

Improved Lower Bounds on Multiflow-Multicut Gaps

Sina Kalantarzadeh 

University of Waterloo, Canada

Nikhil Kumar 

University of Waterloo, Canada

Abstract

Given a set of source-sink pairs, the maximum *multiflow* problem asks for the maximum total amount of flow that can be feasibly routed between them. The minimum *multicut*, a dual problem to multiflow, seeks the minimum-cost set of edges whose removal disconnects all the source-sink pairs. It is easy to see that the value of the minimum multicut is at least that of the maximum multiflow, and their ratio is called the *multiflow-multicut gap*. The classical max-flow min-cut theorem states that when there is only one source-sink pair, the gap is exactly one. However, in general, it is well known that this gap can be arbitrarily large. In this paper, we study this gap for classes of planar graphs and establish improved lower bound results. In particular, we show that this gap is at least $\frac{20}{9}$ for the class of planar graphs, improving upon the decades-old lower bound of 2. More importantly, we develop new techniques for proving such a lower bound, which may be useful in other settings as well.

2012 ACM Subject Classification Theory of computation → Approximation algorithms analysis; Mathematics of computing → Approximation algorithms

Keywords and phrases Approximation Algorithms, Randomized Algorithms, Linear Programming, Graph Algorithms, Scheduling, Multicut, Multiflow

Digital Object Identifier 10.4230/LIPIcs.APPROX/RANDOM.2025.14

Category APPROX

Acknowledgements We would like to thank Joseph Cheriyan for many helpful discussions throughout the course of this project.

1 Introduction

Given an edge-weighted graph with k source-sink pairs, a *multicut* is a set of edges whose removal disconnects all the source-sink pairs. The *minimum multicut problem* seeks a multicut with the minimum total edge weight. This problem generalizes the classical minimum s - t cut problem and has been extensively studied in the past. Computing the minimum multicut is NP-hard, even in highly restricted settings such as trees [11].

A problem closely related to the multicut problem is the multicommodity flow problem (also known as *multiflow*). The goal of this problem is to maximize the total flow that can be routed between the source-sink pairs. If the flow is restricted to take only integer values, the problem is called the *maximum integer multiflow problem*, which generalizes the well-known *edge-disjoint paths* problem. Since any source-sink path must use at least one edge of any multicut, the value of any feasible multicut is at least that of the maximum multicommodity flow. In fact, it turns out that the LP relaxation of multicut problem is the linear programming dual of the multiflow problem. The ratio of the minimum multicut to the maximum multicommodity flow is called the *multiflow-multicut gap*. By the strong duality of linear programming, it follows that the integrality gap of the natural linear programming relaxation for the multicut also provides a bound on the multiflow-multicut gap, and vice versa.



© Sina Kalantarzadeh and Nikhil Kumar;

licensed under Creative Commons License CC-BY 4.0

Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2025).

Editors: Alina Ene and Eshan Chattopadhyay; Article No. 14; pp. 14:1–14:18



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

The famous max-flow min-cut theorem [8] states that the multiflow-multicut gap is exactly 1 when $k = 1$, i.e., when there is exactly one source-sink pair. A well-known theorem by Hu [13] further establishes that the gap remains 1 when $k = 2$. However, this equality does not hold when there are three or more source-sink pairs, even for very simple graphs (see [11] for an example).

Garg, Vazirani, and Yannakakis [10] proved a tight bound of $\Theta(\ln k)$ on the multiflow-multicut gap for any graph G . If G is a tree, then the multiflow-multicut gap is exactly 2 [11]. For K_r -minor-free graphs, Tardos and Vazirani [19] used the decomposition theorem of Klein, Plotkin, and Rao [14] to prove a bound of $\mathcal{O}(r^3)$ on the multiflow-multicut gap. This bound was subsequently improved to $\mathcal{O}(r^2)$ by Fakcharoenphol and Talwar [6], and then to $\mathcal{O}(r)$ by Abraham et al. [1]. A tight bound of $\Theta(\log r)$ was then obtained for graphs of bounded treewidth [7, 9]. Finally, building upon this long sequence of results, Conroy and Filtser [5] recently proved an asymptotically tight bound of $\Theta(\log r)$ on the multiflow-multicut gap for K_r -minor-free graphs. Since planar graphs do not contain K_5 as a minor, it follows that the integrality gap of the minimum multicut problem for planar graphs is $\mathcal{O}(1)$.

The primary motivation behind the works mentioned above was to establish an asymptotic bound on the integrality gap (in terms of r) without optimizing the constants involved. However, for specific graph families, such as planar graphs, the constant obtained from these results is quite large (close to 100). Thus, determining the exact integrality gap remains an intriguing question. It is known that the integrality gap is at least 2 for trees, and consequently for planar graphs as well. Better upper and lower bounds for this problem remain elusive, serving as the primary motivation for this paper.

1.1 Related Work: Demand Multicommodity Flow

In another well studied version of the problem, called the *demand multicommodity flow*, we are given a demand value for each source-sink pair, denoted as d_i for the source-sink pair s_i-t_i . The goal is to determine whether there exists a feasible flow satisfying all the demands. A necessary condition for the existence of a feasible flow is as follows: across every bi-partition $(S, \bar{S})$ of the vertex set, the total demand that must be routed across $(S, \bar{S})$ must not exceed the total capacity of edges crossing $(S, \bar{S})$. This condition is known as the *cut-condition*, and it is a sufficient condition for the existence of flows in trees, outerplanar graphs, and similar graph classes.

In general, however, the cut-condition is not sufficient for the existence of a feasible flow. This leads to a natural question: what is the minimum relaxation of the cut-condition that ensures feasibility? Specifically, what is the smallest $\alpha \geq 1$ such that if the total capacity of edges across every bi-partition is at least α times the demand across the partition, then a feasible flow is guaranteed? In their seminal work, Linial, London, and Rabinovich [16] showed that this gap is $\Theta(\log k)$ for general graphs.

In contrast to the multiflow-multicut gap, our understanding of the flow-cut gap for planar graphs remains limited. Rao [17] showed that the flow-cut gap for planar graphs is $\mathcal{O}(\sqrt{\log n})$. However, the best known lower bound remains just 2 [15, 4], and it is conjectured that the true answer is $\mathcal{O}(1)$ [12].¹

On the other hand, we have a much better understanding of this gap for series-parallel graphs, a subclass of planar graphs. The flow-cut gap is exactly 2 for series-parallel graphs [3, 15]. Given the current state of research, one might be tempted to claim that we understand

¹ This conjecture is widely known as the Planar Embedding Conjecture or the GNRS Conjecture [12].

multiflow-multicut gaps better than flow-cut gaps. However, somewhat surprisingly, the precise multiflow-multicut gap for series-parallel graphs remains unknown, despite the well-understood flow-cut gap. One of the primary motivations of this paper is to bridge this gap in our understanding.

2 Preliminaries

Given a graph G , we denote its vertex and edge sets by $V(G)$ and $E(G)$, respectively. We will use K_r to denote the complete graph on r vertices. In this paper, we will only be concerned with planar graphs. A graph G is *planar* if it does not contain K_5 or $K_{3,3}$ as a minor. Equivalently, a graph is planar if it can be drawn in the plane without any of its edges crossing. Graphs in which every edge is contained in at most one cycle are called *cactus* graphs. Cactus graphs are a subclass of series-parallel and planar graphs, and are arguably the simplest family of planar graphs after trees and cycles. Cactus and series-parallel graphs do not contain K_4 as a minor.

