Multicommodity Flows and Cuts in Polymatroidal Networks*

Chandra Chekuri† Sreeram Kannan‡ Adnan Raja§ Pramod Viswanath¶

University of Illinois, Urbana-Champaign, IL 61801
{chekuri,kannan1,araja2,pramodv}@illinois.edu

January 16, 2013

Abstract

We consider multicommodity flow and cut problems in *polymatroidal* networks where there are submodular capacity constraints on the edges incident to a node. Polymatroidal networks were introduced by Lawler and Martel [36] and Hassin [29] in the single-commodity setting and are closely related to the submodular flow model of Edmonds and Giles [19]; the well-known maxflow-mincut theorem holds in this more general setting. Polymatroidal networks for the multicommodity case have not, as far as the authors are aware, been previously explored. Our work is primarily motivated by applications to information flow in wireless networks. We also consider the notion of undirected polymatroidal networks and observe that they provide a natural way to generalize flows and cuts in edge and node capacitated undirected networks.

We establish flow-cut gap results in several scenarios that have been previously considered in the standard network flow models where capacities are on the edges or nodes [40, 41, 24, 35, 34, 23]. These results have found applications in wireless network information flow [32, 33] and we anticipate others in the future. Our results are based on analyzing the dual of the flow relaxations via continuous extensions of submodular functions, in particular, the Lovász extension. For directed graphs we rely on a simple yet useful reduction from polymatroidal networks to standard networks. For undirected graphs we rely on the interplay between the Lovász extension of a submodular function and line embeddings with low average distortion introduced by Matousek and Rabinovich [44]; this connection is inspired by, and generalizes, the work of Feige, Hajiaghayi and Lee [23] on node-capacitated multicommodity flows and cuts.

1 Introduction

Consider a communication network represented by a directed graph $G = (V, E)$. In the so-called edge-capacitated scenario, each edge e has an associated capacity $c(e)$ that limits the information flowing on it. We consider a more general network model called the *polymatroidal network* introduced by Lawler and Martel [36] and independently by Hassin [29]. This model is closely related to the submodular flow model introduced by Edmonds and Giles [19]. Both models capture as special cases, single-commodity s-t flows in edge-capacitated directed networks, and polymatroid intersection, hence their importance. Moreover the models are known to be equivalent (see Chapter 60 in [52], in particular Section 60.3b). The polymatroidal network flow model is more directly and intuitively related to standard network flows and one can easily generalize it to the multicommodity setting which is the focus in this paper.

*A preliminary version of this paper appeared in *Proc. of Third Innovations in Theoretical Computer Science (ITCS), Jan 2012.*
†Supported in part by NSF grant CCF-1016684.
‡Supported in part by NSF grants CCF 1017430 and CNS 0721652.
§Supported in part by NSF grants CCF 1017430 and CNS 0721652.
¶Supported in part by NSF grants CCF 1017430 and CNS 0721652.

The polymatroidal network flow model differs from the standard network flow model in the following way. Consider a node v in a directed graph G and let $\delta_G^-(v)$ be the set of edges in to v and $\delta_G^+(v)$ be the set of edges out of v. In the standard model each edge (u,v) has a non-negative capacity $c(u,v)$ that is independent of other edges. In the polymatroidal network for each node v there are two associated submodular functions (in fact polymatroids[1]) ρ_v^- and ρ_v^+ which impose joint capacity constraints on the edges in $\delta_G^-(v)$ and $\delta_G^+(v)$ respectively. That is, for any set of edges $S \subseteq \delta_G^-(v)$, the total capacity available on the edges in S is constrained to be at most $\rho_v^-(S)$, similarly for $\delta_G^+(v)$. Note that an edge (u,v) is influenced by ρ_u^+ and ρ_v^-. Lawler and Martel considered the problem of finding a maximum s-t flow in this model. The results in [36, 29] show that various important properties that hold for s-t flows in standard networks generalize to polymatroid networks; these include the classical maxflow-mincut theorem of Ford and Fulkerson (and Menger) and the existence of an integer valued maximum flow when capacities are integral.

The original motivation for the Lawler-Martel model came from an application to a scheduling problem [43]. More recently, there have been several applications of polymatroid network flows, (and submodular flows) and their generalizations such as linking systems [53], to information flow in wireless networks [1, 4, 56, 26, 49, 32]. Our main motivation comes from the study of wireless networks. A node in a wireless network communicates with several nodes over a broadcast medium and hence the channels interfere with each other; this imposes joint capacity constraints on the channels. Several interference scenarios of interest can be modeled by submodular functions.

Most of the work on this topic so far has focused on the case of a single source. In this paper we consider *multicommodity* flows and cuts in polymatroidal networks where several source-sink pairs (s_1,t_1), $(s_2,t_2),\ldots,(s_k,t_k)$ share the capacity of the network. In the communications literature this is referred to as the multiple unicast setting. Our primary motivation is applications to (wireless) network information flow; see companion papers [32, 33] that build on results of this paper. Another motivation is to understand the extent to which techniques and results that were developed for multicommodity flows and cuts in standard networks generalize to polymatroidal networks. We note that polymatroidal networks allow for a common treatment of edge and node capacities; an advantage is that one can define cuts with respect to edge removals while the cost is based on nodes. As far as we are aware, multicommodity flows and cuts in polymatroidal networks have not been studied previously.

Flow-cut gaps in polymatroidal networks: The maxflow-mincut theorem for single commodity flows does not generalize to the multicommodity case even when the number of source-sink pairs is is three or more (two or more in case of directed graphs). See [52] for some special cases where min-max results do hold. Cuts typically upper bound the corresponding flows (in terms of value); the worst-case ratio between the two is referred to as the flow-cut gap. Obtaining tight bounds on flow-cut gaps has been an active and fruitful area of research in theoretical computer science, starting with the seminal work of Leighton and Rao [40]. The initial motivation was approximation algorithms for NP-Hard cut and separator problems. There has been much subsequent work with a tight bound of $O(\log k)$ established for flow-cut gaps in undirected graphs in a variety of settings [24, 41, 7, 23]. It has also been shown that strong lower bounds exist for flow-cut gaps in directed graphs; for instance the gap is $O(\min\{k, n^\delta\})$ between the maximum concurrent flow and the sparsest cut [51, 18] where δ is a fixed constant. However, poly-logarithmic upper bounds on the gaps are known for the case of symmetric demands in directed graphs [35, 20].

The focus of this paper is understanding multicommodity flow-cut gaps in polymatroidal networks. In communication networks, cuts can be used to information theoretically upper-bound achievable rates, while flows allow one to develop lower bounds on achievable rates by combining a variety of routing and coding schemes. Flow-cut gaps are of therefore of much interest in understanding the capacity of communication

[1]A set function $f: 2^N \to \mathbb{R}$ over a finite ground set N is submodular iff $f(A)+f(B) \geq f(A\cap B)+f(A\cup B)$ for all $A,B \subseteq N$; equivalently $f(A\cup\{i\})-f(A) \geq f(B\cup\{i\})-f(B)$ for all $A \subset B$ and $i \notin A$. It is monotone if $f(A) \leq f(B)$ for all $A \subset B$. In this paper a polymatroid refers to a non-negative monotone submodular function with $f(\emptyset)=0$.

networks. We show that several of the flow-cut gap results that have been established in standard networks can be extended to polymatroid networks. In addition to applications to wireless networks [32, 33], our results lead to approximation algorithms for cut problems in polymatroidal networks which could have applications.

Bidirected and undirected polymatroidal networks: As we mentioned already, strong lower bounds exist on flow-cut gaps for directed networks. Positive results in the form of poly-logarithmic upper bounds on flow-cut gaps for standard networks hold when the demands are symmetric or when the supply graph is undirected. A natural model for wireless networks is the *bidirected* polymatroidal network. For two nodes u and v in a wireless network, it is a reasonable approximation to assume that the channel from u to v is similar to that from v to u; hence one can assume that the underlying graph G is bidirected in that if the edge (u, v) is present then so is (v, u). Moreover, we assume that for any node v and $S \subseteq \delta^-(v)$, $\rho_v^-(S) = \rho_v^+(S')$ where $S' \subseteq \delta^+(v)$ is the set of edges that correspond to the reverse of the edges in S. Within a factor of 2, bidirected polymatroidal networks can be approximated by *undirected* polymatroidal networks: we have an undirected graph G and for each node v a single polymatroid ρ_v that constrains the capacity of the edges $\delta_G(v)$, the set of edges incident to v. The main advantage of undirected polymatroid networks is that we can use existing tools and ideas from metric embeddings to understand flow-cut gap results. Undirected polymatroidal networks have not been considered previously. We observe that they allow a natural way to capture both edge and node-capacitated flows in undirected graphs. To capture node-capacitated flows[2] we set $\rho_v(S) = 2c(v)$ for all $\emptyset \neq S \subseteq \delta(v)$ where $c(v)$ is the capacity of v. We mention an advantage of using polymatroidal networks even when considering the special case of node-capacitated flows and cuts: one can define cuts with respect to edges even though the cost is on the nodes. This is in fact quite natural and simplifies certain aspects of the algorithms in [23].

1.1 Overview of results

We do a systematic study of flow-cut gaps in multicommodity polymatroidal networks, both directed and undirected. Let $G = (V, E)$ be a polymatroidal network on n nodes with k source-sink pairs $(s_1, t_1), \ldots, (s_k, t_k)$. We consider two flow problems and their corresponding cut problems: (i) maximum throughput flow and multicut (ii) maximum concurrent flow and sparsest cut. Formal definitions of these terms can be found in Section 2. The main bounds we obtain are summarized below, and in Table 1.

- For directed networks we show a reduction based on the dual that establishes a correspondence between flow-cut gaps in polymatroidal networks and the standard edge-capacitated networks. This allows us to obtain poly-logarithmic upper bounds for flow-cut gaps in directed polymatroidal networks with *symmetric* demands via results in [35, 20]. In particular we obtain an $O(\min\{\log^3 k, \log^2 n \log\log n\})$ gap between the maximum concurrent flow and sparsest cut. The reduction is applicable only to directed graphs.
- We show that line embeddings with low average distortion [44, 48] lead to upper bounds on flow-cut gaps in polymatroidal networks — this connection is inspired by the work in [23] for node-capacitated flows. For undirected polymatroidal networks this leads to an optimal $O(\log k)$ gap between maximum concurrent flow and sparsest cut. We also obtain an optimal $O(\log k)$ gap between throughput flow and multicut. These imply corresponding results for bidirected networks.
- We consider polymatroidal networks that exclude a fixed graph K_h as minor (this includes planar graphs). We show an $O(h^2)$ gap between the maximum throughput flow and minimum multicut for

[2] The factor of 2 is needed since a flow path p through an internal node v uses two edges. On the other hand, it is not needed for the sources and sinks. This technical issue is a minor inconvenience with undirected polymatroidal networks; we note that this also arises in treating node-capacitated multicommodity flows [23].

Setting	Max. Concurrent Flow / Sparsest Cut Gap	Max. Throughput Flow / Multicut Gap
Undirected polymatroidal network	$\Theta(\log k)$	$\Theta(\log k)$
Directed polymatroidal network (symmetric demands)	$O(\min\{\log^2 k, \log n \log\log n\})$	$O(\min\{\log^3 k, \log^2 n \log\log n\})$
Planar undirected polymatroidal network		$\Theta(1)$

Table 1: Summary of Results

these networks. As a corollary, we obtain a constant factor approximation for *node-weighted* multicut in such graphs. Our result is based on interpreting the network decomposition theorem in [34] as a line embedding.

Most of the literature on multicommodity flow-cut gaps is based on analyzing the dual of the linear program for the flow. The dual linear program can be viewed as a fractional relaxation for the corresponding cut problem. A flow-cut gap is established by showing the existence of a cut within some factor of this relaxation. In standard edge-capacitated networks, the dual linear program has length variables on the edges which induce distances on the nodes. The situation is more involved in polymatroidal networks, in particular, the definition of the cost of a cut is somewhat complex and is discussed in more detail in Section 2.2. Our starting point is the use of the Lovász extension of a submodular function [42] to cleanly rewrite the dual of the flow linear programs. This simplifies the constraint structure of the dual at the expense of making the objective a convex function. However, we are able to exploit properties of the Lovász extension in several ways to obtain our results. Our techniques give two new dual-based proofs of the maxflow-mincut theorem for single commodity polymatroid networks that was first established by Lawler and Martel (also Hassin [29]) algorithmically [36] via an augmenting path based approach. We believe that the applicability of embedding based methods for polymatroidal networks is of independent mathematical interest. For the most part we ignore algorithmic issues in this paper although all the flow-cut gap results lead to polynomial-time algorithms for finding approximate cuts.

1.2 Related Work

We have already mentioned several relevant results on multicommodity flows and cuts. We refer the reader to an article by Shmoys [54], and some more recent papers [18, 23, 39]. For several cut problems, approximation algorithms that improve over the flow-cut gap bounds have been obtained via semi-definite programming based relaxations, starting with the seminal work of Arora, Rao and Vazirani [5] — see [6, 3, 23].

Schrijver [52] has extensive treatment of classical results on submodular functions in combinatorial optimization; the equivalence between the submodular flow model of Edmonds and Giles [19] and the polymatroid network flow model of Lawler and Martel [36] can be found there. Federgruen and Groenevelt [22] consider a slight generalization of the Lawler-Martel model to single-source and multiple-sinks which can be reduced to the single commodity case relatively easily. As we already remarked, the multicommodity flows in polymatroidal networks do not appear to have been considered previously. There has been a resurgence of recent interest in submodular functions and their applications. Continuous extension based approaches to optimizing with submodular objectives, for minimization via the Lovász-extension [42], and for maximization via the multilinear extension [10], have led to several new algorithmic results [11, 16, 30, 25, 14, 15]. Our work here demonstrates another application of this approach.

1.3 Organization

The rest of this paper is organized as follows. Formal definitions of multicommodity flows and cuts in polymatroidal networks are described in Section 2. Section 3 describes the convex programming relaxations for cut problems that are equivalent to the dual of the linear programs for the corresponding flow problems. These relaxations are exploited in Section 4 to show flow-cut gaps for directed polymatroidal networks by using a reduction from the polymatroidal network problem to the standard network problem. In Section 5, flow-cut gap bounds are shown for undirected polymatroidal networks via line embeddings. Section 5.3 describes an $O(1)$ bound on the gap between multicut and throughput flow in planar and minor-free graphs.

2 Multicommodity Flows and Cuts in Polymatroidal Networks

We let $G = (V, E)$ represent a graph whether directed or undirected. We use (u, v) for an ordered pair of nodes and uv to denote an unordered pair. In a directed graph G, for a given node v, $\delta^-_G(v)$ and $\delta^+_G(v)$ denote the set of incoming and outgoing edges at v. In undirected graphs we use $\delta_G(v)$ to denote the set of edges incident to v. We omit the subscript G if it is clear from the context. We are interested in multicommodity flows and cuts. In addition to the graph, the input consists of a set of k source-sink pairs $(s_1, t_1), \dots, (s_k, t_k)$ that wish to communicate independently and share the network capacity.

In a directed polymatroidal network, each node $v \in V$ has two associated polymatroids ρ^-_v and ρ^+_v with ground sets as $\delta^-(v)$ and $\delta^+(v)$ respectively. These functions constrain the joint capacity on the edges incident to v as follows. If $S \subseteq \delta^-(v)$, then $\rho^-_v(S)$ upper-bounds the total capacity of the edges in S; similarly, if $S \subseteq \delta^+(v)$, then $\rho^+_v(S)$ upper-bounds the total capacity of the edges in S. We assume that the functions $\rho^-_v(\cdot), \rho^+_v(\cdot), \quad v \in V$, are provided via value oracles. In undirected polymatroidal graphs we have a single function $\rho_v(\cdot)$ at a node v that constrains the capacity of the edges incident to v. Continuous extensions of submodular functions, namely the Lovász extension [42] and the convex closure, are important technical tools in interpreting and analyzing the duals of the linear programs for multicommodity flow in the polymatroid setting. We discuss these in Section 2.2. We first discuss the two flow problems of interest, namely maximum throughput flow and the maximum concurrent flow.

