_j$, then the vertex sets of the connected components of $G[U_i]$ are subsets of the vertex sets of the connected components of $G[U_j]$, so U'_i is either disjoint from U'_j or a subset of U'_j . The case $U_j \subseteq U_i$ is symmetric.

Now $G[U'_i]$ is connected for all i . For $i = 1, \dots, 2|J|$ let $U''_i \subseteq V \setminus U'_i$ be the vertex set of the connected component of $G[V \setminus U'_i]$ for which $|\delta(U''_i) \cap J| = 1$. Then $\delta(U''_i) \subseteq \delta(U'_i) \subseteq \delta(U_i)$. Note that for each i , $G[U'_i]$ and $G[V \setminus U'_i]$ are connected and $|\delta(U''_i) \cap J| = 1$. In other words, $\delta(U''_1), \dots, \delta(U''_{2|J|})$ are J -cuts.

We claim that the sets $U''_1, \dots, U''_{2|J|}$ form a cross-free family \mathcal{L}'' . Let $1 \leq i < j \leq 2|J|$. If $U'_i \subseteq U'_j$, then the vertex sets of the connected components of $G[V \setminus U'_j]$ are subsets of the vertex sets of the connected components of $G[V \setminus U'_i]$, so U''_j is either disjoint from U''_i or a subset of U''_i . The case $U'_j \subseteq U'_i$ is symmetric.

It remains to consider the case $U'_i \cap U'_j = \emptyset$. Suppose $X := V \setminus (U''_i \cup U''_j)$ is nonempty. Since G is connected, we have $\emptyset \neq \delta(X) = \delta(U''_i \cup U''_j) \subseteq \delta(U'_i \setminus U''_j) \cup \delta(U'_j \setminus U''_i)$, where the last step is because all edges with exactly one endpoint in U''_i have the other endpoint in U'_i , and all edges with exactly one endpoint in U''_j have the other endpoint in U'_j . Hence at least one of $U'_i \setminus U''_j$ or $U'_j \setminus U''_i$ is nonempty. Suppose, without loss of generality, $U'_i \setminus U''_j \neq \emptyset$. Because $G[U'_i]$ is connected, U'_i is a subset of the vertex set of a connected component of $G[V \setminus U''_j]$. So $U'_i \cap U''_j = \emptyset$. But then U''_j is a subset of the vertex set of a connected component of $G[V \setminus U'_i]$, and thus it is a subset of U''_i or disjoint from U''_i .

So indeed \mathcal{L}'' is cross-free. To obtain a laminar family, choose $v \in V$ arbitrarily, and for every $U'' \in \mathcal{L}''$ with $v \in U''$, replace U'' by $V \setminus U''$. We obtain a laminar family \mathcal{L}''' of $2|J|$ (not necessarily distinct) vertex sets such that $\delta(U)$ is a J -cut for all $U \in \mathcal{L}'''$ and every edge is contained in at most two of these cuts. It is obvious that all the above steps can be performed in polynomial time. Applying Lemma 6 concludes the proof. \square

To prove Lemma 6, consider the following recursive algorithm:

Algorithm 3: CUTPACKING

Input: a planar graph $G = (V, E)$ and a laminar family \mathcal{L} of (not necessarily distinct) nonempty subsets of V such that $G[U]$ and $G[V \setminus U]$ are connected for all $U \in \mathcal{L}$ and every edge is contained in at most two of the cuts $\delta(U)$, $U \in \mathcal{L}$.

Output: a subset $\mathcal{L}' \subseteq \mathcal{L}$ such that the cuts $\delta(U)$, $U \in \mathcal{L}'$, are pairwise disjoint

```

1 if the elements of  $\mathcal{L}$  are pairwise disjoint then /* case 1 */
2   | Consider the planar graph  $P$  obtained by contracting each  $U \in \mathcal{L}$  into a single
3   | vertex and deleting all other vertices. Compute a 4-coloring on  $P$ .
4   | return a subset of  $\mathcal{L}$  that corresponds to a maximum-cardinality color class
5 if there exist  $U_1, U_2 \in \mathcal{L}$  such that  $U_1 = U_2$  then /* case 2 */
6   | return  $\{U_1\} \cup \text{CUTPACKING}(G, \mathcal{L} \setminus \{U_1, U_2\})$ 
7   /* case 3 */
8 Let  $\bar{U} \in \mathcal{L}$  be a set that is not minimal but all sets  $U_1, \dots, U_l \in \mathcal{L}$  that are proper
9 subsets of  $\bar{U}$  are minimal (inclusionwise).
10 Consider the planar graph  $P$  obtained by contracting each  $U_i$  into a normal vertex,
11 deleting all other vertices in  $\bar{U}$ , and contracting all vertices outside  $\bar{U}$  into a single
12 special vertex.
13 Find a 4-coloring of  $P$  and choose a color class  $\mathcal{K}$  with the largest number of normal
14 vertices. If all color classes contain  $\frac{l}{4}$  normal vertices, choose the one that contains the
15 special vertex. Let  $\mathcal{L}' \subset \mathcal{L}$  be the set of  $U \in \mathcal{L}$  that correspond to normal vertices in  $\mathcal{K}$ .
16 if  $\mathcal{K}$  contains more than  $\frac{l}{4}$  normal vertices then /* case 3a */
17   | return  $\mathcal{L}' \cup \text{CUTPACKING}(\mathcal{L} \setminus \{\bar{U}, U_1, \dots, U_l\})$ 
18 else /* case 3b */
19   | return  $\mathcal{L}' \cup \text{CUTPACKING}(\mathcal{L} \setminus \{U_1, \dots, U_l\})$ 