Let G be a simple undirected graph with edge costs $c: E(G) \rightarrow \mathbb{Q}_{\geq 0}$, and let $\{(s_i, t_i)\}_{i=1}^k$ be the set of source-sink pairs. Let $\mathcal{P}_i$ denote the set of all paths between s_i and t_i in G , and let $\mathcal{P} = \bigcup_{i=1}^k \mathcal{P}_i$. A *multicut* is a set of edges $F \subseteq E(G)$ such that every $P \in \mathcal{P}$ contains at least one edge in F . Equivalently, a multicut is a set of edges whose removal disconnects every source-sink pair. A *multicommodity flow* is an assignment of non-negative real numbers to the paths in $\mathcal{P}$ that respects the capacity constraints of the edges. In the *maximum multiflow problem*, the objective is to find an assignment which maximizes the total value of flow routed.

Given an edge length function $l: E(G) \rightarrow \mathbb{Q}_{\geq 0}$ and $u, v \in V(G)$, we use $l(u, v)$ to denote the shortest path distance between u and v with respect to l . The diameter of G is the maximum distance between a pair of vertices in G , i.e., $\text{diam}(G) = \max_{u, v \in V(G)} l(u, v)$. We use $l(v, e)$ to denote the distance of a vertex v from an edge $e = (x, y)$, i.e., $l(v, e) = \min\{l(v, x), l(v, y)\}$.

For $F \subseteq E(G)$ and $v \in V(G)$, we use $C_F(v)$ to denote the connected component of $G - F$ containing v . We overload notation and also use $C_F(v)$ to denote the set of vertices in the connected component containing v . We define the radius of v with respect to F as the distance of the farthest vertex from v in $C_F(v)$, i.e., $\text{rad}_F(v) = \max_{u \in C_F(v)} l(v, u)$. In addition, the diameter of F is the maximum diameter of a connected component after the removal of F from G , i.e., $\text{diam}(F) = \max_{v \in V(G)} \text{diam}(C_F(v))$. Given $t \in \mathbb{R}_{\geq 0}$ as a parameter, we say that F forms a t -diameter decomposition if $\text{diam}(F) < t$. We denote the set of all t -diameter decompositions of G by $\mathcal{F}_t(G)$. Note that when referring to the distance between two vertices u, v in a component C , $l(u, v)$ denotes their distance in G , rather than in the subgraph induced by C , i.e., $G[C]$.

2.1 Linear Programming Relaxation for the Minimum Multicut Problem

We begin by describing an integer programming (IP) formulation for the minimum multicut problem. For each edge $e \in E(G)$, we introduce an integer variable $x(e) \in \{0, 1\}$, which indicates whether the edge is selected in the multicut. For a given path P , we define $x(P) = \sum_{e \in E(P)} x(e)$. A feasible multicut must include at least one edge from each source-sink path, so we impose the constraint $x(P) \geq 1$ for all $P \in \mathcal{P}$, ensuring that each path is *cut* by at least one edge. We relax the integrality constraints to obtain the linear programming (LP) relaxation of the multicut problem, which is formulated as follows:

$$\begin{aligned}
& \min \sum_{e \in E(G)} c(e) \cdot x(e) \\
& \text{s.t.} \quad x(P) \geq 1 \quad \forall P \in \mathcal{P} \\
& \quad \quad x(e) \geq 0 \quad \forall e \in E(G).
\end{aligned}$$

Even though there are an exponential number of constraints, it is well known that the optimal solution to this LP can be computed in polynomial time [10]. We denote the optimal solutions of the integer and linear programs as OPT_{IP} and OPT_{LP} , respectively. We refer to OPT_{LP} as the minimum fractional multicut. We know that the value of the maximum multiflow is equal to the minimum fractional multicut. Furthermore, a bound on the integrality gap of the LP relaxation for the multicut problem provides the same bound for the multiflow-multicut gap. Therefore, from this point onward, we will focus solely on the integrality gap of the multicut LP.

► **Definition 1.** Let $\mathcal{G}$ be a family of graphs, and let $\mathcal{M}(\mathcal{G})$ be the family of all instances of the minimum multicut problem on $\mathcal{G}$, obtained by assigning arbitrary capacities to the edges and selecting a set of source-sink pairs. The integrality gap $\alpha_{\mathcal{M}(\mathcal{G})}$ of the minimum multicut problem on $\mathcal{M}(\mathcal{G})$ is defined as follows:

$$\alpha_{\mathcal{M}(\mathcal{G})} := \max_{M \in \mathcal{M}(\mathcal{G})} \frac{\text{OPT}_{IP}(M)}{\text{OPT}_{LP}(M)}.$$

From the discussion above, we know that $\alpha_{\mathcal{M}(\text{TREE})} = 2$ [11], where TREE denotes the family of all trees, and $\alpha_{\mathcal{M}(\text{PLANAR})} = \mathcal{O}(1)$ [14], where PLANAR refers to the family of all planar graphs. In this paper, we focus on planar graphs, specifically the class of cactus graphs.

3 Our Results and Techniques

We provide a partial answer to the questions raised above by showing that the integrality gap of the minimum multicut problem for the family of cactus graphs (and therefore for series-parallel graphs and planar graphs) is strictly greater than 2. In particular, we show that the multiflow-multicut gap is at least $\frac{20}{9}$ for the class of *cactus* graphs.

We give two different proofs of this lower bound. In the first proof, we develop a novel technique to argue that the integrality gap of the multicut LP is at least $\frac{20}{9}$. We first observe that the integrality gap of the multicut LP for a class of graphs is α only if any fractional solution to the natural linear programming relaxation of the minimum multicut problem can be approximately written as a convex combination (or equivalently a probability distribution) of feasible multicuts. Furthermore, a feasible multicut can be interpreted in terms of small diameter decompositions (i.e., a set of edges whose removal results in connected components of small diameter) with an appropriate distance function. Therefore, if a graph class admits an integrality gap of at most α , then there exists a set of small diameter decompositions that do not cut any fixed edge too many times. We describe this in detail in Section 4.

Our crucial insight is that if the integrality gap is α for a class of graphs, then there exists a well-structured set of small diameter decompositions that can be used to construct the aforementioned convex combination. These structured decompositions are inspired by the well-known single-source distance-based decomposition algorithms for trees. We also describe this in detail in Section 4. The final step of the proof involves using these structural

insights to argue that there cannot exist a small diameter decomposition with a small value of α for the family of cactus graphs. Note that this proof is non-constructive and does not lead to an explicit example with a large gap. Nevertheless, this proof provides sufficient structural insights into instances with a large integrality gap, allowing us to construct explicit examples of cactus graphs where the gap is at least $\frac{20}{9}$. We emphasize that we attempted to construct these examples through an exhaustive computer search and manual crafting but were unsuccessful. Furthermore, the structural properties established in the first proof hold for very general classes of graphs, specifically those closed under edge subdivision and 1-sum operations, and may prove useful in other settings as well. We now formally state our main theorem, which we prove in Section 7.

► **Theorem 2.** *If $\mathcal{G}$ is the class of cactus graphs, then $\alpha_{\mathcal{M}(\mathcal{G})} \geq \frac{20}{9}$.*

4 Multicut and Small Diameter Decompositions

Let $\mathcal{G}$ be a family of graphs, $\mathcal{M}(\mathcal{G})$ be the family of all instances of the minimum multicut problem on $\mathcal{G}$ and $\alpha = \alpha_{\mathcal{M}(\mathcal{G})}$. In the following theorem, we are going to show that any *feasible* fractional solution to the linear programming relaxation of the minimum multicut problem of an instance $M \in \mathcal{M}(\mathcal{G})$ can be approximately written as a convex combination of feasible multicut solutions of M . The proof of this theorem also follows from the work of Carr and Vempala [2], and for completeness we have included a proof in the appendix A.

