



Approximating Maximum Integral Multiflows on Bounded Genus Graphs

Chien-Chung Huang¹ · Mathieu Mari²  · Claire Mathieu³ · Jens Vygen⁴

Received: 31 May 2021 / Revised: 29 August 2022 / Accepted: 16 September 2022 /

Published online: 9 October 2023

© The Author(s) 2023

Abstract

We devise the first constant-factor approximation algorithm for finding an integral multi-commodity flow of maximum total value for instances where the supply graph together with the demand edges can be embedded on an orientable surface of bounded genus. This extends recent results for planar instances. Our techniques include an uncrossing algorithm, which is significantly more difficult than in the planar case, a partition of the cycles in the support of an LP solution into free homotopy classes, and a new rounding procedure for freely homotopic non-separating cycles.

Keywords Optimization · Approximation algorithms · Multicommodity flow · Bounded-genus graphs

Mathematics Subject Classification 68Q25

Editor in Charge: János Pach

This work was partially funded by the Grant ANR-19-CE48-0016 from the French National Research Agency (ANR).

Chien-Chung Huang
chien-chung.huang@ens.fr

Mathieu Mari
Mathieu.Mari@ideas-ncbr.pl

Claire Mathieu
claire.mathieu@irif.fr

Jens Vygen
vygen@or.uni-bonn.de

¹ CNRS, DI ENS, Université PSL, Paris, France

² University of Warsaw, IDEAS-NCBR, Warsaw, Poland

³ CNRS, IRIF, Université de Paris, Paris, France

⁴ Research Institute for Discrete Mathematics, Hausdorff Center for Mathematics, University of Bonn, Bonn, Germany

1 Introduction

Multi-commodity flows, or *multiflows* for short, are well-studied objects in combinatorial optimization; see, e.g., [38, Part VII]. A multiflow of maximum total value can be found in polynomial time by linear programming. In many applications (see, e.g., [28]), a multiflow must be integral, and then the problem is much harder; the well-known edge-disjoint paths problem is a special case. There is a huge number of works addressing the question when the edge-disjoint paths problem or the maximum integral multiflow problem is solvable in polynomial time or at least admits a constant-factor approximation algorithm; surveys can be found in [11, 28, 33, 38].

Recently, constant-factor approximation algorithms have been found for maximum edge-disjoint paths and integral multiflows in *fully planar instances*, i.e., when $G + H$, the supply graph together with the demand edges, can be embedded in the plane [17, 22] (this setting was first studied by [40]). We generalize these results to surfaces of bounded genus and devise the first constant-factor approximation algorithm for that case. Another motivation of our work is to provide a novel application of recent works in geometric topology.

Beyond using some ideas of [17, 22], we need several new ingredients. Like [17], we start by computing an optimal (fractional) multiflow and “uncross” the cycles in its support as much as possible, but uncrossing is significantly more complicated on general surfaces than in the plane. Next, we need to deal with two cases separately: depending on whether most of the fractional multiflow is on separating cycles (that case is similar to the planar case) or on non-separating cycles. In the latter case we partition the cycles into free homotopy classes and define a cyclic order in each free homotopy class, which is possible due to the uncrossing and allows for a simple greedy algorithm.

1.1 Our Results

The (fractional) *maximum multiflow problem* can be described as follows. An instance consists of two undirected graphs $G = (V, E)$ (the *supply graph*) and $H = (V, D)$ (the *demand graph*) on the same vertex set, as well as a function $u: D \dot{\cup} E \rightarrow \mathbb{Z}_{>0}$. A *demand edge* $d \in D$ with endpoints s and t specifies the demand of a flow from s to t , ideally of value $u(d)$. The demand edges are sometimes referred to as *commodities*. All the flow must be routed in the supply graph. Each *supply edge* $e \in E$ has a capacity $u(e)$, and the total amount of flow routed along e (summed up over all commodities) must not exceed $u(e)$. The goal is to satisfy as much of the demand as possible. More precisely, we ask for an s - t -flow f^d of value at most $u(d)$ for every demand edge $d = \{t, s\}$ such that the total flow on each supply edge is at most its capacity, i.e., $\sum_{d \in D} f^d(e) \leq u(e)$ for all $e \in E$, and the total value of all those flows is maximum. The special case where there is only one demand edge is the classical maximum flow problem [16].

An elegant equivalent description works with the graph $G + H := (V, D \dot{\cup} E)$ that contains both the supply and the demand edges. It is well known that every s - t -flow can be decomposed into flows on s - t -paths and on cycles, and for integral flows there

is an integral decomposition. The cycles in such a decomposition do not contribute to the value of the s - t -flow and can be ignored. An s - t -path in (V, E) together with the demand edge $d = \{t, s\}$ forms a D -cycle: a cycle in $G + H$ that contains exactly one demand edge. If we let \mathcal{C} denote the set of all D -cycles in $G + H$, we can write the maximum multiflow problem equivalently as

$$\max \sum_{C \in \mathcal{C}} f_C \quad \text{such that} \quad \begin{cases} \sum_{C \in \mathcal{C}: C \ni e} f_C \leq u(e) & \text{for all } e \in D \dot{\cup} E, \\ f_C \geq 0 & \text{for all } C \in \mathcal{C}. \end{cases} \quad (1)$$

In some previous works, the problem has been defined with $u(d) = \infty$ for $d \in D$, and this variant is easily seen to be equivalent. We call the linear program (1) the *maximum multiflow LP*. The *maximum integral multiflow problem* is identical, except that the flow must be integral:

$$\max \sum_{C \in \mathcal{C}} f_C \quad \text{such that} \quad \begin{cases} \sum_{C \in \mathcal{C}: C \ni e} f_C \leq u(e) & \text{for all } e \in D \dot{\cup} E, \\ f_C \in \mathbb{Z}_{\geq 0} & \text{for all } C \in \mathcal{C}. \end{cases} \quad (2)$$

The special case where $u(e) = 1$ for every edge $e \in D \dot{\cup} E$ is known as the *maximum edge-disjoint paths* problem. Even that special case is unlikely to have a constant-factor approximation algorithm for general instances (see Sect. 1.2). Our main result is a constant-factor approximation algorithm in the case when $G + H$ can be embedded on an orientable surface of bounded genus.

Theorem 1.1 *There is a polynomial-time algorithm that takes as input an instance (G, H, u) of the maximum integral multiflow problem such that $G + H$ is embedded on an orientable surface of genus g , and which outputs an integral multiflow whose value is at most a factor $O(g^2 \log g)$ smaller than the value of any fractional multiflow.*

See Sect. 3 for an outline of the algorithm and the proof. It is worth pointing out that almost all known hardness results for the maximum edge-disjoint paths problem hold even when G is planar (see Sect. 1.2). Theorem 1.1, along with the two recent papers [17, 22], highlight that for tractability one needs more than the planarity of G alone. The topology of $G + H$ together plays an important role.

The dual LP of (1) is:

$$\min \sum_{e \in D \dot{\cup} E} u(e) y_e \quad \text{such that} \quad \begin{cases} \sum_{e \in C} y_e \geq 1 & \text{for all } C \in \mathcal{C}, \\ y_e \geq 0 & \text{for all } e \in D \dot{\cup} E, \end{cases} \quad (3)$$

and this may be called the *minimum fractional multicut problem*. The *minimum multicut problem* results from replacing the inequality $y_e \geq 0$ in (3) by $y_e \in \{0, 1\}$ for all edges $e \in D \dot{\cup} E$. So a *multicut* is a set of X of edges in $G + H$ such that every D -cycle contains at least one edge of X . The *capacity* of a multicut X is $\sum_{e \in X} u(e)$. Again, many previous works considered the equivalent special case where $u(d) = \infty$

for $d \in D$, in which case no dual variables for demand edges are needed and X must consist of supply edges only. By weak duality, the value of any multiflow is at most the capacity of any multicut. Using Theorem 1.1 and a previous result of [41], we obtain (in Sect. 9):

Corollary 1.2 *For any instance (G, H, u) of the maximum integral multiflow problem such that $G + H$ is embedded on an orientable surface of genus g , the minimum capacity of a multicut is at most $O(g^{3.5} \log g)$ times the maximum value of an integral multiflow.*

In general the integral multiflow-multicut gap, i.e., the ratio of the minimum capacity of a multicut and the maximum value of an integral multiflow,¹ and even the integrality gap of (1), i.e., the ratio of (1) and (2), can be as large as $\Theta(|D|)$. This is true even when G is planar and $G + H$ is embedded in the projective plane [19]; see Sect. 8. In this paper we consider orientable surfaces only. Corollary 1.2 states that the gap becomes constant when $G + H$ has bounded genus. So far very few such constant integral multiflow-multicut gaps are known, for example when G is a tree [19], or when $G + H$ is planar, as recently shown in [17, 22].

Finally, in Sect. 10, we obtain an improved approximation ratio (but not with respect to the LP value):

Theorem 1.3 *There is a polynomial-time algorithm that takes as input an instance (G, H, u) of the maximum integral multiflow problem such that $G + H$ is embedded on an orientable surface of genus g , and which outputs an integral multiflow whose value is at most a factor $O(g^2)$ smaller than the optimum.*

Whether a quadratic dependence on g is necessary remains open. However, we note in Sect. 8 that the integrality gap of the maximum multiflow LP can depend at least linearly on g .

1.2 Related Work

Approximation Algorithms and Hardness for Integral Multiflows Most of the hardness results for the maximum integral multiflow problem follow from the special case of the maximum edge-disjoint paths problem (EDP). The decision version of EDP is one of Karp's original NP-complete problems [23], and remains NP-complete even in many special cases [33], including the case of interest in this paper, namely even when $G + H$ is planar [31]. In terms of approximation, EDP is APX-hard [2]. Assuming that $\text{NP} \not\subseteq \text{DetTIME}(n^{O(\log n)})$, where $n = |V|$, there is no $n^{o(1/\sqrt{\log n})}$ approximation for EDP, even when G is planar and sub-cubic [8]. Assuming that for some positive δ , $\text{NP} \not\subseteq \text{RandTIME}(2^{n^\delta})$, there is no $n^{O(1/(\log \log n)^2)}$ approximation for EDP, even when G is planar and sub-cubic [7]. As far as we know, no stronger hardness result is known for integral multiflows.

¹ There is a closely related, but different, notion of integral flow-cut gap introduced in [6]: they study the smallest constant c such that whenever $u(C \cap E) \geq u(C \cap D)$ for every cut C (the cut condition), there is an integral multiflow satisfying all demands and violating capacities by at most a factor c .

On the positive side, EDP can be solved in polynomial time when the number of demand edges is bounded by a constant [36]. The same holds for integral multiflows when $G + H$ is planar [39]. For exact algorithms in various special cases, see the survey [33]. In general, the best known approximation guarantee for EDP and maximum integral multiflows is $O(\sqrt{n})$ [5]. Approximation algorithms with better approximation ratios for various special cases have been designed. We refer the readers to the survey [11] and to [19, 24, 33] and the references therein.

Recent Work on the Planar Case Recently, [17, 22] gave constant-factor approximation algorithms for maximum integer multiflows when $G + H$ is planar. Both papers proceed by first obtaining a half-integral multiflow and then using the four color theorem to round it to an integral solution (similar to Sect. 6). The main difference between the two works is the way such half-integral multiflows are obtained. In [17], it is constructed by uncrossing a fractional multiflow (see Sect. 5 for a definition) to construct a certain network matrix, which is known to be totally unimodular; in [22], such a half-integral multiflow is obtained by rounding a feasible solution of a related problem in the planar dual graph of $G + H$. Neither approach extends to higher genus graphs in a straightforward way, because the dual of a cycle is no longer a cut in general and cycles cannot always be uncrossed.