2.1 Flows

A multicommodity flow for a given collection of k source-sink pairs $(s_1, t_1), \dots, (s_k, t_k)$ consists of k separate single-commodity flows, one for each pair (s_i, t_i). The flow for the i'th commodity can either be viewed as an edge-based flow $f_i : E \to \mathbb{R}_+$, or as a path-based flow $f_i : \mathcal{P}_i \to \mathbb{R}_+$, where $\mathcal{P}_i$ is the set of all simple paths between s_i and t_i in G. We use path-based flows since they are allow us to treat directed and undirected graphs in a unified fashion, and also for writing the linear programs for flows and cuts in a convenient way. However, it is easier to argue polynomial-time solvability of the linear programs via edge-based flows. Given path-based flows f_i, $i = 1, \dots, k$ for the k source-sink pairs, the total flow on an edge e is defined as $f(e) = \sum_{i=1}^{k} \sum_{p \in \mathcal{P}_i} f_i(p)$. The total flow for commodity i is $R_i = \sum_{p \in \mathcal{P}_i} f_i(p)$, where R_i is interpreted as the rate of commodity flow i. In *directed* polymatroidal networks, the flow is constrained to satisfy the following capacity constraints.

$$\sum_{e \in S} f(e) \quad \leq \quad \rho_v^-(S) \quad \forall S \subseteq \delta^-(v) \tag{1}$$

$$\sum_{e \in S} f(e) \quad \leq \quad \rho_v^+(S) \quad \forall S \subseteq \delta^+(v) \tag{2}$$

The constraints in *undirected* polymatroidal networks are:

$$\sum_{e \in S} f(e) \leq \rho_v(S) \quad \forall v \, \forall S \subseteq \delta(v). \tag{3}$$

A rate tuple $(R_1, ..., R_k)$ is said to be *achievable* if commodities $1, \ldots, k$ can be sent at rates $R_1, \ldots, R_k$ simultaneously between the corresponding source-sink pairs. For a given polymatroidal network and source-sink pairs the set of achievable rate tuples is easily seen from the above constraints to be a polyhedral set. We let $P(G, \mathcal{T})$ denote this rate region where G is the network and $\mathcal{T}$ is the set of given source-sink pairs. In the *maximum throughput multicommodity flow* problem the goal is to maximize $\sum_{i=1}^{k} R_i$ over $P(G, \mathcal{T})$. In the *maximum concurrent multicommodity flow* problem each source-sink pair has an associated demand D_i and the goal is to maximize λ such that the rate tuple $(\lambda D_1, ..., \lambda D_k)$ is achievable, that is the tuple belongs to $P(G, \mathcal{T})$. It is easy to see that both these problems can be cast as linear programming problems. The path-formulation results in an exponential (in n the number of nodes of G) number of variables and we also have an exponential number of constraints due to the polymatroid constraints at each node. However, one can use an edge-based formulation and solve the linear programs in polynomial time via the ellipsoid method and polynomial-time algorithms for submodular function minimization.

Networks with symmetric demands: In directed polymatroidal networks we are primarily interested in *symmetric demands*: node s_i intends to communicate with t_i and node t_i intends to communicate with s_i at the same *rate*. Conceptually one can reduce this to the general setting by having two commodities (s_i, t_i) and (t_i, s_i) for a pair $s_i t_i$ and adding a constraint that ensures their rates are equal. To be technically consistent with previous work we do the following. We will assume that we are given k unordered source-sink pairs $s_1t_1, \ldots, s_kt_k$. Now consider the $2k$ ordered pairs $(s_1, t_1), \ldots, (s_k, t_k), (t_1, s_1), \ldots, (t_k, s_k)$. We are interested in achievable rate tuples of the form $(R_1, \ldots, R_k, R'_1, \ldots, R'_k)$ where $R'_i = R_i$. In the maximum throughput setting we maximize $\sum_{i=1}^{k}(R_i + R'_i)$. Note that even though the rates for (s_i, t_i) and (t_i, s_i) are the same, the flow paths along which they route can be different. In the maximum concurrent flow setting both (s_i, t_i) and (t_i, s_i) have a common demand D_i and we find the maximum λ such that rate tuple $(\lambda D_1, ..., \lambda D_k, \lambda D_1, ..., \lambda D_k)$ is achievable for the pairs $(s_1, t_1), \ldots, (s_k, t_k), (t_1, s_1), \ldots, (t_k, s_k)$.

2.2 Cuts

The multicommodity flow problems have natural dual cut problems associated with them. Given a graph $G = (V, E)$ and a set of edges $F \subseteq E$ we say that the ordered node pair (s, t) is separated by F if there is no path from s to t in the graph $G[E \setminus F]$. In directed graphs F may separate (s, t) but not (t, s). In undirected graphs we say that F separates the unordered node pair st if s and t are in different connected components of $G[E \setminus F]$. In the standard network model the cost of a cut defined by a set of edges F is simply $\sum_{e \in F} c(e)$ where $c(e)$ is the cost of e (capacity in the primal flow network) . In polymatroid networks the cost of F is defined in a more involved fashion. Each edge (u, v) in F is assigned to either u or v; we say that an assignment of edges to nodes $g : F \to V$ is *valid* if it satisfies this restriction. A valid assignment partitions F into sets $\{g^{-1}(v) \mid v \in V\}$ where $g^{-1}(v)$ (the pre-image of v) is the set of edges in F assigned to v by g.

For a given valid assignment g of F the cost of the cut $\nu_g(F)$ is defined as

$$\nu_g(F) := \sum_v \left(\rho_v^-(\delta^-(v) \cap g^{-1}(v)) + \rho_v^+(\delta^+(v) \cap g^{-1}(v)) \right).$$

In undirected graphs the cost for a given assignment is $\sum_v \rho_v(g^{-1}(v))$.

Given a set of edges F we define its cost to be the minimum over all possible valid assignments of F to nodes, the expression for the cost as above. We give a formal definition below.

Definition 1. Cost of edge cut: *Given a directed polymatroid network $G = (V, E)$ and a set of edges $F \subseteq E$, its cost denoted by $\nu(F)$ is*

$$\min_{g:F \to V,\, g \text{ valid}} \sum_v \left(\rho_v^-(\delta^-(v) \cap g^{-1}(v)) + \rho_v^+(\delta^+(v) \cap g^{-1}(v)) \right). \tag{4}$$

In an undirected polymatroid network $\nu(F)$ is

$$\min_{g:F \to V,\, g \text{ valid}} \sum_v \rho_v(g^{-1}(v)). \tag{5}$$

Lemma 1. *Consider a feasible multicommodity flow in a polymatroidal network $G = (V, E)$ where $f(e)$ is the total flow value on edge e over all commodities. Then $\sum_e f(e) \leq \nu(F)$.*

Proof. We consider directed graphs, the argument is similar for undirected graphs. Fix any valid assignment $g : F \to V$. The constraints (1) and (2) that apply to any feasible flow imply that

$$\begin{aligned} \sum_{e \in F} f(e) &= \sum_v \sum_{e \in g^{-1}(v)} f(e) \\ &\leq \sum_v \left(\rho_v^-(\delta^-(v) \cap g^{-1}(v)) + \rho_v^+(\delta^+(v) \cap g^{-1}(v)) \right). \end{aligned}$$

Since the above applies to any valid assignment g, the claim follows. □

The lemma below easily follows from sub-additivity of non-negative submodular functions and Definition 1.

Lemma 2. *The cut cost function is sub-additive, that is, $\nu(F \cup F') \leq \nu(F) + \nu(F')$ for all $F, F' \subseteq E$.*

We now define the two cuts problems of interest.

Definition 2. *Given a collection of source-sink pairs $(s_1, t_1), \ldots, (s_k, t_k)$ in $G = (V, E)$ and associated demand values $D_1, \ldots, D_k$, and a set of edges $F \subseteq E$ the demand separated by F, denoted by $D(F)$, is $\sum_{i:(s_i,t_i) \text{ separated by } F} D_i$. F is a* multicut *if all the given source-sink pairs are separated by F. The* sparsity *of F is defined as $\frac{\nu(F)}{D(F)}$.*

The above definitions extend naturally to undirected graphs. Given the above definitions two natural optimization problems that arise are the following. The first is to find a multicut of minimum cost for a given collection of source-sink pairs. The second is to find a cut of minimum sparsity. These problems are NP-hard even in edge-capacitated undirected graphs and have been extensively studied from an approximation point of view [40, 24, 41, 7, 5, 2]. The lemma follows easily from Lemma 1.

Lemma 3. *Given a multicommodity polymatroidal network instance, the value of the maximum throughput flow is at most the cost of a minimum multicut. The value of the maximum concurrent flow is at most the minimum sparsity.*

Networks with symmetric demands: For a directed network with symmetric demands the notion of a "cut" has to be defined appropriately. We say that a set of edges F separates a pair $s_i t_i$ if it separates (s_i, t_i) *or* (t_i, s_i). With this notion of separation, the definitions of multicut and sparsest cut extend naturally. A multicut is a set of edges F whose removal separates all the given pairs. Similarly for a set of edges F its sparsity is defined to $\nu(F)/D(F)$ where $D(F)$ is the total demand of pairs separated; note that if both (s_i, t_i) and (t_i, s_i) are separated by F we count D_i twice in $D(F)$. This is to be consistent with the definition of flows given earlier. Lemma 3 extends to the symmetric demand case with the definition of flows given for symmetric demands in the previous section.

A key question of interest is to quantify the relative gap between the cut and flow values. These gaps are relatively well-understood in standard networks and the goal of this paper is to obtain results for polymatroid networks.

3 Relaxations for Cuts

Lemma 3 gives a way to lower-bound the value of multicut and sparsest cut via corresponding flow problems. The flow problems can be cast as linear programs. The duals of these linear programs can be directly interpreted as linear programming relaxations for integer programming formulations for the cut problems. Here we take the approach of writing the formulation with a convex objective function and linear constraints; this simplifies and clarifies the constraints and aids in the analysis. For one of the cases we show the equivalence of the formulation with the dual of the corresponding flow linear program. We first discuss continuous extensions of submodular functions.

3.1 Continuous extensions of submodular functions

Given a submodular set function $\rho : 2^N \to \mathbb{R}$ on a finite ground set N, it is useful to *extend* it to a function $\rho' : [0,1]^N \to \mathbb{R}$ defined over the cube in $|N|$ dimensions. That is, we wish to assign a value for each $\mathbf{x} \in [0,1]^N$ such that $\rho'(\mathbf{1}_S) = \rho(S)$ for all $S \subseteq N$ where $\mathbf{1}_S$ is the characteristic vector of the set S. For minimizing submodular functions a natural goal is to find an extension that is convex. We describe two extensions below.

Convex closure: For a set function $\rho : 2^N \to \mathbb{R}$ (not necessarily submodular) its convex closure is a function $\tilde{\rho} : [0,1]^N \to \mathbb{R}$ with $\tilde{\rho}(\mathbf{x})$ defined as the optimum value of the following linear program:

$$\begin{aligned}
\tilde{\rho}(\mathbf{x}) &= \min \sum_{S \subseteq N} \alpha_S \rho(S) \\
\text{s.t.} & \\
\sum_S \alpha_S &= 1 \\
\sum_{S : i \in S} \alpha_S &= x_i \qquad \forall i \in N \\
\alpha_S &\geq 0 \qquad \forall S.
\end{aligned}$$

The function $\tilde{\rho}$ is convex for any ρ. Moreover, when ρ is submodular, for any given $\mathbf{x}$, the linear program above can be solved in polynomial time via submodular function minimization and hence $\tilde{\rho}(\mathbf{x})$ can be computed in polynomial time (assuming a value oracle for ρ). It is known and not difficult to show that if ρ is a polymatroid (monotone and $f(\emptyset) = 0$) the value of the linear program does not change if we drop the constraint that $\sum_S \alpha_S = 1$.

Lovász extension: For a set function $\rho : 2^N \to \mathbb{R}$ (not necessarily submodular) its Lovász extension [42] denoted by $\hat{\rho} : [0,1]^N \to \mathbb{R}$ is defined as follows:

$$\hat{\rho}(\mathbf{x}) = \int_0^1 \rho(\mathbf{x}^\theta) d\theta$$

, where $\mathbf{x}^\theta = \{i \mid x_i \geq \theta\}$. This is not the standard way the Lovász extension is stated but is entirely equivalent to it. The standard definition is the following. Given $\mathbf{x}$ let $i_1, \ldots, i_n$ be a permutation of $\{1, 2, \ldots, n\}$ such that $x_{i_1} \geq x_{i_2} \geq \ldots \geq x_{i_n} \geq 0$. For ease of notation define $x_0 = 1$ and $x_{n+1} = 0$. For $1 \leq j \leq n$ let $S_j = \{i_1, i_2, \ldots, i_j\}$. Then

$$\hat{\rho}(\mathbf{x}) = (1 - x_{i_1})\rho(\emptyset) + \sum_{j=1}^{n} (x_{i_j} - x_{i_{j+1}})\rho(S_j).$$

It is typical to assume that $\rho(\emptyset) = 0$ and omit the first term in the right hand side of the preceding equation. Note that it is easy to evaluate $\hat{\rho}(\mathbf{x})$ given a value oracle for ρ.

We state some well-known facts.

Lemma 4. *For a submodular set function ρ, $\tilde{\rho}(\mathbf{x}) = \hat{\rho}(\mathbf{x})$ for any $\mathbf{x} \in [0,1]^N$. Therefore the convex closure coincides with the Lovász extension and $\hat{\rho}(\cdot)$ is convex.*

Proposition 1. *For a monotone submodular function ρ and $\mathbf{x} \leq \mathbf{x}'$ (coordinate-wise), $\hat{\rho}(\mathbf{x}) \leq \hat{\rho}(\mathbf{x}')$.*

The equivalence of $\tilde{\rho}$ and $\hat{\rho}$ also implies that an optimum solution to the linear program defining $\tilde{\rho}(\mathbf{x})$ is obtained by a solution $\bar{\alpha}$ where the support of $\bar{\alpha}$ is a chain on N (a laminar family whose tree representation is a path). In fact we have the following. Given $\mathbf{x} \in [0,1]^N$ consider the ordering of the coordinates and the associated sets as in the definition of the $\hat{\rho}(\mathbf{x})$. One can verify that $\alpha_{S_j} = x_{i_j} - x_{i_{j-1}}$ for $1 \leq j \leq n$, $\alpha_\emptyset = (1 - x_{i_n})$, and $\alpha_S = 0$ for all other sets S is an optimum solution to the linear program that defines $\tilde{\rho}(\mathbf{x})$. We will use this fact later.

3.2 Multicut

We now consider the multicut problem. Recall that we wish to find a subset $F \subseteq E$ such that F separates all the given source-sink pairs so as to minimize the cost $\nu(F)$. The only difference between the polymatroid networks and standard networks is in the definition of the cost. We first focus on expressing the constraint that F is a feasible set for separating the pairs. For each edge e we have a variable $\ell(e) \in [0,1]$ in the relaxation that represents whether e is cut or not. For feasibility of the cut we have the condition that for any path p from s_i to t_i (that is $p \in \mathcal{P}_i$) at least one edge in p is cut; in the relaxation this corresponds to the constraint that $\sum_{e \in p} \ell(e) \geq 1$. In other words $\text{dist}_\ell(s_i, t_i) \geq 1$ where $\text{dist}_\ell(u, v)$ is the distance between u and v with edge lengths given by $\ell(e)$ values.