► **Theorem 3.** *Suppose we are given an instance of multicut $M \in \mathcal{M}(\mathcal{G})$. Let $\mathcal{F} \subseteq 2^{E(G)}$ be the set of all the feasible multicuts for instance M and x be a feasible fractional solution to the LP relaxation. Then there exists a probability distribution y over $\mathcal{F}$ such that:*

$$\sum_{F: e \in F, F \in \mathcal{F}} y_F \leq \alpha \cdot x(e) \quad \forall e \in E(G).$$

► **Corollary 4.** *Let G be a graph with edge length function l and let $t > 0$ be a given parameter. Then there exists a probability distribution y over $\mathcal{F}_t(G)$ such that:*

$$\sum_{F: e \in F, F \in \mathcal{F}_t(G)} y_F \leq \frac{\alpha \cdot l(e)}{t} \quad \forall e \in E(G).$$

Proof. We will define a multicut instance on G and then use Theorem 3 to complete the proof. Let $S = \{(u, v) \in V(G) \times V(G) \mid l(u, v) \geq t\}$ be the set of source-sink pairs. Let x be a fractional solution such that $x(e) = \frac{l(e)}{t}$ for all $e \in E(G)$. It is easy to verify that this is feasible for the multicut instance. By Theorem 3, there exists a probability distribution over feasible multicuts that satisfies the statement of the corollary, i.e., an edge is cut with probability at most $\alpha \cdot \frac{l(e)}{t}$. The proof then follows by noting that a set of edges is a feasible multicut for this instance if and only if it forms a t -diameter decomposition. ◀

In other words, if the integrality gap of any multicut instance for a class of graphs $\mathcal{G}$ is at most α , then for any $G \in \mathcal{G}$ equipped with an edge length function l , there exists a distribution over t -diameter decompositions such that an edge e is included in no more than $\frac{\alpha \cdot l(e)}{t}$ fraction of the decompositions. This formulation of the integrality gap eliminates the notion of source-sink pairs, which could be arbitrarily situated, and provides a more uniform way of analysis. Furthermore, we will work with a family of graphs that are *closed under edge subdivision*, i.e., the operation of replacing an edge by a path of arbitrary length. By scaling and subdividing edges, we can assume that all edge lengths are 1. This simplifies the presentation significantly, and we will make this assumption from here on.

5 Small Diameter Decomposition for Trees

In this section, we provide a detailed description of a well-known construction of a probability distribution over t -diameter decompositions for trees and highlight some useful properties of this distribution. This will serve as a foundation for developing intuition regarding the construction of structured small-diameter decompositions in the next section. As mentioned earlier, we will assume without loss of generality that all edge lengths are set to 1. Additionally, for reasons that will become clear shortly, we will assume that $t = 2w$ for some positive even integer w .

We know that $\alpha = \alpha_{\mathcal{M}(\text{TREE})} = 2$ for the family of trees [11]. Given a tree T , Corollary 4 implies that there exists a probability distribution over $2w$ -diameter decompositions of T such that each edge $e \in E$ is part of the decomposition with probability at most $\frac{\alpha}{2w} = \frac{2}{2w} = \frac{1}{w}$. In the following, we provide an explicit construction of this distribution.

► **Theorem 5.** *Let T be a tree and $\mathcal{F}_{2w}(T)$ be the family of all $2w$ -diameter decompositions of T . Then there exists a probability distribution y over $\mathcal{F}_{2w}(T)$ such that:*

$$\sum_{F \in \mathcal{F}_{2w}(T), e \in F} y_F \leq \frac{1}{w} \quad \forall e \in E(T). \quad (1)$$

Proof. We root the tree T at an arbitrary vertex $r \in V$. Let

$$F_i = \{e \in E \mid l(r, e) = i + kw \text{ where } k \in \mathbb{Z}_{\geq 0}\} \text{ for } i = 0, \dots, w-1.$$

We set $y_{F_i} = \frac{1}{w}$ for $i = 0, \dots, w-1$, and $y_F = 0$ otherwise. It is easy to see that F_i 's partition the edge set $E(T)$, i.e. $E(T) = \cup_{i=0}^{w-1} F_i$ and $F_i \cap F_j = \emptyset$ for $i \neq j$. Thus,

$$\sum_{F \in \mathcal{F}_{2w}(T), e \in F} y_F = \sum_{i=0}^{w-1} y_{F_i} = \frac{1}{w} \quad \forall e \in E(T).$$

To complete the proof of the theorem, we need to show that F_i is a $2w$ -diameter decomposition for $i = 0, \dots, w-1$. Fix a F_i . Let (u, v) be a pair of vertices such that $l(u, v) \geq 2w$, and let q be the lowest common ancestor (LCA) of u and v . The unique path $u - v$ between u and v can be partitioned into two subpaths: one from u to q , and the other from q to v . Since the length of the $u - v$ path is at least $2w$, one of the paths $u - q$ or $q - v$ must have length at least w . Without loss of generality, assume that the $q - v$ path has length at least w . Denote this path by $Q = e_0, e_1, \dots, e_p$.

Since q is an ancestor of v , the length of the path from any vertex r to an edge e_i is given by $l(r, e_i) = l(r, e_{i-1}) + 1$ for $i = 1, \dots, p$. This implies that there exists at least one edge in Q , say e_j , such that $j = i + kw$ for some $k \in \mathbb{Z}_{\geq 0}$. This implies that $e_j \in F_i$, and hence, every u, v pair with distance at least $2w$ is disconnected after removing F_i . This completes the proof that F_i is a valid $2w$ -diameter decomposition. ◀

We now point out a useful property of the $2w$ -diameter decompositions constructed above, which will be helpful later.

► **Observation 6.** *The following holds for the $2w$ -diameter decompositions $F_0, \dots, F_{w-1}$ described in the proof of Theorem 5:*

1. $\text{rad}_{F_i}(r) \leq i \leq w-1$ for every $i = 0, \dots, w-1$. Equivalently,

$$\sum_{F: \text{rad}_F(r) \leq w-1} y_F = 1.$$

2. For all $1 \leq k \leq w$, $\text{rad}_{F_i}(r) \leq k - 1$ with sufficiently high probability. More precisely,

$$\sum_{F_i: \text{rad}_{F_i}(r) \leq k-1} y_{F_i} = \sum_{i=0}^{k-1} y_{F_i} = \frac{k}{w} = 1 - \frac{2}{2w}(w - k) = 1 - \frac{\alpha_{\mathcal{M}(\text{TREE})}}{2w}(w - k).$$


This observation shows that not only does the distribution claimed in Theorem 5 satisfy the condition in equation (1), but it also fulfills additional constraints. These results are succinctly captured in the following corollary. In Section 6, we will show that a similar distribution exists for the families of graphs which are closed under 1-sum operation.

► **Corollary 7.** *Let T be a tree, and let $r \in V(T)$ be an arbitrary vertex of T . Let $\mathcal{F}_{2w}(T)$ be the family of all $2w$ -diameter decompositions. Then, there exists a probability distribution y over $\mathcal{F}_{2w}(T)$ such that:*

$$\begin{aligned} \sum_{F \in \mathcal{F}_{2w}(T), e \in F} y_F &\leq \frac{1}{w} \quad \forall e \in E(T), \\ \sum_{F: \text{rad}_F(r) \leq w-1} y_F &= 1, \\ \sum_{F: \text{rad}_F(r) \leq k-1} y_F &\geq 1 - \frac{\alpha_{\mathcal{M}(\text{TREE})}}{2w}(w - k) \quad \forall k = 1, \dots, w. \end{aligned}$$

6 A Structural Result for Small Diameter Decompositions

We now define the 1-sum operation on graphs, which will play a crucial role going forward. Let $G_1, \dots, G_k$ be non-empty graphs, and let $r_i \in V(G_i)$ for $i = 1, \dots, k$. The graph G^S is obtained by taking the disjoint union of $G_1, G_2, \dots, G_k$, and identifying the vertices $r_1, r_2, \dots, r_k$. We say that G^S is obtained by performing the 1-sum of the G_i 's at the vertices r_i 's. The vertex $r = r_1 = \dots = r_k$ is called the main vertex of G^S . See figure 1 for an illustration.