Minimum Multicut Problem The minimum multicut problem is NP-hard even when there are only three demand edges [12]. In general, assuming that the Unique Games conjecture holds, there is no $O(1)$ -approximation [4], but an $O(\log |D|)$ -approximation algorithm [18]. Better approximations also have been shown for special cases; see [19, 41] and the references therein. In particular, when $G + H$ is planar, [26] gave an approximation scheme. When G has genus g , an FPT-approximation scheme with parameters of g and $|D|$ has been proposed [9].

Tools from Topology The design of multiflows on surfaces is closely related to the properties of sets of curves on a surface. In a recent breakthrough, Przytycki [34] proved that the maximum number of essential curves on a closed surface of genus g such that no two of them are freely homotopic or intersect more than once is $O(g^3)$, improving on the previous exponential upper bound by [30]. Very recently, this number was shown to be $O(g^2 \log g)$ by [20], which almost matches the lower bound $\Omega(g^2)$ on the size of such sets [30]. We will use this result in Sect. 7.

2 Preliminaries

Consider an instance (G, H, u) of the maximum integral multiflow problem, and let $G + H = (V, E \dot{\cup} D)$ be the graph whose edge set is the disjoint union of the edge sets of the supply graph $G = (V, E)$ and the demand graph $H = (V, D)$. Throughout the paper, we assume that the graph $G + H$ is connected, otherwise, we can run the algorithm on each of its connected components.

Graphs on Surfaces Surfaces are either orientable or non-orientable; in this paper we only consider closed orientable surfaces. A closed orientable surface of genus g can be seen as a connected sum of g tori, or equivalently a sphere with g handles attached

on it, where g is called the *genus* of the surface. Given an integer $g \geq 0$, all closed surfaces with genus g are mutually homeomorphic, and we refer to any one of them as \mathbb{S}_g . For instance, \mathbb{S}_0 is the sphere and \mathbb{S}_1 is the torus.

A (multi)graph has *genus* g or is a *genus- g graph*, if it can be drawn on \mathbb{S}_g without edge crossings, but not on \mathbb{S}_{g-1} . A genus- g graph may have several non-equivalent embeddings on \mathbb{S}_g , but all of them satisfy the same invariant, called the *Euler characteristic*: $\# \text{Faces} - \# \text{Edges} + \# \text{Vertices} = 2 - 2g$. A simple application of Euler’s formula gives the following upper bound on the coloring number of genus- g graphs, when $g \geq 1$.

Theorem 2.1 (map color theorem) *A genus- g graph can be colored in polynomial time with at most*

$$\chi_g \leq \left\lfloor \frac{7 + \sqrt{1 + 48g}}{2} \right\rfloor$$

colors.

For $g = 0$, this is an algorithmic version of the 4-color theorem [35]. For $g \geq 1$, the coloring is obtained in polynomial time by a simple recursive algorithm that removes a vertex of minimum degree and colors the remaining graph [21]. For additional details and results about graphs on surfaces see e.g. [10, 32].

Combinatorial Embeddings Given a graph, let $\delta(v)$ denote the set of edges incident to a vertex v , and $\delta(U)$ the set of edges with exactly one endpoint in vertex set U . Given an embedding of a graph on an orientable surface, and an arbitrary orientation of this surface, for each vertex v , a clockwise cyclic order can be defined on the edges of $\delta(v)$. Note that contracting an edge $e = \{u, v\}$ results in removing e from $\delta(u)$ and from $\delta(v)$ and concatenating the orders to obtain the clockwise cyclic order of the edges around the vertex created by the contraction. Using these orders together with the incidence relation between edges and faces, embeddings become purely combinatorial objects. For additional details see, e.g., [32, Chap. 4].

Graph Duality Given an embedding of a genus- g graph G on \mathbb{S}_g , there exists a uniquely defined dual graph, denoted as G^* . This graph can be embedded on the same surface as G . There exists a bijection between the faces of G and the vertices of G^* , a bijection between the vertices of G and the faces of G^* , and a bijection between the edge sets of G and of G^* . Moreover, the embeddings of G and G^* are consistent: with this bijection, every edge only crosses its dual edge, and every face only contains its corresponding dual vertex and reciprocally. For notational simplicity, the latter bijection is implicit.

Cycles and Cuts A *path* in a graph G is a sequence $(v_0, e_1, v_1, \dots, e_k, v_k)$ for some $k \geq 0$, where v_0, \dots, v_k are distinct vertices and $e_i = \{v_{i-1}, v_i\}$ is an edge for all $i = 1, \dots, k$. A *cycle* in a graph G is a sequence $(v_0, e_1, v_1, \dots, e_k, v_k)$ such that v_1, \dots, v_k are distinct vertices, $\{v_{i-1}, v_i\}$ is an edge for all $i = 1, \dots, k$, and $v_0 = v_k$. Sometimes we view cycles as edge sets or as graphs. A *cut* is an edge set $\delta(U)$ for some proper subset $\emptyset \neq U \subset V$. A cut $\delta(U)$ is *simple* if both U and $V \setminus U$ are connected. We say that an edge set F in a graph is a (simple) *dual cut* if the corresponding set of

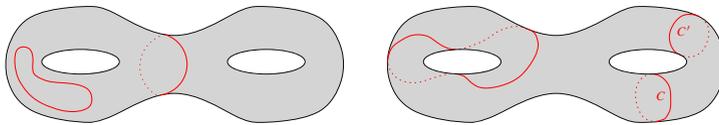


Fig. 1 Some cycles on an orientable surface of genus 2. On the left, two *separating* cycles. On the right, three *non-separating* cycles. C and C' are *freely homotopic* and their union disconnects the surface

edges F^* in the dual is a (simple) cut. A cycle C in G is called *separating* if it is a dual cut, and *non-separating* otherwise. Note that every separating cycle is a simple dual cut.

Homotopy Given a surface \mathbb{S} , a (simple) *topological cycle* is a continuous injective map γ from the unit cycle $S^1 := \{z \in \mathbb{C} : \|z\| = 1\}$ to \mathbb{S} . Two topological cycles γ_1 and γ_2 are *freely homotopic* if there exists a continuous function $\varphi: [0, 1] \times S^1 \rightarrow \mathbb{S}$ such that $\varphi(0, \cdot) = \gamma_1$ and $\varphi(1, \cdot) = \gamma_2$. Intuitively, cycle γ_1 is transformed into cycle γ_2 by continuously moving it on the surface. Free homotopy is an equivalence relation. Given an embedding of the graph $G + H$ on \mathbb{S} , we say that a cycle C in $G + H$ is *represented*² by a topological cycle γ of \mathbb{S} if the image of γ is the embedding of C on \mathbb{S} . Two cycles in $G + H$ are freely homotopic if and only if they can be represented by two freely homotopic topological cycles.

In the sequel, we use the following well-known fact.

Fact 1 *If two cycles C and C' are freely homotopic, then their symmetric difference is a dual cut. If C and C' are additionally disjoint and non-separating, then their union is a simple dual cut.*

Intuitively, the image of the continuous homotopy function from C to C' on the surface forms an annulus [13]. See Fig. 1 for an illustration.

3 Overview

In this section, we give an overview of our constant-factor approximation algorithm for the maximum integral multiflow problem when $G + H$ is embedded on an orientable surface \mathbb{S}_g of genus g , where g is bounded by a constant (Theorem 1.1). Again, without loss of generality, we assume that $G + H$ is connected. Here is the main algorithm. Steps 1, 2, 3, 4 will be described in detail in Sects. 4, 5, 6, 7, respectively.

1. Solve the linear program (1) to obtain a (fractional) multiflow f^* .
2. Construct another multiflow \bar{f} such that any two cycles in the support of \bar{f} cross at most once (Lemma 5.3). See Definition 5.1 for the definition of “crossing.”
3. If at least half of the total value of \bar{f} is contributed by separating cycles, these cycles now form a laminar family. Construct a half-integral multiflow f^{half} (Theorem 6.2), and from there, using the map color theorem (Theorem 2.1), compute an integral multiflow f' (Theorem 6.3), which is the output.

² Topological cycles are considered up to orientation-preserving reparameterization. Therefore, a cycle in $G + H$ may be represented by a topological cycle from two classes, one for each orientation: the class of γ and the class of γ' where $\gamma'(e^{i\theta}) = \gamma(e^{-i\theta})$.

4. Otherwise, partition the non-separating cycles in the support of \bar{f} into free homotopy classes. Pick the class \mathcal{H} with the largest total flow value. Remove the flow on all other cycles and greedily construct an integral multiflow (Lemmas 7.7 and 7.4), which is the output.

It can be proved that we only lose a constant factor at every step of the algorithm: see Sect. 8 for the analysis of the above algorithm, proving Theorem 1.1.

4 Finding a Fractional Multiflow (Step 1)

A feasible solution f to the maximum multiflow LP (1) will be simply called a *multiflow*. Recall that \mathcal{C} denotes the set of all D -cycles, i.e., all cycles in $G + H$ that contain precisely one demand edge. We denote by $|f| = \sum_{C \in \mathcal{C}} f_C$ the *value* of f , and by $\mathcal{C}(f) := \{C \in \mathcal{C} : f_C > 0\}$ the *support* of f . Although formulation (1) has an exponential number of variables, it is well known that it can be reformulated by polynomially many flow variables and constraints (see, e.g., [16, Chap. 1.2] or [1, Chap. 3.5]) and thereby solved in polynomial time:

Proposition 4.1 *There is an algorithm that finds an optimal solution f^* to the maximum multiflow LP (1) such that $|\mathcal{C}(f^*)| \leq |D| \cdot |E|$. Its running time is polynomial in the size of the input graph.*

Proof We reformulate the linear program as follows. Let D^{\rightarrow} contain an arbitrary orientation of each demand edge, and let E^{\leftrightarrow} contain both orientations of each supply edge. Introduce nonnegative flow variables x_e^d for all $d \in D$ and $e \in \{d^{\rightarrow}\} \dot{\cup} E^{\leftrightarrow}$, where d^{\rightarrow} denotes the chosen orientation of D . We maximize $\sum_{d \in D} x_{d^{\rightarrow}}^d$ subject to the constraints that x^d is a nonnegative circulation (i.e., obeying the flow conservation constraint at every vertex) for each $d \in D$, obeying the capacity constraints $x_{d^{\rightarrow}}^d \leq u(d)$ for all $d \in D$, and $\sum_{d \in D} (x_{e^{\leftarrow}}^d + x_{e^{\rightarrow}}^d) \leq u(e)$ for all $e \in E$, where e^{\leftarrow} and e^{\rightarrow} denote the two orientations of e . This is a linear program of polynomial size.

For a D -cycle $C \in \mathcal{C}$, let C^{\rightarrow} denote the orientation of C as a directed cycle such that $d^{\rightarrow} \in C^{\rightarrow}$. For every solution f to (1), $x_e^d := \sum_{C \in \mathcal{C}: d^{\rightarrow}, e \in C^{\rightarrow}} f_C$ defines a feasible solution of the new LP with the same objective function value. Conversely, given an optimum solution x to the new LP, each circulation x^d can be decomposed into flows on at most $|E|$ cycles, resulting in a solution f to (1) with $x_e^d = \sum_{C \in \mathcal{C}: d^{\rightarrow}, e \in C^{\rightarrow}} f_C$ for all d and e . Hence this f is an optimum solution, and it has support at most $|D| \cdot |E|$. □

Later we will restrict a multiflow to subsets of D -cycles. For $\mathcal{C}' \subseteq \mathcal{C}$ we define a multiflow f' by $f'_C := f_C$ for $C \in \mathcal{C}'$ and $f'_C := 0$ for $C \in \mathcal{C} \setminus \mathcal{C}'$, and write $f(\mathcal{C}') := f'$.