```

Proof. (Lemma 6) We show the assertion by induction on $|\mathcal{L}|$. We say here that a subset $\mathcal{L}' \subseteq \mathcal{L}$ is *stable* if all $\delta(U)$ for $U \in \mathcal{L}'$ are pairwise disjoint. We apply Algorithm 3. Note that the graphs P in step 2 and step 7 are indeed planar because we only contract connected vertex sets. By Theorem 3, the algorithm runs in polynomial time. We start with the base case of our induction.

Case 1: The elements of \mathcal{L} are pairwise disjoint. Let $\mathcal{L}' \subseteq \mathcal{L}$ be the solution returned. \mathcal{L}' corresponds to an independent set in P , so it is stable and has size at least $|V(P)|/4 = |\mathcal{L}|/4$. Notice that when $\mathcal{L} = \emptyset$, case 1 applies, and the empty set is returned.

Case 2: Some $U_1, U_2 \in \mathcal{L}$ are equal. We assume by the induction hypothesis that $\text{CUTPACKING}(G, \mathcal{L} \setminus \{U_1, U_2\})$ is stable and $\text{CUTPACKING}(G, \mathcal{L} \setminus \{U_1, U_2\})$ has size at least $(|\mathcal{L}| - 2)/4$. First, the solution returned has size $\frac{|\mathcal{L}|-2}{4} + 1 = \frac{|\mathcal{L}|+2}{4} > \frac{|\mathcal{L}|}{4}$. To prove that $\{U_1\} \cup \text{CUTPACKING}(G, \mathcal{L} \setminus \{U_1, U_2\})$ is stable, we only have to check that $\delta(U_1) \cap \delta(U) = \emptyset$ for all $U \in \mathcal{L} \setminus \{U_1, U_2\}$. This holds since any edge in $\delta(U_1)$ is also contained in $\delta(U_2)$ and is contained in at most two of the cuts.

Case 3: There exists a set $\bar{U} \in \mathcal{L}$ that is not minimal but all sets $U_1, \dots, U_l \in \mathcal{L}$ that are proper subsets of \bar{U} are minimal. Therefore the sets U_i ($i = 1, \dots, l$) are pairwise disjoint.




Figure 2: The different cases in the proof of Lemma 6.

Then consider the graph P constructed in step 7 of the algorithm and a partition of its vertex set into four independent sets. Now we consider two subcases.

If one of these independent sets contains at least $\frac{l+1}{4}$ normal vertices (case 3a), then by induction hypothesis, the solution has size at least $\frac{l+1}{4} + \frac{|\mathcal{L}|-(l+1)}{4} = \frac{|\mathcal{L}|}{4}$. Otherwise (case 3b) all these independent sets contain exactly $\frac{l}{4}$ normal vertices and we break ties by taking the one that contains the special vertex. Using the induction hypothesis again, the solution returned has size at least $\frac{l}{4} + \frac{|\mathcal{L}|-l}{4} = \frac{|\mathcal{L}|}{4}$.

After using induction hypothesis, this solution is stable because any edge in $\delta(U_i) \cap \delta(U)$ for $U \in \mathcal{L}$ such that $U \subseteq V \setminus \bar{U}$ would also be contained in $\delta(\bar{U})$, but no edge is contained in three of the cuts. If $\bar{U} \in \text{CUTPACKING}(\mathcal{L} \setminus \{U_1, \dots, U_l\})$, then the special vertex is in the independent set \mathcal{K} , and hence $\delta(\bar{U}) \cap \delta(U_i) = \emptyset$ for all $U_i \in \mathcal{L}'$. \square