■ **Figure 1** An illustration of the 1-sum operation.

Let $\mathcal{G}$ be a family of graphs. We say that $\mathcal{G}$ is closed under the 1-sum operation if for any $G_1, \dots, G_k \in \mathcal{G}$ and $r_i \in G_i$, the graph obtained by taking 1-sum of $G_1, \dots, G_k$ at $r_1, r_2, \dots, r_k$ is a graph in $\mathcal{G}$. Many natural classes of family are closed under the 1-sum operation, such as trees, cactus graphs and planar graphs. Note that 1-sum is a special case of a well known and a more general notion of clique-sums.

We note down a few more definitions before stating the main theorem of this section. Let G be a graph and $r \in V(G)$ be an arbitrary vertex. Recall that $\mathcal{F}_{2w}(G)$ denotes the set of all $2w$ -diameter decompositions of G . For $k \in \{1, \dots, w\}$, we use $\mathcal{F}_{2w}^k(G, r)$ to denote the set of all $2w$ -diameter decompositions of G such that every vertex in the connected component containing r is within distance strictly less than k from it. More precisely,

$$\mathcal{F}_{2w}^k(G, r) = \{F \in \mathcal{F}_{2w}(G) \mid \text{rad}_F(r) < k\}.$$

► **Definition 8.** Let G be a graph, and let $\mathcal{F}_{2w}(G)$ be the family of $2w$ -diameter decompositions of G . We say that $y = \{y_F \mid F \in \mathcal{F}_{2w}(G)\}$ is a p -load distribution if the following conditions hold:

$$\begin{aligned} \sum_{F \in \mathcal{F}_{2w}(G)} y_F &= 1, \\ \sum_{F \in \mathcal{F}_{2w}(G), e \in F} y_F &\leq p \quad \forall e \in E(G), \\ y_F &\geq 0 \quad \forall F \in \mathcal{F}_{2w}(G). \end{aligned}$$

If such a p -load distribution exists, we say that G accepts a p -load distribution. We also say that a family $\mathcal{G}$ accepts a p -load distribution if every graph $G \in \mathcal{G}$ accepts a p -load distribution.

Note that if we have a p -load distribution, then we can use this to sample a $2w$ -diameter decomposition of G such that each edge is picked in the decomposition with probability at most p .

► **Corollary 9.** Let $\mathcal{G}$ be a family of graphs. Let $p = \frac{\alpha_{\mathcal{M}(\mathcal{G})}}{2w}$, then $\mathcal{G}$ accepts a p -load distribution.

Proof. Recall that we assumed that all edges have unit length. Corollary 4 gives a direct proof by setting $t = 2w$. ◀


Suppose that $\mathcal{G}$ is a family of graphs closed under 1-sum for the rest of the section. Let $p = \frac{\alpha}{2w}$, where $\alpha = \alpha_{\mathcal{M}(\mathcal{G})}$. Corollary 9 implies that $\mathcal{G}$ accepts a p -load distribution. In Theorem 10, we show that if $\mathcal{G}$ is closed under the 1-sum operation and accepts a p -load distribution, then for any $G \in \mathcal{G}$ and $r \in V(G)$, there exists a distribution over $2w$ -diameter decompositions $\mathcal{F}_{2w}(G)$ such that if we sample a decomposition $F \in \mathcal{F}_{2w}(G)$ from this distribution, we are guaranteed that $\text{rad}_F(r) \leq w - 1$, i.e. $F \in \mathcal{F}_{2w}^w(G, r)$. Note that this is similar to the first item of Observation 6.

► **Theorem 10.** Suppose that $\mathcal{G}$ is closed under 1-sum and accepts a p -load distribution. Let $G \in \mathcal{G}$ and $r \in V(G)$ be an arbitrary vertex. Then there exists a p -load distribution $y = \{y_F \mid F \in \mathcal{F}_{2w}(G)\}$ such that $\sum_{F \in \mathcal{F}_{2w}^w(G, r)} y_F = 1$.

Proof. Let $\mathcal{F}_{2w} = \mathcal{F}_{2w}(G)$ and $\mathcal{F}_{2w}^w(r) = \mathcal{F}_{2w}^w(G, r)$ for simplicity. It is sufficient to show that the following LP is feasible and has optimal value 0.

$$\begin{aligned} \min \quad & \sum_{F \in \mathcal{F}_{2w} \setminus \mathcal{F}_{2w}^w(r)} y_F \\ \sum_{F \in \mathcal{F}_{2w}, e \in F} y_F &\leq p \quad \forall e \in E(G) \\ \sum_{F \in \mathcal{F}_{2w}} y_F &= 1 \\ y_F &\geq 0 \quad \forall F \in \mathcal{F}_{2w} \end{aligned}$$

The above LP is feasible since $\mathcal{G}$ accepts a p -load distribution. For the sake of contradiction, assume that the optimal value of the above LP is $z > 0$. Let $m > \frac{1}{z}$ be a natural number. Let $G_1, \dots, G_m$ be m disjoint copies of G and r_i be the vertex of G_i which corresponds to r . Let G' be formed by taking 1-sum of $G_1, \dots, G_m$ at $r_1, \dots, r_m$. See figure 2 for an illustration.



■ **Figure 2** The construction of G' .

Note that $G' \in \mathcal{G}$ since $\mathcal{G}$ is closed under the 1-sum operation. Let $\mathcal{F}'_{2w} = \mathcal{F}_{2w}(G')$ be the set of all $2w$ -diameter decompositions of G' . By the assumption of the theorem, every graph in $\mathcal{G}$ accepts a p -load distribution. Let $\{g_{F'}\}_{F' \in \mathcal{F}'_{2w}}$ denote such a distribution for G' . Let $G_i = (V_i, E_i)$ and $\mathcal{F}_{2w}(G_i)$ be the set of all $2w$ -diameter decompositions of G_i for $i = 1, 2, \dots, m$.² The distribution g over $\mathcal{F}'_{2w}$ defines a distribution g^i over $\mathcal{F}_{2w}(G_i)$ as follows:

$$g_F^i = \sum_{F' \in \mathcal{F}'_{2w}, F' \cap E_i = F} g_{F'} \quad \text{for all } F \in \mathcal{F}_{2w}(G_i).$$

Since g^i is a projection of g onto G_i and g is a p -load distribution, it follows that g^i is also a p -load distribution (see Appendix B for a short proof). Furthermore, since G_i is an identical copy of G and g^i is a feasible solution to the LP mentioned above, we have that,

$$\sum_{F \in \mathcal{F}_{2w}(G_i) \setminus \mathcal{F}_{2w}^w(G_i, r)} g_F^i \geq z \quad \text{for } i = 1, 2, \dots, m.$$

Recall that $\mathcal{F}_{2w}^w(G_i, r)$ denotes the set of all $2w$ -diameter decompositions of G_i in which the distance of every vertex in the connected component containing r is at most $w - 1$ from it. Let T_i be the event that, when sampling $F' \in \mathcal{F}'_{2w}$ according to the distribution g , the intersection $F' \cap E_i$ does not belong to $\mathcal{F}_{2w}^w(G_i, r)$. From the above discussion, it follows that $\Pr[T_i] \geq z$. Since $m > \frac{1}{z}$, we have

$$\sum_{i=1}^m \Pr[T_i] \geq z \cdot m > 1.$$

This implies that the events $T_1, \dots, T_m$ are not disjoint, and there exist indices i, j such that $\Pr[T_i \cap T_j] \neq \emptyset$. Therefore, there exists a $F' \in \mathcal{F}'_{2w}$, and vertices $u \in V_i, v \in V_j$ such that:

1. $g_{F'} > 0$,
2. u and v are in the connected component containing r in $G' - F'$,
3. the distance of u and v from r is at least w .

