




A tight max-flow min-cut duality theorem for nonlinear multicommodity flows

Matthew Broussard¹ · Bala Krishnamoorthy² 

Accepted: 21 February 2024 / Published online: 11 April 2024

© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2024

Abstract

The Max-Flow Min-Cut theorem is the classical duality result for the MAX-FLOW problem, which considers flow of a single commodity. We study a multiple commodity generalization of MAX-FLOW in which flows are composed of real-valued k -vectors through networks with arc capacities formed by regions in \mathbb{R}^k . Given the absence of a clear notion of ordering in the multicommodity case, we define the generalized max flow as the feasible region of all flow values. We define a collection of concepts and operations on flows and cuts in the multicommodity setting. We study the *mutual capacity* of a set of cuts, defined as the set of flows that can pass through all cuts in the set. We present a method to calculate the mutual capacity of pairs of cuts, and then generalize the same to a method of calculation for arbitrary sets of cuts. We show that the mutual capacity is exactly the set of feasible flows in the network, and hence is equal to the max flow. Furthermore, we present a simple class of the multicommodity max flow problem where computations using this tight duality result could run significantly faster than default brute force computations. We also study more tractable special cases of the multicommodity max flow problem where the objective is to transport a maximum real or integer multiple of a given vector through the network. We devise an augmenting cycle search algorithm that reduces the optimization problem to one with m constraints in at most $\mathbb{R}^{(m-n+1)k}$ space from one that requires mn constraints in \mathbb{R}^{mk} space for a network with n nodes and m edges. We present efficient algorithms that compute ϵ -approximations to both the ratio and the integer ratio maximum flow problems.

Keywords Multicommodity flows · Max-flow min-cut · Sheaf theory

✉ Bala Krishnamoorthy
kbala@wsu.edu

Matthew Broussard
matthewbrouss@gmail.com

¹ TD Bank, N.A., Cherry Hill, USA

² Washington State University, Vancouver, USA

1 Introduction

Network flows are widely used to model transportation systems of commodities or people subject to capacity constraints. Ford and Fulkerson (1956) introduced one of the first motivating examples for network flows in 1956: a system of rail lines connecting cities. Each city is represented by a node in the network and the rail lines connecting them form the arcs. Each arc is then labeled with the maximum number of people its associated rail line can transport in, say, a day—this is the capacity of the arc. Another canonical example is a network of pipes transporting water. Pipes form the arcs of the network and junctions form the nodes. The circumference of a pipe determines the flow rate it permits, which is specified as the capacity of the arc associated with the pipe.

In such examples, we often wish to find the maximum amount of a commodity we can transport. We might wish to find out how many people could possibly travel from New York to Chicago by train, or the maximum flow rate we can achieve through a network of pipes leading to a sink. This class of problems is referred to as the MAX-FLOW problem. A concept central to the study of MAX-FLOW is that of *cuts*, in particular the collections of arcs that, when removed, separate the starting point from our destination. Also called *disconnecting sets* (of arcs), the concept of cuts is dual to flows. The *value* $v(D)$ of a disconnecting set D is the sum of the capacities of all (forward) arcs in D . Ford and Fulkerson (1956) presented the classical theorem linking cuts to flow values, as well as a method for finding the maximum flow by finding a minimal cut.

Theorem 1.1 (Max-Flow Min-Cut Theorem, Ford and Fulkerson (1956)) *The maximal flow value obtainable in a network N is the minimum of $v(D)$ taken over all disconnecting sets D .*

Noting that a disconnecting set with a minimum possible value is automatically a cut, it follows that the maximum flow value is exactly the minimum cut value. The authors described an algorithm to identify the minimum cut in polynomial time by finding augmenting paths.

One generalization of MAX-FLOW is the MIN-COST FLOW problem (Ahuja et al. 1993) where each arc has both a capacity and a cost for each unit of the commodity transported. A subset of the nodes are source (supply) nodes with given amounts of the commodity and another subset are sink (demand) nodes with specific required amounts. The problem is to route the commodity from source nodes to sink nodes such that demands at the sink nodes are met, capacities of arcs are honored, and the total cost is minimized. A further generalization of MIN-COST FLOW to handle multiple types of commodities termed multicommodity minimum cost flows has been well studied (Ahuja et al. 1993, Chap. 17).

However, the direct generalization of MAX-FLOW to the problem with multiple types of commodities without the consideration of costs, supply, and demand has not received as much attention. While this generalizes the standard max-flow problem (Krishnan 2021), the single commodity algorithms may not generalize easily. It is not clear how to generalize the classical result equating maximum flows with minimum cuts (Theorem 1.1) to the multicommodity case. In fact, specifying the boundary of

feasible regions for flows of two or more commodities is already challenging. We focus on this problem (termed **MULT- COM MAX- FLOW**) with the goal of finding the set of feasible flows through the a network.

A more tractable version of this problem could seek a maximum multiple of a certain ratio of commodities that could be transported through a network. This problem is termed **RATIO MAX- FLOW**, and we could also consider a special case termed **INT RATIO MAX- FLOW** which allows only integer multiples. As these problems search for a single value, we may be able to modify some of the single-commodity algorithms to search for the maximum (integer) multiple. Both **RATIO MAX- FLOW** and **INT RATIO MAX- FLOW** could model practical situations. Suppose a factory manufacturing cars uses tires transported from one depot and chassis transported from another. In this case, the desired ratio would be 4 : 1 tires to chassis, and combinations which are not integer multiples of this ratio do not form a full car and are thus not usable. This is a situation that **INT RATIO MAX- FLOW** models. On the other hand, we could model a dye dispensary which mixes a dark green dye from one part yellow and two parts blue using **RATIO MAX- FLOW**, since we can make partial units of the green dye.

1.1 Our contributions

We investigate a general case of **MULT- COM MAX- FLOW** with k commodities where the capacity restrictions on arcs are defined as regions of \mathbb{R}^k . With a goal toward developing a tight multicommodity max-flow min-cut theorem, we define several new concepts related to flows and cuts in this setting. As it is not clear how one would define a *maximum* flow in the multicommodity setting, we analyze the set of all feasible flows instead. We study *pseudoflows* that could violate capacity constraints while honoring flow balance at nodes, and *local flows* over a cut that assign feasible flows for all arcs in the cut (see Sect. 2.2). We also define the concept of *mutual capacity* as the correct generalization of cut capacity in the single commodity case to our general setting (Sect. 3). Our main theorem presents a tight duality result in the multicommodity setting:

Theorem 4.4 *The set of feasible flows through a network is precisely the mutual capacity of all cuts in the network.*

Furthermore, our proof is a constructive algorithm, so calculation of the max flow, which generalizes to the feasible region of flows, gives us the set of all flows that have a particular flow value in the feasible region.