6 Max-Multiflow Min-Multicut Ratio (Proof of Theorem 2)

6.1 Linear Programming Relaxations and Integrality Gaps

EDGEDISJOINTPATHS consists of finding as many D -cycles as possible. If \mathcal{C} denotes the set of all D -cycles, then we look for a maximum subset of \mathcal{C} whose elements are pairwise disjoint. The natural fractional relaxation is (along with its linear programming dual):

$$\begin{aligned}
 \max \quad & \sum_{C \in \mathcal{C}} f_C \\
 \sum_{C \in \mathcal{C}: e \in C} f_C & \leq 1 \quad \forall e \in D \dot{\cup} E \\
 f_C & \geq 0 \quad \forall C \in \mathcal{C}.
 \end{aligned}
 \qquad
 \begin{aligned}
 \min \quad & \sum_{e \in D \dot{\cup} E} y_e \\
 \sum_{C \in \mathcal{C}} y_e & \geq 1 \quad \forall C \in \mathcal{C} \\
 y_e & \geq 0 \quad \forall e \in D \dot{\cup} E
 \end{aligned}$$

Figure 3: The D -cycle packing LP and its dual, the D -cycle covering LP.

If we go to the planar dual graph and let \mathcal{U} denote the set of vertex sets inducing D -cuts,

these LPs become

$$\begin{array}{ll}
\max \sum_{U \in \mathcal{U}} f_U & \min \sum_{e \in D \dot{\cup} E} y_e \\
\sum_{U \in \mathcal{U}: e \in \delta(U)} f_U \leq 1 & \forall e \in D \dot{\cup} E \\
f_U \geq 0 & \forall U \in \mathcal{U}. \\
\sum_{e \in \delta(U)} y_e \geq 1 & \forall U \in \mathcal{U} \\
y_e \geq 0 & \forall e \in D \dot{\cup} E
\end{array}$$

Figure 4: The D -cut packing LP and its dual, the D -cut covering LP.

It turns out that the D -cut packing LP is equivalent to the nonnegative cycle LP. We need the following well-known generalization of Theorem 5 (see, e.g., [32] and the references therein):

Theorem 7. *For every graph $G = (V, E)$ with capacities $u : E \rightarrow \mathbb{R}_{\geq 0}$ and every set J of edges such that $u(C \cap J) \leq u(C \setminus J)$ for every cycle C , one can compute in strongly polynomial time a laminar family \mathcal{L} of vertex sets with weights $w : \mathcal{L} \rightarrow \mathbb{R}_{> 0}$ such that $|\delta(U) \cap J| = 1$ for all $U \in \mathcal{L}$, $\sum_{U \in \mathcal{L}} w(U) = \sum_{e \in J} u(e)$, and $\sum_{U \in \mathcal{L}: e \in \delta(U)} w(U) \leq u(e)$ for all $e \in E$. If u is integral, then w can be chosen half-integral.*

Using this, we can show:

Lemma 7. *The D -cut packing LP is equivalent to the nonnegative cycle LP (5). Their values are the same, from an optimum solution to the former we can get an optimum solution to the latter in polynomial time, and vice versa.*

Proof. For any feasible solution f of the D -cut packing LP define a vector x by $x_e := \sum_{U \in \mathcal{U}: e \in \delta(U)} f_U$ for $e \in D$. Then x is a feasible solution to the nonnegative cycle LP because for every cycle C in $G + H$ we have

$$\sum_{e \in C \cap D} x_e = \sum_{e \in C \cap D} \sum_{\substack{U \in \mathcal{U}: \\ e \in \delta(U)}} f_U = \sum_{U \in \mathcal{U}} f_U |C \cap D \cap \delta(U)| \leq \sum_{U \in \mathcal{U}} f_U |C \cap E \cap \delta(U)| = \sum_{e \in C \cap E} \sum_{\substack{U \in \mathcal{U}: \\ e \in \delta(U)}} f_U \leq |C \cap E|.$$

Here the first inequality holds because a cycle C and a cut $\delta(U)$ intersect in an even number of edges, and for $U \in \mathcal{U}$ at most one edge in the intersection belongs to D . We also have $\sum_{e \in D} x_e = \sum_{U \in \mathcal{U}} f_U$.

Conversely, let x be a feasible solution to the nonnegative cycle LP. Define $u(e) := x_e$ for $e \in D$ and $u(e) := 1$ for $e \in E$. By Theorem 7 (applied to $G + H$ and $J := D$), one can compute a laminar family \mathcal{L} of vertex sets with weights $w : \mathcal{L} \rightarrow \mathbb{R}_{> 0}$ such that $|\delta(U) \cap D| = 1$ for all $U \in \mathcal{L}$, $\sum_{U \in \mathcal{L}} w(U) = \sum_{e \in D} u(e)$, and $\sum_{U \in \mathcal{L}: e \in \delta(U)} w(U) \leq u(e)$ for all $e \in D \dot{\cup} E$. Set $f_U := w(U)$ for $U \in \mathcal{L}$ and $f_U := 0$ otherwise. Then f is a feasible solution to the D -cut packing LP with $\sum_{U \in \mathcal{U}} f_U = \sum_{U \in \mathcal{L}} w(U) = \sum_{e \in D} u(e) = \sum_{e \in D} x_e$. \square

This implies:

Corollary 1. *If $G + H$ is planar, the integrality gap of the D -cut packing LP (and hence the integrality gap of the D -cycle packing LP) is at least 2 and at most 32.*

Proof. We showed in the proof of Theorem 1 that there is an algorithm that computes an integral solution of the D -cut packing LP of at least $\frac{1}{32}$ times the value of the nonnegative cycle LP. By Lemma 7 this implies the upper bound.