We now consider the cost of the cut. Note that $\nu(F)$ is defined by valid assignments of F to the nodes, and submodular costs on the nodes. In the relaxation we model this as follows. For an edge $e = (u, v)$ we have variables $\ell(e, u)$ and $\ell(e, v)$ which decide whether e is assigned to u or v. We have a constraint $\ell(e, u) + \ell(e, v) = \ell(e)$ to model the fact that if e is cut then it has to be assigned to either u or v. Now consider a node v and the edges in $\delta^+(v)$. The variables $\ell(e, v), e \in \delta^+(v)$ in the integer case give the set of edges $S \subseteq \delta^+(v)$ that are assigned to v and in that case we can use the function $\rho_v^+(S)$ to model the cost. However, in the fractional setting the variables lie in the real interval $[0,1]$ and here we use the extension approach to obtain a convex programming relaxation; we can rewrite the convex program as an equivalent linear program via the definition of $\tilde{\rho}$. Let $\mathbf{d}_v^-$ be the vector consisting of the variables $\ell(e, v)$, $e \in \delta^-(v)$ and similarly $\mathbf{d}_v^+$ denote the vector of variables $\ell(e, v)$, $e \in \delta^+(v)$. The relaxation for the directed case is

$$\begin{aligned}
&\min \sum_v (\hat{\rho}_v^-(\mathbf{d}_v^-) + \hat{\rho}_v^+(\mathbf{d}_v^+)) \\
&\ell(e,u) + \ell(e,v) = \ell(e) \qquad e = (u,v) \in E \\
&\text{dist}_\ell(s_i,t_i) \geq 1 \qquad 1 \leq i \leq k \\
&\ell(e), \ell(e,u), \ell(e,v) \geq 0 \qquad e = (u,v) \in E.
\end{aligned}$$

$$\begin{aligned}
&\min \sum_v \hat{\rho}_v(\mathbf{d}_v) \\
&\ell(e,u) + \ell(e,v) = \ell(e) \qquad e = uv \in E \\
&\text{dist}_\ell(s_i,t_i) \geq 1 \qquad 1 \leq i \leq k \\
&\ell(e), \ell(e,u), \ell(e,v) \geq 0 \qquad e = uv \in E.
\end{aligned}$$

Figure 1: Lovász-extension based relaxations for multicut in directed and undirected polymatroidal networks

$$\begin{aligned}
&\min \sum_v \hat{\rho}_v^-(\mathbf{d}_v^-) + \hat{\rho}_v^+(\mathbf{d}_v^+)) \\
&\ell(e,u) + \ell(e,v) = \ell(e) \qquad e = (u,v) \in E \\
&\sum_{i=1}^k D_i \cdot \text{dist}_\ell(s_i,t_i) = 1 \\
&\ell(e), \ell(e,u), \ell(e,v) \geq 0 \qquad e = (u,v) \in E.
\end{aligned}$$

$$\begin{aligned}
&\min \sum_v \hat{\rho}_v(\mathbf{d}_v) \\
&\ell(e,u) + \ell(e,v) = \ell(e) \qquad e = uv \in E \\
&\sum_{i=1}^k D_i \cdot \text{dist}_\ell(s_i,t_i) = 1 \\
&\ell(e), \ell(e,u), \ell(e,v) \geq 0 \qquad e = (u,v) \in E.
\end{aligned}$$

Figure 2: Relaxations for sparsest cut in directed and undirected polymatroidal networks

formally described in Fig 1 in the box on the left. For the symmetric demands case the relaxation is similar, but since we need to separate either (s_i, t_i) or (t_i, s_i) the constraint $\text{dist}_\ell(s_i, t_i) \geq 1$ is replaced by the constraint $\text{dist}_\ell(s_i, t_i) + \text{dist}_\ell(t_i, s_i) \geq 1$.

For the undirected case we let $\mathbf{d}_v$ denote the vector of variables $\ell(e, v), e \in \delta(v)$ and the resulting relaxation is shown on the right in Fig 1.

One can replace $\hat{\rho}_v$ in the above convex programming relaxations by $\tilde{\rho}_v$ the convex closure; further, one can use the definition of $\tilde{\rho}_v$ via a linear program to convert the convex program into an equivalent linear program. The resulting linear program can be shown to be equivalent to the dual of the maximum throughput flow problem. See Section A for a formal proof.

3.3 Sparsest cut

Now we consider the sparsest cut problem. In the sparsest cut problem we need to decide which pairs to disconnect and then ensure that we pick edges whose removal separates the chosen pairs. Moreover we are interested in the ratio of the cost of the cut to the demand separated. We follow the known formulation in the edge-capacitated case with the main difference, again, being in the cost of the cut. There is a variable y_i which determines whether pair i is separated or not. We again have the edge variables $\ell(e), \ell(e, u), \ell(e, v)$ to indicate whether $e = (u, v)$ is cut and whether e's cost is assigned to u or v. If pair i is to be separated to the extent of y_i we ensure that $\text{dist}_\ell(s_i, t_i) \geq y_i$. To express sparsity, which is defined as a ratio, we normalize the demand separated to be 1. Fig 2 has a formal description on the left for the directed case. For the symmetric demands case we have essentially the same relaxation; the constraint $\sum_i D_i \text{dist}_\ell(s_i, t_i) = 1$ is replaced by the constraint $\sum_i D_i (\text{dist}_\ell(s_i, t_i) + \text{dist}_\ell(t_i, s_i)) = 1$.

The relaxation for the undirected case is shown on the right in Fig 2 where $\mathbf{d}_v$ is the vector of variables $\ell(e, v), e \in \delta(v)$.

4 Flow-Cut Gaps in Directed Polymatroidal Networks

In this section we consider flow-cut gaps in directed polymatroidal networks. We show via a reduction that these gaps can be related to corresponding gaps in directed edge-capacitated networks that have been well-studied. We note that this reduction is specific to directed graphs and does not apply to undirected polymatroidal networks. The embedding based approach for the undirected case that we discuss in Section 5 is also applicable to directed graphs.

The reduction is similar at a high level for both gap questions of interest and is based on the relaxations for the two cut problems that we described in Section 3. We take a feasible fractional solution for relaxation of the cut problem in question and produce an instance of a cut problem in an edge-capacitated network and a feasible fractional solution to the corresponding cut problem. We also provide a correspondence between feasible integer solutions to the edge-capacitated network instance and the original problem such that the cost of the solution is preserved. These correspondences allow us to translate known gap results for the edge-capacitated networks to polymatroidal networks.

4.1 Details of the reduction

Let $G = (V, E)$ be a directed graph and let $\ell : E \to \mathbb{R}_+$ be a length function on the edges. We let $\text{dist}_\ell(u, v)$ be the shortest path distance from u to v in G with edge lengths ℓ. Moreover, for each edge (u, v) let $\ell(e, u)$ and $\ell(e, v)$ be two non-negative numbers such that $\ell(e) = \ell(e, u) + \ell(e, v)$. For a node v let $\mathbf{d}_v^+$ be the vector of $\ell(e, v)$ values for all edges $e \in \delta^+(v)$ and similarly $\mathbf{d}_v^-$ is the vector of $\ell(e, v)$ values for edges in $\delta^-(v)$. In the polymatroidal setting the cost induced by the edge length variables is given by $\sum_{v \in V}(\hat{\rho}^-(\mathbf{d}_v^-) + \hat{\rho}^+(\mathbf{d}_v^+))$. Note that for multicut we have that $\text{dist}_\ell(s_i, t_i) \geq 1$ for each demand pair (s_i, t_i) while in sparsest cut we are interested in the ratio of the cost to $\sum_i D_i \cdot \text{dist}_\ell(s_i, t_i)$. We now describe the construction of a graph $H = (V_H, E_H)$ where $V_H = V \uplus V'$ (that is the nodes of G are also in H) and an edge length function $\ell' : E_H \to \mathbb{R}_+$ such that $\text{dist}_\ell(u, v) = \text{dist}_{\ell'}(u, v)$ for all $u, v \in V$; that is the distances between nodes in V are the same in G and H. We also create an edge-cost (or capacity in the primal sense) function $c : E_H \to \mathbb{R}_+$. The construction will also establish the correspondence of cuts in G and H and their costs.

The graph $H = (V \uplus V', E_H)$ is constructed as follows. To aid the reader we first describe the idea of the construction at a high-level. Consider a node $v \in V$ and the in-coming edges $\delta^-(v)$ and out-going edges $\delta^+(v)$. In H we have nodes of V and build an in-tree T_v^- and an out-tree T_v^+ that are rooted at v. The leaves of T_v^- are the edges in $\delta^-(v)$ the leaves of T_v^+ are the edges in $\delta^+(v)$. Note that an edge (u, v) will thus participate in T_u^+ and T_v^-. Now for the formal details. The nodes of H, denoted by V_H, consist of the nodes V of G and additional nodes V'. V' has two types of nodes. First, for each edge $e \in E$ there is a node γ_e. Second, for each node $v \in V$ we create two sets of nodes $N^-(v)$ and $N^+(v)$ where $|N^-(v)| = n_v^- = |\delta_G^-(v)|$ and $|N^+(v)| = n_v^+ = |\delta_G^+(v)|$; thus one node for each edge in $\delta^-(v) \cup \delta^-(v)$; these will be the internal nodes of the trees T_v^- and T_v^+ respectively. For notational convenience we refer to the j'th node in $N^-(v)$ as v_j^- and similarly v_j^+ for the j'th node in $N^+(v)$.

Now we describe the edge set E_H of the graph H, the edge length function $\ell' : E_H \to \mathbb{R}_+$, and the cost function $c : E_H \to \mathbb{R}_+$. The edge set is essentially prescribed by specifying the trees T_v^- and T_v^+ for each $v \in V$. Consider the vector $\mathbf{d}^-(v)$ of values $\ell(e, v)$ for $e \in \delta_G^-(v)$. Recall the definition of the Lovász extension $\hat{\rho}^-(\mathbf{d}_v^-)$. We order the edges in $\delta^-(v)$ as $e_1, e_2, \ldots, e_{n_v^-}$ where $\ell(e_j, v) \geq \ell(e_{j+1}, v)$ for $1 \leq j < n_v^-$ and then $\hat{\rho}^-(\mathbf{d}_v^-) = \sum_j(\ell(e_j, v) - \ell(e_{j+1}, v))\rho_v^-(S_j)$ where $S_j = \{e_1, \ldots, e_j\}$. We associate the node v_j^- with the set S_j. The edge set of T_v^- is defined as follows. For ease of notation we let $v_{n_v^-+1}^-$ represent the node v. We create a directed path $v_1^- \to v_2^- \to \ldots \to v_{n_v^-}^- \to v_{n_v^-+1}^- = v$ with edge lengths $\ell'(v_1^-, v_2^-) = \ell(e_1, v) - \ell(e_2, v), \ell'(v_2^-, v_3^-) = \ell(e_2, v) - \ell(e_3, v), \ldots, \ell'(v_{n_v^-}^-, v) = \ell(e_{n_v^-}, v) - 0$. The costs of these edges are defined as follows: $c(v_j^-, v_{j+1}^-) = \rho_v^-(S_j)$ for $1 \leq j \leq n_v^-$. For each j we add the edge

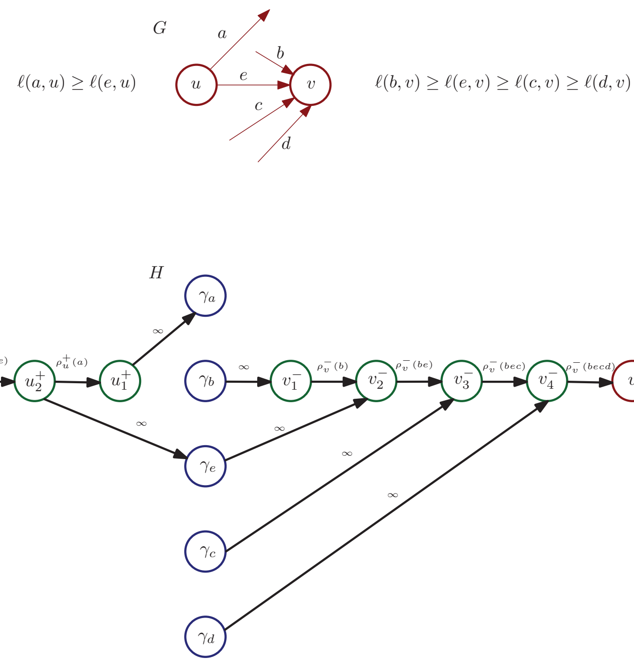

Figure 3: Illustration of the reduction. Only $\delta_G^+(u)$ and $\delta_G^-(v)$ are shown. The costs on edges in H are shown but not their lengths. The lengths of the infinite cost edges is 0 and $\ell'(u_2^+, u_1^+) = \ell(a,u) - \ell(e,u)$ and $\ell'(v_3^-, v_4^-) = \ell(c,v) - \ell(d,v)$.

(γ_{e_j}, v_j^-) with length 0 and cost ∞ (for computational purpose a sufficiently large number M would do); this connects the node γ_{e_j} corresponding to the edge e_j to v_j^- that corresponds to S_j. See Fig 3.

The construction of T_v^+ is quite similar except that the edge directions are reversed; assuming that the edges in $\delta^+(v)$ are ordered such that $\ell(e_1, v) \geq \ell(e_2, v) \geq \ldots \geq \ell(e_{n_v^+}, v)$, we create a path $v \to v_{n_v^+}^+ \to \ldots v_2^+ \to v_1^+$ with edge lengths $\ell(e_{n_v^+}, v) - 0, \ldots, \ell(e_j, v) - \ell(e_{j+1}, v), \ldots, \ell(e_1, v) - \ell(e_2, v)$. The costs for the edges in this path are set to $\rho_v^+(S_{n_v^+}), \ldots, \rho_v^+(S_1)$ where $S_j = \{e_1, \ldots, e_j\}$. For each j we add an edge (v_j^+, γ_{e_j}) with length 0 and cost ∞. This finishes the description of H. We now describe various properties of the graph H. Several of these properties are straightforward from the description of the construction and we omit proofs of the easy claims.

The proposition below asserts the cost of the fractional solution in the edge-capacitated network H is the same as the cost of the fractional solution in the polymatroidal network G.

Proposition 2. $\sum_{e \in E_H} c(e) \cdot \ell'(e) = \sum_{v \in V} (\hat{\rho}^-(\mathbf{d}_v^-) + \hat{\rho}^+(\mathbf{d}_v^+))$.

Proposition 3. *For any edge $e \in \delta_G^-(v)$ the length of the unique path in T_v^- from the node γ_e to v is equal to $\ell(e, v)$. Similarly for $e \in \delta_G^+(v)$, the length of the unique path in T_v^+ from the node v to the node γ_e is equal to $\ell(e, v)$.*

We now establish a correspondence between paths in G and H that connect nodes in V. Let $e = (u, v)$ be an edge in G. We obtain a canonical path $q(u, v)$ from u to v in H as follows: concatenate the unique path from u to γ_e in T_v^+ with the unique path from γ_e to v in T_v^-. For any two nodes $s, t \in V$ let $\mathcal{P}_G(s, t)$ be the set of (simple) s-t paths on G and similarly $\mathcal{P}_H(s,t)$ be the paths in H. We create a map $g : \mathcal{P}_G(s,t) \to \mathcal{P}_H(s,t)$ as follows. Consider a path $p \in \mathcal{P}_G(s,t)$; we obtain a path $p' \in \mathcal{P}_H(s,t)$ corresponding to p as follows. We replace each edge $(u,v) \in p$ by the canonical path $q(u,v)$.

Lemma 5. *The map g is a bijection. Moreover, for any two nodes $u, v \in V$, $dist_{\ell'}(u,v) = dist_{\ell}(u,v)$.*

Now we establish a correspondence between cuts in G and H. For a given set of edges $F \subseteq E$ let $\text{sep}_G(F)$ be set of node pairs in $V \times V$ separated by F in the graph G. Similarly for a set of edges $F' \subseteq E_H$ let $\text{sep}_H(F')$ be the set of node pairs in $V \times V$ separated by F' in the graph H. We say that a set of edges F is minimal with respect to separating node pairs if there is no proper subset of F that separates the same node pairs as F.