While we address a less general version of the max flow problem than Krishnan (2021) as we restrict ourselves to subsets of \mathbb{R} for our fields of coefficients, we use the restriction to give more detailed descriptions of the behavior of network flows. We present a method for calculating the max flow for any network with capacities in \mathbb{R}^k . Additionally, we prove our results from first principles rather than using the machinery of category theory as done by Krishnan (2021).

We use our duality result to construct an algorithm to solve the general **MULT- COM MAX- FLOW** problem. In addition, we present a class of instances of the problem on which our algorithm calculates the max flow region exponentially faster than the default brute force approach.

Motivated by the work of Leighton and Rao (1999), we also consider more tractable cases of the MULT-COM MAX-FLOW problem that aim to send a maximum multiple of a particular ratio of commodities. We present efficient algorithms for two problems that take advantage of this optimization step. We show that when the capacities of the network are bounded and reducible, these algorithms return accurate results and also terminate in finite time. Furthermore, these algorithms reduce both the size and the dimension of the optimization approach by using the structure inherent in the network. We improve the complexity of the optimization problem from requiring mn constraints in \mathbb{R}^{mk} to m constraints in at most $\mathbb{R}^{(m-n+1)k}$ when there are n nodes and m arcs.

1.2 Previous work

While multicommodity *minimum cost* flow problems are well studied in the context of linear and integer programming (Ahuja et al. 1993, Chap. 17), multicommodity maximum flow problems have not received as much attention. Leighton and Rao (1999) studied a version of the multicommodity max flow problem where a network with scalar capacities on each arc is given for multiple commodities, each with its own source and sink, as well as a desired ratio of commodities. (In contrast, motivated by Krishnan's work (Krishnan 2021), we consider more general classes of capacities.) The max flow problem here was to maximize the ratio transported from source to sink subject to the restriction that the sum of all commodities through an arc cannot exceed its scalar capacity. They also defined an analogue to the min cut called the sparsest cut, which is a cut that minimizes the ratio of the capacity of the cut to the sum of the demands for which the source and sink are separated by the cut.

They summarized several results for this problem, including ones from previous work, showing that the max flow and sparsest cut are equal for two-commodity flows, that the max flow and sparsest cut are equal provided that the graph constructed with edges (s_i, t_i) for each commodity i has no set of three disjoint edges and no triangle along with a disjoint edge, and further information on special cases. Their main result was a general proof that the max flow is within a $\Theta(\log_2 n)$ -factor of the sparsest cut. They also developed an algorithm for determining an approximate max flow.

Ghrist and Krishnan (2013) and then Krishnan (2021) used sheaf theory to develop a notion of directed homology and cohomology. This framework allowed the characterization of flows and cuts as directed homology and cohomology classes. The "min cut" was defined as the intersection over all cuts of the Minkowski sum of arc capacities in cuts, and the max flow was replaced by the feasible region of flows. Poincaré duality was then used to establish a general max-flow min-cut theorem. However, the general theorem presented by Ghrist and Krishnan (2013) had duality gaps in that the region given as the set of feasible flows could include some flow values that cannot be achieved.

The manuscript titled *Flow-Cut Dualities for Sheaves on Graphs* (Krishnan 2021) ends with two cases for which the max flow is equal to the min cut. If all paths from source to sink pass through a minimum, i.e., have an arc which has a capacity that is a subset of each capacity along a path, or if the capacities come from a lattice-ordered semi-module, then the min cut and feasible region for flows are equal.

However, this approach does not give a general tractable solution to the multi-commodity flow problem since the model given for multicommodity flow, which uses Minkowski sums and intersections, does not form a lattice order (in general, $(A +_M B) \cap (A +_M C) \neq A +_M (B \cap C)$). Furthermore, there exist networks where some paths do not pass through a minimum, and indeed transforming the bounds on arcs so as to make all paths pass through a minimum modifies the real network capacity. Our approach develops a tight bound on max flow using arguments from first principles. With the bound established, we show how our results fit into Krishnan's framework by modifying certain definitions.

Even et al. (1976) studied timetable and integral multicommodity flow problems, showing both directed and undirected versions of the multicommodity flow problems are NP-complete. In the undirected case, the flows are assigned a direction (\pm) even though the edges themselves are undirected. The capacity functions considered in this work were restricted to natural numbers as upper bounds on the simple sums of flows of all commodities. Our multicommodity flow problem allows highly general nonlinear capacity functions, and hence includes the version studied by Even, Itai, and Shamir as a special case. As such, our multicommodity flow problem turns out to be NP-complete as well.

2 Problems and definitions

We first formally define the three multicommodity max-flow problems of interest. We then show that the main problem turns out to be NP-complete, as it covers as special cases a class of 2-commodity flow problems that are known to be NP-complete. We then list various definitions related to networks and flows that we use throughout the paper. Several of them are standard concepts defined in previous work, but are listed here for the sake of completeness. Definitions without citations are our original contributions.

Following standard convention (Ahuja et al. 1993), we let n and m denote the number of nodes and arcs in the network, and k denote the number of commodities. We denote the capacity of arc a by C_a , and will drop the subscript when the arc is clear from context. Furthermore \mathcal{C} denotes a set of cuts, and is also referred to as a cut set.

Definition 2.1 (*Multicommodity Max Flow* Krishnan (2021)) Given a network with capacities in k commodities specified as subsets of \mathbb{R}^k , MULT-COM MAX-FLOW aims to find the set of feasible flows through the network.

Definition 2.2 (*Ratio Max Flow* Leighton and Rao (1999)) Given a network with real capacities in k commodities, RATIO MAX-FLOW aims to find the flow that is the maximum multiple of a desired ratio of the k commodities.

Definition 2.3 (*Integer Ratio Max Flow*) Given a network with real capacities in k commodities, INT RATIO MAX-FLOW aims to find the flow that is the maximum integer multiple of a desired ratio of the k commodities.

Note that the arc capacities we consider (see Definition 2.10) are fairly general—in particular, we do not assume they are convex, or even compact, subsets of \mathbb{R}^k . Even, Itai, and Shamir studied versions of the integer multicommodity flow problem (Even et al. 1976), and their connections to certain timetable scheduling problems. They showed these classes of problems are NP-complete for both directed and undirected networks. These problems turn out to be special cases of MULT-COM MAX-FLOW.

More precisely, D2CIF is the directed 2-commodity integer flow problem where the directed network has a pair (s_i, t_i) of origin and destination nodes for $i = 1, 2$, and integer requirements $R_i \geq 0$ on the amount of outflow of commodity i from s_i . The two source nodes s_1, s_2 could be identical, and so could be the destination nodes t_1, t_2 . The capacity restrictions on the edges are nonnegative integers that present upper bounds on the simple sum of flows of both capacities. The goal of D2CIF is to decide if a nonnegative integer feasible flow exists that satisfies the requirements. The related U2CIF considers the same problem on *undirected* networks, but with the flows still given a positive or negative direction. The capacity restrictions are now applied to the sum of absolute values of directed flows of the two commodities.