The lower bound of 2 is attained for $G + H = K_4$ (the complete graph on 4 vertices) if D is a perfect matching in $G + H$, see Figure 5 (a). \square

For the nonnegative cycle LP the integrality gap could be smaller. We know that it is between $\frac{3}{2}$ (shown by $G + H = K_4$ and $D = \delta(u)$ for an arbitrary vertex u , see Figure 5 (b)) and 16 (shown by Theorem 4).

We now observe that the dual LPs have integrality gap at most 2.

Lemma 8. *The integrality gap of the D -cut covering LP (and hence the integrality gap of the D -cycle covering LP if $G + H$ is planar) is at most 2, and it is at least $\frac{3}{2}$ even if $G + H$ is planar.*

Proof. The lower bound follows from an example by Cheriyan, Karloff, Khandekar, and Köneemann [5] for the tree augmentation problem (let D contain the tree edges and E contain the links of this example, then their LP is equivalent to the D -cut covering LP). See Figure 5 (c).

We now show the upper bound. The D -cut covering LP is equivalent to

$$\begin{aligned}
 & \min \sum_{e \in D \dot{\cup} E} y_e \tag{7} \\
 & \sum_{e \in \delta(U)} y_e + \sum_{e \in \delta(U) \cap D} z_e \geq 2 \quad \forall U \subseteq V \text{ with } D \cap \delta(U) \neq \emptyset \\
 & y_e \geq 0 \quad \forall e \in D \dot{\cup} E \\
 & 0 \leq z_e \leq 1 \quad \forall e \in D
 \end{aligned}$$

because in (7) we can assume that $z_e = 1$ for all $e \in D$, and then it is exactly the same as the D -cut covering LP. Now (7) is a survivable network design LP of the type

$$\begin{aligned}
 & \min \sum_{e \in E'} c(e)x_e \\
 & \sum_{e \in \delta(U)} x_e \geq f(U) \quad \forall U \subseteq V \\
 & 0 \leq x_e \leq 1 \quad \forall e \in E'
 \end{aligned}$$

for the graph whose edge set E' consists of an edge e of cost $c(e) = 1$ for each $e \in D \dot{\cup} E$ and another parallel edge $e' = \{v, w\}$ with $c(e') = 0$ for each $\{v, w\} \in D$. The requirement function $f : 2^V \rightarrow \mathbb{Z}_{\geq 0}$ is given by $f(U) = 2$ if $D \cap \delta(U) \neq \emptyset$ and $f(U) = 0$ otherwise. This function f is easily seen to be proper (and thus weakly supermodular), and hence Jain's iterative rounding theorem [21] tells that there is an integral feasible solution (y, z) to (7) of cost at most twice the LP value. Then y is an integral feasible solution to the D -cut covering LP of the same cost. \square




Figure 5: Examples for lower bounds of integrality gaps. Solid black edges belong to E , dashed red edges belong to D .

We remark that even if just G is planar (not $G + H$) the integrality gap of the D -cycle covering LP is bounded by a (large) constant [35].

6.2 Multiflows and Multicuts

A natural generalization of EDP takes as input, in addition to G and H , a capacity function $u : D \dot{\cup} E \rightarrow \mathbb{Z}_{>0}$ and asks for a maximum number of D -cycles such that each edge $e \in D \dot{\cup} E$ belongs to at most $u(e)$ of them. If $u(e) = 1$ for all $e \in D \dot{\cup} E$, this is EDP.

This generalization reduces to EDP by replacing each edge e by $u(e)$ parallel edges of unit capacity. In the planar dual, this corresponds to replacing e by a path with $u(e)$ edges. If u is given in binary representation, this reduction should of course not be performed explicitly. Nevertheless our algorithm can easily be generalized to run in polynomial time for general capacities. A dual of a weighted planar graph can be coded by a weighted planar graph where an edge e of weight w_e represents a path of length w_e . All the algorithms involved in proving Theorem 4 can be easily extended to weighted instances. Once we have a weighted feasible solution J that respects the non-negative cycle condition, we apply Theorem 7 to get a (weighted) laminar family that satisfies the requirements of Lemma 6. Then, the final feasible solution of EDP can be represented by an integral flow of a certain value for some demand edges. These flows can be decomposed into (and hence represented by) at most $|E|$ paths by standard flow decomposition.