Proposition 4. *Let $F' \subseteq E_H$ be minimal with respect to separating node pairs in $V \times V$ and of finite cost. Then for any $v \in V$, F' contains at most one edge from T_v^- and at most one edge from T_v^+.*

Proof. Consider a node v and edge-sets $F' \cap T_v^-$ and $F' \cap T_v^+$. For an edge $e \in E$ there is a node $\gamma_e \in V_H$ and there is exactly one edge coming into γ_e and exactly one edge going out of γ_e and both are of infinite cost. Therefore, if F' is of finite cost, $F' \cap T_v^-$ consists of some edges in the path $v_1^- \to v_2^- \ldots \to v_{n_v^-}^- \to v$ contained in T_v^-. Since the only way to reach v is through T_v^- it follows that if F' contains an edge (v_j^-, v_{j+1}^-) then it is redundant to remove an edge (v_i^-, v_{i+1}^-) for $i < j$. Thus minimality of F' implies F' contains exactly one edge from T_v^-. The reasoning for T_v^+ is similar. □

Lemma 6. *Let $F' \subseteq E_H$ be minimal with respect to separating node pairs in $V \times V$ and of finite cost. There exists a set of edges $F \subseteq E$ such that $sep_G(F) \supseteq sep_H(F')$ and $\nu(F) \leq c(F')$.*

Proof. Given a minimal F' we obtain a set of edges $F \subseteq E$ as follows. From the proof of Proposition 4 we see that for any node v, F' contains at most one edge from T_v^- and in particular if it contains an edge then it is an edge (v_j^-, v_{j+1}^-) for some $1 \leq j \leq n_v^-$ (for simplicity we identify v with $v_{n_v^-+1}^-$). Suppose there is such an edge $e' = (v_j^-, v_{j+1}^-)$ in F'. Note that e' corresponds to the set $S_j = \{e_1, \ldots, e_j\}$ of edges in $\delta_G^-(v)$ ordered in increasing order by $\ell(e,v)$ values. We add S_j to F and assign these edges to v in upper bounding $\nu(F)$: by construction $c(e') = \rho_v^-(S_j)$. We do a similar procedure if $e' \in F \cap T_v^+$. It follows that the edge set F that we construct satisfies the property that $\nu(F) \leq c(F')$.

We now show that $\text{sep}_G(F) \supseteq \text{sep}_H(F')$. Consider a pair (s,t) such that s is separated from t by F' in H. Suppose (s,t) is not separated by F in G. Let p be an s-t path that remains in $G \setminus F$. From Proposition 3 there is a unique path $g(p) \in \mathcal{P}_H(s,t)$. For every edge $e = (u,v) \in p$ consider the canonical path $q(u,v)$ in H. Since e is not in F it implies that u can reach γ_e in $H \setminus F'$ and that γ_e can reach v in $H \setminus F'$. This means that $q(u,v)$ exists in $H \setminus F'$. This would imply that $g(p)$ exists in $H \setminus F'$ contradicting that assumption that (s,t) is separated by F'. □

We summarize the properties of the reduction. We assume that we have a polymatroidal network $G = (V,E)$ with k demand pairs $(s_i,t_i), \ldots, (s_k,t_k)$ with associated demand values $D_1, \ldots, D_k$. For all the cut problems of interest, the relaxations in Section 3 produce a length function $\ell : E \to \mathbb{R}_+$ and for each $e = (u,v)$ associated non-negative values $\ell(e,u)$ and $\ell(e,v)$ such that $\ell(e) = \ell(e,u) + \ell(e,v)$. As before we use $\mathbf{d}_v^-$ and $\mathbf{d}_v^+$ to denote the vector of $\ell(e,v)$ values for the incoming and outgoing edges at v. The reduction produces an edge-capacitated network $H = (V_H, E_H)$ with the following properties:

- each node of V is a node in V_H
- for all $u, v \in V$, $\text{dist}_\ell(u,v) = \text{dist}_{\ell'}(u,v)$
- $\sum_{e \in E_H} c(e)\ell'(e) = \sum_{v \in V}(\hat{\rho}_v^-(\mathbf{d}_v^-) + \hat{\rho}_v^+(\mathbf{d}_v^+))$

- for any set of edges $F' \subseteq E_H$ there is a corresponding set $F \subseteq E$ such that $\text{sep}_G(F) \supseteq \text{sep}_H(F')$ and $\nu(F) \leq c(F')$.

We also note that the reduction can be carried out in polynomial time. Moreover, given a set $F' \subseteq E_H$ a set $F \subseteq E$ that satisfies the last property in the list above can be found in polynomial time.

We build on the reduction to obtain flow-cut gap results, all of which are based on using the relaxations from Section 3 which are dual to the corresponding flow problems. We argue via the reduction and known results on edge-capacitated networks that there exist integral cuts within some factor α of the fractional solution.

4.2 Multicut

We consider the multicut problem for arbitrary demand pairs as well as symmetric demands. The relaxation satisfies the constraint that $\text{dist}_\ell(s_i, t_i) \geq 1$ for each demand pair (s_i, t_i). The reduction from the preceding section produces a graph $H = (V_H, E_H)$ and a fractional solution $\ell' : E_H \to \mathbb{R}_+$ such that $\text{dist}_{\ell'}(s_i, t_i) \geq 1$. We note that ℓ' is a feasible solution for the standard distance based relaxation for multicut in edge-capacitated networks which is the dual for the maximum throughput multicommodity flow problem. The integrality gap of this relaxation has been studied and several results are known. Let $\beta = \sum_{e \in E_H} c(e)\ell'(e)$ be the fractional solution value. Then one can obtain an integral multicut F' with cost $c(F')$ that can be bounded in terms of β. We summarize the known results.

- The single commodity case corresponds to $k = 1$. In this case the classical maxflow-mincut theorem for standard edge-capacitated graphs implies that there is a cut F' separating s and t of cost at most β.
- Cheriyan, Karloff and Rabani [17] showed that there exists an F' such that $c(F') \leq O(1) \cdot \beta^3$; this was improved by Gupta [27] to show the existence of a multicut F' such that $c(F') \leq O(1) \cdot \beta^2$. These results hold under the assumption that $c(e) \geq 1$ for all e.
- Agrawal, Alon and Charikar [2] improving the results in [17, 27] showed the existence of a cut F' such that $c(F') = \tilde{O}(n^{11/23}) \cdot \beta$. Here n is the number of nodes in the graph.
- Saks, Samorodnitsky and Zosin [51] showed that there exist instances on which every integral multicut has a value $\Omega(k) \cdot \beta$.
- Chuzhoy and Khanna [18] showed that there exist instances on which every multicut has a value $\tilde{\Omega}(n^{1/7}) \cdot \beta$. Further, they showed that the multicut problem is hard to approximate to within a factor of $\Omega(2^{\log^{1-\epsilon} n})$ unless $NP \subseteq ZPP$.

Since polymatroidal networks generalize edge-capacitated networks it follows that all the lower bounds in the above hold for the polymatroidal network case as well. The reduction also allows us to obtain an upper-bound for polymatroidal networks. We have to be careful when using bounds that depend on the number of nodes in the graph. The reduction takes G with n nodes and m edges and produces an edge-capacitated graph H with $n + 2m$ nodes. In the worst case H has $\Omega(n^2)$ nodes. We thus obtain the following theorem.

Theorem 1. *In a directed polymatroidal network G on n nodes, for any given multicommodity flow instance with k pairs, if β is the maximum throughput multicommodity flow then:*

- *If $k = 1$ then there is a feasible cut separating s_1 and t_1 of cost at most β.*
- *There is a feasible multicut F' such that $\nu(F') \leq O(1) \cdot \beta^2$ assuming that ρ_v^+ and ρ_v^- are integer valued for all $v \in V$.*

- *There is a feasible multicut F' such that $\nu(F') \leq \tilde{O}(n^{22/23}) \cdot \beta$.*

Moreover, there exist polynomial-time algorithms to find multicuts guaranteed as above.

Remark 1. *In the preceding theorem the bound $\nu(F') \leq \tilde{O}(n^{22/23}) \cdot \beta$ is obtained via a black box application of the result in [2] and our reduction that blows up the number of nodes. A closer examination of the proof in [2] may lead to a improved bound.*

Symmetric demands: We now consider the symmetric demand case when a multicut corresponds to separating (s_i, t_i) or (t_i, s_i) for a given demand pair $s_i t_i$. The relaxation for this has a constraint that $\text{dist}_\ell(s_i, t_i) + \text{dist}_\ell(t_i, s_i) \geq 1$. In contrast to the strong negative results for the general multicut problem, poly-logarithmic upper bounds on flow-cut gaps are known for symmetric demands in standard networks. In particular Klein et al. [35] show that if β is the cost of a fractional solution then there exists an integral multicut of cost $O(\log^2 k) \cdot \beta$. Even et al. [20] showed the existence of a multicut of cost $O(\log n \log \log n) \cdot \beta$. Note that these bounds are incomparable in that depending on the relationship between k and n one is better than the other. It is also known that there exist instances on which the gap is at least $\Omega(\log n)$. Via the reduction we obtain the following.

Theorem 2. *In a directed polymatroidal network G on n nodes, for any given multicommodity flow instance with symmetric demands on k pairs, the minimum multicut is $O(\min\{\log^2 k, \log n \log \log n\}) \cdot \beta$ where β is maximum throughput multicommodity flow for the symmetric demands.*

Remark 2. *The flow-cut gap in polymatroidal networks for multiterminal flows[3] can be shown to be 2 via the reduction and the result of Naor and Zosin [45].*

4.3 Sparsest cut

Now we consider the sparsest cut problem where the goal is to find a set of edges F to minimize $\nu(F)/D(F)$ where $D(F)$ is the total demand of the pairs separated by F. The relaxation corresponds to finding edge length variables ℓ to minimize the fractional cost subject to the constraint that $\sum_i D_i \cdot \text{dist}_\ell(s_i, t_i) = 1$. Via the reduction we produce an edge-capacitated network H such that $\sum_i D_i \cdot \text{dist}_{\ell'}(s_i, t_i) = 1$ and with the fractional cost preserved. In edge-capacitated networks there is a generic strategy that translates the flow-cut gap for multicut into a flow-cut gap for sparsest cut at an additional loss of an $O(\log \sum_i D_i)$ factor due to Kahale [31] (see also [54]); this has been refined via a more intricate analysis in [47] to lose only an $O(\log k)$ factor although one needs to apply it carefully. In [2] a simple reduction that loses an $O(\log n)$ factor is given (this builds on [31]). For directed graphs the known-gaps for sparsest cut are essentially based on using the corresponding gap for multicut and translating via the above mentioned schemes. We thus obtain the following results.

Theorem 3. *In a directed polymatroidal network G on n nodes, for any given multicommodity flow instance with k pairs, if β is the value of the maximum concurrent flow then there is a cut of sparsity at most $\tilde{O}(n^{22/23}) \cdot \beta$.*

Theorem 4. *In a directed polymatroidal network G on n nodes, for any given multicommodity flow instance with symmetric demands on k pairs, there is a cut of sparsity $O(\min\{\log^3 k, \log^2 n \log \log n\}) \cdot \beta$ where β is maximum concurrent flow.*

[3] In multiterminal flows we have a set of k terminals $\{s_1, s_2, \ldots, s_k\}$ and flow can be sent between any pair of terminals; the goal is to maximize the total flow. The corresponding cut is referred to as multiterminal cut or multiway cut in which the goal is to remove a minimum-cost set of edges to disconnect every (ordered) pair of terminals.

5 Flow-Cut Gaps in Undirected Polymatroidal Networks

In this section we consider flow-cut gaps in undirected polymatroidal networks. As we already noted, node-capacitated flows are a special case of polymatroidal flows. We show that line embeddings with low average distortion introduced by Matousek and Rabinovich [44] (and further studied in [48]) are useful for bounding the gap between the maximum concurrent flow and sparsest cut; we are inspired to make this connection from [23] who considered node-capacitated flows. For multicut we show that the region growing technique from [40] that was used in [24] for edge-capacitated multicut can be adapted to the polymatroidal setting.

5.1 Maximum concurrent flow and sparsest cut

We start with the definition of line embeddings and average distortion.

Let (V, d) be a finite metric space. A map $g : V \to \mathbb{R}$ is an embedding of V into a line; it is a *contraction* (also called 1-Lipschitz) if for all $u, v \in V$,

$$|g(u) - g(v)| \leq d(u, v).$$

Given a demand function $w : V \times V \to \mathbb{R}_+$ and a contraction $g : V \to \mathbb{R}$, its *average distortion* with respect to w is defined as

$$\text{avgd}_w(g) = \frac{\sum_{u,v \in V} w(u, v) \cdot d(u, v)}{\sum_{u,v \in V} w(u, v) \cdot |g(u) - g(v)|}$$

The following theorem is implicit in [8]; see [23] for a sketch.

Theorem 5 (Bourgain [8]). *For every n-point metric space (V, d) and every weight function $w : V \times V \to \mathbb{R}_+$ there is a polynomial-time computable contraction $g : V \to \mathbb{R}$ such that $\text{avgd}_w(g) = O(\log n)$. Moreover, if the support of w is k there is a map g such that $\text{avgd}_w(g) = O(\log k)$.*

Using the above we prove the following.

Theorem 6. *In undirected polymatroidal networks, for any given multicommodity flow instance with k pairs, the ratio between the value of the sparsest cut and the value of the maximum concurrent flow is $O(\log k)$. Moreover, there is an efficient algorithm to compute an $O(\log k)$ approximation to the sparsest cut problem.*

Recall the relaxation for the sparsest cut from Section 3.3 and the associated notation. To prove the theorem we consider an optimum solution to the relaxation and show the existence of a cut whose sparsity is $O(\log k)$ times the value of the relaxation. Let (V, d) be the metric induced on V by shortest path distances in the graph with edge lengths given by $\ell : E \to \mathbb{R}_+$ from the optimum fractional solution. Let $g : V \to \mathbb{R}$ be line embedding guaranteed by Theorem 5 with respect to d and the weight function given by the demands D_i; that is $w(s_i, t_i) = D_i$ for a demand pair and is 0 for any pair of nodes that do not correspond to a demand. Without loss of generality we can assume that g maps V to the interval $[0, \beta]$ for some $\beta > 0$. For $\theta \in (0, \beta)$ let $S_\theta = \{u \mid g(u) \leq \theta\}$. We show that there is a θ such that $\delta(S_\theta)$ is an approximately good sparse cut. Let $D(\delta(S_\theta))$ be the total demand of pairs separated by S_θ, that is $D(\delta(S_\theta)) = \sum_{i:S_\theta \text{ separates } s_it_i} D_i$.

Lemma 7.

$$\int_0^\beta D(\delta(S_\theta))d\theta = \Omega\left(\frac{1}{\log k}\right).$$

Proof. From the definition of $D(\delta(S_\theta))$,

$$\begin{aligned} \int_0^\beta D(\delta(S_\theta))d\theta &= \int_0^\beta (\sum_{i:S_\theta \text{ separates } s_it_i} D_i)d\theta \qquad (6) \\ &= \sum_{i=1}^k D_i \cdot \int_0^\beta \mathbf{1}_{S_\theta \text{ separates } s_it_i} d\theta = \sum_{i=1}^k D_i \cdot |g(s_i) - g(t_i)|. \end{aligned}$$

From the properties of g,

$$\frac{\sum_i D_i \cdot d(s_i, t_i)}{\sum_i D_i \cdot |g(s_i) - g(t_i)|} \le O(\log k).$$

We have the constraint $\sum_i D_i \cdot d(s_i, t_i) = 1$ from the LP relaxation; this combined with the above inequality proves the lemma. □

The main insight in the proof is the following lemma. A version of the lemma also holds for directed graphs that we address in a remark following the proof.