Given an instance of D2CIF (or U2CIF), we can create an equivalent instance of MULT-COM MAX-FLOW as follows. We add extra nodes s, s', t , and edges $(s, s'), (s', s_1), (s', s_2)$, as well as $(t_1, t), (t_2, t)$ (directed or undirected as needed). The capacity restriction for edge (s, s') is set as $(f_1, f_2) \geq (R_1, R_2)$ and $f_1, f_2 \in \mathbb{Z}$, the set of integers, i.e., the given requirements are set as elementwise lower bounds on the flow of the two commodities that are also required to be integral. All other added arcs are set as uncapacitated. It follows that the answer to D2CIF is Y_{eS} if and only if a solution for the instance of MULT-COM MAX-FLOW exists, i.e., the set of feasible flows is non-empty.

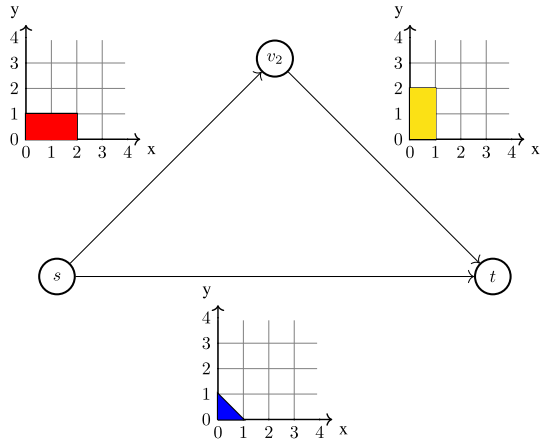
Since D2CIF and U2CIF were shown to be NP-complete (Even et al. 1976), it turns out that MULT-COM MAX-FLOW is NP-complete as well. While we focus on instances of MULT-COM MAX-FLOW where the capacities are fairly general, specific instances may admit efficient solution approaches, e.g., when the capacities are all closed hyperrectangles specifying simple upper bounds on flows of individual commodities.

2.1 Definitions on networks

Definition 2.4 (*Single commodity Network* Ahuja et al. (1993)) A single commodity network N is a directed graph with node set V and arc set E such that each arc $a \in E$ has exactly one capacity C_a . These capacities are regions in \mathbb{R} ranging from 0 to a maximum value. We denote this region by noting only the maximum value. The network has a distinguished source node s and terminus node t .

Definition 2.5 (*k-commodity Network*) A k -commodity network N for $k \geq 2$ is a directed graph with node set V and arc set E such that each arc $a \in E$ has exactly one capacity C_a . These capacities are regions embedded in \mathbb{R}^k , where each dimension corresponds to a particular commodity. Arcs are directed as a way to fix an orientation with regards to the capacity region, i.e., flow in the direction of the arc is considered positive flow along the arc and flow in the opposite direction is considered negative.

Fig. 1 Example of a 2-commodity network. Capacities of the arcs are the regions shown in red, yellow, and blue



We can reverse the orientation of the arc if we multiply all vectors in the arc’s capacity by -1 . The network has a distinguished source node s and terminus node t .

Note that it is possible for flow of one commodity to be in the positive direction while another commodity flows in the negative direction. For instance, suppose we have the (oriented) arc (a, b) in the network and its capacity includes the vector $(2, -3)$. Hence the network can transport two units of commodity 1 from a to b so long as it also transports three units of commodity 2 from b to a . We could also consider the arc in the orientation (b, a) , whose capacity includes the vector $(-2, 3)$. This represents the same flow in the network, i.e., two units of commodity 1 from a to b and three units of commodity 2 from b to a .

Definition 2.6 (*Opposite Orientation*) An arc (a, b) with capacity $C \subseteq \mathbb{R}^k$ is associated with the arc (b, a) with capacity $-C \subseteq \mathbb{R}^k$. These two arcs are said to have opposite orientations.

Since our focus is on multicommodity flows, we use the terms “network” and “ k -commodity network” interchangeably. We say “single commodity network” when referring to cases where there is only one commodity.

Example 2.7 Consider a 2-commodity network with $V = \{s, v_2, t\}$, $E = \{(s, v_2), (v_2, t), (s, t)\}$, and capacities on the three arcs as shown in Fig. 1.

Note that we have chosen the arc capacities in most of the examples as simple, yet nontrivial, sets in order to keep the examples accessible and yet insightful. For instance, the capacities of the arcs in the network in Example 2.7 (and in the related Example 3.1) are simplex convex polygons. Capacities in Figs. 3, 4, 5 are finite sets of 2D integer points (from $\mathbb{Z}_{\geq 0}^2$). And in Example 3.4, the capacity of arc $(2, 3)$ is specified as a disjunction. At the same time, the definitions and results apply to most general networks—no particular structure is assumed to be satisfied by the arc capacities unless specified otherwise.

Definition 2.8 (*Enhanced Network* Krishnan (2021)) An enhanced network N^E is a network N together with an arc e directed from t to s with the capacity constraints $x_i \geq 0$ for all commodities i , where x_i is the flow value of commodity i in the arc e .

We assume that all networks are equipped with this arc e by default, and use the term “network” to refer to enhanced networks. Likewise, we omit e from diagrams and examples for the brevity of illustration. We also denote the arc set of the enhanced network by E .

Definition 2.9 (*Boundary Functions* Krishnan (2021)) The functions $\partial_-, \partial_+ : E \rightarrow V$ provide the tail or head of an arc, respectively. That is, $\partial_-(u, v) = u$ and $\partial_+(u, v) = v$ for all arcs $(u, v) \in E$.

Note we often use a or e to denote arcs rather than listing them as explicit ordered pairs.

Definition 2.10 (*Capacity* Krishnan (2021)) The capacity of an arc is defined as the set of vectors in \mathbb{R}^k which are permissible for that arc, i.e., all possible amounts of the k commodities that we can transport across that arc at the same time.

Capacities can be quite general sets of vectors in \mathbb{R}^k . However, we restrict ourselves to two-dimensional polytopes or integer points within the same in our example networks for ease of calculation and visualization. Our results do hold for the most general forms of capacity unless noted otherwise.

Definition 2.11 (*Cut*) A cut in a k -commodity network is a partition of nodes in V into two sets S and T such that $s \in S$ and $t \in T$. We associate with each cut all arcs that originate at nodes in S and terminate at nodes in T , i.e., forward arcs, in their original orientation. Similarly, we associate all arcs that originate at nodes in T and terminate at nodes in S , i.e., backward arcs, in their opposite orientation. The capacity of the cut is $v(\{a_i\})$ where $\{a_i\}$ is the set of appropriately oriented arcs in the cut.