If $u(e) = \infty$ (or large enough) for all $e \in D$, the problem has been called the maximum integer multiflow problem. (Note that $u(e) = \infty$ for all $e \in D$ can in fact be assumed without loss of generality because a demand edge $e = \{v, w\}$ with capacity $u(e)$ can be replaced equivalently by a demand edge $\{v, x\}$ of infinite capacity and a supply edge $\{x, w\}$ of capacity $u(e)$, where x is a new vertex.)

If we do not require integrality, the maximum multiflow problem is of course a linear program:

$$\begin{array}{ll}
\max \sum_{C \in \mathcal{C}} f_C & \min \sum_{e \in E} u(e) y_e \\
\sum_{C \in \mathcal{C}, C \ni e} f_C \leq u(e) & \forall e \in E \\
f_C \geq 0 & \forall C \in \mathcal{C}. \\
\sum_{e \in C \cap E} y_e \geq 1 & \forall C \in \mathcal{C} \\
y_e \geq 0 & \forall e \in E
\end{array}$$

Figure 6: The multiflow LP and its dual, the multicut LP.

The feasible solutions to the multiflow LP are called multiflows (here we use the form after flow decomposition). The integral feasible solutions to the dual LP correspond to edge sets $F \subseteq E$ such that deleting F destroys all D -cycles. They are called *multicuts*. The *capacity* of a multicut is the total capacity of its edges. Of particular interest is the worst ratio of the minimum capacity of a multicut and the maximum value of an integer multiflow. The (integer version of the) famous max-flow min-cut theorem says that this ratio is 1 if $|D| = 1$. However in general, even if G is sub-cubic and planar the ratio can be as large as $\Theta(|D|)$ [18]. Besides when G is a tree (then the ratio is 2 [18]), when G is planar and has bounded tree-width [3] or when $G + H$ is series-parallel [11] (then the ratio is 1), very few cases are known where the ratio can be bounded by a constant. We can now show a constant upper bound when $G + H$ is planar. This is Theorem 2, which we restate here:

Theorem 2. *If $G + H$ is planar, the minimum cardinality of a multicut is at most 64 times the maximum value of an integer multiflow.*

Proof. The integrality gap of the multiflow LP equals the integrality gap of the D -cycle packing LP, as the two problems can be reduced to each other: the gap cannot be smaller because given an instance of the D -cycle packing LP, as said above, we can replace a demand edge $\{v, w\}$ by a demand edge $\{v, x\}$ of infinite (or very large) capacity and a supply edge $\{x, w\}$ of capacity 1 to form a multiflow instance.

Conversely, the gap also cannot be larger since we can reduce the multiflow problem to D -cycle packing by replacing every edge e by $u(e)$ parallel edges with unit capacity. Note that both reductions preserve planarity. The integrality gap of the multicut LP equals the integrality gap of the D -cycle covering LP by the same argument. Now Corollary 1 and Lemma 8 imply the assertion. \square

Determining the exact integrality gaps remains an open question. Without the planarity assumption, all the LPs except for the D -cut covering LP have unbounded integrality gap (for the D -cycle packing LP, a well-known example in [18] shows that the gap is in the order of $\Omega(\sqrt{n})$, even when G is planar, subcubic, and all demand pairs lie in the boundary of the outer face of G ; for the D -cut packing LP and the nonnegative cycle LP, consider $G + H = K_n$ and $D = \delta(v)$ for some vertex v ; for the D -cycle covering LP, let G be a bounded-degree expander graph with n vertices and D the set of $\frac{n^2}{4}$ vertex pairs with largest distance [17]).




Figure 7: The reduction from vertex cover in planar graphs. On the left-hand side picture, the vertices surrounded by a circuit form a minimum vertex cover X in G . On the right-hand side, the supply edges are shown in black, and the demand edges are colored (red or blue). The solution to NNC corresponding to X is the set of thicker edges. It is obtained by taking the red edges in each gadget associated to a vertex in X and the blue edges otherwise.

7 NP-completeness of NONNEGATIVECYCLES

In this section we prove that NONNEGATIVECYCLES is NP-hard. In fact, we consider the NONNEGATIVECYCLES decision problem, which takes as input an instance of NONNEGATIVECYCLES and an integer k , and asks whether there exists a solution of cardinality k .