5 Making a Fractional Flow Minimally Crossing (Step 2)

In this section we show that for a given embedding, we can “uncross” a multiflow in such a way that any two D -cycles in the support cross at most once. While doing this we will lose only an arbitrarily small fraction of the multiflow value.

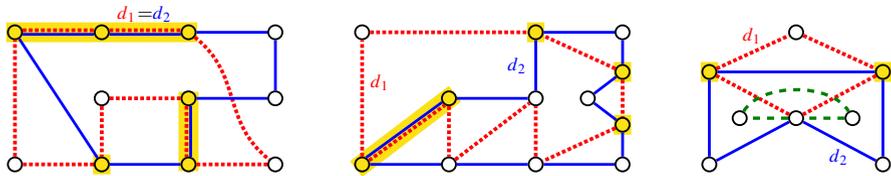


Fig. 2 Each of the two figures on the left show two D -cycles, C_1 (red, dotted) and C_2 (blue, solid). The edges belonging to D are marked as d_1 and d_2 . Edges are arranged at every vertex in the order of their embedding. Crossings are marked by yellow shade. The two D -cycles on the left cross three times. The two D -cycles in the middle cross four times. The figure on the right shows two D -cycles C_1 and C_2 that cross twice, and a third D -cycle C_3 (green, dashed) that crosses neither C_1 nor C_2 . Uncrossing C_1 and C_2 here generates a crossing of C_3 with a new D -cycle (namely with the triangle containing d_2)

Uncrossing is a well-known technique in combinatorial optimization, but in most cases it is applied to families of subsets of a ground set U . Such a family is said to be cross-free if, for any two of its sets, A and B , at least one of the four sets $A \setminus B$, $B \setminus A$, $A \cap B$, and $U \setminus (A \cup B)$ is empty. Here we want to uncross D -cycles in the topological sense, and this can be reduced to the above (with some extra care) only if all these cycles are separating (which, for example, is always the case if $G + H$ is planar; cf. [17]).

Definition 5.1 We say that two D -cycles C_1 and C_2 *cross* if there exists a path P (possibly a single vertex), which is a subpath of both C_1 and C_2 , and such that in the embedding, after contracting the edges of P , the vertex v thus obtained is incident to two edges of C_1 and to two edges of C_2 , all distinct, and in the embedding the restriction of the cyclic order of $\delta(v)$ to those four edges alternates between an edge of C_1 and an edge of C_2 .

Remark 5.2 A set of pairwise non-crossing D -cycles can be equivalently characterized in a topological manner as follows. Cycles C_1, \dots, C_r are pairwise non-crossing if and only if for all N_1, \dots, N_r such that N_i is a neighbourhood of the embedding of C_i on \mathbb{S} , $i = 1, \dots, r$, there are topological cycles C'_i in N_i , freely homotopic to C_i , for $i = 1, \dots, r$, that are pairwise disjoint. We use this equivalence in Sects. 6 and 7.

Two cycles may cross multiple times. We denote by $\text{cr}(C, C')$ the number of times that C and C' cross. See Fig. 2 for three examples. In contrast to the planar case, it is possible that two cycles cross exactly once and cannot be uncrossed. The third example in Fig. 2 shows another difficulty: when uncrossing two D -cycles it might be necessary to generate new crossings with other cycles.

Lemma 5.3 Let $\epsilon > 0$ be fixed. Given a multifold f whose support has size at most $|E| \cdot |D|$, there is a polynomial-time algorithm to construct another multifold \bar{f} , of value at least $|\bar{f}| \geq (1 - \epsilon)|f|$, and such that any two cycles in the support of \bar{f} cross at most once.

Proof First we discretize the multifold, losing an ϵ fraction in value; then we iteratively modify it, without changing its value, to reduce the number of crossings or the total

amount of flow on all edges; finally, we analyze the process and argue that the number of iterations is polynomially bounded.

Discretization The statement is trivial if $|f| = 0$. Otherwise, before uncrossing, we round down the flow on every D -cycle to integer multiples of $\epsilon|f|/(|E| \cdot |D|)$. That is, we define

$$f'_C := \frac{\epsilon|f|}{|E| \cdot |D|} \left\lfloor \frac{|E| \cdot |D| f_C}{\epsilon|f|} \right\rfloor$$

for all $C \in \mathcal{C}$. Note that f' is a multiflow. We claim that $|f'| \geq (1 - \epsilon)|f|$. Indeed,

$$|f'| = \sum_{C \in \mathcal{C}} f'_C \geq \sum_{C \in \mathcal{C}(f)} \left(f_C - \frac{\epsilon|f|}{|E| \cdot |D|} \right) = |f| - |\mathcal{C}(f)| \frac{\epsilon|f|}{|E| \cdot |D|} \geq |f| - \epsilon|f|.$$

The discretized multiflow f' can be represented by a multi-set \mathcal{S} of unweighted D -cycles: if $f'_C = k\epsilon|f|/(|E| \cdot |D|)$, then k identical copies of cycle C are added to \mathcal{S} . The number of cycles in \mathcal{S} (counting multiplicities) is at most $|E| \cdot |D|/\epsilon$ because

$$|\mathcal{S}| = \sum_{C \in \mathcal{C}} f'_C \frac{|E| \cdot |D|}{\epsilon|f|} \leq \sum_{C \in \mathcal{C}} f_C \frac{|E| \cdot |D|}{\epsilon|f|} = \frac{|E| \cdot |D|}{\epsilon}.$$

Uncrossing To construct \bar{f} , we perform a sequence of transformations of the multiflow. We will modify \mathcal{S} while maintaining the following invariants:

- (a) The number of elements of \mathcal{S} (counting multiplicities) remains constant.
- (b) For every $e \in D \dot{\cup} E$, the number of elements of \mathcal{S} (counting multiplicities) that contain e never increases.

Thanks to (b), at any stage, \bar{f} is a multiflow, where \bar{f} is defined by $\bar{f}_C = k\epsilon|f|/(|E| \cdot |D|)$ for $C \in \mathcal{C}$, where k is the multiplicity of C in \mathcal{S} . Initially $\bar{f} = f'$. Thanks to (a), the value of the multiflow is preserved. In the following we work only with \mathcal{S} .

While there exist two cycles C_1 and C_2 in \mathcal{S} that cross at least twice, do the following *uncrossing* operation (on one copy of C_1 and one copy of C_2). Let d_1 be the edge in $C_1 \cap D$, and let d_2 be the edge in $C_2 \cap D$. Let P and Q be two of the paths where C_1 and C_2 cross (cf. Definition 5.1), such that Q contains only edges of E . Fix an orientation of C_1 as follows. If P contains d_1 , then orient C_1 arbitrarily. Otherwise, orient C_1 so that in that orientation, d_1 is located between P and Q . Let \mathbf{C}_1 denote the resulting directed cycle. Let a be the first vertex on P in the orientation of \mathbf{C}_1 , and let b be an arbitrary vertex on Q . Vertices a and b partition \mathbf{C}_1 into a path \mathbf{C}_1^+ from a to b that contains d_1 and a path \mathbf{C}_1^- from b to a that does not contain d_1 .

Case 1 P contains an edge of D . Then this edge is $d_1 = d_2$. We orient C_2 so that the orientation on P agrees with the orientation of \mathbf{C}_1 on P . Let \mathbf{C}_2 denote the resulting directed cycle. Then the vertices a and b also partition \mathbf{C}_2 into a path \mathbf{C}_2^+ from a to b that contains d_2 and a path \mathbf{C}_2^- from b to a that does not contain d_2 .

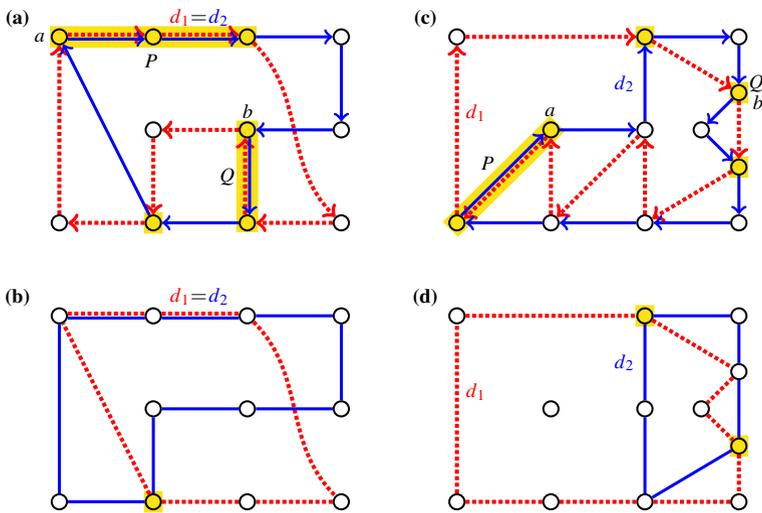


Fig. 3 Uncrossing the pairs of D -cycles from Fig. 2. (a) and (b) show an example for Case 1, (c) and (d) an example for Case 2. The initial situation (C_1 red, dotted, and C_2 blue, solid) and a possible choice of P, Q, a, b and the resulting orientation is shown in (a) and (c). As the result of the uncrossing operation, shown in (b) and (d), we have the new D -cycles C'_1 (red, dotted) and C'_2 (blue, solid) with fewer crossings among each other

Case 2 P contains edges of E only. Then we orient C_2 so that in that orientation, d_2 is located between P and Q . Let \mathbf{C}_2 denote the directed cycle. With that orientation, vertices a and b also partition \mathbf{C}_2 into a path C_2^+ from a to b that contains d_2 and a path C_2^- from b to a that does not contain d_2 .

To obtain C'_1 , we concatenate C_1^+ and C_2^- , remove any closed walk (allowing repeated vertices) that does not contain d_1 , and remove the orientation. To obtain C'_2 , we concatenate C_2^+ and C_1^- , remove any closed walk that does not contain d_2 , and remove the orientation. Note that C'_1 and C'_2 are D -cycles because each of C_1^+ and C_2^+ contains exactly one demand edge, and C_1^- and C_2^- contain no demand edge. See Fig. 3 for two examples, one for each case.

Analysis From the construction it follows that C'_1 and C'_2 are D -cycles and no edge is contained more often in C'_1 and C'_2 than in C_1 and C_2 . Hence removing one copy of C_1 and C_2 from \mathcal{S} and adding one copy of C'_1 and C'_2 to \mathcal{S} maintains the invariants (a) and (b).

To show that after a polynomial number of uncrossing operations any pair of cycles in \mathcal{S} crosses at most once, we consider the total number of edges $\Phi_1 = \sum_{C \in \mathcal{S}} |C|$ (counting multiplicities) and the total number of crossings $\Phi_2 = \sum_{C, C' \in \mathcal{S}} \text{cr}(C, C')$ (where we again count multiplicities). Note that $|\mathcal{S}|$ remains constant by invariant (a), and Φ_1 never increases by invariant (b). Moreover, $0 \leq \Phi_1 \leq |V| \cdot |\mathcal{S}|$ and $0 \leq |\Phi_2| \leq |V| \cdot |\mathcal{S}|^2$.

Claim 5.4 Each uncrossing operation either decreases Φ_1 or leaves Φ_1 unchanged and decreases Φ_2 .