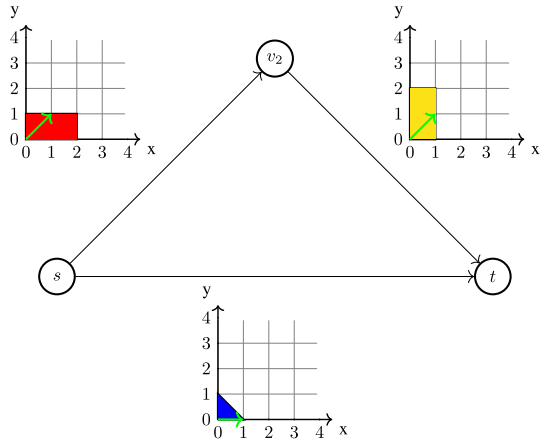
Note that we study only s, t -cuts, i.e., partitions that separate s and t . More generally, any partition of V may be referred to as a cut in some literature (Ahuja et al. 1993). We further note that, unlike in the single commodity case, forward and backward arcs can contribute to the capacity of a cut. Also note that e is a backward arc over all cuts in our setting.

Definition 2.12 (*Arc Set*) The arc set $E(c)$ of a cut c is the collection of forward and backward arcs over the cut, i.e., arcs from S to T or vice-versa, along with their orientations. The arc set of a set of cuts \mathcal{C} is similarly defined as $E(\mathcal{C}) = \cup_{c \in \mathcal{C}} E(c)$.

Definition 2.13 (*Value Function*) We denote the capacity of a set of arcs $\{a_i\}_{i \in I}$ by $v(\{a_i\})$ and define it as the Minkowski sum of the capacities of the individual arcs. We refer to v as the value function. When $\{a_i\}$ is the arc set of a cut and the arcs are oriented from s to t , we call $v(\{a_i\})$ the capacity of the cut.

In practice, we exclude the capacity of the arc e in calculations of a network’s capacity. Recall that e by definition cannot transport commodities from s to t as it takes only nonnegative values from t to s . Hence the inclusion of e in the value function does not give any useful information on the network’s capacity.

Fig. 2 Example of a multicommodity flow in the network shown in Fig. 1. The green vectors denote flow values for the two commodities (x, y) on each arc



Definition 2.14 (*Total Capacity* Ghrist and Krishnan (2013)) The *total capacity* of a set of cuts is the intersection of their capacities.

Definition 2.15 (*Mutual Capacity*) The *mutual capacity* $U_m(\mathcal{C})$ of a set of cuts \mathcal{C} is the set of flow values f for which there exists an assignment of flow values to the arcs in the arc sets of \mathcal{C} that respect the capacity constraints and give a total value of f over each cut in \mathcal{C} .

Note that the mutual capacity must be a subset of total capacity, but the two need not be equal. See Example 3.1 for an illustration.

Definition 2.16 (*Reducible Capacity*) A capacity C for an arc is *reducible* if for every flow value $f \in C$ and every flow value $0 \leq f' \leq f$ (element wise), $f' \in C$.

Working with reducible capacities gives us the ability to speed up searches for a max flow since we can iteratively search increasing flows. We can then terminate our search when we find a flow with no realization.

2.2 Flows and flow-like objects

Definition 2.17 (*Flow* Krishnan (2021)) A *flow* on a network is a function $\phi : E \rightarrow \mathbb{R}^k$ subject to the following restrictions:

- conservation: for $v \in V, \sum_{a \in \partial_-^{-1}(v)} \phi(a) = \sum_{a \in \partial_+^{-1}(v)} \phi(a)$;
- capacity constraints: for $a \in E, \phi(a) \in C_a$ and $\phi_e \geq 0$ for the edge $e = (t, s)$.

Note that $\phi(v) = 0$ for all nodes v in an enhanced network N^E .

Example 2.18 Figure 2 shows a flow in the 2-commodity network presented in Example 2.7.

Definition 2.19 (*Flow Value on a Network* Krishnan (2021)) The *flow value* of a given flow on a network is the net value assigned to the arcs in a cut set for that network. Equivalently, a flow on an enhanced network has a value equal to the value assigned to arc e .

We prove in Lemma 4.1 that a flow takes the same total value across each cut, so the choice of the cut in the above definition is immaterial. For example, the flow value of flow shown in Fig. 2 is the net outflow of s , which is $(1, 1) + (1, 0) = (2, 1)$.

Definition 2.20 (*Flow Value on an Arc* Krishnan (2021)) The *flow value on an arc* is the value assigned to that arc by a flow.

We define a pseudoflow by loosening the capacity constraints rather than the conservation constraints.

Definition 2.21 (*Pseudoflow*) A *pseudoflow* on a network is a function $\phi : E \rightarrow \mathbb{R}^k$ subject to flow conservation, i.e., for $v \in V$, $\sum_{a \in \partial_-^{-1}(v)} \phi(a) = \sum_{a \in \partial_+^{-1}(v)} \phi(a)$.

A pseudoflow can be considered as an ideal transportation situation where transports have no limits. In our approach, we will find realizations of particular flow values by transforming pseudoflows into flows.

We are also interested in assigning values to a subset of arcs. Our first definition gives terminology to discuss a particular assignment of flow values to a particular cut.

Definition 2.22 (*Local Flow*) Given a set of cuts \mathcal{C} , a *local flow* $\phi_L : E(\mathcal{C}) \rightarrow \mathbb{R}^k$ is an assignment of values to the arcs in the cut set $E(\mathcal{C})$ subject to the following restrictions:

- capacity constraints: for $a \in E(\mathcal{C})$, $\phi_L(a) \in C_a$;
- flow value: for c in \mathcal{C} , $\phi_L(e)$ is the sum $\sum_{a \in F} \phi_L(a) - \sum_{a \in B} \phi_L(a)$, where F and B are the sets of forward and backward arcs in c , respectively.

The value assigned to e by ϕ_L is called the *flow value* of the local flow.

Note that since $e \in E(\mathcal{C})$ because e is in all cuts, e has a value assigned to it by any local flow. Also, a local flow assigns values to the arcs in a set of cuts in such a way that net flow over each cut is equal. When reading the following definitions, recall that each local flow is defined on its given set of cuts.

Definition 2.23 (*Compatible*) A set of local flows are *compatible* if for each arc a in the union of arc sets of the local flows, $\phi_L(a)$ is equal for each local flow ϕ_L that contains a in its arc set.

Note that compatibility requires the local flows to assign the same value to e , so a set of local flows must each have the same flow value to be compatible. For example, Fig. 4 presents a local flow compatible with that in Fig. 3.

Definition 2.24 (*Gluing*) The *gluing* G of a compatible collection of local flows $(\{\phi_L\})$ is another local flow. The arc set of G is the union of arc sets of local flows in $\{\phi_L\}$, and likewise the cut set of G is the union of the cut sets of the local flows. For each arc a in the arc set of G , it holds that $G(a) = \phi_L(a)$ for any ϕ_L that has a in its arc set. Since all ϕ_L are compatible, this is a unique definition of $G(a)$.