Theorem 8. *The NONNEGATIVECYCLES decision problem is NP-complete even when $G + H$ is planar.*

To prove membership in NP, here is a polynomial-time algorithm that, given an instance $G = (V, E)$ and $H = (V, D)$ and a subset $D' \subseteq D$ of cardinality k , verifies whether D' is a solution: As in the proof of Lemma 1, assign weight 1 to edges of E , -1 to edges of D' , and 0 to edges of $D \setminus D'$; compute a minimum-weight \emptyset -join; then D' is a solution if and only if that minimum weight is 0.

To prove NP-completeness, recall the vertex cover problem: given a graph G and an integer k , it asks whether G has a vertex cover of size k . This problem is well-known to be NP-complete even in planar graphs [16]. We give a polynomial-time transformation from that problem.

Let $G = (V, E)$ be a planar graph. Let $G_2 := (V, E \dot{\cup} E)$ denote the (multi)graph where each edge of G is duplicated. We construct an instance $\tilde{G} + \tilde{H} = (\tilde{V}, \tilde{E} \dot{\cup} \tilde{D})$ of NONNEGATIVECYCLES as follows, illustrated in Figure 7.

For each vertex $v \in V$ of degree $d = d_G(v)$ in G , there is a *gadget* $(\tilde{V}_v, \tilde{E}_v \dot{\cup} \tilde{D}_v)$ built as follows. The vertex set \tilde{V}_v consists of $4d$ vertices $z_v, b_{v,1}, r_{v,1}, b_{v,2}, \dots, r_{v,2d-1}, b_{v,2d}$. The supply edge set \tilde{E}_v consists of edges $\{b_{v,i}, r_{v,i}\}$ and $\{r_{v,i}, b_{v,i+1}\}$, for $1 \leq i \leq 2d - 1$. The demand edge set \tilde{D}_v consists of $\tilde{B}_v \cup \tilde{R}_v$, where $\tilde{B}_v := \{\{z_v, b_{v,i}\} \mid 1 \leq i \leq 2d\}$ and $\tilde{R}_v := \{\{z_v, r_{v,i}\} \mid 1 \leq i \leq 2d - 1\}$.

To complete the construction of (\tilde{G}, \tilde{H}) , starting from a planar embedding of G_2 , replace each vertex v by the corresponding gadget, and each edge $e = \{u, v\}$ in G_2 by a supply edge $\{b_{u,i}, b_{v,j}\}$ in \tilde{E} , where i and j are chosen so that $\tilde{G} + \tilde{H}$ is planar.

The construction of $\tilde{G} + \tilde{H}$ can be done in polynomial time. Hence the following lemma completes the proof of Theorem 8.

Lemma 9. *Let G be a graph. There exists a vertex cover in G of size k if and only if there exists a subset $D' \subseteq \tilde{D}$ of size $4|E| - k$ that is a feasible solution to NONNEGATIVECYCLES in $\tilde{G} + \tilde{H}$.*

Proof. (\Rightarrow) Let X be a vertex cover of size k in G . Define $D' := (\bigcup_{v \in X} \tilde{R}_v) \cup (\bigcup_{v \notin X} \tilde{B}_v)$. It is easy to check that D' has size $4|E| - k$. We show now that D' is a feasible solution for NONNEGATIVECYCLES. Let \tilde{I} denote the set of $|D'|$ vertices naturally associated to D' , taking the endpoint of each edge of D' that is not z_v for any $v \in V$. By construction and since X is a vertex cover in G , \tilde{I} is an independent set in $\tilde{G} + \tilde{H}$. Hence $\{\delta(\{\tilde{v}\}) \mid \tilde{v} \in \tilde{I}\}$ is a D' -cut packing. For each cycle \tilde{C} in $\tilde{G} + \tilde{H}$, and for each $\tilde{v} \in \tilde{I}$, we have $|\tilde{C} \cap \delta(\{\tilde{v}\}) \cap D'| \leq |\tilde{C} \cap \delta(\{\tilde{v}\}) \cap \tilde{E}|$ and thus

$$|\tilde{C} \cap D'| = \sum_{\tilde{v} \in \tilde{I}} |\tilde{C} \cap \delta(\{\tilde{v}\}) \cap D'| \leq \sum_{\tilde{v} \in \tilde{I}} |\tilde{C} \cap \delta(\{\tilde{v}\}) \cap \tilde{E}| \leq |\tilde{C} \cap \tilde{E}|.$$

(\Leftarrow) Assume now that we are given a feasible solution D' for NONNEGATIVECYCLES in $\tilde{G} + \tilde{H}$ of size $4|E| - k$. We define $X := \{v \in V \mid D' \cap \tilde{D}_v \neq \tilde{B}_v\}$. We show that X is a vertex cover in G with size at least k . If there was an edge $\{u, v\} \in E$ with $u \notin X$ and $v \notin X$, then there would be a cycle, induced by vertices $\{z_u, b_{u,i}, b_{v,j}, z_v, b_{v,j+1}, b_{u,i+1}\}$ in $\tilde{G} + \tilde{H}$ for some indices i, j , that contains four edges in D' but only two supply edges. We now prove that $|X| \leq k$.