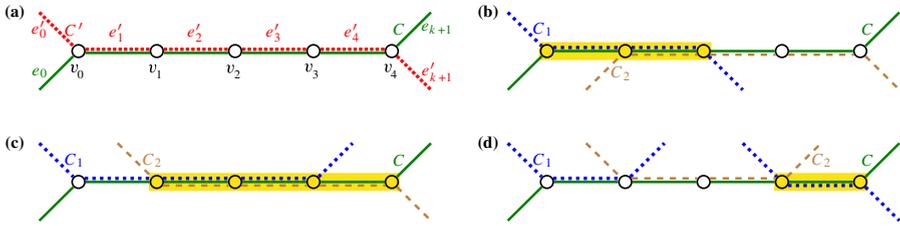


Fig. 4 For each crossing of C with a new cycle $C' \in \{C'_1, C'_2\}$ at a path R there is a crossing of C with one of the old cycles C_1 and C_2 at a subpath of R . This crossing is marked with yellow shade in the three examples

To prove Claim 5.4, consider an uncrossing operation that replaces C_1 and C_2 by C'_1 and C'_2 , and suppose that Φ_1 remains the same, so C'_1 consists of C_1^+ plus C_2^- , and C'_2 consists of C_2^+ plus C_1^- . We first observe that $\text{cr}(C'_1, C'_2) < \text{cr}(C_1, C_2)$. Indeed, the crossings at P and at Q go away, and no new crossing arises.

Finally we need to show that for any cycle $C \in \mathcal{C}$,

$$\text{cr}(C, C'_1) + \text{cr}(C, C'_2) \leq \text{cr}(C, C_1) + \text{cr}(C, C_2). \tag{4}$$

To show (4), consider a crossing of C and $C' \in \{C'_1, C'_2\}$ at a path R . Let $e'_1 = \{v_0, v_1\}, \dots, e'_k = \{v_{k-1}, v_k\}$ be the edges of R ($k \geq 0$), and let $e_0, e_{k+1}, e'_0, e'_{k+1}$ be edges such that $e_0, e'_1, \dots, e'_k, e_{k+1}$ are subsequent on C and $e'_0, e'_1, \dots, e'_k, e'_{k+1}$ are subsequent on C' . After contracting R , the incident edges $e_0, e'_0, e_{k+1}, e'_{k+1}$ are embedded in this cyclic order. (Note that $e_0 = e_{k+1}$ or $e'_0 = e'_{k+1}$ is possible if $k \geq 1$, then contracting R yields a loop.) See Fig. 4(a).

Now e'_0 belongs to C_1 or C_2 , say C_1 . If R contains neither a nor b , then e'_0, \dots, e'_{k+1} all belong to C_1 , and C_1 crosses C at R . If R contains either a or b , say at v_i , then e'_0, \dots, e'_i belong to C_1 and $e'_{i+1}, \dots, e'_{k+1}$ belong to C_2 . Moreover, C_1 and C_2 cross at a path containing v_i , so either C_1 crosses C at a subpath of R (Fig. 4(b)) or C_2 crosses C at a subpath of R (Fig. 4(c)). Finally, if R contains a and b , say at v_i and v_j for $0 \leq i < j \leq k$, then e'_0, \dots, e'_i and $e'_{j+1}, \dots, e'_{k+1}$ belong to C_1 and e'_{i+1}, \dots, e'_j belong to C_2 (Fig. 4(d)). Again, C_1 or C_2 crosses C at a subpath of R . This concludes the proof of Claim 5.4.

We can now conclude the proof of Lemma 5.3 because Φ_1 decreases at most $|V| \cdot |S|$ times, and while Φ_1 is constant, Φ_2 decreases at most $|V| \cdot |S|^2$ times, so the total number of uncrossing operations is at most $|V|^2 |S|^3 \leq |V|^2 |E|^3 |D|^3 / \epsilon^3$. \square

6 Separating Cycles: Routing an Integral Flow (Step 3)

Let \bar{f} be the flow resulted from Proposition 4.1 and Lemma 5.3, and let \mathcal{C}^{sep} denote the set of separating cycles in the support of \bar{f} . We now consider the case when the separating cycles contribute at least half to the total flow value, i.e., $|\bar{f}(\mathcal{C}^{\text{sep}})| \geq |\bar{f}|/2$. Note that $|\mathcal{C}^{\text{sep}}| \leq |D| \cdot |E|$. This branch of our algorithm consists of two steps:

1. given $\bar{f}(\mathcal{C}^{\text{sep}})$, construct a half-integral multiflow f^{half} of value at least $|\bar{f}(\mathcal{C}^{\text{sep}})|/2$;
2. given f^{half} , construct an integral multiflow of value at least $|f^{\text{half}}|/\Theta(\sqrt{g})$.

6.1 Obtaining a Half-Integral Multiflow

To obtain a half-integral multiflow, we follow the technique used by [17] for the case where $G + H$ is planar. By the Jordan curve theorem, any cycle in a planar graph is separating. As for the plane, the following property is easy to check for higher genus surfaces.

Proposition 6.1 *If C and C' are two cycles embedded on a surface, and C' is a separating cycle, then C and C' must cross an even number of times.*

Proof C' is separating the surface into two sides. While walking along C from a vertex v , we go from one side to the other each time we cross C' . When we return at v , we are on the same side where we started, so the number of crossing is even. \square

Since any pair of cycles in the support of \bar{f} crosses at most once, \mathcal{C}_{sep} must be a non-crossing family by Proposition 6.1. In particular, we can show that \mathcal{C}_{sep} have a laminar structure.

We say that a family of subsets of the dual vertex set V^* is laminar if any two members either are disjoint or one contains the other. Let us take any face of $G + H$ that we call ∞ . For any cycle $C \in \mathcal{C}_{\text{sep}}$ we define $\text{in}(C)$ and $\text{out}(C)$ to be the two connected components of $(G + H)^* \setminus C^*$, such that $\infty \in \text{out}(C)$, where C^* denotes the set of dual edges corresponding to edges in C . We claim that the family $\mathcal{L} := \{\text{in}(C) : C \in \mathcal{C}_{\text{sep}}\}$ is laminar.

Indeed, take any two cycles C and C' in \mathcal{C}_{sep} . Since they do not cross, either (i) $(C' \setminus C)^* \subseteq \text{in}(C)$ or, (ii) $(C' \setminus C)^* \subseteq \text{out}(C)$. In case (i) we must have $\text{in}(C') \subseteq \text{in}(C)$. In case (ii), we have either (ii.a) $\text{in}(C) \subseteq \text{in}(C')$ or (ii.b) $\text{in}(C) \cap \text{in}(C') = \emptyset$, hence laminarity.

Using the terminology in [17], we say a multiflow f is *laminar* if $\{C^* : C \in \mathcal{C}, f_C > 0\} = \{\delta(U) : U \in \mathcal{L}\}$ where \mathcal{L} is a laminar family (of subsets of V^*). Thus, $\bar{f}(\mathcal{C}_{\text{sep}})$ is laminar and we can apply the following result to get f^{half} .

Theorem 6.2 [17] *If f is a laminar multiflow, then there exists a laminar half-integral multiflow f' of value $|f'| \geq |f|/2$ such that $\mathcal{C}(f') \subseteq \mathcal{C}(f)$. Such a multiflow can be computed in polynomial time.*

6.2 Obtaining an Integral Multiflow

In this section we show the following result, which is an extension of a result from [17, 22], who proved it for planar graphs.

Theorem 6.3 *Let (G, H, u) be an instance of the maximum multiflow problem such that $G + H$ has genus g , and let f^{half} be a laminar half-integral multiflow whose support $\mathcal{C}(f^{\text{half}})$ contains only separating cycles, and $|\mathcal{C}(f^{\text{half}})| \leq |D| \cdot |E|$. Then there exists an integral multiflow f' of value $|f'| \geq 2|f^{\text{half}}|/\chi_g$ (such that $\mathcal{C}(f') \subseteq \mathcal{C}(f^{\text{half}})$). Such a multiflow can be found in polynomial time.*

In the rest of this section, we prove this theorem. We will first reduce to the case where all cycles in $\mathcal{C}(f^{\text{half}})$ have flow value $1/2$ and every edge has capacity 1. It turns

out that constructing a large integral multiflow for this new instance is equivalent to finding a large independent set in the intersection graph of cycles in $\mathcal{C}^{\text{half}}$. The rest of our proof follows the same outline as the proof of Theorem 1 of Fiorini et al. [15]: with Lemmas 6.5 and 6.6, we show that this intersection graph has genus at most g , and can therefore be colored using χ_g colors. Finally, the integral multiflow output corresponds to the largest color class.

Proposition 6.4 *If Theorem 6.3 holds for the special case when $u(e) = 1$ for all $e \in D \dot{\cup} E$ and $f_C^{\text{half}} \in \{0, 1/2\}$ for all $C \in \mathcal{C}$, then it holds in general.*

Proof Let (G, H, u) be a general instance of the maximum multiflow problem such that $G + H$ has genus g , and let f^{half} be a laminar half-integral multiflow whose support $\mathcal{C}^{\text{half}} := \mathcal{C}(f^{\text{half}})$ contains only separating cycles, and with $|\mathcal{C}^{\text{half}}| \leq |D| \cdot |E|$. Define the multiflows f^{int} and \widehat{f} such that for each $C \in \mathcal{C}^{\text{half}}$ we have $f_C^{\text{int}} = \lfloor f_C^{\text{half}} \rfloor$ and $\widehat{f}_C = f_C^{\text{half}} - f_C^{\text{int}} \in \{0, 1/2\}$. The multiflow f^{int} is integral, and we have $f^{\text{half}} = f^{\text{int}} + \widehat{f}$.

Define a new instance $\widehat{G} + \widehat{H}$ obtained from $G + H$ by replacing each edge $e \in D \dot{\cup} E$ by $\min \{|\mathcal{C}^{\text{half}}|, u(e) - \sum_{C \in \mathcal{C}: e \in C} f^{\text{int}}(C)\}$ parallel edges in $\widehat{G} + \widehat{H}$, each of unit capacity, embedded in a sufficiently small neighbourhood of the embedding of e . The graph $\widehat{G} + \widehat{H}$ has the same genus as $G + H$, and its size is polynomial in the size of $G + H$ because $|\mathcal{C}^{\text{half}}| \leq |D| \cdot |E|$. We can now associate to \widehat{f} a feasible half-integral multiflow in $\widehat{G} + \widehat{H}$ (that we call also \widehat{f}) of the same value, by replacing each edge of each cycle $C \in \mathcal{C}(\widehat{f})$ by one of its parallel edges. To see that this can be done while ensuring that the resulting multiflow is still feasible and laminar, recall Remark 5.2: we can associate to each cycle $C \in \mathcal{C}(\widehat{f})$ a topological cycle C' in an arbitrarily small neighbourhood such that these cycles C' for $C \in \mathcal{C}(\widehat{f})$ are pairwise disjoint. Now, for any edge $e \in D \dot{\cup} E$, let e_1, \dots, e_r denote the corresponding parallel edges, ordered according to their embedding in the surface; and let C_1, \dots, C_t denote the set of cycles in $\mathcal{C}(\widehat{f})$ that contain e , ordered according to the embedding of the associated pairwise disjoint cycles C'_1, \dots, C'_t . Then, for each $i \in \{1, \dots, t\}$ we replace the edge e in C_i by its parallel edge $e_{\lceil i/2 \rceil}$. Since $\widehat{f}(C) = 1/2$ for all $C \in \mathcal{C}(\widehat{f})$, we have $t \leq 2r$, and every edge in $\widehat{G} + \widehat{H}$ is used at most twice, which guarantees feasibility. Finally, laminarity follows from the fact that the cycles C' for $C \in \mathcal{C}(\widehat{f})$ are pairwise disjoint.