But then, the diameter of F' is at least $2w$, which contradicts the fact that $F' \in \mathcal{F}'_{2w}$. This implies that $z = 0$, and completes the proof of the theorem. ◀

We say that a graph G and one of its vertex r accepts a (r, p) -load distribution if there exists a p -load distribution y over $2w$ -diameter decompositions $\mathcal{F}_{2w}(G)$ such that $\sum_{F \in \mathcal{F}_{2w}^w(G, r)} y_F = 1$.

² Note that $\bigcap_{i=1}^m V_i = \{r\}$.

14:10 Improved Lower Bounds on Multiflow-Multicut Gaps

We say that a graph class $\mathcal{G}$ accepts a *strong- p -load distribution* if for every $G \in \mathcal{G}$ and $r \in V(G)$, there exists a (r, p) -load distribution.

In Theorem 10, we proved that if a class of graphs $\mathcal{G}$, closed under the 1-sum operation, accepts a p -load distribution, then it also accepts a strong- p -load distribution. In Theorem 11, we generalize Theorem 10 to arbitrary radii. To prove this theorem, we assume, in addition to the graph class being closed under the 1-sum operation, that the class also includes K_2 , the complete graph on two vertices, i.e. a single edge.


► **Theorem 11.** *Suppose that $\mathcal{G}$ is closed under 1-sum such that $K_2 \in \mathcal{G}$ and accepts a strong- p -load distribution. Let $G \in \mathcal{G}$, $r \in V(G)$ and $k \in \{1, \dots, w\}$. Then there exists a (r, p) -load distribution $y = \{y_F \mid F \in \mathcal{F}_{2w}(G)\}$ such that,*

$$\sum_{F \in \mathcal{F}_{2w}^k(G, r)} y_F \geq 1 - p(w - k).$$

Proof. Let $\mathcal{F}_{2w} = \mathcal{F}_{2w}(G)$, $\mathcal{F}_{2w}^w(r) = \mathcal{F}_{2w}^w(G, r)$ and $\mathcal{F}_{2w}^k(r) = \mathcal{F}_{2w}^k(G, r)$ for simplicity. It is sufficient to show that the following LP is feasible and has optimal value 0.

$$\begin{aligned} \min \quad & \sum_{F \in \mathcal{F}_{2w} \setminus \mathcal{F}_{2w}^w(r)} y_F \\ & \sum_{F \in \mathcal{F}_{2w}^k(r)} y_F \geq 1 - p(w - k) \\ & \sum_{F \in \mathcal{F}_{2w}, e \in F} y_F \leq p \quad \forall e \in E(G) \\ & \sum_{F \in \mathcal{F}_{2w}} y_F = 1 \\ & y_F \geq 0 \quad \forall F \in \mathcal{F}_{2w} \end{aligned}$$

For now, assume that the LP is feasible and its optimal value is $z > 0$. The proof that LP is feasible will also follow from the discussion below. Let $m > \frac{1}{z}$ and $G_1, \dots, G_m$ be m disjoint copies of G . Let r_i be the vertex of G_i which corresponds to r and G' be formed by taking 1-sum of $G_1, \dots, G_m$ at $r_1, \dots, r_m$. We construct H by adding a path of length $w - k$ to G' at r . Let $P = \{e_1, e_2, \dots, e_{w-k}\}$ be the set of edges on this path, where $e_1 = (r, v_1), e_2 = (v_1, v_2), \dots, e_{w-k} = (v_{w-k-1}, v)$. See figure 3 for an illustration.



■ **Figure 3** The construction of H .

Since $\mathcal{G}$ is closed under taking 1-sum and $K_2 \in \mathcal{G}$, we can conclude that $H \in \mathcal{G}$. By Theorem 10, there exists a probability distribution $x = \{x_{F'} \geq 0 \mid F' \in \mathcal{F}_{2w}^w(H, v)\}$, such that

$$\sum_{F' \in \mathcal{F}_{2w}^w(H, v)} x_{F'} = 1, \text{ and, } \sum_{F' \in \mathcal{F}_{2w}^w(H, v), e \in F'} x_{F'} \leq p \text{ for all } e \in E(H).$$

Let $A = \{F' \in \mathcal{F}_{2w}^w(H, v) \mid F' \cap E(P) \neq \emptyset\}$ and $B = \{F' \in \mathcal{F}_{2w}^w(H, v) \mid F' \cap E(P) = \emptyset\}$. Let $A_i = \{F' \in A \mid e_i \in F'\}$ for $i = 1, \dots, w - k$. Note that $\sum_{F' \in A} x_{F'} + \sum_{F' \in B} x_{F'} = 1$. Since $\sum_{F' \in A_i} x_{F'} \leq p$, and $A = \bigcup_{i=1}^{w-k} A_i$, it follows that:

$$\sum_{F' \in A} x_{F'} \leq (w - k)p \Rightarrow \sum_{F' \in B} x_{F'} \geq 1 - (w - k)p.$$

Let $G_i = (V_i, E_i)$ and $\mathcal{F}_{2w}(G_i)$ be the set of all $2w$ -diameter decompositions of G_i for $i = 1, 2, \dots, m$. Recall that $\mathcal{F}_{2w}^w(G_i, r)$, $\mathcal{F}_{2w}^k(G_i, r)$ is the set of $2w$ -diameter decompositions of G_i in which the distance of every vertex in the connected component containing r is at most $w - 1$ and $k - 1$ from it, respectively. For $F \in \mathcal{F}_{2w}(G_i)$, let

$$y_F^i = \sum_{F' \in \mathcal{F}_{2w}^w(H, v), F' \cap E_i = F} x_{F'}.$$

In other words, y_F^i is the projection of x to the graph G_i . The next claim shows that the projection of the distribution x onto G_i is a feasible solution to the LP 6.

▷ **Claim 12.** $y^i = \{y_F^i \mid F \in \mathcal{F}_{2w}(G_i)\}$ is a feasible solution to the LP 6.

Proof. Appendix B implies that y^i is a p -load distribution for G_i . Let $B_i = \{F' \cap E_i \mid F' \in B\}$. Observe that $B_i \in \mathcal{F}_{2w}(G_i)$. Let $F \in B_i$. Then there exists $F' \in B$ such that $F = F' \cap E_i$. Furthermore, for any $F' \in B$, we have $F' \in \mathcal{F}_{2w}^w(H, v)$ and $F' \cap E(P) = \emptyset$. Hence we can conclude that $F \in \mathcal{F}_{2w}^k(G_i, r)$. Thus,

$$\sum_{F \in \mathcal{F}_{2w}^k(G_i, r)} y_F^i \geq \sum_{F \in B_i} y_F^i = \sum_{F \in B_i} \sum_{F' \in B, F' \cap E_i = F} x_{F'} = \sum_{F' \in B} x_{F'} \geq 1 - (w - k) \cdot p. \quad \triangleleft$$

Since G_i is an identical copy of G and $\mathcal{F}_{2w}(G_i)$ is also an identical copy of $\mathcal{F}_{2w} = \mathcal{F}_{2w}(G)$, Claim 6 shows that y^i is a feasible solutions for LP 6 for G_i , it follows that LP 6 is feasible. Since z is the optimal solution of LP 6, for each y^i , we have,

$$\sum_{F \in \mathcal{F}_{2w}(G_i) \setminus \mathcal{F}_{2w}^w(G_i, r)} y_F^i \geq z.$$

This implies that,

$$\sum_{i=1}^m \sum_{F \in \mathcal{F}_{2w}(G_i) \setminus \mathcal{F}_{2w}^w(G_i, r)} y_F^i \geq mz > 1.$$