Lemma 8.

$$\int_0^\beta \nu(\delta(S_\theta))d\theta \le 2\sum_u \hat{\rho}_u(\mathbf{d}_u).$$

Proof. Consider an edge $uv \in \delta(S_\theta)$ and for simplicity assume $g(u) < g(v)$. The length of e in the embedding is $\ell'(e) = |g(v) - g(u)| \le \ell(e)$. The edge $(u, v) \in \delta(S_\theta)$ iff θ is in the interval $[g(u), g(v)]$. Note that the cost $\nu(\delta(S_\theta))$ is in general a complicated function to evaluate. We upper bound $\nu(\delta(S_\theta))$ by giving an explicit way to assign $e = uv$ to either u or v as follows. Recall that in the relaxation $\ell(e) = \ell(e, u) + \ell(e, v)$ where $\ell(e, u)$ and $\ell(e, v)$ are the contributions of u and v to e. Let $r = \frac{\ell(e,u)}{\ell(e)}$ and let $\ell'(e, u) = r\ell'(e)$ and $\ell'(e, v) = (1-r)\ell'(e)$. We partition the interval $[g(u), g(v)]$ into $[g(u), g(u)+\ell'(e, u))$ and $[g(u) + \ell'(e, u), g(v)]$; if θ lies in the former interval we assign e to u, otherwise we assign e to v. This assignment procedures describes a way to upper bound $\nu(\delta(S_\theta))$ for each θ. Now we consider the quantity $\int_0^\beta \nu(\delta(S_\theta))d\theta$ and upper bound it as follows.

Consider a node u and let $L_u = \{uv \in \delta(u) \mid g(v) < g(u)\}$ be the set of edges uv that go from u to the left of u in the embedding g. Similarly $R_u = \{uv \in \delta(u) \mid g(v) \ge g(u)\}$. Note that L_u and R_u partition $\delta(u)$. Let $\mathbf{d}'_u$ be the vector of dimension $|\delta(u)|$ consisting of the values $\ell'(e, u)$ for $e \in \delta(u)$. We obtain $\mathbf{d}^L_u$ from $\mathbf{d}'_u$ by setting the values for $e \in R_u$ to 0 and similarly $\mathbf{d}^R_u$ from $\mathbf{d}'_u$ by setting the values for $e \in L_u$ to 0. Since $0 \le \ell'(e, u) \le \ell(e, u)$ for each $e \in \delta(u)$ we see that $\mathbf{d}'_u \le \mathbf{d}_u$ and (component wise) and hence $\mathbf{d}^L_u \le \mathbf{d}_u$ and $\mathbf{d}^R_u \le \mathbf{d}_u$. Since ρ_u is monotone we have that $\hat{\rho}_u(\mathbf{d}^L_u) \le \hat{\rho}_u(\mathbf{d}_u)$ and $\hat{\rho}_u(\mathbf{d}^R_u) \le \hat{\rho}_u(\mathbf{d}_u)$ (see Proposition 1).

We claim that

$$\int_0^\beta \nu(\delta(S_\theta))d\theta \le \sum_{u \in V}(\hat{\rho}_u(\mathbf{d}^L_u) + \hat{\rho}_u(\mathbf{d}^R_u)),$$

which would prove the lemma.

To see the claim consider some fixed θ and $\nu(\delta(S_\theta))$. Fix a node u and consider the edges in $\delta(u) \cap S_\theta$ assigned to u by the procedure we described above; call this set $A_{\theta,u}$. First assume that $\theta < g(u)$. Then the edges assigned to u by the procedure, denoted by $A_{\theta,u} = \{e \in L_u \mid \theta > g(u) - \ell'(e, u)\}$. Similarly, if $\theta > g(u)$, $A_{\theta,u} = \{e \in L_u \mid \theta < g(u) + \ell'(e, u)\}$. From these definitions we have

$$\int_0^\beta \nu(\delta(S_\theta))d\theta \le \sum_{u \in V} \int_0^\beta \rho_u(A_{\theta,u})d\theta.$$

For a fixed node u,

$$\int_0^\beta \rho_u(A_{\theta,u})d\theta = \int_0^{g(u)} \rho_u(A_{\theta,u})d\theta + \int_{g(u)}^\beta \rho_u(A_{\theta,u})d\theta.$$

Let $L_u = \{e_1, e_2, \ldots, e_h\}$ where $0 \le \ell'(e_1, u) \le \ell'(e_2, u) \le \ldots \le \ell'(e_h, u)$. Then

$$\int_0^{g(u)} \rho_u(A_{\theta,u})d\theta = \sum_{j=1}^h (\ell'(e_j, u) - \ell'(e_{j-1}, u))\rho(\{e_1, e_2, \ldots, e_j\}).$$

The right-hand side of the above, is, by construction and the definition of the Lovász extension, equal to $\hat{\rho}_u(\mathbf{d}_u^L)$. Similarly, $\int_{g(u)}^{\beta} \rho_u(A_{\theta,u})d\theta = \hat{\rho}_u(\mathbf{d}_u^R)$. □

We now finish the proof of Theorem 6 via the preceding two lemmas.

$$\begin{aligned}\min_{\theta\in(0,\beta)} \frac{\nu(\delta(S_\theta))}{D(\delta(S_\theta))} &\leq \frac{\int_0^\beta \nu(\delta(S_\theta))d\theta}{\int_0^\beta D(\delta(S_\theta))d\theta} \\ &\leq 2\sum_u \hat{\rho}_u(\mathbf{d}_u)\cdot O(\log k) = O(\log k)\sum_u \hat{\rho}_u(\mathbf{d}_u).\end{aligned}$$

The above shows that the sparsity of S_θ for some θ is at most $O(\log k)$ times $\sum_u \hat{\rho}_u(\mathbf{d}_u)$ which is the value of the relaxation. Given a line embedding g there are only $n-1$ distinct cuts of interest and one can try all of them to find the one with the smallest sparsity. The efficiency of the algorithm therefore depends on complexity of the solving the fractional relaxation and the complexity of finding a line embedding guaranteed by Theorem 5. Since both have polynomial time algorithms, one can find an $O(\log k)$ approximation to the sparsest cut in polynomial time.

Remark 3. *Node-weighted flows and cuts/separators can be cast as special cases of flows and cuts in polymatroid networks. Our algorithm produces edge-cuts from line embeddings in a simple way even for node-weighted problems — the ν cost of the edge-cut automatically translates into an appropriate node-weighted cut. In contrast, the algorithm in [23] has to solve several instances of s-t separator problems in auxiliary graphs obtained from the line embedding.*

A remark on directed polymatroidal networks: An examination of the proof of Lemma 8 explains the factor of 2 on the right hand side; the edges in $\delta(v)$ can be both to the left and right of v in the line embedding and each side contributes $\hat{\rho}_u(\mathbf{d}_v)$ to the cost. This is related to the technical issue about undirected polymatroid networks where the flow through v takes up capacity on two edges incident to v. For directed graphs the same proof outline can be used to show a related statement. Let $g : V \to [0,\beta]$ be an embedding of the nodes in to the interval $[0,\beta]$ such that the following property is true: if $g(u) < g(v)$ and (u,v) is an edge then $g(v) - g(u) \leq \ell(u,v)$. For $\theta in[0,\beta]$ let $\delta^+(S_\theta)$ be the set of edges leaving S_θ. Then,

$$\int_0^\beta \nu(\delta^+(S_\theta))d\theta \leq \sum_u (\hat{\rho}_u^-(\mathbf{d}_u^-) + \hat{\rho}_u^+(\mathbf{d}_u^+)). \tag{7}$$

Notice that there is no factor of 2 since one treats the incoming and outgoing edges separately. The above statement gives an embedding proof of the maxflow-mincut theorem for single-commodity directed polymatroidal networks as follows.

Consider the relaxation in Section 3.2 with $k = 1$, that is, there is a single pair (s,t). Let $\ell(e), e \in E$ be the edge lengths given by an optimum solution. Define a line embedding $g : V \to [0,\beta]$ where $g(v)$ is the shortest path distance from s to v according to the edge lengths ℓ; note that the distance from s to v is not necessarily the same as the distance from v to s since the graph is now directed. Since ℓ is a feasible solution to the relaxation $g(t) = 1$ and it is not hard to see that in an optimum solution $g(v) \leq 1$ for all v. We can now apply (7) to find a $\theta \in (0,1)$ such that $\nu(\delta^+(S_\theta)) \leq \sum_u (\hat{\rho}_u^-(\mathbf{d}_u^-) + \hat{\rho}_u^+(\mathbf{d}_u^+))$. This is a valid s-t cut and its cost is at most the value of the relaxation.

Sparsest bi-partition cut: We have so far worked with general edge cuts, but for certain applications, it is necessary to work with a special type of edge cut called the bi-partition cut. In an undirected polymatroidal network, an edge-cut F is said to be a *bi-partition cut* if there exists a set $S \subseteq V$ such that $F := \{e =$

$uv : u \in S, v \in S^c$ or $v \in S, u \in S^c\}$; we denote such an edge cut by F_S. In the case of edge-capacitated undirected networks, it is well known that a sparsest bi-partition cut is also a sparsest edge cut[4]. While this no longer continues to be true for polymatroidal networks, a factor 2 gap can indeed be shown between the sparsest cut and the sparsest cut restricted to only bi-partition cuts. This is captured in the theorem below whose proof can be found in Section C.

Theorem 7. *Given any edge cut for an undirected polymatroidal network, there exists a bi-partition cut whose sparsity is at most* 2 *times the sparsity of the edge cut. Furthermore this factor is tight.*

Now, Theorem 6 and Theorem 7 together imply a logarithmic gap between maximum concurrent flow and sparsest bi-partition cut. This is formally stated in the following corollary.

Corollary 1. *In undirected polymatroidal networks, for any given multicommodity flow instance with k pairs, the ratio between the value of the* sparsest bi-partition cut *and the value of the maximum concurrent flow is* $O(\log k)$.

5.2 Maximum throughput flow and multicut

We prove the following theorem in this section.

Theorem 8. *In undirected polymatroidal networks, for any given multicommodity flow instance with k pairs, the ratio between the value of the minimum multicut and the value of the maximum throughput flow is* $O(\log k)$. *Moreover, there is an efficient algorithm to compute an* $O(\log k)$ *approximation to the minimum multicut problem.*

We recall the relaxation for the minimum multicut problem from Section 3.2. Consider an optimum solution to the relaxation given by edge lengths $\ell(e), e \in E$ and the partition of $\ell(e)$ for each $e = uv$ between u and v given by the variables $\ell(e, u)$ and $\ell(e, v)$. We will show that there exists a multicut $F \subseteq E$ for the given pairs such that $\nu(F) = O(\log k)(\sum_v \hat{\rho}_v(\mathbf{d}_v))$.

By slightly generalizing the proof of Lemma 8 we obtain the following.

Lemma 9. *Let* $g : V \to [0, \beta]$ *be a contraction, let* $0 \le a_0 \le a < b \le b_0 \le \beta$ *and* $S_\theta = \{u \mid g(u) < \theta\}$. *Suppose for every edge* $e = uv \in \cup_{\theta \in [a,b]} \delta(S_\theta)$, $g(u)$ *and* $g(v)$ *are both in* $[a_0, b_0]$. *Then,*

$$\int_a^b \nu(\delta(S_\theta)) d\theta \le 2 \sum_{v: g(v) \in [a_0, b_0]} \hat{\rho}_v(\mathbf{d}_v).$$

The proof is very similar to the proof of Lemma 8, except that to upper bound the left-hand side in the statement of the lemma, we only need to consider edges that are in the set $\cup_{\theta \in [a,b]} \delta(S_\theta)$. The condition in the lemma assures us that any node that is involved in $\delta(S_\theta)$ has to lie within the interval $[a_0, b_0]$. Thus, it is sufficient to consider the set of nodes $v : g(v) \in [a_0, b_0]$ in the integral on the right hand side. A formal proof with all details can be found in Sec. B of the appendix.

Given a graph G with edge lengths $\ell : E \to \mathbb{R}_+$, a node v and radius r, let $B_G^\ell(v, r) = \{u \mid \text{dist}_\ell(v, u) \le r\}$ denote the ball of radius r around v according to edge lengths ℓ. We omit ℓ and G if they are clear from the context. For a set of nodes $X \subseteq V$ we let $\text{vol}(X) = \sum_{v \in X} \hat{\rho}_v(\mathbf{d}_v)$ denote the total contribution of the nodes in X to the objective function.

Lemma 10. *Let* $\delta < 1$ *and suppose* $\ell(e) < \frac{\delta}{2 \log k}$ *for all e. Then, for any given node s and* $k \ge 2$ *there exists a* $r \in [0, \delta)$ *such that* $\nu(\delta(B(s, r)) \le a \log k \cdot \frac{1}{\delta}(\text{vol}(B(s, r)) + \text{vol}(V)/k)$, *with* $a = 28$.

[4] It is also known that this does not hold for directed networks.

Proof. For simplicity we assume here that $\log k$ is an integer multiple of 3. Order the nodes in increasing order of distance from s: this produces a line embedding $g_s : V \to \mathbb{R}_+$. For integer $i \geq 0$ define $r_i = \frac{i \cdot \delta}{2 \log k}$. Define $\alpha_0 = \text{vol}(V)/k$ and for $i \geq 1$ let $\alpha_i = \alpha_0 + \text{vol}(B(s, r_i))$.

Consider any $1 \leq j \leq 2 \log k$. We apply Lemma 9 to the embedding g_s and the interval $[r_{j-1}, r_j]$; note that $\ell(e) < \frac{\delta}{2 \log k}$ which implies that we can indeed apply the lemma. Also any edge $e \in \cup_{\theta \in [r_{j-1}, r_j]} \delta(S_\theta)$ satisfies the property that $g(u) \in [r_{j-2}, r_{j+1}]$ and $g(v) \in [r_{j-2}, r_{j+1}]$ since $\ell(e) < \frac{\delta}{2 \log k}$. Thus

$$\begin{aligned} \int_{r_{j-1}}^{r_j} \nu(\delta(B(s, \theta))d\theta &\leq 2 \sum_{v: g_s(v) \in [r_{j-2}, r_{j+1}]} \hat{\rho}_v(\mathbf{d}_v) \\ &\leq 2(\alpha_{j+1} - \alpha_{j-2}). \end{aligned} \tag{8}$$

We claim that there is some $1 \leq j < 2 \log k$ such that $\alpha_{j+1} \leq 8\alpha_{j-2}$. Suppose not, then $\alpha_{3i} > 8\alpha_{3(i-1)}$ for all $1 \leq i \leq \frac{2 \log k}{3}$. This implies that $\alpha_{3i} > 8^i \alpha_0 = 2^{3i}\alpha_0$. Therefore, with $i = \frac{2 \log k}{3}$, this implies that $\alpha_{2 \log k} > 2^{2 \log k} \frac{\text{vol}(V)}{k} > 4\text{vol}(V)$ which is impossible.

Thus there exists a j such that $\alpha_{j+1} \leq 8\alpha_{j-2}$. Consider that j, equation (8) implies that

$$\begin{aligned} \int_{r_{j-1}}^{r_j} \nu(\delta(B(s, \theta))d\theta &\leq 2(\alpha_{j+1} - \alpha_{j-2}) \\ &\leq 2(7\alpha_{j-2}). \end{aligned}$$

If we pick r uniformly at random from the interval $[r_{j-1}, r_j]$, where satisfies the above property, the expected cost of $\nu(\delta(B(s, r)))$ is

$$\frac{1}{r_j - r_{j-1}} \int_{r_{j-1}}^{r_j} \nu(\delta(B(s, \theta))d\theta \leq \frac{28 \log k}{\delta} \alpha_{j-2},$$

from the preceding inequality and the fact that $r_j - r_{j-1} = \frac{\delta}{2 \log k}$. Hence there exists an $r \in [r_{j-1}, r_j]$ such that $\nu(\delta(B(s, r))) \leq \frac{28 \log k}{\delta} \alpha_{j-2}$. Since $\alpha_{j-2} - \alpha_0 \leq \text{vol}(B(s, r))$, the lemma follows. □

Now we consider the following algorithm for finding a multicut from a given fractional solution.

- Let $F \leftarrow \{e \mid \ell(e) \geq \frac{1}{4 \log k}\}$.
- $G' \leftarrow G[E \setminus F]$.
- Until there exists a pair $s_i t_i$ connected in G', do the following:
  - Let $s_j t_j$ be a pair connected in G'.
  - Via Lemma 10 with $\delta = 1/2$ find $r < 1/2$ such that $\nu(\delta_{G'}(B_{G'}(s_j, r))) \leq 2a \log k \cdot (\text{vol}(B_{G'}(s_j, r)) + \text{vol}(V)/k)$.
  - $F \leftarrow F \cup \delta_{G'}(B_{G'}(s_j, r))$.
  - Remove the vertices $B_{G'}(s_j, r)$ and edges incident to them from G'.
- Output F as the multicut.