Now we are in the special case: suppose that there exists a polynomial-time algorithm that computes an integral multiflow \widehat{f}' in $\widehat{G} + \widehat{H}$ of value $|\widehat{f}'| \geq 2|\widehat{f}|/\chi_g$, and such that $\mathcal{C}(\widehat{f}') \subseteq \mathcal{C}(\widehat{f})$. Then, \widehat{f}' induces the integral multiflow $f' := f^{\text{int}} + \widehat{f}'$ in $G + H$, which satisfies the conditions of Theorem 6.3. Indeed, $f' := f^{\text{int}} + \widehat{f}'$ is integral and feasible (for the initial capacities u), its support is contained in $\mathcal{C}^{\text{half}}$, and its value is

$$|f'| = |f^{\text{int}}| + |\widehat{f}'| \geq |f^{\text{int}}| + \frac{2|\widehat{f}|}{\chi_g} \geq \frac{2(|f^{\text{int}}| + |\widehat{f}|)}{\chi_g} = \frac{2|f^{\text{half}}|}{\chi_g},$$

as required. □

Therefore, to prove Theorem 6.3 it is enough to consider the special case of unit capacities and $f_C^{\text{half}} \in \{0, 1/2\}$ for all $C \in \mathcal{C}$.

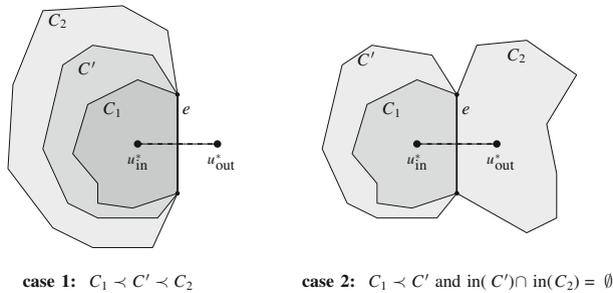


Fig. 5 Proof of Lemma 6.5

Recall that cycles in $\mathcal{C}^{\text{half}} \subseteq \mathcal{C}^{\text{sep}}$ are separating and do not cross each other, so the family $\{\text{in}(C) : C \in \mathcal{C}^{\text{half}}\}$ is laminar. We partially order $\mathcal{C}^{\text{half}}$ with the following relation: $C \prec C'$ if $\text{in}(C) \subset \text{in}(C')$. We have the following simple property:

Lemma 6.5 *Let $\mathcal{C}^{\text{half}}$ be a set of separating cycles such that $\{\text{in}(C) : C \in \mathcal{C}^{\text{half}}\}$ is laminar and no edge is contained in more than three of these cycles. If $C_1, C_2, C' \in \mathcal{C}^{\text{half}}$ are such that $C_1 \prec C'$ and $C_2 \not\prec C'$, then C_1 and C_2 are edge-disjoint.*

Proof Assume, for a contradiction, that C_1 and C_2 share an edge e . Let $e^* = \{u_{\text{in}}^*, u_{\text{out}}^*\}$ denote its dual edge, such that $u_{\text{in}}^* \in \text{in}(C_1)$ and $u_{\text{out}}^* \in \text{out}(C_1)$. Since $C_2 \not\prec C'$, by laminarity either $C' \prec C_2$ or $\text{in}(C') \cap \text{in}(C_2) = \emptyset$. In the first case we have $C_1 \prec C' \prec C_2$ and then

$$u_{\text{in}}^* \in \text{in}(C_1) \subseteq \text{in}(C') \subseteq \text{in}(C_2) \quad \text{and} \quad u_{\text{out}}^* \in \text{out}(C_2) \subseteq \text{out}(C'),$$

so $e \in C'$. In the second case we have $C_1 \prec C'$ and $\text{in}(C') \cap \text{in}(C_2) = \emptyset$ and then

$$u_{\text{in}}^* \in \text{in}(C_1) \subseteq \text{in}(C') \subseteq \text{out}(C_2) \quad \text{and} \quad u_{\text{out}}^* \in \text{in}(C_2) \subseteq \text{out}(C'),$$

so $e \in C'$. See Fig. 5. Thus in both cases e belongs to C' as well as to C_1 and C_2 . Since these three D -cycles are in the support of a half-integral multiflow, this implies that the flow along this edge is at least $3/2$, contradicting feasibility. \square

Our goal is to get a large subset $\mathcal{C}' \subseteq \mathcal{C}^{\text{half}}$ such that any two cycles in \mathcal{C}' are edge-disjoint. This is equivalent to finding a large independent set in a properly defined graph $\text{Int}(\mathcal{C}^{\text{half}})$ with vertex set $\mathcal{C}^{\text{half}}$ and such that two cycles are adjacent if they share at least one edge. Using Lemma 6.5 we can show:

Lemma 6.6 *Let G be a graph embedded in \mathbb{S}_g , and let $\mathcal{C}^{\text{half}}$ be a set of separating cycles in G such that $\{\text{in}(C) : C \in \mathcal{C}^{\text{half}}\}$ is laminar and no edge is contained in more than three of these cycles. Let $\text{Int}(\mathcal{C}^{\text{half}})$ be the graph with vertex set $\mathcal{C}^{\text{half}}$ and such that two cycles are adjacent if they share at least one edge. Then $\text{Int}(\mathcal{C}^{\text{half}})$ is a genus- g graph.*

Proof We prove the statement by induction on $g + |\mathcal{C}^{\text{half}}|$. When $g + |\mathcal{C}^{\text{half}}| \leq 2$, it is trivial. Otherwise let G be a connected genus- g graph, embedded on \mathbb{S}_g , and $\mathcal{C}^{\text{half}}$ a family as described above.

Suppose first that $\{\text{in}(C) : C \in \mathcal{C}^{\text{half}}\}$ are pairwise disjoint. Then, contract in G^* each set $\text{in}(C)$ into a single node. Two cycles C and C' share an edge if and only if in this contracted graph, the nodes corresponding to $\text{in}(C)$ and $\text{in}(C')$ are adjacent. This means that $\text{Int}(\mathcal{C}^{\text{half}})$ is a minor of G^* , and in particular has genus less than or equal to the genus of G^* . The case where there is one cycle \bar{C} such that $C < \bar{C}$ for all $C \in \mathcal{C}^{\text{half}} \setminus \bar{C}$ and $\{\text{in}(C) : C \in \mathcal{C}^{\text{half}} \setminus \bar{C}\}$ are pairwise disjoint works similarly; here we contract out \bar{C} .

Otherwise there exists a triple $C_1, C_2, C \in \mathcal{C}^{\text{half}}$ such that $C_1 < C$ and $C_2 \not< C$. The separating cycle C divides \mathbb{S}_g into two sides. Each side can be closed—by identifying the boundary of a disk with the boundary formed by C —so that they are homeomorphic to $\mathbb{S}_{g_{\text{in}}}$ and $\mathbb{S}_{g_{\text{out}}}$, respectively. The connected sum of these two surfaces is homeomorphic to \mathbb{S}_g , and in particular we have $g_{\text{in}} + g_{\text{out}} = g$. This equality can easily be checked with Euler’s formula.

Let G_{in} (resp. G_{out}) be the subgraph of G induced by the vertices embedded on the side corresponding to $\mathbb{S}_{g_{\text{in}}}$ (resp. $\mathbb{S}_{g_{\text{out}}}$), such that both contain C . The embedding of G in \mathbb{S}_g induces an embedding of G_{in} in $\mathbb{S}_{g_{\text{in}}}$ and an embedding of G_{out} in $\mathbb{S}_{g_{\text{out}}}$. Thus, $\text{genus}(G_{\text{in}}) + \text{genus}(G_{\text{out}}) \leq g$.

Now we define $\mathcal{C}_{\leq C}^{\text{half}} := \{C' \in \mathcal{C}^{\text{half}} : C' < C\} \cup \{C\}$ and $\mathcal{C}_{\not< C}^{\text{half}} := \{C' \in \mathcal{C}^{\text{half}} : C' \not< C\} \cup \{C\}$. The choice of C implies that these two families are proper subsets of $\mathcal{C}^{\text{half}}$. Since the cycles in $\mathcal{C}^{\text{half}}$ do not cross, we have $\{C \in \mathcal{C}^{\text{half}} : C \subseteq G_{\text{in}}\} = \mathcal{C}_{\leq C}^{\text{half}}$ and $\{C \in \mathcal{C}^{\text{half}} : C \subseteq G_{\text{out}}\} = \mathcal{C}_{\not< C}^{\text{half}}$.

By the induction hypothesis, $\text{Int}(\mathcal{C}_{\leq C}^{\text{half}})$ and $\text{Int}(\mathcal{C}_{\not< C}^{\text{half}})$ can be embedded on $\mathbb{S}_{g_{\text{in}}}$ and $\mathbb{S}_{g_{\text{out}}}$, respectively. By Lemma 6.5, the graph $\text{Int}(\mathcal{C}^{\text{half}})$ arises from $\text{Int}(\mathcal{C}_{\leq C}^{\text{half}})$ and $\text{Int}(\mathcal{C}_{\not< C}^{\text{half}})$ by identifying the two vertices that correspond to C .

Finally we prove that $\text{Int}(\mathcal{C}^{\text{half}})$ can be embedded on a surface genus $g_{\text{in}} + g_{\text{out}} \leq g$. To see that, remove small disks D_{in} and D_{out} in $\mathbb{S}_{g_{\text{in}}}$ and $\mathbb{S}_{g_{\text{out}}}$, respectively, around the point that corresponds to vertex C and that intersects only edges incident to C , and glue them together by identifying boundaries of D_{in} and D_{out} . The surface obtained is homeomorphic to $\mathbb{S}_{g_{\text{in}}+g_{\text{out}}}$. It is easy to see that C , and the edges incident to C , can be re-embedded in this surface without intersecting any other edges. This terminates the proof of Lemma 6.6. □

Now the proof of Theorem 6.3 (in the special case that is sufficient by Proposition 6.4) is easy. We apply Lemma 6.6 to $G + H$ and the support $\mathcal{C}^{\text{half}}$ of a laminar half-integral multiflow. Using Theorem 2.1, one can compute in polynomial time a subset $C' \subseteq \mathcal{C}^{\text{half}}$ of at least $|\mathcal{C}^{\text{half}}|/\chi_g$ pairwise edge-disjoint D -cycles. From this set, we define an integral multiflow by setting $f'_C = 1$ for $C \in C'$ and $f'_C = 0$ for $C \in \mathcal{C} \setminus C'$. It is easy to check that f' is a multiflow that satisfies the properties of Theorem 6.3. This concludes the proof of Theorem 6.3.

7 Non-Separating Cycles: Routing an Integral Multiflow (Step 4)

If the separating cycles contribute less than half to the total value of the multiflow \bar{f} obtained by Lemma 5.3, we consider the non-separating cycles in the support of \bar{f} . We first partition them into *free homotopy classes*. The following theorem by Greene [20] gives an upper bound on the number of such classes.

Theorem 7.1 [20] *Let \mathbb{S}_g be an orientable surface of genus g . Then there are at most $O(g^2 \log g)$ topological cycles such that any two of them are in different free homotopy classes and cross each other at most once.*

Corollary 7.2 *The D -cycles in the support of \bar{f} can be partitioned into $O(g^2 \log g)$ free homotopy classes in polynomial time.*

Proof Take pairs of cycles in the support of \bar{f} and check whether they are freely homotopic, for example as in [14, 29]. \square

7.1 Greedy Algorithm

Let \mathcal{H} be a free homotopy class of non-separating cycles whose total flow value $|\bar{f}(\mathcal{H})|$ is largest. We will run the following simple greedy algorithm (Algorithm 1) on \mathcal{H} to get an integral multiflow.