Since D' is a feasible solution, we have $|D' \cap \tilde{D}_v| \leq 2d_G(v)$ for all $v \in V$, for otherwise one could find a triangle in $\tilde{G}_v + \tilde{H}_v$ with two edges in D' . Moreover, $|D' \cap \tilde{D}_v| = 2d_G(v)$ only if $D' \cap \tilde{D}_v = \tilde{B}_v$, i.e., only if $v \notin X$. This implies

$$|X| \leq \sum_{v \in V} (2d_G(v) - |D' \cap \tilde{D}_v|) \leq 4|E| - |D'| = k.$$

□

8 Conclusion

We designed the first constant-factor approximation algorithm for EDGEDISJOINTPATHS if $G + H$ is planar.

If all demand edges lie on a single face of a planar embedding of G , then the planar dual H^* has only one nontrivial connected component (and isolated vertices). In this case, any instance satisfying the cut condition has a complete solution [24]. Then a 3-coloring of the outerplanar graph $G[L]$ in Algorithm 1 and a small modification of Algorithm 2 (reducing the upper bounds on y from 1 and 2 to $\frac{1}{2}$ and 1) yields a 4-approximation. If we could solve NONNEGATIVECYCLES in this case, we would even obtain an exact algorithm. This motivates the following open question: can NONNEGATIVECYCLES be solved in polynomial time if H has only one nontrivial connected component?

One attempt, also to reduce the constant factor that we lose when rounding a solution to the nonnegative cycle LP (5) in general, would be to find a characterisation of optimal LP




Figure 8: This instance shows that the nonnegative cycle LP (5) does not always have half-integral optimal solutions. Black (solid) edges are supply edges, red (dashed) edges represent demand edges together with their value in the unique optimum fractional solution.

solutions. However, this seems to be difficult. Figure 8 shows that there are instances in which the unique optimum LP solution is not half-integral.

References

- [1] M. ANDREWS, J. CHUZHUY, S. KHANNA, AND L. ZHANG, *Hardness of the undirected edge-disjoint paths problem with congestion*, in Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2015, pp. 226–244.
- [2] Y. AUMANN AND Y. RABANI, *Improved bounds for all optical routing*, in Proceedings of the 6th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 1995, pp. 567–576.
- [3] C. BENTZ, *Résolution exacte et approchée de problèmes de multiflot entier et de multicoûte : algorithmes et complexité*, Ph.D. thesis, Conservatoire National des Arts et Métiers, 2006.
- [4] C. CHEKURI, S. KHANNA, AND F. B. SHEPHERD, *An $O(\sqrt{n})$ approximation and integrality gap for disjoint paths and unsplittable flow*, Theory of Computing, 2 (2006), pp. 137–146.

- [5] J. CHERIYAN, H. KARLOFF, R. KHANDEKAR, AND J. KÖNEMANN, *On the integrality ratio for tree augmentation*, Operations Research Letters, 36 (2008), pp. 399–401.
- [6] J. CHUZHOUY, *Routing in undirected graphs with constant congestion*, in Proceedings of the 44th Annual ACM Symposium on Theory of Computing Conference (STOC), 2012, pp. 855–874.
- [7] J. CHUZHOUY AND D. H. K. KIM, *On approximating node-disjoint paths in grids*, in Proceedings of the Workshop on Approximation, Randomization, and Combinatorial Optimization (APPROX), N. Garg, K. Jansen, A. Rao, and J. D. P. Rolim, eds., vol. 40 of LIPIcs, 2015, pp. 187–211.
- [8] J. CHUZHOUY, D. H. K. KIM, AND S. LI, *Improved approximation for node-disjoint paths in planar graphs*, in Proceedings of the 48th Annual ACM Symposium on Theory of Computing Conference (STOC), D. Wicks and Y. Mansour, eds., 2016, pp. 556–569.
- [9] J. CHUZHOUY, D. H. K. KIM, AND R. NIMAVAT, *New hardness results for routing on disjoint paths*, in Proceedings of the 49th Annual ACM Symposium on Theory of Computing Conference (STOC), H. Hatami, P. McKenzie, and V. King, eds., 2017, pp. 86–99.
- [10] —, *Almost polynomial hardness of node-disjoint paths in grids*, in Proceedings of the 50th Annual ACM Symposium on Theory of Computing Conference (STOC), 2018, pp. 1220–1233.
- [11] D. CORNAZ, *Max-multiflow/min-multicut for $G+H$ series-parallel*, Discrete Mathematics, 311 (2011), pp. 1957–1967.
- [12] J. EDMONDS AND E. L. JOHNSON, *Matching, Euler tours and the Chinese postman*, Mathematical Programming, 5 (1973), pp. 88–124.
- [13] S. FORTUNE, J. HOPCROFT, AND J. WYLLIE, *The directed subgraph homeomorphism problem*, Theoretical Computer Science, 10 (1980), pp. 111–121.
- [14] A. FRANK, *Packing paths, circuits and cuts – a survey*, in Paths, Flows, and VLSI-Layout, B. Korte, L. Lovász, H. J. Prömel, and A. Schrijver, eds., Springer, 1990, pp. 47–100.
- [15] —, *A survey on T -joins, T -cuts, and conservative weightings*, in Combinatorics, Paul Erdős is Eighty, D. Miklós, V. T. Sós, and T. Szőnyi, eds., vol. 2, Springer, 1996, pp. 213–252.
- [16] M. R. GAREY, D. S. JOHNSON, AND L. STOCKMEYER, *Some simplified NP-complete graph problems*, Theoretical Computer Science, 1 (1976), pp. 237–267.
- [17] N. GARG, V. VAZIRANI, AND M. YANNAKAKIS, *Approximate max-flow min-(multi)cut theorems and their applications*, SIAM Journal on Computing, 25 (1996), pp. 235–251.
- [18] —, *Primal-dual approximation algorithms for integral flow and multicut in trees*, Algorithmica, 18 (1997), pp. 3–20.