Lemma 11. *The set of edges F output by the algorithm is a feasible multicut for the given instance.*

Proof. By induction on the number of steps in the while loop. We consider the first step. The diameter of the ball $B_{G'}(s_j, r)$ is $2r < 1$ and hence the end points of any pair cannot both be inside this ball. We remove the edges $\delta(B_{G'}(s_j, r))$ and by the preceding observation there is no need to recurse on this ball. The algorithm recurses on the remaining graph $G' - B_{G'}(s_j, r)$, and by induction separates any pair with both end points in that graph. □

Now we argue about the cost of the set F output by the algorithm. Let $F_0 \leftarrow \{e \mid \ell(e) \geq \frac{1}{4\log k}\}$ be the initial set of edges added to F and let F_i be the set of edges added in the i'th iteration of the while loop.

Lemma 12. $\nu(F_0) \leq 8 \log k \cdot \sum_v \hat{\rho}_v(\mathbf{d}_v)$.

Proof. For $v \in V$ let $A_v = \{e \in \delta(v) \cap F_0 \mid \ell(e, v) \geq \frac{1}{8\log k}\}$. We can upper bound $\nu(F_0)$ by $\sum_v \rho_v(A_v)$ since the latter term counts each edge $uv \in F_0$ in at least one of A_u and A_v since $\ell(e, u) + \ell(e, v) = \ell(e) \geq \frac{1}{4\log k}$. From the definition of the Lovász extension

$$\hat{\rho}_v(\mathbf{d}_v) = \int_0^1 \rho_v(\mathbf{d}_v^\theta) d\theta \geq \int_0^{1/(8\log k)} \rho_v(\mathbf{d}_v^\theta) d\theta \geq \frac{1}{8\log k} \rho_v(A_v),$$

where we used non-negativity of ρ_v for the first inequality above and monotonicity for the second. □

Lemma 13. $\sum_{i \geq 1} \nu(F_i) \leq 4a \log k \sum_v \hat{\rho}_v(\mathbf{d}_v)$.

Proof. From the algorithm description, $F_i = \delta(B_{G'}(s_j, r))$ for some terminal s_j and radius $r < 1/2$ where G' is the remaining graph in iteration i. Moreover, $\nu(F_i) \leq 2a \log k \cdot (\text{vol}(B_{G'}(s_j, r)) + \text{vol}(V)/k)$. Since the nodes in $B_{G'}(s_j, r)$ are removed from the graph, a node u is charged only once inside a ball. Hence

$$\sum_i \nu(F_i) \leq \sum_i 2a \log k \cdot \text{vol}(V)/k + 2a \log k \sum_v \hat{\rho}_v(\mathbf{d}_v) \leq 4a \log k \sum_v \hat{\rho}_v(\mathbf{d}_v),$$

since there are at most k iterations of the while loop; each iteration separates at least one pair. □

Since ν is subadditive (see Lemma 2)

$$\nu(F) \leq \nu(F_0) + \sum_{i \geq 1} \nu(F_i) \leq (8 + 4a) \log k \sum_v \hat{\rho}_v(\mathbf{d}_v).$$

This finishes the proof of Theorem 8.

5.3 Max throughput flow and multicut in planar and minor-free graphs

We now consider the flow-cut gaps in undirected *planar* polymatroidal networks[5] and more generally networks (equivalently graphs) that exclude the complete graph K_h as a minor[6] for some fixed h. Klein, Plotkin and Rao [34] proved an important network decomposition theorem for such graphs that leads to two results in edge capacited graphs. First, it gives an $O(1)$ bound on the gap between concurrent flow and sparsest cut for product multicommodity flows. Second, as shown in [55], it leads to an $O(1)$ bound on the gap between throughput flow and multicut. Gupta et al. [28] conjectured that the concurrent flow-sparsest cut gap is $O(1)$ for these networks. This is still an importan open problem but some non-trivial results have been shown in support of this conjecture. Rao [50] proved an upper bound of $O(\sqrt{\log n})$ thereby improving upon the gap for general graphs which can be $\Omega(\log n)$ in the worst case. The gap for series parallel graphs is known to be

[5] By a planar polymatroidal network we simply mean that the underlying graph G is planar.

[6] A graph H is called a *minor* of a graph G if H can be obtained by G by a sequence of edge deletions, vertex deletions and contraction of edges (i.e., collapsing two nodes connected by an edge into a single node).

2 [12], and the gap for k-outerplanar graphs is known to be $2^{O(k)}$ [13] — see [39] for further results. Much less is known for node-capacitated planar and minor-free graphs; the only result that we are aware of is that of Brinkman, Karagiozova and Lee [9] that shows an $O(\sqrt{\log n})$ gap for series-parallel graphs. In fact the $O(1)$ gap between throughput flow and multicut has not been generalized to node-capacitated graphs. Here, we show that an $O(1)$ bound for the throughput flow-multicut gap in planar and minor-free polymatroidal graphs.

Our result, not surprisingly, is based on the KPR network decomposition theorem [34]. Rabinovich [48] used the KPR theorem to give a line embedding theorem for planar and minor-free graphs with $O(1)$ average distortion when restricted to product multicommodity flows; this interpretation gives an $O(1)$ bound on concurrent flow and sparsest cut for node-capacitated case [23], and, from the discussion in Section 5.1, also for the polymatroidal case. The line embedding theorem does not directly lead to a bound for the gap between throughput flow and multicut. We observe that the KPR decomposition is based on $O(1)$ iterations, each of which can be thought of as providing a line embedding. We use these iterative line embeddings to derive our result that is formally stated below.

Theorem 9. *Let G be an undirected polymatroidal network such that the underlying graph excludes K_h as a minor. Then, for any multicommodity instance on G, the minimum multicut is within a factor $O(h^2)$ of the maximum throughput flow.*

As an easy corollary we obtain the following result.

Corollary 2. *There is an $O(h^2)$-approximation for finding a minimum* node-weighted *multicut in a graph that excludes K_h as a minor.*

The rest of this section is dedicated to proving Theorem 9. We prove a weaker bound of $O(h^3)$ via the KPR network decomposition theorem, and then indicate how the bound can be improved to $O(h^2)$ via the result from [21]. Consider an optimum solution to the relaxation for the minimum multicut problem from Section 3.2; let $\ell(e), e \in E$ be the edge lengths given by the solution, and for $e = uv$, $\ell(e, u)$ and $\ell(e, v)$ are the values such that $\ell(e, u) + \ell(e, v) = \ell(e)$. The goal is to show that there exists a multicut $F \subseteq E$ such that F separates each source from its corresponding sink and $\nu(F) = O(h^3) \sum_v \hat{\rho}_v(\mathbf{d}_v)$.

Chopping operation: We describe a chopping operation, which is used to partition the network. We use the terminology of [38] to describe this process.

Given a connected graph H, a special node $v_0 \in V(H)$, positive numbers τ and γ, and a metric ℓ on the nodes, we define a partitioning operation, called τ-chop of H rooted at v_0 with offset γ, as follows. Consider a line embedding of the nodes $V(H)$, induced by the shortest path distance from v_0 using the metric ℓ; i.e., $g : V \to \mathbb{R}_+$ is defined as

$$g(u) = \text{dist}_\ell(u, v_0) \quad \forall u \in V(H).$$

Since the graph is connected, $g(u)$ is bounded, and therefore define

$$d_{\max} = \max_{u \in V(H)} g(u).$$

The τ-chop partitioning operation divides V into partitions V_i defined as follows:

$$V_i = \{v \in V(H) : \gamma + (i-1)\tau \leq g(v) < \gamma + i\tau\}, \quad i = 1, 2, .., \frac{\lceil d_{\max} \rceil}{\tau}.$$

Clearly $V(H) = \biguplus_i V_i$. This partitioning operation disconnects the edges,

$$F := \{e = uv \in E(H) : \exists i \neq j \text{ s.t. } u \in V_i, v \in V_j\}.$$

Thus we can think of F as the cut associated with the τ-chop. The cost of the τ-chop is equal to the cut cost $\nu(F)$. More generally, a τ-chop on a disconnected graph is defined as the result of performing a τ-chop on each of its connected components. When we perform a sequence of τ-chops, the i-th chop performs partitioning individually on each of the partitions created by the $i-1$-th chop.

Fig. 4 shows an example of a graph with distances and the application of two successive τ-chops. In the figure, the root node for the chop is shown in a transparent circle, whereas the other nodes are shown as filled circles. Observe that in each iteration, for each connected component, we use a different line embedding depending upon the root node selected.

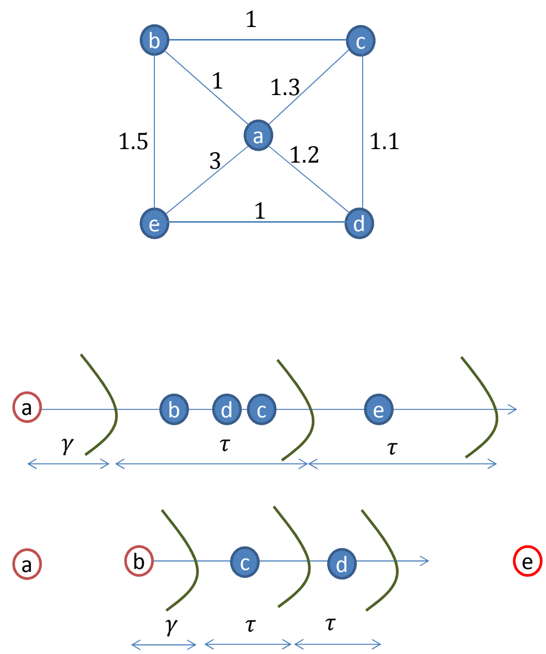

Figure 4: Example of a weighted graph and two successive τ-chop operations.

We will show that there exists a "good" offset γ, such that the cost of the cut is within a constant factor of the dual cost.

Lemma 14. *Given a graph $G=(V,E)$, a distance metric ℓ satisfying $\ell(e)\le\tau\ \forall e\in E$, any root node $v_0\in V$ and a positive number τ, let the offset γ be uniformly random in $[0,\tau]$ and F_γ be the random cut corresponding to the τ-chop rooted at v_0 with offset γ. Then the expected value of the random cut F_γ is*

$$\mathbb{E}\left[\nu(F_\gamma)\right]\le\frac{2}{\tau}\sum_v\hat{\rho}_v(\mathbf{d}_v).$$

Proof. We consider the case when the graph is comprised of a single connected component. The case of a disconnected (partitioned) graph can be dealt with by dealing with each of the connected components (partitions) separately.

We begin by considering the line embedding $g(u)$ induced by the shortest path distance from v_0 using distances ℓ, i.e., $g(u)=\mathrm{dist}_\ell(u,v_0),\quad\forall u\in V(H)$. The length of edge $e=uv$ in the embedding is given by $\ell'(e):=|g(v)-g(u)|\le\ell(e)$ where we have used the triangle inequality. While the cost $\nu(F_\gamma)$ is in general complicated to evaluate, we can upper-bound $\nu(F_\gamma)$ by using one particular way to assign every edge

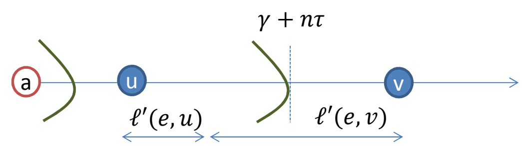

Figure 5: Charging an edge

$e = uv \in F_\gamma$ to either of the nodes u or v, i.e., by charging the edge to the submodular constraint on either u or v. To do this, we critically use the finer grain information contained in the dual variables, $\ell(e,u)$ and $\ell(e,v)$, which add up to give $\ell(e)$ in the relaxation. Let $r = \frac{\ell(e,u)}{\ell(e)}$ and let $\ell'(e,u) = r\ell'(e)$ and $\ell'(e,v) = (1-r)\ell'(e)$. For $g(u) < g(v)$, we partition the interval $[g(u), g(v)]$ into $[g(u), g(u) + \ell'(e,u))$ and $[g(u) + \ell'(e,u), g(v)]$. If the τ-chop cuts the edge $e = uv$ in the former interval, we assign it to u , else we assign it to v, i.e., charge it to the submodular constraint at v. This is illustrated in Fig. 5, where the cut is charged to node v.

In order to state a formal bound on $\nu(F_\gamma)$, we need some definitions. Consider a node u and let $L_u = \{uv \in \delta(u) \mid g(v) < g(u)\}$ be the set of edges uv that go from u to the left of u in the embedding g. Similarly $R_u = \{uv \in \delta(u) \mid g(v) \geq g(u)\}$. Note that L_u and R_u partition $\delta(u)$. Let $\mathbf{d}'_u$ be the vector of dimension $|\delta(u)|$ consisting of the values $\ell'(e,u)$ for $e \in \delta(u)$. We obtain $\mathbf{d}^L_u$ from $\mathbf{d}'_u$ by setting the values for $e \in R_u$ to 0 and similarly $\mathbf{d}^R_u$ from $\mathbf{d}'_u$ by setting the values for $e \in L_u$ to 0. Since $0 \leq \ell'(e,u) \leq \ell(e,u)$ for each $e \in \delta(u)$ we see that $\mathbf{d}'_u \leq \mathbf{d}_u$ and (component wise) and hence $\mathbf{d}^L_u \leq \mathbf{d}_u$ and $\mathbf{d}^R_u \leq \mathbf{d}_u$. Since ρ_u is monotone, the extension $\hat{\rho}$ is also monotone and we have $\hat{\rho}_u(\mathbf{d}^L_u) \leq \hat{\rho}_u(\mathbf{d}_u)$ and $\hat{\rho}_u(\mathbf{d}^R_u) \leq \hat{\rho}_u(\mathbf{d}_u)$.

We start with

$$\mathbb{E}\left[\nu(F_\gamma)\right] = \frac{1}{\tau}\int_{\gamma=0}^{\tau} \nu(F_\gamma), \tag{9}$$

and upper-bound $\nu(F_\gamma)$, for any fixed γ using the assignment formalized below. Define

$$A^L_{\gamma,u} = \left\{ \begin{array}{ll} e = uv \in L_u : \exists n \in \mathbb{N} : & g(u) > \gamma + n\tau \\ & g(u) - \ell'(e,u) < \gamma + n\tau \end{array} \right\}, \tag{10}$$

and similarly define

$$A^R_{\gamma,u} = \left\{ \begin{array}{ll} e = uv \in R_u : \exists n \in \mathbb{N} : & g(u) < \gamma + n\tau \\ & g(u) + \ell'(e,u) > \gamma + n\tau \end{array} \right\}. \tag{11}$$

Since $\ell(e) \leq \tau$, we note that $A^L_{\gamma,u} \cup A^R_{\gamma,u} = \emptyset$, i.e., only one of the sets is active for a given γ. We can write the upper bound on $\nu(F_\gamma)$ using these sets as

$$\nu(F_\gamma) \leq \sum_{u \in V} \rho_u(A^L_{\gamma,u}) + \rho_u(A^R_{\gamma,u}). \tag{12}$$

For a fixed node u, let us consider $\mathbb{E}\left[\rho_u(A^L_{\gamma,u})\right]$. To compute this, let us order the set L_u as $L_u = \{e_1, ..., e_h\}$ such that $\ell'(e_1,u) \geq \ell'(e_2,u) \geq \ldots \geq \ell'(e_h,u) \geq 0$. When we take a random $\gamma \in [0,\tau]$, the probability that edge e_i is cut and assigned to node u, is given by $\frac{\ell'(e,u)}{\tau}$. Furthermore, we observe that whenever edge e_i is cut and assigned to u, all the edges $e_1, ..., e_i$ are also cut and assigned to u. Thus the set of edges $e_1, ..., e_i$ is assigned to node u with probability $\frac{\ell'(e_i,u) - \ell'(e_{i-1},u)}{\tau}$. This gives us the equality

$$\mathbb{E}\left[\rho_u(A^L_{\gamma,u})\right] = \frac{1}{\tau}\sum_{j=1}^{h}(\ell'(e_j,u) - \ell'(e_{j-1},u))\rho(\{e_1, e_2, \ldots, e_j\}). \tag{13}$$

By the definition of Lovász extension, the right-hand side of this equation is equal to $\frac{1}{\tau}\hat{\rho}_u(\mathbf{d}_u^L)$. We can perform a symmetric calculation for the expected value of the second term in (12),

$$\mathbb{E}\left[\rho_u(A_{\gamma,u}^L)\right] = \frac{1}{\tau}\hat{\rho}_u(\mathbf{d}_u^L). \tag{14}$$

Thus (12) implies that

$$\begin{aligned} \mathbb{E}\left[\nu(F_\gamma)\right] &\leq \frac{1}{\tau}\sum_{u\in V}\hat{\rho}_u(\mathbf{d}_u^L) + \hat{\rho}_u(\mathbf{d}_u^R) & (15)\\ &\leq \frac{2}{\tau}\sum_{u\in V}\hat{\rho}_u(\mathbf{d}_u), & (16)\end{aligned}$$

where the second inequality follows due to the fact that $\hat{\rho}(\mathbf{d}_u^L) \leq \hat{\rho}(\mathbf{d}_u)$ and $\hat{\rho}(\mathbf{d}_u^R) \leq \hat{\rho}(\mathbf{d}_u)$. □

We use the following lemma from [34] that shows that if a graph excludes K_h as a minor, then a sequence of $h-1$ τ-chops will yield components with diameter $O(h\tau)$.