Using same argument as in the proof of Theorem 10, we can show that there exists $1 \leq i \neq j \leq m$ and $F' \in \mathcal{F}_{2w}^w(H, v)$ such that

$$y_{F'} > 0, F' \cap E_i \notin \mathcal{F}_{2w}^w(G_i, r) \text{ and } F' \cap E_j \notin \mathcal{F}_{2w}^w(G_j, r).$$

This means that there exists a vertex $a \in V_i, b \in V_j$ such that the distance of a and b from r is at least w , and they are both included in the connected component containing r in $H - F'$. Hence u and v are at least $2w$ distance apart, and this contradicts the fact that F' is a $2w$ -diameter decomposition. Hence $z = 0$ and this completes the proof of the theorem. ◀

14:12 Improved Lower Bounds on Multiflow-Multicut Gaps

So far, we have shown that probability distributions over $2w$ -diameter decompositions for any class of graphs closed under 1-sum are consistent with those observed for trees, for a fixed value of k^3 . We now state a simple claim that will be useful in proving the lower bound. From this point onward, by Theorem 10 and Theorem 11, we can restrict ourselves to $\mathcal{F}_{2w}^w(G, r)$ instead of $\mathcal{F}_{2w}(G)$ for a given graph G and vertex $r \in V(G)$.

► **Lemma 13.** *Let $G \in \mathcal{G}$, and $r \in V(G)$ be an arbitrary vertex. Let $\mathcal{F}_{2w}^w(G, r)$ be defined as before. Let P be an arbitrary shortest path with length w , starting at r . Denote the other endpoint of P as r' . If $y = \{y_F \mid F \in \mathcal{F}_{2w}^w(G, r)\}$ is a (r, p) -load distribution then $F \cap E(P) \neq \emptyset$ for all $F \in \mathcal{F}_{2w}^w(G, r)$, and,*

$$\sum_{F \in \mathcal{F}_{2w}^w(G, r), |F \cap E(P)| \geq 2} y_F \leq (w \cdot p) - 1.$$

Proof. If there exists $F \in \mathcal{F}_{2w}^w(G, r)$ such that $F \cap E(P) = \emptyset$, then r, r' are within the same connected component in $G - F$, which contradicts the fact that $F \in \mathcal{F}_{2w}^w(G, r)$. Let,

$$A = \{F \in \mathcal{F}_{2w}^w(G, r) \mid |F \cap E(P)| \geq 2\} \text{ and } B = \{F \in \mathcal{F}_{2w}^w(G, r) \mid |F \cap E(P)| = 1\}.$$

Note that A, B forms a partition of $\mathcal{F}_{2w}^w(G, r)$. Using the definition of A and B , and the fact that $\sum_{F \in A} y_F + \sum_{F \in B} y_F = 1$, we can derive the statement of the theorem as follows:

$$\sum_{F \in A} y_F + 1 = 2 \cdot \sum_{F \in A} y_F + \sum_{F \in B} y_F \leq \sum_{e \in E(P)} \sum_{e \in F, F \in \mathcal{F}_{2w}^w(G, r)} y_F \leq \sum_{e \in E(P)} p \leq w \cdot p.$$

Note that the first inequality can be derived by showing that y_F appears at least twice in the right hand side if $F \in A$, and exactly once if $F \in B$. ◀


7 Lower Bound For Cactus Graphs

We are now ready to prove our main theorem. We will apply the tools developed in the previous sections to the family of cactus graphs. Let $\mathcal{G}$ be the family of cactus graphs, and define $\alpha = \alpha_{\mathcal{M}(\mathcal{G})}$. Let $q = \frac{\alpha}{2w}$. First, since $\mathcal{G}$ is closed under 1-sum, Corollary 9 implies that $\mathcal{G}$ accepts a q -load distribution and Theorem 10 ensures that $\mathcal{G}$ also accepts a strong q -load distribution. In Theorem 14, we show that if $\mathcal{G}$ accepts a p -load distribution, then $w \cdot p \geq \frac{10}{9}$. This implies that $w \cdot q = w \cdot \frac{\alpha}{2w} \geq \frac{10}{9}$, and hence $\alpha \geq \frac{20}{9}$. This concludes the proof of our main theorem, i.e. Theorem 2.

► **Theorem 14.** *Let $\mathcal{G}$ be the family of cactus graphs. If $\mathcal{G}$ accepts a p -load distribution, then $w \cdot p \geq \frac{10}{9}$.*

Proof. Let H be a cycle of length $2w$. Let r, u, r' , and v be four vertices of the cycle in anti-clockwise order, such that $l(r, u) = l(u, r') = l(r', v) = l(v, r) = \frac{w}{2}$. We construct G from H by attaching paths $u - u'$ and $v - v'$ of length $\frac{w}{2}$ from u and v , respectively. We denote the path of length $\frac{w}{2}$ from r to u by P_u , from u to r' by P_a , from r' to v by P_b , and from v to r by P_v . We denote the path from u to u' by P_c and the path from v to v' by P_d . See figure 4 for an illustration. Note that $G \in \mathcal{G}$. For the sake of contradiction, assume

³ Note that for trees, this property holds for all values of k simultaneously; however, here we can only guarantee it for a fixed value of k .



■ **Figure 4** Construction of G .

that $w \cdot p < \frac{10}{9}$. Let $k = \frac{w}{2}$, $\mathcal{F} = \mathcal{F}_{2w}^w(G, r)$ and $A = \mathcal{F}_{2w}^k(G, r)$. Since $K_2 \in \mathcal{G}$ and $\mathcal{G}$ is closed under 1-sum, we can use Theorem 10 and Theorem 11 to conclude that there exists (r, p) -load distribution $y = \{y_F \mid F \in \mathcal{F}_{2w}^w(G, r)\}$ such that,

$$\sum_{F \in A} y_F = \sum_{F \in \mathcal{F}_{2w}^k(r)} y_F \geq 1 - (w - k) \cdot p = 1 - \frac{w}{2} \cdot p = 1 - \frac{w \cdot p}{2}.$$

Suppose that $F \in A$. Since $l(u, r) = k = \frac{w}{2}$ and $l(v, r) = k = \frac{w}{2}$, we have that $F \cap E(P_v) \neq \emptyset$ and $F \cap (P_u) \neq \emptyset$. The next claim shows that under the assumption that $w \cdot p < \frac{10}{9}$, there exists $F \in A$ which does not pick any edges from P_a, P_b, P_c, P_d . This will lead to a contradiction, as u' and v' are $2w$ distance apart, and since F is a $2w$ -diameter decomposition, they should have been separated by F . This implies that $w \cdot p \geq \frac{10}{9}$ and completes the proof of the theorem.

▷ **Claim 15.** There exists $F \in A$ such that $F \cap E(P_a), F \cap E(P_b), F \cap E(P_c), F \cap E(P_d) = \emptyset$.