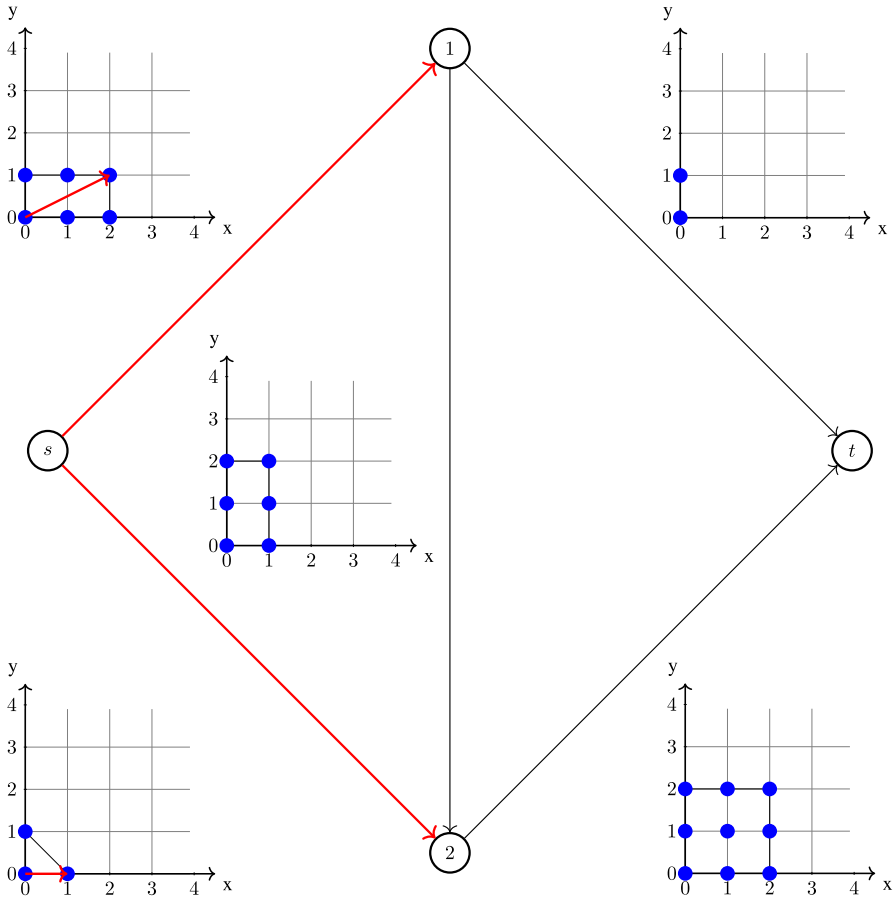


Fig. 3 Example of a local flow over a cut (in red) with arc set $\{(s, 1), (s, 2)\}$ (Color figure online)

For instance, the local flows in Figs. 3 and 4 can be glued together as they are compatible local flows. This gluing is a local flow over a pair of cuts shown in Fig. 5.

A gluing is local flow over a broad set of cuts built from local flows over smaller subsets of those cuts in a way that preserves the assignments given by the local flows that are glued together. We require consistency so that no two local flows try to assign different values to the same arc. Further, since e takes a value under each local flow, we also ensure that the local flows have the same net value over every cut in the cut set. We will build global flows as gluings of local flows that fit together in this manner.

2.3 Operations

We define the operations we use to consider and compare capacities of arcs, which are defined as regions of k -vectors. The basic operation is the Minkowski sum.

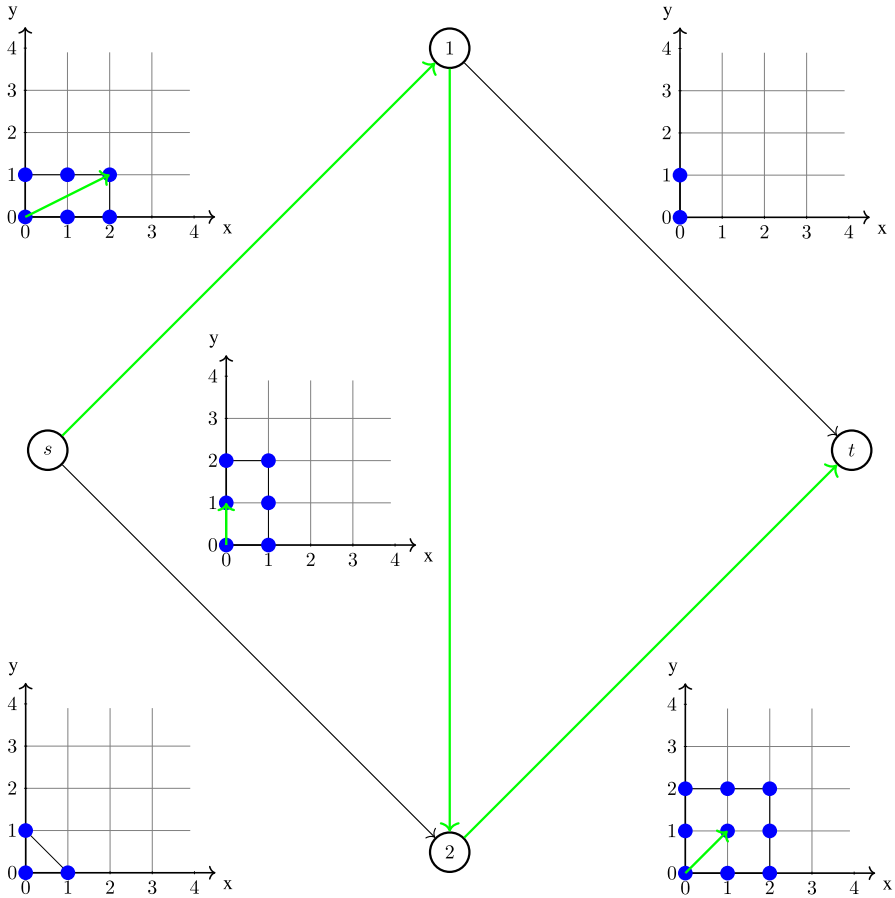


Fig. 4 A second example of a local flow on the same network used in Fig. 3. Arc set of the single cut is shown in green. Note that arc $(1, 2)$ is a backward arc of this cut, and hence subtracts from the local flow's flow value (Color figure online)

Definition 2.25 (Minkowski Sum Delos and Teissandier (2014)) The *Minkowski sum* C of two sets of vectors A and B (denoted by $+_M$) is the set of vectors that can be written as the sum of a vector from A and a vector from B .