Algorithm 1: Greedy algorithm for integral multiflows

Input: a sequence C_1, \dots, C_k of D -cycles of $\mathcal{C}(\bar{f})$.

Output: an integral multiflow f .

$f \leftarrow$ the all-zero multiflow;

for $i = 1$ **to** k **do**

\lfloor Set f_{C_i} to be the greatest integer such that f remains feasible.

The value of the integral multiflow returned by this algorithm depends on the order of the D -cycles in the input. If they are ordered according to the following definition, then we show that we lose only a constant fraction of the flow value.

Definition 7.3 A family of cycles $\{C_1, C_2, \dots, C_k\}$ is *cyclically ordered*, or has a *cyclic order* if, whenever two cycles C_a and C_b share an edge, where $a < b$, then this edge is:

- shared by all cycles $C_a, C_{a+1}, \dots, C_{b-1}, C_b$, or
- shared by all cycles $C_b, C_{b+1}, \dots, C_k, C_1, \dots, C_{a-1}, C_a$.

The following lemma establishes the approximation ratio of Algorithm 1 on cyclically ordered input.

Lemma 7.4 *Let \bar{f} be a multiflow and $\mathcal{H} = \{C_1, C_2, \dots, C_k\}$ a cyclically ordered family of $\mathcal{C}(\bar{f})$. Then Algorithm 1 returns in polynomial time an integral multiflow of value at least $\lfloor \bar{f}(\{C_1, \dots, C_k\}) \rfloor / 2$.*

Proof Let \bar{f} be a multiflow and $\mathcal{H} = \{C_1, C_2, \dots, C_k\}$ a cyclically ordered family of $\mathcal{C}(\bar{f})$. It is clear that Algorithm 1 runs in $O(k|E|)$ time and returns an integral multiflow. Let f be this flow. We show that its value is at least $|\bar{f}(\mathcal{H})|/2$.

Let us define $\mathcal{H}_{a,b} = \{C_a, C_{a+1}, \dots, C_{b-1}\}$ and $\mathcal{H}_{b,a} = \{C_b, C_{b+1}, \dots, C_k, C_1, \dots, C_{a-1}\}$ for all $1 \leq a \leq b \leq k$. Additionally, for all edges $e \in \bigcup_{C \in \mathcal{H}} C$, we define $\mathcal{H}^e := \{C \in \mathcal{H} : e \in C\}$. Since we assumed that \mathcal{H} is cyclically ordered, we know that for each $e \in \bigcup_{C \in \mathcal{H}} C$, there are indexes $1 \leq a, b \leq k$, such that $\mathcal{H}^e = \mathcal{H}_{a,b}$. We call i_0 the smallest index $1 \leq i \leq k$ such that there exists an edge $e \in C_i$ such that $f(\mathcal{H}_{1,i+1})(e) = u(e)$ and $\mathcal{H}_{i,1} \subseteq \mathcal{H}^e$, where $f(\mathcal{H}_{1,i+1})(e)$ denotes the total amount of flow that $f(\mathcal{H}_{1,i+1})$ ships along edge e . Remark that in particular, for all $i > i_0$, we must have $f_{C_i} = 0$, and thus $|f| = |f(\mathcal{H}_{1,i_0+1})|$.

We first show by induction that for all $1 \leq i < i_0$ we have $|f(\mathcal{H}_{1,i+1})| \geq |\bar{f}(\mathcal{H}_{1,i+1})|$. For $i = 1$, we have $|f(\mathcal{H}_{1,2})| = |f_{C_1}| = \min\{u(e) : e \in C_1\} \geq \bar{f}_{C_1} = |\bar{f}(\mathcal{H}_{1,2})|$.

Assume now that at some iteration $1 < i < i_0$ of the algorithm we set $f_{C_i} = x$. By the choice of x , we know that there is an edge $e \in C_i$ such that $u(e) = f(\mathcal{H}_{1,i+1})(e)$. In particular, notice that $|f(\mathcal{H}^e)| = |f(\mathcal{H}^e \cap \mathcal{H}_{1,i+1})| = u(e)$. By feasibility of \bar{f} , we have

$$|f(\mathcal{H}^e \cap \mathcal{H}_{1,i+1})| = u(e) \geq |\bar{f}(\mathcal{H}^e)|. \tag{5}$$

Now, let a, b be the two indices such that $\mathcal{H}_{a,b} = \mathcal{H}^e$. Since we assumed that $i < i_0$, we must have $i < b \leq k$. There are two cases: either $1 \leq a \leq i < b$ or $1 < i < b < a$. If $1 \leq a \leq i < b$, then (5) becomes $|f(\mathcal{H}_{a,i+1})| \geq |\bar{f}(\mathcal{H}^e)| \geq |\bar{f}(\mathcal{H}_{a,i+1})|$. Together with the induction hypothesis we obtain

$$|f(\mathcal{H}_{1,i+1})| = |f(\mathcal{H}_{1,a})| + |f(\mathcal{H}_{a,i+1})| \geq |\bar{f}(\mathcal{H}_{1,a})| + |\bar{f}(\mathcal{H}_{a,i+1})| = |\bar{f}(\mathcal{H}_{1,i+1})|.$$

Otherwise if $1 < i < b < a$, then $\mathcal{H}_{1,i+1} \subseteq \mathcal{H}^e$, and thus the inequality claimed follows directly from (5). We have established the induction. In particular, we have proved that $|f| = |f(\mathcal{H}_{1,i_0+1})| \geq |f(\mathcal{H}_{1,i_0})| \geq |\bar{f}(\mathcal{H}_{1,i_0})|$. To conclude the proof of Lemma 7.4, it remains to show that $|f| \geq |\bar{f}(\mathcal{H}_{i_0,1})|$.

By definition of i_0 , we know that there exists an edge $e \in C_{i_0}$ such that $f(e) = u(e)$ and such that $\mathcal{H}_{i_0,1} \subseteq \mathcal{H}^e$. By feasibility of \bar{f} , we deduce that $|\bar{f}(\mathcal{H}_{i_0,1})| \leq u(e) = f(e) \leq |f|$. This concludes the proof. \square

Remark 7.5 The analysis of Algorithm 1 for cyclically ordered inputs is tight. To see this, imagine that $\mathcal{H} = \{C_1, \dots, C_{2k-1}\}$, and there are two edges e_1, e_2 , both of capacity k , such that $\{C \in \mathcal{H} : e_1 \in C\} = \{C_1, \dots, C_k\}$ and $\{C \in \mathcal{H} : e_2 \in C\} = \{C_{k+1}, \dots, C_{2k-1}, C_1\}$. Then Algorithm 1 may only set $f_{C_1} = k$ while \bar{f} could be such that $\bar{f}_C = 1$ for all $C \in \mathcal{H}$, for a total value $2k - 1$.

7.2 Computing a Cyclic Order

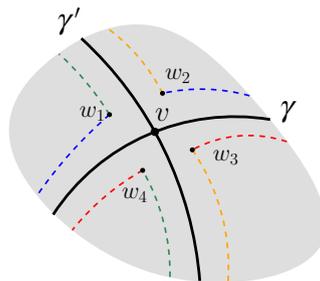
Lemma 7.7, the second main result of the section, states that a family \mathcal{H} of pairwise freely homotopic cycles crossing at most once can be cyclically ordered in polynomial

time. One ingredient in the proof is that cycles in \mathcal{H} are pairwise non-crossing. This fact uses the assumption that the surface is orientable. In a non-orientable surface, two freely homotopic cycles may cross exactly once. One may easily prove that freely homotopic cycles cross an even number of times using classic tools from algebraic topology. Here we give an elementary and self-contained proof. Recall that \bar{f} denotes the minimally-crossing multiflow obtained by Lemma 5.3.

Lemma 7.6 *Two freely homotopic cycles in $\mathcal{C}(\bar{f})$ do not cross.*

Proof By construction of \bar{f} , if two cycles C and C' in $\mathcal{C}(\bar{f})$ cross, then they cross at exactly one path P . To simplify, let us take two topological cycles γ and γ' , freely homotopic to C and C' , that are in a small neighborhood around C and C' , respectively, and such that γ and γ' only cross at a single point v of the surface. We show that $\gamma \cup \gamma'$ do not disconnect the orientable surface. By Fact 1 this implies that C and C' are not freely homotopic.

To see that $\gamma \cup \gamma'$ do not disconnect the surface, pick four points w_1, w_2, w_3, w_4 in a small neighborhood of v , each one of them being on a different of the four sections of this neighborhood delimited by $\gamma \cup \gamma'$. If $(w_i)_{1 \leq i \leq 4}$ are in clockwise order around v , then w_i and w_{i+1} are still connected for $i = 1, \dots, 4$ (where $w_5 := w_1$), because we can walk all along γ (or γ'). Notice that here we use the property that the surface is orientable (otherwise, w_i might be connected to w_{i+2} instead of w_{i+1}). By transitivity, we conclude that $\gamma \cup \gamma'$ do not disconnect the surface.



Now we can show³ (see Fig. 6 for an illustration):

Lemma 7.7 *A family of non-separating, pairwise non-crossing and freely homotopic cycles of a graph embedded in an orientable surface can be cyclically ordered. Such a cyclic order can be found in polynomial time.*

Proof Let \mathcal{H} be a set of non-separating, pairwise freely homotopic and non-crossing cycles. We first order the cycles in \mathcal{H} and then prove that this is a cyclic order. We assume that $|\mathcal{H}| \geq 3$, otherwise any order on \mathcal{H} is a cyclic order. By Remark 5.2, we can associate to each cycle $C \in \mathcal{H}$ a freely homotopic topological cycle C' in an arbitrarily small neighbourhood of C , such that the set \mathcal{H}' of resulting topological

³ This result holds more generally for a family of *non-contractible* (cycles that are not freely homotopic to a point on the surface), pairwise non-crossing and freely homotopic cycles. For simplicity, we only consider the special case of non-separating cycles, which is sufficient for our main result.

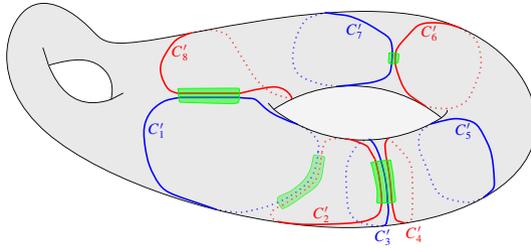


Fig. 6 A family of non-separating, pairwise non-crossing and freely homotopic cycles can be cyclically ordered. If some cycles share common edges or common vertices (green regions), we can slightly move these cycles along the surface in order to obtain a family of pairwise disjoint cycles. Then, each of the connected components obtained after cutting the surface along these new cycles are bordered by exactly two cycles (Claim 7.8). These connected components naturally induce a cyclic ordering of the original cycles

cycles are pairwise disjoint. By transitivity, the cycles in \mathcal{H}' are also pairwise freely homotopic. To order the cycles in \mathcal{H}' (and thus also those in \mathcal{H}), we consider the connected components obtained after cutting the surface along cycles in \mathcal{H}' .