Proof. Let P'_a, P'_b be the two paths between r, r' containing P_a, P_b respectively. Also, let P'_u, P'_v be the unique shortest paths starting at r and ending at u', v' , respectively. Let $A_a = \{F \in A \mid F \cap E(P_a) \neq \emptyset\}$ and $F \in A_a$. This implies that $|F \cap E(P'_a)| \geq 2$. Since P'_a is a shortest path with length w from r , using Lemma 13, we have,

$$\sum_{F \in \mathcal{F}, |F \cap E(P'_a)| \geq 2} y_F \leq (w \cdot p) - 1 \Rightarrow \sum_{F \in A_a} y_F \leq (w \cdot p) - 1.$$

Similarly, we define,

$$A_b = \{F \in A \mid F \cap E(P_b) \neq \emptyset\}, A_c = \{F \in A \mid F \cap E(P_c) \neq \emptyset\}, A_d = \{F \in A \mid F \cap E(P_d) \neq \emptyset\},$$

and doing the same argument for P'_b, P'_u, P'_v , we obtain,

$$\sum_{F \in A_b} y_F \leq (w \cdot p) - 1; \sum_{F \in A_c} y_F \leq (w \cdot p) - 1; \sum_{F \in A_d} y_F \leq (w \cdot p) - 1.$$

Let $A^* = A \setminus (A_a \cup A_b \cup A_c \cup A_d)$. From the discussion above, it follows that,

$$\begin{aligned} \sum_{F \in A^*} y_F &= \sum_{F \in A} y_F - \sum_{F \in A_a \cup A_b \cup A_c \cup A_d} y_F \geq \left(1 - \frac{w \cdot p}{2}\right) - 4 \cdot (w \cdot p - 1) \\ &= 5 - \frac{9 \cdot w \cdot p}{2} > 0. \end{aligned}$$

This shows that $A^* \neq \emptyset$ and completes the proof of the claim. ◁


◀

8 An Explicit Construction

Inspired by the proof presented in the previous sections, we now give an explicit example on a cactus graph where the integrality gap is at least $\frac{20}{9}$. In particular, given an $\epsilon > 0$, we are going to give an instance of the minimum multicut problem M on a cactus graph G such that,


$$\frac{\text{OPT}_{IP}(M)}{\text{OPT}_{LP}(M)} \geq \frac{20}{9} - \epsilon.$$

We will construct the graph G by using two 1-sum operations. Let H be the graph depicted in figure 5.



■ **Figure 5** Construction of H .

Let k be a sufficiently large natural number, and let $H_1, \dots, H_k$ be k distinct copies of H . Let v_1^i be the vertex corresponding to v_1 in H . We construct H' by taking 1-sum of $H_1, H_2, \dots, H_k$ at $v_1^1, v_1^2, \dots, v_1^k$, respectively. Let $v_1 = v_1^1 = v_1^2 = \dots = v_1^k$. We obtain H'' by adding an edge (v_1, v_0) to H' . See figure 6 for an illustration.




■ **Figure 6** Construction of H'' .

Let $H_1'', \dots, H_k''$ be k disjoint copies of H'' . Let v_0^i be the vertex of H_i'' corresponding to the vertex v_0 of H'' . Let $G = (V, E)$ be the graph obtained by taking 1-sum of $H_1'', \dots, H_k''$ at $v_0^1, \dots, v_0^k$, respectively. Let $v_0 = v_0^1 = \dots = v_0^k$. Let v_i be the unique neighbor of v_0 in H_i'' . See figure 7 for an illustration.

Let $l(e) = 1$ for all $e \in E$. We partition the set of edges E based on their distance from v_0 as follows:

$$E_1 = \{e \in E \mid l(v_0, e) = 0\}, \quad E_2 = \{e \in E \mid l(v_0, e) = 1\}, \quad E_3 = \{e \in E \mid l(v_0, e) = 2\}.$$

Note that E_1 is the set of incident edges to v_0 , and E_1, E_2, E_3 is a partition of E . Now we are going to define an instance of the minimum multicut problem on G . We first assign costs to edges as follows:



■ **Figure 7** Construction of G .

$$c(e) = \begin{cases} k & e \in E_1 \\ 2 & e \in E_2 \\ 1 & e \in E_3 \end{cases}$$

Let $S = \{(u, v) \in V \times V \mid l(u, v) \geq 4\}$ be the set of source-sink pairs. We will denote this multicut instance by M . It is easy to see that $x = \{x_e = \frac{1}{4} \mid e \in E\}$ is a feasible fractional solution to M with cost

$$\frac{k \cdot k + 2 \cdot k^2 \cdot 2 + 4k^2 \cdot 1}{4} = \frac{9 \cdot k^2}{4}.$$

We will now show that the cost of any feasible multicut (i.e. an integral solution) is at least $5k^2 - 9k$. This will imply that the integrality gap for this instance is at least $\frac{20}{9} - \frac{4}{k}$, which can be arbitrarily close to $\frac{20}{9}$. More precisely, for any $\epsilon > 0$, we can set $k > \frac{4}{\epsilon}$ to obtain a lower bound of $\frac{20}{9} - \epsilon$.

Recall that for a graph H , we use $V(H)$ and $E(H)$ to denote the set of vertices and edges in H , and for $E' \subseteq E(H)$, we use $c(E')$ to denote the total cost of edges in E' . Let F be a feasible multicut solution to M . Let $t = |E_1 \cap F|$. For now assume that $0 < t < k$. Without loss of generality, assume that $(v_0, v_i) \in F$ for $i = 1, \dots, t$.

▷ **Claim 16.** $c(F \cap E(H''_i)) \geq 4 \cdot (k - 1) + k$ for $i = 1, 2, \dots, t$.

Proof. We will prove the claim for $i = 1$. The proof for other values of i is identical. Denote $H_1, \dots, H_k$ as the k copies of H incident at v_1 . For each $1 \leq i \leq k$, we call H_i good if $\text{rad}_F(v_1) \leq 1$, and we call it bad otherwise. It is not too difficult to see that there is at most one bad H_i . Suppose that H_1, H_2 are bad graphs, then there exists $a \in V(H_1), b \in V(H_2)$ such that $l(a, v_1), l(b, v_1) \geq 2$. But this is a contradiction since a and b are within the same component as v_1 after the removal of F , and are at a distance 4 apart, i.e. they are a source-sink pair in the multicut instance M . Thus, there are at least $k - 1$ good graphs attached to v_1 . By doing a simple case analysis, it can be verified that $c(F \cap E(H_i)) \geq 4$ if H_i is good. Combining the above with the fact that the edge between v_0, v_1 is included in F , we obtain the statement of the claim. ◁

Note that $v_{t+1}, \dots, v_k$ are within the same component as v_0 after the removal of F . For each $t + 1 \leq j \leq k$, we call H''_j good if $\text{rad}_F(v_0) \leq 1$ in H''_j , otherwise we call it bad. Using a similar argument as in the proof of Claim 8, one can show that there at most 1 bad H''_j . Thus, we have at least $k - t - 1$ good H''_j 's.

▷ Claim 17. $c(F \cap E(H_j'')) \geq 5 \cdot k$ if H_j'' is good for $t + 1 \leq j \leq k$.

Proof. Since the edge between v_0, v_j is not included in F , $E_2 \cap E(H_j'') \subseteq F$ or equivalently, all the edges with cost 2 of H_j'' are included in F . On the other hand, let H_i be one of the attached copies of H to v_j . Note that we have already showed that $E_2 \cap E(H_i) \in F$. If $E_3 \cap E(H_i) \cap F = \emptyset$, then there is a source-sink pair at distance 4 in H_i which is not disconnected. Therefore, in each H_i , F picks edges of total weight at least 5. ◁

Therefore we obtain,

$$c(F) \geq t \cdot (5k - 4) + (k - 1 - t) \cdot (5k) = 5k^2 - 5k - 4t \geq 5k^2 - 9k.$$

Even when $t = 0, k$, one can use the same arguments as above to obtain the same lower bound on the cost of F .

9 Conclusions and Future Work

We improve upon a decade-old lower bound on the multiflow-multicut gap for planar graphs and, in doing so, develop new techniques. The main question our work raises is whether tight gap results can be obtained, even for the class of cactus and series-parallel graphs, and more generally, for planar graphs. Proving such a result likely requires new techniques, making it an interesting and challenging problem.