It makes sense to consider the capacity of a cut to be the Minkowski sum of the capacities of its arcs. After all, we can transport $\mathbf{x} = (x_1, x_2, \dots, x_k)$ goods over a cut if and only if we can split it up into one portion for each arc in the cut which fits through that arc. That is, we can transport \mathbf{x} only if it is the element-wise vector sum of the capacities of arcs in the cut, which is the Minkowski sum.

However, when we try to compare capacities for different cuts, the Minkowski sum may give misleading results if a particular flow value requires an arc a to take one value for one of the cuts it is involved with for that cut to reach a value f , but another cut requires a different value on a to reach f (see Example 3.1).

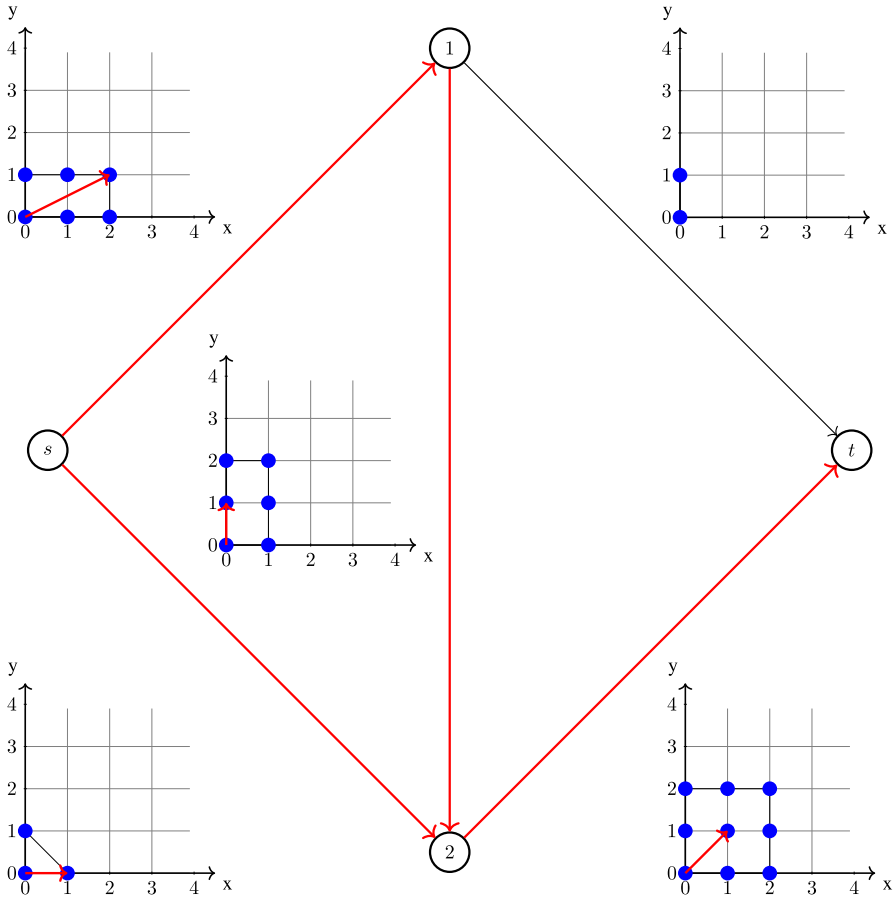


Fig. 5 A local flow over a pair of cuts. This flow is obtained by gluing (see Definition 2.24) together the flows in Figs. 3 and 4

To address this discrepancy, we introduce an extension of the Minkowski sum that not only tracks the total flow value across a cut as the standard Minkowski sum does, but also records the arcs contributing to the flow and the flow values assigned to each arc. Subsequently we define an intersection that accurately detects and discards such “double booking”.

Definition 2.26 (Enriched Minkowski Sum) The enriched Minkowski sum of the capacities of a collection $\{a_i\}$ of arcs in a cut is the Minkowski sum $+_M : \prod C_{a_i} \rightarrow \mathbb{R}^k$ of those capacities along with the level set under $+_M$ in $\prod C_{a_i}$ of each vector in the Minkowski sum. That is, we find all vectors that can be produced as an element-wise vector sum from the arc capacities. We further include for each such vector the set of all legal vector assignments to arcs in the cut which produce that vector.

For instance, in Fig. 6 the enriched Minkowski sum of each point on the RHS of the two equations (corresponding to the two cuts) will have each vector v associated

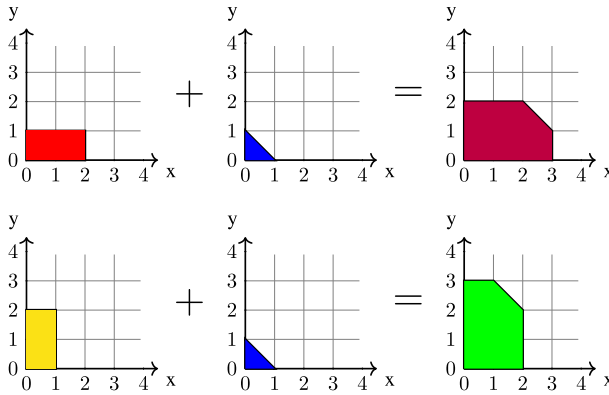


Fig. 6 Capacities of the two cuts C_1, C_2 in Fig. 1. $U(a_1)+U(a_3) = U(C_1)$ (first row) and $U(a_2)+U(a_3) = U(C_2)$ (second row)

with the set of all pairs of vectors from a_1 and a_3 or a_2 and a_3 , respectively, that sum to v . Note that the Enriched Minkowski sum gives the set of all local flows over a cut indexed by their flow values.

We consider a generalization of intersection for enriched Minkowski sums of cuts. This generalization finds flow values in the capacities of a collection of cuts, but only so long as there is a consistent assignment of flow values to arcs which produces the desired flow value across each compared cut—i.e., provided no cuts require double (or more) booking an arc. We show that taking this intersection over the enriched Minkowski sums of all cuts returns the region of feasible flows along with all realizations of flows that produce any desired flow value.

Definition 2.27 (Enriched Intersection) The *enriched intersection* of enriched Minkowski sums over a collection of cuts \mathcal{C} is the intersection of the Minkowski sum of capacities over each cut, along with, for each vector v in the intersection of Minkowski sums, the set of gluings of consistent local flows with flow value v for each $c \in \mathcal{C}$.

We will use intersection in place of enriched intersection and Minkowski sum (or even simply sum) in place of enriched Minkowski when the context is clear. Likewise, when it is clear that we are taking the sum of a set of vectors, we will use $+$ rather than $+_M$ for the Minkowski sum. We also use the \cap symbol for the general intersection when it is clear that we are taking the general intersection of enriched Minkowski sums.

3 Mutual capacity: towards an MFMC result

Our goal is to define a capacity restriction to replace the straight intersection of cut capacities as used by Krishnan (2021) which allows us to reduce (or eliminate) the duality gap. To this end, we first introduce a tighter notion of combining capacities of two cuts, and then generalize to higher number of cuts.