Lemma 15 ([34]). *If $G = (V, E)$ with distances $\ell(e), e \in E$ excludes K_h as a minor, then for any $\tau \geq 1$, any sequence of $h-1$ iterated τ-chops on V results in a partition $V = S_1 \cup S_2 \cup \cdots \cup S_m$ such that $\mathrm{diam}(S_i) \leq O(h^2\tau)$, where* diam *refers to the diameter in G using the shortest path distance $dist_\ell$.*

Algorithm for finding a multicut:

- Compute the optimal solution to the relaxation. This can be done efficiently using the ellipsoidal algorithm, since the separation oracle for the dual is a simple shortest path computation.
- Initialize $F \leftarrow F_0 := \{e \mid \ell(e) \geq \tau\}$, i.e., remove all edges greater than length τ.
- Set $G' \leftarrow G[E \setminus F]$ with distance function $\ell(e), e \in E(G')$.
- Perform $(h-1)$ τ-chops on G' as follows. For the i-th chop, choose an arbitrary node in each connected component as the corresponding root node and use uniformly independently chosen offsets $\gamma \in [0, \tau]$. Let F_i be the cut associated with the i-th τ-chop. For each $i = 1, 2, ..., h-1$, update

$$F \leftarrow F \cup F_i. \tag{17}$$

- Output F as the multicut.

Since the graph avoids K_h as a minor, by Lemma 15, the diameter of every component will be smaller than $O(h\tau)$. By setting $\tau = \frac{\triangle}{Ch^2}$, with C large enough, the diameter of every component will be smaller than $\triangle$. We set $\triangle = \frac{1}{2}$, which implies that for any i, s_i and t_i are never in the same component due to the fact that $\mathrm{dist}_\ell(s_i, t_i) \geq 1$ and the triangle inequality which implies that $\mathrm{dist}_\ell(s_i, v_0) + \mathrm{dist}_\ell(t_i, v_0) \geq 1$.

Theorem 10. *The algorithm outputs a multicut F such that*

$$\mathbb{E}\left[\nu(F)\right] \leq O(h^3)\sum_{v}\hat{\rho}_v(\mathbf{d}_v).$$

Proof. We compute the cost of the multicut F as follows:

$$\mathbb{E}\left[\nu(F)\right] \leq \nu(F_0) + \sum_{i=1}^{h-1} \mathbb{E}\left[\nu(F_i)\right], \tag{18}$$

since the cost function $\nu(.)$ is subadditive (this follows from the fact that ρ_v is a polymatroid, and hence is subadditive).

We first compute the cost $\nu(F_0)$ as follows. Since for each edge $e = uv \in E$, $\ell(e) \geq \tau$, either $\ell(e,u) \geq \frac{\tau}{2}$ or $\ell(e,v) \geq \frac{\tau}{2}$ as $\ell(e) = \ell(e,u) + \ell(e,v)$. Define for $v \in V$, $A_v = \{e \in \delta(v) \cap F_0 \mid \ell(e,v) \geq \frac{\tau}{2}\}$. We can upper-bound $\nu(F_0)$ by $\sum_v \rho_v(A_v)$ since the latter term counts each edge $uv \in F_0$ in at least one of A_u or A_v. From the definition of the Lovász extension,

$$\hat{\rho}_v(\mathbf{d}_v) = \int_0^1 \rho_v(\mathbf{d}_v^\theta) d\theta \geq \int_0^{\tau/2} \rho_v(\mathbf{d}_v^\theta) d\theta \geq \frac{\tau}{2} \rho_v(A_v),$$

where we used non-negativity of ρ_v for the first inequality above. The second inequality follows from the fact that $A_v \subseteq \mathbf{d}_v^\theta$, whenever $\theta \leq \frac{\tau}{2}$ and the monotonicity of ρ_v. Thus, we get

$$\nu(F_0) \leq \sum_v \rho_v(A_v) \leq \frac{2}{\tau} \sum_v \hat{\rho}_v(\mathbf{d}_v). \tag{19}$$

By Lemma 14, we get that, for the i-th τ-chop, the expected cost is

$$\mathbb{E}\left[\nu(F_i)\right] \leq \frac{1}{\tau} \sum_v \hat{\rho}_v(\mathbf{d}_v). \tag{20}$$

Substituting this into (18), we get

$$\mathbb{E}\left[\nu(F\right] \leq \frac{h+1}{\tau} \sum_v \hat{\rho}_v(\mathbf{d}_v) = \frac{Ch^2(h+1)}{\triangle} \hat{\rho}_v(\mathbf{d}_v) \tag{21}$$

$$= O(h^3) \sum_v \hat{\rho}_v(\mathbf{d}_v), \tag{22}$$

using the choice $\triangle = \frac{1}{2}$, which concludes the proof of the theorem. $\square$

Theorem 10 implies a weaker version of Theorem 9 with a bound of $O(h^3)$. Now we sketch how the result of [21] that implies the claimed bound of $O(h^2)$. The algorithm in [34], in each iteration, chooses the root node v_0 in each connected component of the current graph, in an arbitrary fashion. In [21] it is shown that one can choose the root nodes in a careful fashion such that after $h-1$ iterations the diameter of each connected component is at most $O(h\tau)$. The multicut algorithm is now modified to use the choice of the roots given by the algorithm from [21] and τ can be chosen to be C/h rather than C/h^2. This gives the desired improvement.

Proof of Corollary 2: A multicut in a node-weighted graph G can therefore be modeled by a multicut in a polymatroidal network G' obtained from G as follows. For each v with weight $w(v)$ we define the function ρ_v as : $\rho_v(S) = w(v)$ for each $S \subseteq \delta(v)$, $S \neq \emptyset$. Note that the multicut in the polymatroidal network G' is defined as a set of edges F but its cost $\nu(F)$ takes into account the minimum weight set of nodes whose removal ensures that all edges of F are removed. For instance if an edge $uv \in F$ is assigned to u in the evaluation of $\nu(F)$ then the node u will be part of the multicut in the original graph G.

6 Conclusions

We considered multicommodity flows and cuts in polymatroidal networks and derived flow-cut gap results in several settings. These results generalize some existing results for the well-studied edge and node-capacitated networks. We briefly mention two results that can be obtained via the line embeddings technique that we did not include in this paper. A multicommodity flow instance in an undirected network $G = (V, E)$ is a product multicommodity flow instance if there there is a non-negative weight function $\pi : V \to \mathbb{R}_+$ and the demand D_{uv} between u and v is $\pi(u) \cdot \pi(v)$. The associated cut problem is interesting because it corresponds to finding sparse separators in graphs which in turn can be used to find balanced separators; these have several applications. Arora, Rao and Vazirani [5] gave an $O(\sqrt{\log n})$-approximation, via a semi-definite programming relaxation, for the sparsest cut problem in an undirected edge-capacitated network. Note that this is not a traditional flow-cut gap result since the SDP-based relaxation used is strictly stronger than the dual of the multicommodity flow relaxation. By interpreting the main technical result in [5] as a line-embedding theorem, [23] obtained an $O(\sqrt{\log n})$-approximation for sparsest cut in node-capacitated graphs; this can also be extended to the polymatroidal setting via the techniques in Section 5.1. It may also be possible to extend the results of Agarwal et al. [3] on $O(\sqrt{\log n})$ approximation for directed cut problems to the polymatroidal setting.

Flow-cut gap questions for node-capacitated problems are less well-understood than the corresponding questions for edge-capacitated problems; line-embeddings provide a tool to obtain upper bounds on the gap but they do not provide a tight characterization as ℓ_1-embeddings do for the edge-capacitated case. We hope that polymatroidal networks and their applications to network information flow provide a new impetus for understanding these questions. Recently, partially motivated by our work, Lee, Mendel and Moharrami [37] obtained results for node-capacitated and polymatroidal versions of the well-known Okamura-Seymour theorem [46].

References

[1] A. S. Avestimehr, S. N. Diggavi, and D. N. C. Tse. Wireless Network Information Flow: A Deterministic Approach. *IEEE Trans. Info. Theory.*, 57(4):1872–1905, 2011.

[2] A. Agarwal, N. Alon and M. Charikar. Improved approximation for directed cut problems. *Proc. of ACM STOC*, 671–680, 2007.

[3] A. Agarwal, M. Charikar, K. Makarychev, and Y. Makarychev. $O(\sqrt{\log n})$ approximation algorithms for min UnCut, min 2CNF deletion, and directed cut problems. *Proc. of ACM STOC*, 2005.

[4] A. Amadruz, and C. Fragouli. Combinatorial algorithms for wireless information flow. *Proc. of ACM-SIAM SODA*, Jan. 2009.

[5] S. Arora, S. Rao and U. Vazirani. Expander Flows, Geometric Embeddings, and Graph Partitionings. *JACM*, 56(2), 2009. Preliminary version in *Proc. of ACM STOC*, 2004.

[6] S. Arora, J. Lee and A. Naor. Euclidean distortion and the Sparsest Cut. *J. of the AMS*, to appear. Preliminary version in *Proc. of ACM STOC*, 2005.

[7] Y. Aumann and Y. Rabani. An $O(\log k)$ approximate min-cut max-flow theorem and approximation algorithm. *SIAM Journal on Computing*, 27(1):291–301, 1998.

[8] J. Bourgain. On Lipschitz embedding of finite metric spaces in Hilbert space. *Israeli J. Math.*, 52:46–52, 1985.

[9] B. Brinkman, A. Karagiozova and J. Lee. Vertex cuts, random walks, and dimension reduction in series-parallel graphs. *Proc. of ACM STOC*, 2007.

[10] G. Calinescu, C. Chekuri, M. Pál and J. Vondrák. Maximizing a submodular set function subject to a matroid constraint. *Proc. of IPCO*, 182–196, 2007.

[11] G. Calinescu, C. Chekuri, M. Pál and J. Vondrák. Maximizing a submodular set function subject to a matroid constraint. To appear in *SIAM Journal on Computing, special issue for STOC 2008*.

[12] A. Chakrabarti, A. Jaffe, J. Lee and J. Vincent. Local moves and lossy invariants in planar graph embeddings. *Proc. of IEEE FOCS*, 2008.

[13] C. Chekuri, A. Gupta, I. Newman, Y. Rabinovich, and A. Sinclair. Embedding k-outerplanar graphs into ℓ_1. *SIAM J. on Disc. Math.*, 20(1):119–136, 2006. Preliminary version in *Proc. of ACM-SIAM SODA*, 2003.

[14] C. Chekuri and A. Ene. Submodular Cost Allocation Problems and Applications. *Proc. of ICALP*, 2011.

[15] C. Chekuri and A. Ene. Approximation Algorithms for Submodular Multiway Partition. *Proc. of IEEE FOCS*, 2011.

[16] C. Chekuri, J. Vondrak and R. Zenklusen. Submodular Function Maximization via the Multilinear Relaxation and Contention Resolution Schemes. *Proc. of ACM STOC*, 2011.

[17] J. Cheriyan, H. Karloff and Y. Rabani. Approximating Directed Multicuts. *Combinatorica*, 25(3): 251–269, 2005.

[18] J. Chuzhoy and S. Khanna. Polynomial Flow-Cut Gaps and Hardness of Directed Cut Problems. *JACM*, Volume 56, issue 2, Article 6, 2009. Preliminary version in *Proc. of the ACM STOC*, 2007.

[19] J. Edmonds and E. Giles A min-max relation for submodular functions on graphs. *Annals of Discrete Mathematics*, vol.1, 185–204, 1977.

[20] G. Even, J. Naor, S. Rao and B. Schieber, Divide-and-conquer approximation algorithms via spreading metrics. *JACM*, Vol. 47 (2000), pp. 585-616.

[21] J. Fakcharoenphol and K. Talwar. An Improved Decomposition Theorem for Graphs Excluding a Fixed Minor. *Proc. Random-Approx*, Springer Verlag, 2003.

[22] A. Federgruen and H. Groenevelt, Polymatroidal flow network models with multiple sinks. *Networks,*, 18(4): 285–302, 1988.

[23] U. Feige, M.T. Hajiaghayi, and J. Lee. Improved approximation algorithms for minimum-weight vertex separators. *SIAM J. on Computing*, 38(2): 629–657, 2008. Preliminary version in *Proc. of ACM STOC*, 563–572, 2005.

[24] N. Garg, V. Vazirani, and M. Yannakakis. Approximate Max-Flow Min-(Multi)Cut Theorems and Their Applications. *SIAM J. Comput.*, 25(2): 235-251, 1996.

[25] Gagan Goel, Chinmay Karande, Pushkar Tripathi, and Lei Wang. Approximability of combinatorial problems with multi-agent submodular cost functions. *Proc. of IEEE FOCS*, 755–764, 2009.

[26] M. X. Goemans, S. Iwata, and R. Zenklusen. A flow model based on polylinking system. *Math. Programming Ser A*, to appear. Preliminary version in *Proc. Allerton Conf*, Sep. 2009

[27] Anupam Gupta. Improved results for directed multicut. *Proc. of ACM-SIAM SODA*, 454–455, 2003.

[28] A. Gupta, I. Newman, Y. Rabinovich, and A. Sinclair. Cuts, trees and ℓ_1-embeddings of graphs. *Combinatorica*, 24 (2004), pp. 233–269. Preliminary version in *Proc. of IEEE FOCS*, 1999.

[29] R. Hassin. On Network Flows. Ph.D Dissertation, Yale University, 1978.

[30] Satoru Iwata and Kiyohito Nagano. Submodular function minimization under covering constraints. *Proc. of IEEE FOCS*, 671–680, 2009.

[31] N. Kahale. On reducing the cut ratio to the multicut problem. Technical Report TR-93-78, DIMACS 1993.

[32] S. Kannan, A. Raja and P. Viswanath. Local Phy + Global Flow: A Layering Principle for Wireless Networks. *Proc. of IEEE ISIT* , Aug. 2011.

[33] S. Kannan and P. Viswanath. Multiple-Unicast in Fading Wireless Networks: A Separation Scheme is Approximately Optimal. *Proc. of IEEE ISIT*, Aug. 2011.

[34] P. Klein, S. Plotkin and S. Rao. Planar graphs, multicommodity flow, and network decomposition. *Proc. of ACM STOC*, 1993.

[35] P. N. Klein, S. A. Plotkin, S. Rao, and E. Tardos. Approximation Algorithms for Steiner and Directed Multicuts. *J. Algorithms*, 22(2):241–269, 1997.

[36] E. L. Lawler and C. U. Martel, Computing maximal "Polymatroidal" network flows. *Math. Oper. Res.,*, 7(3):334–347, 1982.