References

- 1 Ittai Abraham, Cyril Gavoille, Anupam Gupta, Ofer Neiman, and Kunal Talwar. Cops, robbers, and threatening skeletons: Padded decomposition for minor-free graphs. *SIAM Journal on Computing*, 48(3):1120–1145, 2019. doi:10.1137/17M112406.
- 2 Robert Carr and Santosh Vempala. Randomized metarounding. *Probabilistic Methods in Combinatorial Optimization*, 20:342–352, 2002. doi:10.1002/RSA.10033.
- 3 Amit Chakrabarti, Alexander Jaffe, James R Lee, and Justin Vincent. Embeddings of topological graphs: lossy invariants, linearization, and 2-sums. In *49th Annual IEEE Symposium on Foundations of Computer Science, 2008*, pages 761–770. IEEE, 2008. doi:10.1109/FOCS.2008.79.
- 4 Chandra Chekuri, F Bruce Shepherd, and Christophe Weibel. Flow-cut gaps for integer and fractional multiflows. *Journal of Combinatorial Theory, Series B*, 103(2):248–273, 2013. doi:10.1016/J.JCTB.2012.11.002.
- 5 Jonathan Conroy and Arnold Filtser. How to protect yourself from threatening skeletons: Optimal padded decompositions for minor-free graphs. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*, pages 2281–2292, 2025. doi:10.1145/3717823.3718252.
- 6 Jittat Fakcharoenphol and Kunal Talwar. An improved decomposition theorem for graphs excluding a fixed minor. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 36–46. Springer, 2003. doi:10.1007/978-3-540-45198-3_4.
- 7 Arnold Filtser, Tobias Friedrich, Davis Issac, Nikhil Kumar, Hung Le, Nadym Mallek, and Ziena Zeif. Optimal padded decomposition for bounded treewidth graphs. *arXiv preprint arXiv:2407.12230*, 2024. doi:10.48550/arXiv.2407.12230.
- 8 Lester Randolph Ford and Delbert R Fulkerson. Maximal flow through a network. *Canadian journal of Mathematics*, 8:399–404, 1956.
- 9 Tobias Friedrich, Davis Issac, Nikhil Kumar, Nadym Mallek, and Ziena Zeif. Approximate max-flow min-multicut theorem for graphs of bounded treewidth. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, pages 1325–1334, 2023. doi:10.1145/3564246.3585150.

- 10 Naveen Garg, Vijay V Vazirani, and Mihalis Yannakakis. Approximate max-flow min-(multi) cut theorems and their applications. *SIAM Journal on Computing*, 25(2):235–251, 1996. doi:10.1137/S0097539793243016.
- 11 Naveen Garg, Vijay V. Vazirani, and Mihalis Yannakakis. Primal-dual approximation algorithms for integral flow and multicut in trees. *Algorithmica*, 18(1):3–20, 1997. doi:10.1007/BF02523685.
- 12 Anupam Gupta, Ilan Newman, Yuri Rabinovich, and Alistair Sinclair. Cuts, trees and l_1 -embeddings of graphs. *Combinatorica*, 24(2):233–269, 2004.
- 13 T Chiang Hu. Multi-commodity network flows. *Operations research*, 11(3):344–360, 1963.
- 14 Philip Klein, Serge A Plotkin, and Satish Rao. Excluded minors, network decomposition, and multicommodity flow. In *Proceedings of the twenty-fifth annual ACM symposium on Theory of computing*, pages 682–690. ACM, 1993. doi:10.1145/167088.167261.
- 15 James R Lee and Prasad Raghavendra. Coarse differentiation and multi-flows in planar graphs. *Discrete & Computational Geometry*, 43(2):346–362, 2010. doi:10.1007/S00454-009-9172-4.
- 16 Nathan Linial, Eran London, and Yuri Rabinovich. The geometry of graphs and some of its algorithmic applications. *Combinatorica*, 15(2):215–245, June 1995. doi:10.1007/BF01200757.
- 17 Satish Rao. Small distortion and volume preserving embeddings for planar and euclidean metrics. In *Proceedings of the Fifteenth Annual Symposium on Computational Geometry*, pages 300–306. ACM, 1999. doi:10.1145/304893.304983.
- 18 Schrijver. *Theory of Linear and Integer Programming*. Wiley, 1998.
- 19 Eva Tardos and Vijay V Vazirani. Improved bounds for the max-flow min-multicut ratio for planar and $K_{r,r}$ -free graphs. *Information Processing Letters*, 47(2):77–80, 1993. doi:10.1016/0020-0190(93)90228-2.

A Proof of Theorem 3

Proof. Suppose the statement does not hold. Then the following linear system is infeasible.

$$\begin{aligned} \sum_{F \in \mathcal{F}} y_F &= 1 \\ \sum_{F \in \mathcal{F}, e \in F} y_F &\leq \alpha \cdot x(e) \quad \forall e \in E(G) \\ y_F &\geq 0. \end{aligned}$$

This implies that the following system is infeasible as well. The reason is that if the system below is feasible, then we can scale down the feasible solution appropriately and obtain a feasible solution for the system above.

$$\sum_{F \in \mathcal{F}} y_F \geq 1 \tag{2}$$

$$\sum_{F \in \mathcal{F}, e \in F} y_F \leq \alpha \cdot x(e) \quad \forall e \in E(G) \tag{3}$$

$$y_F \geq 0. \tag{4}$$

By reversing the inequality 2, we obtain that the following system is also infeasible,

$$\begin{aligned} \sum_{F \in \mathcal{F}} -y_F &\leq -1 \\ \sum_{F \in \mathcal{F}, e \in F} y_F &\leq \alpha \cdot x(e) \quad \forall e \in E(G) \\ y_F &\geq 0. \end{aligned}$$

Now, we use the following variant of Farkas Lemma (see [18] for a proof).

14:18 Improved Lower Bounds on Multiflow-Multicut Gaps

► **Lemma 18.** $\{x \in \mathbb{R}^n \mid Ax \leq b, x \geq 0\} = \emptyset$ iff there exists a vector u such that $A^T u \geq 0, u \geq 0$ and $b^T u < 0$.

For a feasible multicut F , let $\chi_F \in \{0, 1\}^E$ denote its indicator vector. By Lemma 18, there exists $u \geq 0$ and $c \geq 0$ such that $c^T \chi_F - u \geq 0$ for all $F \in \mathcal{F}$ and $-u + \alpha \cdot c^T x < 0$. This means that $c^T \chi_F \geq u$ for all $F \in \mathcal{F}$, and $\alpha \cdot c^T x < u$. Thus, $c^T \chi_F \geq u > \alpha \cdot c^T x$ for all $F \in \mathcal{F}$. Therefore with respect to the cost function c , $\text{OPT}_{LP} \leq \frac{u}{\alpha}$ and $\text{OPT}_{IP} > u$. This implies that integrality gap of the multicut instance M is $> \alpha$, a contradiction. ◀

B Projection of p -load Distribution

▷ **Claim 19.** Let G be a graph and let x be a p -load distribution of G . Let H be a subgraph of G , and for any $F \in \mathcal{F}_{2w}(H)$, let

$$y_F = \sum_{F': F' \in \mathcal{F}_{2w}(G), F' \cap E(H) = F} x_{F'}.$$

y is a p -load distribution for H .

Proof. For an edge $e \in E(H)$, we have:

$$\sum_{e \in F} y_F = \sum_{e \in F} \sum_{F' \in \mathcal{F}_{2w}(G), F' \cap E(H) = F} x_{F'} = \sum_{e \in F, F' \in \mathcal{F}_{2w}(G)} x_{F'} \leq p.$$

Finally, note that $y \geq 0$, and y forms a distribution:

$$\sum_{F \in \mathcal{F}_{2w}(H)} y_F = \sum_{F \in \mathcal{F}_{2w}(H)} \sum_{F' \cap E(H) = F, F' \in \mathcal{F}_{2w}(G)} x_{F'} = \sum_{F' \in \mathcal{F}_{2w}(G)} x_{F'} = 1. \quad \triangleleft$$