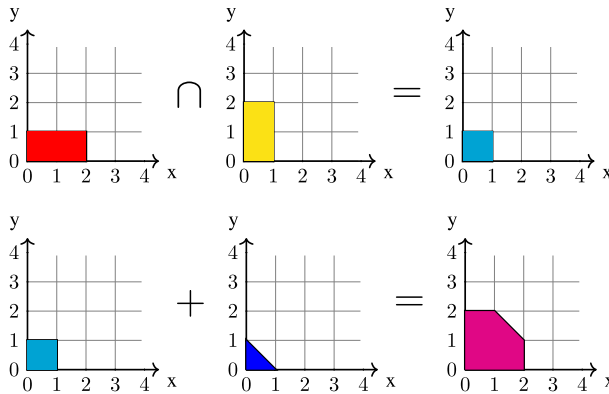


Fig. 7 Pairwise capacity of two cuts C_1, C_2 in Fig. 1. First row shows $v(a_1) \cap v(a_2)$. Second row shows the pairwise capacity as $(v(a_1) \cap v(a_2)) + v(\{a_3\})$. Note that $C_1 \cap C_2 = \{a_3\}$

3.1 Capacity of two cuts: pairwise capacity

We first illustrate the tighter notion of combining capacities of two cuts on a simple example.

Example 3.1 Consider the 2-commodity network introduced in Fig. 1. Naming the arcs as $a_1 = (s, v_2), a_2 = (v_2, t)$, and $a_3 = (s, t)$, we get two cuts $C_1 = \{a_1, a_3\}$ and $C_2 = \{a_2, a_3\}$ with capacities shown in Fig. 6.

Definition 3.2 (Pairwise Capacity) The pairwise capacity of two sets of arcs $A = \{a_i\}$ and $B = \{b_j\}$ is defined as $U_P(A, B) = v(A \cap B) + (v(A \setminus B) \cap v(B \setminus A))$. In words, we take the capacity of the set of arcs in the intersection of the two sets, then add the intersection of the capacity of the arcs unique to each set.

We refer to pairwise capacity also as *two-cut mutual capacity*, as it is a special case of mutual capacity. For instance, the pairwise capacity $U_P(C_1, C_2)$ for the cuts in our example requires us to find $C_1 \cap C_2 = \{a_3\}$. Also note that C_1 has the unique arc a_1 and C_2 has the unique arc a_2 . Hence we get $U_P(C_1, C_2) = v(\{a_3\}) + v(a_1) \cap v(a_2)$, as shown in Fig. 7.

We note that $A + (B \cap C) \subseteq (A + B) \cap (A + C)$ for the Minkowski sum, but the two sets are not equal in general. We also note that for a set of arcs A partitioned into sets B and C , $v(A) = v(B) + v(C)$. This is true because any flow a in $v(A)$ can be written as $a = b + c$ where $b \in v(B)$ and $c \in v(C)$, and any flow that can be written as a sum of a flow in $v(B)$ plus one in $v(C)$ lies in $v(A)$.

It follows that $U_P(A, B) \subseteq v(A) \cap v(B)$. Applied to a family of cuts \mathcal{C} , we can then say that $\bigcap_{c_i, c_j \in \mathcal{C}} U_P(c_i, c_j) \subseteq \bigcap_{c_l \in \mathcal{C}} C_l$. And since the total capacity is given by $\bigcap_{c_l \in \mathcal{C}} C_l$, we can say that $\bigcap_{c_i, c_j \in \mathcal{C}} U_P(c_i, c_j)$ is in general a tighter constraint than that given by Ghrist and Krishnan (2013).

We demonstrate that total capacity may include some flow values that lack a feasible realization in Figs. 8 and 9. For the network in Fig. 1, the total capacity of the two

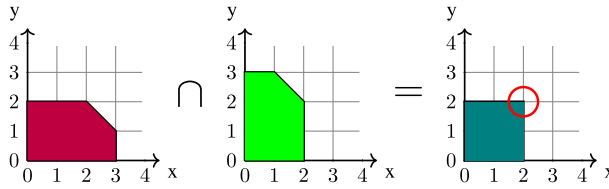


Fig. 8 Capacities $U(C_1)$ (left), $U(C_2)$ (center), and total capacity (right). Figure 6 shows the computation of $U(C_1)$ and $U(C_2)$. Note the vector $(2, 2)$ lies in total capacity

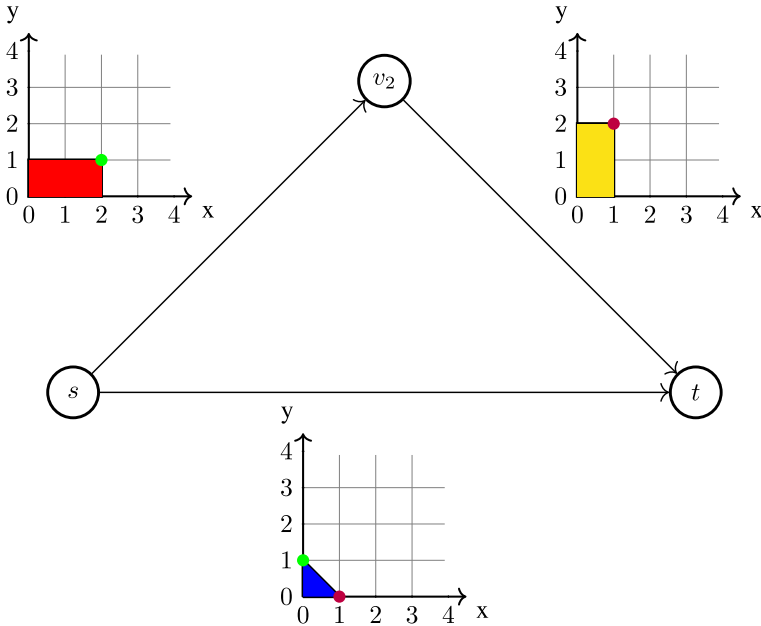


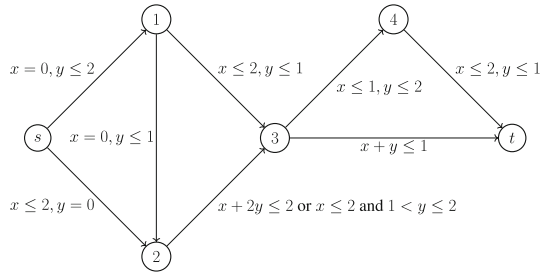
Fig. 9 Demonstration of gap in total capacity. Arc flows needed to realize the flow $(2, 2)$, which is contained in the total capacity as shown in Fig. 8, are shown as green dots for cut C_1 and pink dots for cut C_2 . But both flows (green and pink dots) cannot be *simultaneously* realized in a_3 (Color figure online)

cuts C_1, C_2 is shown in Fig. 8. At the same time, the corresponding pairwise capacity shown in Fig. 7 is strictly smaller than the total capacity.