[37] J. Lee, M. Mendel, and M. Moharrami. A node-capacitated Okamura-Seymour theorem. *ArXiv preprint*, 1209.2744v1, 2012.

[38] J. Lee and C. Rashtchian. A Simpler Proof Of The KPR Theorem. *Available online:* `http://tcsmath.wordpress.com/2012/01/11/a-simpler-proof-of-the-kpr-theorem/`

[39] J. R. Lee and A. Sidiropoulos. On the geometry of graphs with a forbidden minor. *Proc. of ACM STOC*, 2009.

[40] F. T. Leighton and S. Rao, Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *JACM*, 46(6):787–832, 1999. Preliminary version in *Proc. of ACM STOC*, 1988.

[41] N. Linial, E. London, and Y. Rabinovich The Geometry of Graphs and Some of its Algorithmic Applications. *Combinatorica*, 15(2):215–245, 1995.

[42] L. Lovász. Submodular functions and convexity. *Mathematical programming: the state of the art*, pages 235–257, 1983.

[43] C. Martel. Preemptive Scheduling with Release Times, Deadlines, and Due Times. *JACM*, 29(3):812–829, 1982.

[44] J. Matousek and Y. Rabinovich. On dominated ℓ_1 metric. *Israeli J. of Mathematics*, 123:285–301, 2001.

[45] J. Naor and L. Zosin. A 2-approximation algorithm for the directed multiway cut problem. *SIAM J. on Computing*, 31(2):477–482, 2001. Preliminary version in *Proc. of IEEE FOCS*, 1997.

[46] H. Okamura and P.D. Seymour. Multicommodity flows in planar graphs. *Journal of Combinatorial Theory, Series B*, 31 (1981) 75–81.

[47] S. Plotkin and E. Tardos. Improved Bounds on the Max-Flow Min-Cut Ratio for Multicommodity Flows. *Combinatorica*, 15(3): 425–434, 1995.

[48] Y. Rabinovich. On Average Distortion of Embedding Metrics into the Line. *Discrete & Computational Geometry*, 39(4): 720–733, 2008. Preliminary version in *Proc. of ACM STOC*, 2003.

[49] A. Raja and P. Viswanath. Compress-and-Forward Scheme for a Relay Network: Approximate Optimality and Connection to Algebraic Flows *Proc. of IEEE ISIT*, Aug. 2011.

[50] S. Rao. Small Distortion and Volume Preserving Embeddings for Planar and Euclidean Metrics. *Proc. of ACM Symposium on Computational Geometry*, 1999: 300-306.

[51] M. Saks, A. Samorodnitsky and L. Zosin. A Lower Bound On The Integrality Gap For Minimum Multicut In Directed Networks. *Combinatorica*, 24(3): 525–530, 2004.

[52] A. Schrijver. Combinatorial Optimization: Polyhedra and Efficiency. Springer-Verlag, 2003.

[53] A. Schrijver. Matroids and linking systems. *Journal of Combinatorial Theory B*, 26:349–369, 1979.

[54] D. Shmoys. (Approximation algorithms for) Cut problems and their application to divide-and-conquer. *Approximation Algorithms for NP-hard Problems*, (D.S. Hochbaum, ed.) PWS, 1997, 192-235.

[55] E. Tardos and 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.

[56] S. M. S. Yazdi and S. A. Savari. A Max-Flow/Min-Cut Algorithm for Linear Deterministic Relay Networks. *IEEE Trans. on Info. Theory*, 57(5):3005–3015, 2011. Preliminary version in *Proc. of ACM-SIAM SODA*, 2010.

A Proof of Lemma 16

Lemma 16. *For a polymatroidal network, the dual of the maximum throughput flow problem is equivalent (in terms of value) to the program given in Fig. 1.*

Proof. We will show the proof for the undirected case, the proof for the directed case is similar. The program for maximum throughput flow is given by:

$$\max \quad \sum_i \sum_{p \in \mathcal{P}_{(s_i,t_i)}} f(p)$$

s.t.

$$\begin{aligned} \sum_{e:e\in S} \sum_{p:e\in p} f(p) &\leq \rho_v(S) \quad \forall S \subseteq \delta(v)\, \forall v \in V \\ f(p) &\geq 0 \quad \forall p \in \mathcal{P}_{(s_i,t_i)}, \forall i = 1 \ldots k. \end{aligned}$$

The dual of the flow linear program can now be written. Let the dual variables $d_v(S_v)$ correspond to the non-trivial constraint in the above linear program. Then the dual linear program is:

$$\mathcal{P}_d := \min \sum_{v \in V} \sum_{S \subseteq \delta(v)} d_v(S) \rho_v(S)$$

s.t.

$$\sum_{e=uv: e \in p} \left(\sum_{S \subseteq \delta(u): e \in S} d_u(S) + \sum_{S \subseteq \delta(v): e \in S} d_v(S) \right) \geq 1 \quad \forall p \in \mathcal{P}_{(s_i, t_i)} \text{ where } \quad e = uv$$

$$d_u(S) \geq 0 \; \forall u \in V \; \forall S \subseteq \delta(u).$$

This can be rewritten equivalently as

$$\begin{aligned}
\mathcal{P}_d &:= \min \sum_{v \in V} \sum_{S \subseteq \delta(v)} d_v(S) \rho_v(S) \\
&\text{s.t.} \\
\ell(e) &:= \left(\sum_{S \subseteq \delta(u): e \in S} d_u(S) + \sum_{S \subseteq \delta(v): e \in S} d_v(S) \right) \\
\text{dist}_\ell(s_i, t_i) &\geq 1 \quad 1 \leq i \leq k \\
d_u(S) &\geq 0 \; \forall u \in V \; \forall S \subseteq \delta(u).
\end{aligned}$$

Let us define new variables $\ell(e, u)$, $\ell(e, v)$ for each edge $e = uv$, and rewrite the linear program:

$$\min \sum_{v \in V} \sum_{S \subseteq \delta(v)} d_v(S) \rho_v(S)$$

s.t.

$$\begin{aligned}
\ell(e) &:= \ell(e, u) + \ell(e, v), \text{ where } e = uv \\
\ell(e, u) &= \sum_{S \subseteq \delta(u): e \in S} d_u(S) \; \forall e \in E, e = uv \\
\ell(e, v) &= \sum_{S \subseteq \delta(v): e \in S} d_v(S) \quad \forall e \in E, e = uv \\
\text{dist}_\ell(s_i, t_i) &\geq 1 \quad 1 \leq i \leq k \\
d_u(S) &\geq 0 \\
\ell(e, u), \ell(e, v) &\geq 0 \quad \forall u \in V \; \forall S \subseteq \delta(u).
\end{aligned}$$

The minimization is over the variables $\ell(e, u)$ and $d_v(S)$. Observe for any fixed v the variables $d_v(S), S \subseteq \delta(v)$ influence only the variable $\ell(e, v), e \in \delta(v)$. Hence, for any v and a fixed assignment set of values $\ell(e, v), e \in \delta(v)$ the optimal choice of variables $d_v(S), S \subseteq \delta(v)$ can be obtained by solving the following linear program:

$$\min \sum_{S \subseteq \delta(v)} d_v(S) \rho_v(S)$$

s.t.

$$\begin{aligned}
\sum_{S \subseteq \delta(v): e \in S} d_v(S) &= \ell(e, v) \quad \forall e \in E, e = uv \\
d_u(S) &\geq 0, \quad S \subseteq \delta(v), \; \forall v \in V.
\end{aligned}$$

Recalling the definition of the convex closure of a function, one sees that the value of the above linear program is equal to $\tilde{\rho}_v(\mathbf{d}_v)$; note that for polymatroids we can drop the constraint $\sum_S d_v(S) = 1$ in the linear program for the convex closure. Since the convex closure is equval to the Lovász extension we obtain the desired equivalence of the formulations. $\square$

B Proof of Lemma 9

Lemma 17. *Let $g : V \to [0, \beta]$ be a contraction, let $0 \le a_0 \le a < b \le b_0 \le \beta$ and $S_\theta = \{u \mid g(u) < \theta\}$. Suppose for every edge $e = uv \in \cup_{\theta \in [a,b]} \delta(S_\theta)$, $g(u)$ and $g(v)$ are both in $[a_0, b_0]$. Then,*

$$\int_a^b \nu(\delta(S_\theta)) d\theta \le 2 \sum_{v : g(v) \in [a_0, b_0]} \hat{\rho}_v(\mathbf{d}_v).$$

Proof. Consider an edge $uv \in \delta(S_\theta)$ and for simplicity assume $g(u) < g(v)$. The length of e in the embedding is $\ell'(e) = |g(v) - g(u)| \le \ell(e)$. The edge $(u, v) \in \delta(S_\theta)$ iff θ is in the interval $[g(u), g(v)]$. Also by the conditions of the theory for every such (u, v), $g(u) \in [a_0, b_0]$ and $g(v) \in [a_0, b_0]$. Note that the cost $\nu(\delta(S_\theta))$ is in general a complicated function to evaluate. We upper bound $\nu(\delta(S_\theta))$ by giving an explicit way to assign $e = uv$ to either u or v as follows. Recall that in the relaxation $\ell(e) = \ell(e, u) + \ell(e, v)$ where $\ell(e, u)$ and $\ell(e, v)$ are the contributions of u and v to e. Let $r = \frac{\ell(e,u)}{\ell(e)}$ and let $\ell'(e, u) = r\ell'(e)$ and $\ell'(e, v) = (1-r)\ell'(e)$. We partition the interval $[g(u), g(v)]$ into $[g(u), g(u) + \ell'(e, u))$ and $[g(u) + \ell'(e, u), g(v)]$; if θ lies in the former interval we assign e to u, otherwise we assign e to v. This assignment procedures describes a way to upper bound $\nu(\delta(S_\theta))$ for each θ. Now we consider the quantity $\int_a^b \nu(\delta(S_\theta)) d\theta$ and upper bound it as follows.

Consider a node u and let $L_u = \{uv \in \delta(u) \mid g(v) < g(u)\}$ be the set of edges uv that go from u to the left of u in the embedding g. Similarly $R_u = \{uv \in \delta(u) \mid g(v) \ge g(u)\}$. Note that L_u and R_u partition $\delta(u)$. Let $\mathbf{d}'_u$ be the vector of dimension $|\delta(u)|$ consisting of the values $\ell'(e, u)$ for $e \in \delta(u)$. We obtain $\mathbf{d}^L_u$ from $\mathbf{d}'_u$ by setting the values for $e \in R_u$ to 0 and similarly $\mathbf{d}^R_u$ from $\mathbf{d}'_u$ by setting the values for $e \in L_u$ to 0. Since $0 \le \ell'(e, u) \le \ell(e, u)$ for each $e \in \delta(u)$ we see that $\mathbf{d}'_u \le \mathbf{d}_u$ and (component wise) and hence $\mathbf{d}^L_u \le \mathbf{d}_u$ and $\mathbf{d}^R_u \le \mathbf{d}_u$. Since ρ_u is monotone we have that $\hat{\rho}_u(\mathbf{d}^L_u) \le \hat{\rho}_u(\mathbf{d}_u)$ and $\hat{\rho}_u(\mathbf{d}^R_u) \le \hat{\rho}_u(\mathbf{d}_u)$ (see Proposition 1).

We claim that

$$\int_a^b \nu(\delta(S_\theta)) d\theta \le \sum_{u \in V : g(u) \in [a_0, b_0]} (\hat{\rho}_u(\mathbf{d}^L_u) + \hat{\rho}_u(\mathbf{d}^R_u)),$$

which would prove the lemma.

To see the claim consider some fixed θ and $\nu(\delta(S_\theta))$. Fix a node u and consider the edges in $\delta(u) \cap S_\theta$ assigned to u by the procedure we described above; call this set $A_{\theta,u}$. First assume that $\theta < g(u)$. Then the edges assigned to u by the procedure, denoted by $A_{\theta,u} = \{e \in L_u \mid \theta > g(u) - \ell'(e, u)\}$. Similarly, if $\theta > g(u)$, $A_{\theta,u} = \{e \in L_u \mid \theta < g(u) + \ell'(e, u)\}$. From these definitions we have

$$\begin{aligned}
\nu(\delta(S_\theta)) &\le \sum_{u \in V : g(u) \in [a_0, b_0]} \rho_u(A_{\theta,u}) \\
\Rightarrow \int_a^b \nu(\delta(S_\theta)) d\theta &\le \sum_{u \in V : g(u) \in [a_0, b_0]} \int_a^b \rho_u(A_{\theta,u}) d\theta \\
&\le \sum_{u \in V : g(u) \in [a_0, b_0]} \int_0^\beta \rho_u(A_{\theta,u}) d\theta.
\end{aligned}$$

For a fixed node u,

$$\int_0^\beta \rho_u(A_{\theta,u})d\theta \quad = \quad \int_0^{g(u)} \rho_u(A_{\theta,u})d\theta + \int_{g(u)}^\beta \rho_u(A_{\theta,u})d\theta$$

Let $L_u = \{e_1, e_2, \ldots, e_h\}$ where $0 \le \ell'(e_1, u) \le \ell'(e_2, u) \le \ldots \le \ell'(e_h, u)$. Then

$$\int_a^{g(u)} \rho_u(A_{\theta,u})d\theta = \sum_{j=1}^{h} (\ell'(e_j, u) - \ell'(e_{j-1}, u))\rho(\{e_1, e_2, \ldots, e_j\})$$

The right hand side of the above, is by construction and the definition of the Lovász extension, equal to $\hat{\rho}_u(\mathbf{d}_u^L)$. Similarly, $\int_{g(u)}^\beta \rho_u(A_{\theta,u})d\theta = \hat{\rho}_u(\mathbf{d}_u^R)$. □

C Proof of Theorem 7

Proof. Let F be a set of edges that corresponds to an edge cut and let $V_1, V_2, \ldots, V_h$ be the nodes of the connected components in $G - F$. The sparsity of F is $\nu(F)/D(F)$ where $D(F)$ is the sum of the demands of pairs that are separated by F.

Construct an undirected graph H with nodes $\hat{v}_1, ..., \hat{v}_h$ and edges $\hat{v}_i\hat{v}_j$ with weight w_{ij} equal to the demand between partition V_i and V_j in the original graph G. For graph H, there exists a *weighted max-cut*, whose value is greater than half the sum of all the weights (since a random bi-partition of H where each edge gets cut with probability half has expected weight equal to half the sum of all weights). Let this max-cut partition H into sets A and $V \setminus A$ and let $S = \cup_{i:\hat{v}_i \in A} V_i$. Now consider the bi-partition cut F_S consisting of the edges between S and $V \setminus S$. From the construction of S, we have $D(F_S) \ge D(F)/2$. Moreover, since $F_S \subseteq F$ we have $\nu(F_S) \le \nu(FS)$. Thus, the sparsity of F_S is at most twice that of F.

To see that this factor is tight, consider a polymatroidal network with $n + 1$ nodes v_0, v_1,...,v_n, with edge e_i between v_0 and v_i, for each $i \in \{1, 2, \ldots, n\}$ and assume for simplicity that n is even; the network is a star with center v_0. The only capacity constraint is a polymatroidal constraint at node v_0, which constrains the total capacity of every non-empty subset of $\{e_1, ..., e_n\}$ to be 1 (in effect this simulates a node capacity of 1 at v_0). The demand graph is a complete graph on $v_1, \ldots, v_n$ with each demand value set to 1.

Now consider an edge cut F which removes all the edges: $\nu(F) = 1$ and $D(F) = \binom{n}{2}$, and hence the sparsity is $\frac{2}{n(n-1)}$. Consider a bi-partition cut $(S, V \setminus S)$ such that S does not contain v_0 and let $F_S = \delta(S)$. We have $\nu(F_S) = 1$ and $D(F_S) = |S|(n - |S|)$; the sparsity is minimized when $|S| = \frac{n}{2}$ and is given by $\frac{4}{n^2}$. Thus the sparsity of the best bi-partition cut is a factor of $\frac{2(n-1)}{n}$ bigger than the sparsity of the best edge cut. This factor approaches 2 as n approaches ∞. □