Claim 7.8 *The set $\mathbb{S}_g \setminus \bigcup_{C' \in \mathcal{H}'} C'$ has exactly $|\mathcal{H}'|$ connected components, and each of these connected components is bordered by exactly two cycles in \mathcal{H}' .*

We prove this claim by induction on the size of \mathcal{H}' . The base case $|\mathcal{H}'| = 2$ follows directly from Fact 1. Now consider a set \mathcal{H}' of pairwise disjoint and freely homotopic non-separating cycles on \mathbb{S}_g , and let C'_0 be any cycle in \mathcal{H}' . By induction hypothesis, we assume that the set $\mathbb{S}_g \setminus \bigcup_{C' \in \mathcal{H}', C' \neq C'_0} C'$ has exactly $|\mathcal{H}'| - 1$ connected components, and each of these connected components is bordered by exactly two cycles in $\mathcal{H}' \setminus \{C'_0\}$. Since cycles in \mathcal{H}' are pairwise disjoint, C'_0 is strictly contained in one of the connected components of $\mathbb{S}_g \setminus \bigcup_{C' \in \mathcal{H}', C' \neq C'_0} C'$. Let K be this component and C'_1 and C'_2 the cycles in $\mathcal{H}' \setminus \{C'_0\}$ that are bordering K . We claim that $K \setminus C'_0$ has exactly two connected components, the first one bordered by C'_0 and C'_1 and the second one bordered by C'_0 and C'_2 . Indeed, assume for a contradiction that $K \setminus C'_0$ is connected. Let K' be the other connected component of $\mathbb{S}_g \setminus (C'_1 \cup C'_2)$. Then K' is also bordered by C'_1 and C'_2 . We have $\mathbb{S}_g \setminus (C'_0 \cup C'_1) = \mathbb{S}_g \setminus (C'_0 \cup C'_1 \cup C'_2) \cup C'_2 = (K \setminus C'_0) \cup K' \cup C'_2$, and thus $\mathbb{S}_g \setminus (C'_0 \cup C'_1)$ is connected, which contradicts Fact 1. Therefore, C'_0 separates K in two connected components. None of them is bordered only by C'_0 because C'_0 is non-separating. Thus, one of these two connected components is bordered by C'_0 and C'_1 , and the other by C'_0 and C'_2 . We have proved the claim.

We now define a (cyclic) order $\mathcal{H} = \{C_1, \dots, C_k\}$ as follows: C_i and C_j are consecutive in this order, i.e., $j = i + 1 \pmod k$ or $j = i - 1 \pmod k$ whenever their corresponding cycles C'_i and C'_j in \mathcal{H}' are bordering the same connected component of $\mathbb{S}_g \setminus \bigcup_{C' \in \mathcal{H}'} C'$. To prove that this order satisfies the property of Definition 7.3, consider any edge e in $G + H$. Consider the set of cycles in \mathcal{H} that contain e , and let $\mathcal{H}'(e) \subseteq \mathcal{H}'$ denote the corresponding set of pairwise disjoint topological cycles. Then there exists a neighborhood $N \subset \mathbb{S}_g$ around e that intersects among \mathcal{H}' only the cycles $\mathcal{H}'(e)$, so these are indeed consecutive in the cyclic order.

To compute this cyclic order in time polynomial in the size of $G + H$, one can use the dual $(G + H)^*$ of $G + H$. Given any family \mathcal{H} of cycles in $G + H$, the connected components of $\mathbb{S}_g \setminus \bigcup_{C \in \mathcal{H}} C$ correspond to the connected components of $(G + H)^* \setminus \bigcup_{C \in \mathcal{H}} C^*$, where C^* denotes the dual edges of edges in C . The connected components of a graph can be computed in polynomial time. This concludes the proof of Lemma 7.7. \square

8 Proof of Theorem 1.1

The algorithm works as described in Sect. 3. By construction, the output of the algorithm is a feasible solution. We now analyze the value of the output. Since (1) is a relaxation of the maximum integral multiflow problem, $|f^*| \geq \text{OPT}$. By Lemma 5.3, $|\bar{f}| \geq (1 - \epsilon)|f^*|$. For $\epsilon = 1/2$ we have $|\bar{f}| \geq |f^*|/2$. Moreover, $|\mathcal{C}(\bar{f})| \leq |\mathcal{C}(f^*)| \leq |D| \cdot |E|$ by Proposition 4.1.

Consider the multiflow restricted to separating cycles, \bar{f}_{sep} . If $|\bar{f}_{\text{sep}}| \geq |\bar{f}|/2$, then by Theorems 6.2 and 6.3 we obtain an integral flow of value at least $|\bar{f}_{\text{sep}}|/\Theta(\sqrt{g}) = |f^*|/\Theta(\sqrt{g})$. Otherwise, by Theorem 7.1 there exists a free homotopy class \mathcal{H} of non-separating cycles such that $|\bar{f}(\mathcal{H})| \geq |\bar{f}|/\Theta(g^2 \log g)$. Construct a cyclic order (Lemma 7.7) and apply Algorithm 1. By Lemma 7.4 the output has value at least $|f^*|/\Theta(g^2 \log g)$.

Finally, we analyze the running time. As observed in Sect. 4, an optimum fractional multiflow f^* can be found in polynomial time (Proposition 4.1). (Discretizing and) uncrossing is done in time polynomial in $|E| \cdot |D|$ time by Lemma 5.3. Partitioning into free homotopy classes is done by Corollary 7.2. Finally, the operations of Theorems 6.2, 2.1, and 6.3, and Lemmas 7.4 and 7.7 can all be done in polynomial time, hence polynomial running time overall. This concludes the proof of Theorem 1.1.

Lower Bound on the Integrality Gap We claim that the gap between an integral and a fractional multiflow can depend at least linearly on g . For any $n \geq 1$, we define a graph G_n as in [3]. This graph consists of n concentric cycles (rings) and $4n$ radial line segments that intersect each cycle, and each has endpoint s_i or t_i on the outer face, for $1 \leq i \leq 2n$. This circular ordering of these endpoints is $s_1, \dots, s_{2n}, t_1, \dots, t_{2n}$. See Fig. 7. We now define the set of demand edges $H_n = \{\{s_i, t_i\} : 1 \leq i \leq 2n\}$. The graph $G_n + H_n$ can be embedded in the projective plane but cannot be embedded in an orientable surface of genus smaller than n ; see [3] for a proof.

Now, to obtain a large integrality gap, we define a new graph G'_n by splitting each degree-4 vertex of G_n into two vertices, joined by a new edge, such that two of the four incident edges are incident to each of the two new vertices (similarly as in an example of [19]). All edges have unit capacity. We have the following properties:

- (i) $G'_n + H_n$ has orientable genus at least n . This holds since $G_n + H_n$ is a minor of $G'_n + H_n$.
- (ii) In an integral solution we can satisfy only one demand. Indeed, any s_i - t_i -path must share a vertex with any s_j - t_j -path, for $j \neq i$. This holds more generally for any planar graph: if four vertices s, s', t, t' are on the outer face in this order, then any path from s to t disconnects s' from t' . Moreover, since all vertices of

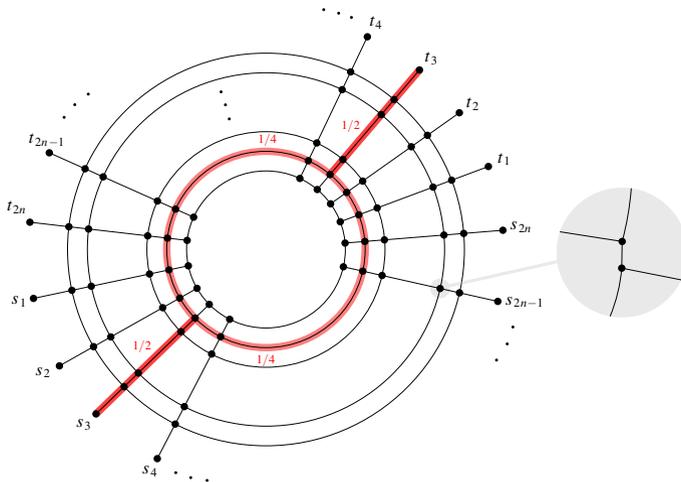


Fig. 7 The graph G_n . To obtain the graph G'_n , we split each degree-4 vertex of G_n into two vertices, joined by a new edge, such that two of the four incident edges are incident to each of the two new vertices. To illustrate the fractional multiflow described in (iii) (at the end of Sect. 8), we show the portion serving the commodity $\{s_3, t_3\}$ in red

G'_n have degree at most three, any two paths that share a vertex must also share an edge.

- (iii) A (fractional) multiflow f of value n exists. For each $i \in \{1, \dots, 2n\}$, there is a flow of value $1/2$ from s_i to t_i . The flow follows the radial line segment corresponding to s_i , until it meets the $\lceil i/2 \rceil$ -th ring, where it is divided in two parts, each with value $1/4$ following a different direction of the ring. Then, the flow merges again when it meets the radial line segment corresponding to t_i , and follows this radial line until t_i . See Fig. 7. Each ring is used by two source-sink pairs, and thus the total multiflow is feasible.

9 Proof of Corollary 1.2

In this section, we observe how Corollary 1.2 follows from Theorem 1.1 and the following result by Tardos and Vazirani [41] (based on work by Klein et al. [27]).

Theorem 9.1 [41] *Let (G, H, u) be a multiflow instance and $\gamma > 1$ such that the supply graph G does not have a $K_{\gamma, \gamma}$ minor. Then the minimum capacity of a multicut is $O(\gamma^3)$ times the maximum value of a (fractional) multiflow.*

The following is well known.

Claim 9.2 *If a graph G has genus at most g , where $g \geq 1$, then it has no $K_{\gamma, \gamma}$ minor for any $\gamma > 2(\sqrt{g} + 1)$.*

Proof Suppose that such a minor $K_{\gamma, \gamma}$ exists in G . As the three operations for obtaining a minor (deleting edges/vertices and contracting edges) do not increase the genus,

$K_{\gamma,\gamma}$ has genus at most g . Furthermore, $K_{\gamma,\gamma}$ has 2γ vertices, γ^2 edges, and at most $\gamma^2/2$ faces (since there is no odd cycle in a bipartite graph). By Euler's formula, $2 - 2g \leq 2\gamma - \gamma^2 + \gamma^2/2$, which implies $\gamma \leq 2(\sqrt{g} + 1)$. \square

By Claim 9.2 and Theorem 9.1, the ratio between the minimum capacity of a multicut and the maximum value of a (fractional) multiflow is $O(g^{1.5})$. This, combined with Theorem 1.1, proves Corollary 1.2.

10 An Improved Approximation Ratio (Proof of Theorem 1.3)

Theorem 1.1 yields an approximation ratio of $O(g^2 \log g)$ for the maximum integer multiflow problem for instances where $G + H$ is embedded on an orientable surface of genus g . Here we show how to improve this ratio to $O(g^2)$, proving Theorem 1.3.

Namely, after applying Corollary 7.2, consider the $O(g^2 \log g)$ free homotopy classes of the non-separating cycles in the support of our uncrossed multiflow, and take a representative cycle in each class. Let I be the graph whose vertices are these free homotopy classes and whose edges correspond to pairs of classes whose representative cycles cross. This definition does not depend on the choice of the representative cycles. Now a theorem of Przytycki [34] says that this graph has maximum degree $O(g^2)$.

Theorem 10.1 [34] *There is a universal constant $\beta > 0$ such that the following is true. Let $g \geq 1$ and let Γ be a family of simple curves on \mathbb{S}_g such that any two of them are not freely homotopic and cross at most once. Then, the maximum degree of the intersection graph of Γ is at most βg^2 .*

Hence we can color the vertices of this graph I greedily with $O(g^2)$ colors so that the color classes are stable sets, i.e., sets of cycles that do not cross. Hence there is a color class \mathcal{K} whose cycles support an $\Omega(1/g^2)$ fraction of the total flow value. Next, we throw away all cycles outside \mathcal{K} and apply the greedy algorithm of Sect. 7.1 to each free homotopy class of this color class \mathcal{K} separately, but before, in each free homotopy class of \mathcal{K} , we reduce the capacity of every edge in the two *extreme cycles* to its total flow value in this class, rounded down.