- [19] M. GRÖTSCHEL, L. LOVÁSZ, AND A. SCHRIJVER, *The ellipsoid method and its consequences in combinatorial optimization*, *Combinatorica*, 1 (1981), pp. 169–197.
- [20] M. GUAN, *Graphic programming using odd and even points*, *Chinese Journal of Mathematics*, 1 (1962), pp. 273–277.
- [21] K. JAIN, *A factor 2 approximation algorithm for the generalized Steiner network problem*, *Combinatorica*, 21 (2001), pp. 39–60.
- [22] R. M. KARP, *On the computational complexity of combinatorial problems*, *Networks*, 5 (1975), pp. 45–68.
- [23] K. KAWARABAYASHI AND Y. KOBAYASHI, *An $O(\log n)$ -approximation algorithm for the edge-disjoint paths problem in Eulerian planar graphs*, *ACM Transactions on Algorithms*, 9 (2013), pp. 16:1–16:13.
- [24] E. KORACH AND M. PENN, *Tight integral duality gap in the Chinese postman problem*, *Mathematical Programming*, 55 (1992), pp. 183–191.
- [25] L. LOVÁSZ, *2-matchings and 2-covers of hypergraphs*, *Acta Mathematica Academiae Scientiarum Hungaricae*, 26 (1975), pp. 433–444.
- [26] M. MIDDENDORF AND F. PFEIFFER, *On the complexity of the disjoint paths problem*, *Combinatorica*, 13 (1993), pp. 97–107.
- [27] B. MOHAR AND C. THOMASSEN, *Graphs on Surfaces*, Johns Hopkins series in the mathematical sciences, Johns Hopkins University Press, 2001.
- [28] G. NAVES AND A. SEBŐ, *Multiflow feasibility: an annotated tableau*, in *Research Trends in Combinatorial Optimization*, W. Cook, L. Lovász, and J. Vygen, eds., Springer, 2009, pp. 261–283.
- [29] N. ROBERTSON, D. SANDERS, P. SEYMOUR, AND R. THOMAS, *The four-colour theorem*, *Journal of Combinatorial Theory, Series B*, 70 (1997), pp. 2–44.
- [30] N. ROBERTSON AND P. SEYMOUR, *Graph minors. XIII. The disjoint paths problem*, *Journal of Combinatorial Theory, Series B*, 63 (1995), pp. 65–110.
- [31] A. SCHRIJVER, *Combinatorial Optimization: Polyhedra and Efficiency*, Springer, 2003.
- [32] A. SEBŐ, *Potentials in undirected graphs and planar multiflows*, *SIAM Journal of Computing*, 26 (1997), pp. 582–603.
- [33] L. SEGUIN-CHARBONNEAU AND F. B. SHEPHERD, *Maximum edge-disjoint paths in planar graphs with congestion 2*, in *Proceedings of the 52nd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, R. Ostrovsky, ed., 2011, pp. 200–209.
- [34] P. D. SEYMOUR, *On odd cuts and plane multicommodity flows*, *Proceedings of the London Mathematical Society*, s3-42 (1981), pp. 178–192.

- [35] É. TARDOS AND V. V. VAZIRANI, *Improved bounds for the max-flow min-multicut ratio for planar and $K_{r,r}$ -free graphs*, Information Processing Letters, 47 (1993), pp. 77–80.
- [36] V. V. VAZIRANI, *Approximation Algorithms*, Springer, 2001.