Lemma 10.2 *Each free homotopy class \mathcal{H}_i in \mathcal{K} has two extreme cycles C_i^+ and C_i^- such that any cycle C of another homotopy class in \mathcal{K} that shares an edge with a cycle in \mathcal{H}_i also shares an edge with C_i^+ or C_i^- . The set of extreme cycles can be computed in polynomial time.*

Intuitively, for each class, the extreme cycles correspond to the pair of cycles that delimits the maximal annulus among all pairs in this class. Notice that when a class consists of a single cycle C , we have $C_i^+ = C_i^- = C$.

Proof We can assume that $g \geq 2$, otherwise \mathcal{K} has at most one free homotopy class and the statement is trivially true. Additionally, if \mathcal{H}_i contains exactly one cycle, the statement is also trivially true. Then, let \mathcal{H} be a free homotopy class of size at least two. Cutting along cycles in \mathcal{H}_i might separate the surface into several components that are all homeomorphic to annuli or disks except one component K that has genus at least one. Its boundary is contained in the union of two cycles, which we call C_i^+ and C_i^- .

All other cycles in $\bigcup \mathcal{K} \setminus \mathcal{H}_i$ are contained in K . Thus, if a cycle in $\bigcup \mathcal{K} \setminus \mathcal{H}_i$ shares an edge e with a cycle in \mathcal{H}_i , this edge must be on K 's boundary, and in particular $e \in C_i^+ \cup C_i^-$. \square

Thus, for each homotopy class \mathcal{H}_i in \mathcal{K} and each edge e that is contained in an extreme cycle of \mathcal{H}_i , we reduce its capacity to $\lfloor f(\mathcal{H}_i)(e) \rfloor$. This is sufficient to make the multiflow problems of the free homotopy classes independent of each other because any edge that lies on two cycles from two distinct classes must also lie on one of the extreme cycles of the corresponding classes. The rounding down loses an additive constant of at most $2|\mathcal{K}|$ (at most two per free homotopy class); by Corollary 7.2, this is $O(g^2 \log g)$. Losing this additive constant can be afforded since this loses only a constant factor unless the optimum value is $\text{OPT} = O(g^2 \log g)$.

To cover this case, we can guess the value of an optimum integral flow $f^{\text{OPT}}(d)$ through each demand edge $d \in D$. For each guess, we create an instance of the edge-disjoint paths problem by replacing each demand edge $d \in D$ by $f^{\text{OPT}}(d)$ parallel demand edges (of unit capacity), and each supply edge $e \in E$ by $\min\{u(e), \text{OPT}\}$ parallel supply edges (of unit capacity). Since $\text{OPT} = O(g^2 \log g)$, this new graph has polynomial size. Since the number of demand edges in the edge-disjoint paths instance is bounded by a constant, we can apply the polynomial-time algorithm by Robertson and Seymour [36] (whose running time $O(n^3)$ was later improved to $O(n^2)$ in [25], with n referring to the number of vertices in the graph) to decide whether this instance is feasible or not. Since we need to enumerate only $|D|^{O(g^2 \log g)}$ guesses, we can compute an optimal solution f^{OPT} to the original maximum integral multiflow instance in polynomial time, assuming that $\text{OPT} = O(g^2 \log g)$. This concludes the proof of Theorem 1.3.

However, due to the last step of this algorithm, this does not imply a stronger bound on the integrality gap shown in Theorem 1.1 or the integral multiflow-multicut gap shown in Corollary 1.2.

Note Added in Proof

Our techniques have recently been used to solve more general cycle packing problems in bounded-genus graphs [37].

Acknowledgements The authors would like to thank Arnaud de Mesmay and an anonymous reviewer for useful suggestions.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Ahuja, R.K., Magnanti, T.L., Orlin, J.B.: *Network Flows*. Prentice Hall, Englewood Cliffs (1993)
2. Andrews, M., Chuzhoy, J., Guruswami, V., Khanna, S., Talwar, K., Zhang, L.: Inapproximability of edge-disjoint paths and low congestion routing on undirected graphs. *Combinatorica* **30**(5), 485–520 (2010)
3. Auslander, L., Brown, T.A., Youngs, J.W.T.: The imbedding of graphs in manifolds. *J. Math. Mech.* **12**(4), 629–634 (1963)
4. Chawla, S., Krauthgamer, R., Kumar, R., Rabani, Y., Sivakumar, D.: On the hardness of approximating multicut and sparsest-cut. *Comput. Complex.* **15**(2), 94–114 (2006)
5. Chekuri, C., Khanna, S., Shepherd, F.B.: An $O(\sqrt{n})$ approximation and integrality gap for disjoint paths and unsplittable flow. *Theory Comput.* **2**, 137–146 (2006)
6. Chekuri, C., Shepherd, F.B., Weibel, C.: Flow-cut gaps for integer and fractional multiflows. *J. Comb. Theory Ser. B* **103**(2), 248–273 (2013)
7. Chuzhoy, J., Kim, D.H.K., Nimavat, R.: Almost polynomial hardness of node-disjoint paths in grids. *Theory Comput.* **17**(6), # 6 (2021)
8. Chuzhoy, J., Kim, D.H.K., Nimavat, R.: New hardness results for routing on disjoint paths. *SIAM J. Comput.* **51**(2), STOC17-189–STOC17-237 (2022)
9. Cohen-Addad, V., Colin de Verdière, É., de Mesmay, A.: A near-linear approximation scheme for multicut of embedded graphs with a fixed number of terminals. *SIAM J. Comput.* **50**(1), 1–31 (2021)
10. Colin de Verdière, É.: Computational topology of graphs on surfaces. In: *Handbook of Discrete and Computational Geometry*, pp. 605–636 (chapter 23). CRC Press, Boca Raton (2018)
11. Costa, M.-C., Létocart, L., Roupin, F.: Minimal multicut and maximal integer multiflow: a survey. *Eur. J. Oper. Res.* **162**(1), 55–69 (2005)
12. Dahlhaus, E., Johnson, D.S., Papadimitriou, C.H., Seymour, P.D., Yannakakis, M.: The complexity of multiterminal cuts. *SIAM J. Comput.* **23**(4), 864–894 (1994)
13. Epstein, D.B.A.: Curves on 2-manifolds and isotopies. *Acta Math.* **115**, 83–107 (1966)
14. Erickson, J., Whittlesey, K.: Transforming curves on surfaces redux. In: *24th Annual ACM-SIAM Symposium on Discrete Algorithms (New Orleans 2013)*, pp. 1646–1655. SIAM, Philadelphia (2013)
15. Fiorini, S., Hardy, N., Reed, B., Vetta, A.: Approximate min-max relations for odd cycles in planar graphs. *Math. Program.* **110**(1), 71–91 (2007)
16. Ford, L.R., Jr., Fulkerson, D.R.: *Flows in Networks*. Princeton University Press, Princeton (1962)
17. Garg, N., Kumar, N., Sebő, A.: Integer plane multiflow maximisation: flow-cut gap and one-quarter-approximation. In: *21st International Conference on Integer Programming and Combinatorial Optimization. Lecture Notes in Computer Science*, vol. 12125, pp. 144–157. Springer, Cham (2020)
18. Garg, N., Vazirani, V.V., Yannakakis, M.: Approximate max-flow min-(multi)cut theorems and their applications. *SIAM J. Comput.* **25**(2), 235–251 (1996)
19. Garg, N., Vazirani, V.V., Yannakakis, M.: Primal-dual approximation algorithms for integral flow and multicut in trees. *Algorithmica* **18**(1), 3–20 (1997)
20. Greene, J.E.: On loops intersecting at most once. *Geom. Funct. Anal.* **29**(6), 1828–1843 (2019)
21. Heawood, P.J.: Map colour theorem. *Q. J. Math.* **24**, 332–338 (1890)
22. Huang, C.-C., Mari, M., Mathieu, C., Schewior, K., Vygen, J.: An approximation algorithm for fully planar edge-disjoint paths. *SIAM J. Discret. Math.* **35**(2), 752–769 (2021)
23. Karp, R.M.: On the computational complexity of combinatorial problems. *Networks* **5**(1), 45–68 (1975)
24. Kawarabayashi, K., Kobayashi, Y.: An $O(\log n)$ -approximation algorithm for the edge-disjoint paths problem in Eulerian planar graphs. *ACM Trans. Algorithms* **9**(2), # 16 (2013)
25. Kawarabayashi, K., Kobayashi, Y., Reed, B.: The disjoint paths problem in quadratic time. *J. Comb. Theory Ser. B* **102**(2), 424–435 (2012)
26. Klein, P.N., Mathieu, C., Zhou, H.: Correlation clustering and two-edge-connected augmentation for planar graphs. In: *32nd International Symposium on Theoretical Aspects of Computer Science (Garching 2015)*. Leibniz Int. Proc. Inform., vol. 30, pp. 554–567. Leibniz-Zent. Inform., Wadern (2015)
27. Klein, P., Plotkin, S.A., Rao, S.: Excluded minors, network decomposition, and multicommodity flow. In: *25th Annual ACM Symposium on Theory of Computing (San Diego 1993)*, pp. 682–690. ACM, New York (1993)
28. Korte, B., Lovász, L., Prömel, H.J., Schrijver, A. (eds.): *Paths, Flows, and VLSI-Layout. Algorithms and Combinatorics*, vol. 9. Springer, Berlin (1990)

29. Lazarus, F., Rivaud, J.: On the homotopy test on surfaces. In: 53rd Annual Symposium on Foundations of Computer Science (New Brunswick 2012), pp. 440–449. IEEE Computer Soc., Los Alamitos (2012)
30. Malestein, J., Rivin, I., Theran, L.: Topological designs. *Geom. Dedicata* **168**, 221–233 (2014)
31. Middendorf, M., Pfeiffer, F.: On the complexity of the disjoint paths problem. *Combinatorica* **13**(1), 97–107 (1993)
32. Mohar, B., Thomassen, C.: *Graphs on Surfaces*. Johns Hopkins Series in the Mathematical Sciences. Johns Hopkins University Press, Baltimore (2001)
33. Naves, G., Sebő, A.: Multiflow feasibility: an annotated tableau. In: *Research Trends in Combinatorial Optimization*, pp. 261–283. Springer, Berlin (2009)
34. Przytycki, P.: Arcs intersecting at most once. *Geom. Funct. Anal.* **25**(2), 658–670 (2015)
35. Robertson, N., Sanders, D., Seymour, P., Thomas, R.: The four-colour theorem. *J. Comb. Theory Ser. B* **70**(1), 2–44 (1997)
36. Robertson, N., Seymour, P.D.: Graph minors. XIII. The disjoint paths problem. *J. Comb. Theory Ser. B* **63**(1), 65–110 (1995)
37. Schlomberg, N., Thiele, H., Vygen, J.: Packing cycles in planar and bounded-genus graphs. In: 34th Annual ACM-SIAM Symposium on Discrete Algorithms (Florence 2023), pp. 2069–2086. SIAM, Philadelphia (2023)
38. Schrijver, A.: *Combinatorial Optimization: Polyhedra and Efficiency*. Algorithms and Combinatorics, vol. 24. Springer, Berlin (2003)
39. Sebő, A.: Integer plane multiflows with a fixed number of demands. *J. Comb. Theory Ser. B* **59**(2), 163–171 (1993)
40. Seymour, P.D.: On odd cuts and plane multicommodity flows. *Proc. Lond. Math. Soc.* **42**(1), 178–192 (1981)
41. Tardos, É., Vazirani, V.V.: Improved bounds for the max-flow min-multicut ratio for planar and $K_{r,r}$ -free graphs. *Inf. Process. Lett.* **47**(2), 77–80 (1993)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.