We observe that flow $(2, 2)$ is contained in the total capacity. To achieve a flow of $(2, 2)$ through the network, cut C_1 requires a flow of $(2, 1)$ through a_1 and $(0, 1)$ through a_3 . At the same time, cut C_2 requires a flow of $(1, 2)$ through a_2 and $(1, 0)$ through a_3 . Since a_3 cannot take a flow of $(0, 1)$ and $(1, 0)$ at the same time, the network cannot transport $(2, 2)$.

The example demonstrates that pairwise capacity could avoid gaps allowed by total capacity. To prove that pairwise capacity still bounds multicommodity flows, we must show that any valid flow must lie within the given capacity bounds. The following lemma presents this result.

Fig. 10 A network for which pairwise capacity reduces but does not eliminate the duality gap



Lemma 3.3 Given a pair of cuts C_i, C_j , any valid flow must fall within

$$U_P(C_i, C_j) = v(C_i \cap C_j) + (v(C_i \setminus C_j) \cap v(C_j \setminus C_i)),$$

i.e., the sum of the values of all arcs common to the cuts in the collection plus the intersection of cut values not included in the intersection.

Proof For any pair of cuts, flow can be partitioned into a portion through the common arcs and flow through the remaining arcs. Since the flow can assign only one value to each of the common arcs, we must be able to write the value of flow as $a + b$ and $a + c$ (in standard vector addition) where a is in v (the common arcs), b is in v (the arcs unique to C_i), and c is in v (the arcs unique to C_j). Since the flow value across each cut must be equal, $b = c$. Hence $a \in v(C_i \cap C_j)$ and $b = c \in (v(C_i \setminus C_j) \cap v(C_j \setminus C_i))$. Thus the feasible flow values through C_i and C_j must fall within $U_P(C_i, C_j)$. \square

We further note that since the feasible flow values through any pair of cuts (C_i, C_j) is bounded by $U_P(C_i, C_j)$, the feasible flow values through the network must be bounded by $\bigcap_{i,j} U_P(C_i, C_j)$ taken over all pairs of cuts.

3.1.1 Duality gap of pairwise capacity

While pairwise capacity gives us a tighter approximation than total capacity, it may still not capture the exact feasible region. We present an example that illustrates this gap.

Example 3.4 We present in Fig. 10 a network for which taking the intersection over all $s-t$ cuts gives a duality gap with the true maximum flow.

The intersection over all cuts gives the capacity $(x \leq 2, y \leq 2)$ for this network. However, we make the following observations.

- We cannot get any x to node 1, so we can rewrite the capacity of $(1, 3)$ as $(x = 0, y \leq 1)$.
- We cannot get $(x = 2, y = 2)$ to node 2, but we can get up to two units of x and one unit of y , so the capacity can be rewritten as $(x + 2y \leq 2)$.

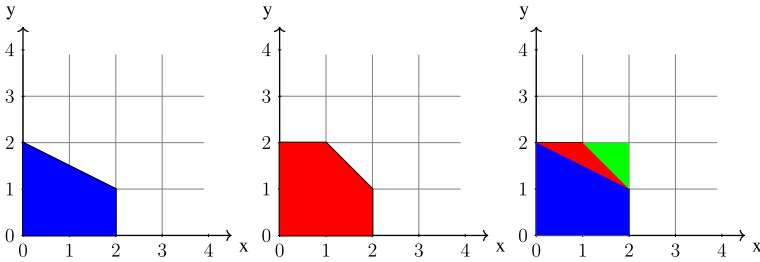


Fig. 11 Feasible regions: true feasible region in blue (left), the region from pairwise capacities in red (middle), and the total capacity in green shown along with the other two regions for easy comparison (right) (Color figure online)

We can determine the true feasible region using these observations as shown in Fig. 11. For instance, it is impossible to find a flow with flow value $(1, 2)$. Since $(s, 1)$ can transport only commodity y and $(s, 2)$ can transport only commodity x , we must route two units of y through $(s, 1)$. Likewise, we must route one unit of x through $(s, 2)$. Then since inflow equals outflow, we must send one unit of y through each of $(1, 2)$ and $(1, 3)$. But this means that we must send one unit each of x and y along $(2, 3)$. However, this is outside the capacity of $(2, 3)$. Hence $(1, 2)$ is outside the feasible region of flow values even though it lies in all pairwise capacities in this network.

To find the mutual capacity of the cuts $\{(s, 2), (1, 2), (1, 3)\}$ and $\{(s, 1), (s, 2)\}$, we take the capacity of the common arc $(s, 2)$, i.e., $(x \leq 2, y = 0)$, and add the intersection of the sum of capacities of arcs unique to each cut, i.e.,

$$(x = 0, y \leq 2) \cap ((x = 0, y \leq 1) + (x \leq 2, y \leq 1)) = (x = 0, y \leq 2),$$

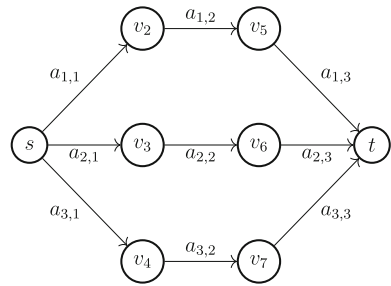
for a total bound of $(x \leq 2, y \leq 2)$. Hence for the portion of the network with nodes $\{s, 1, 2, 3\}$, we get an approximation of the actual capacity given by $(x \leq 2, y \leq 2)$.

For the portion of the network with nodes $\{3, 4, t\}$, we get the region in the center in Fig. 11 (this calculation is identical to the one shown in Fig. 7). As this region is a subset of $(x \leq 2, y \leq 2)$, the intersection over all two-way mutual capacities gives us the same region shown in the center of Fig. 11. Note how the true capacity of the network is contained in the approximation from pairwise capacities, which in turn is contained in the total capacity.

3.2 Mutual capacity of fully disjoint networks

Before we consider the generalization of pairwise capacity to the case of combining capacities of multiple cuts in the most general setting, we present the result for a class of networks with specific structure. The reader will find our arguments for this special case easier to follow before transitioning to the general case. We say a network N is *fully disjoint* if all of its $s-t$ paths are node disjoint except at s and t . We illustrate some new definitions and results on the network in Fig. 12, where $a_{i,j}$ has capacity $c_{i,j}$.

Fig. 12 Example network for restrictions



To use mutual capacity to find the true capacity of a fully disjoint network, we develop a few more nuanced results. First, we expand our idea of mutual capacity.

Definition 3.5 (Restriction) A restriction on a network is a constraint on the flow capacities that can be written as a sum over all possible intersections of arc capacities.

All cuts are restrictions, as are the mutual capacities of pairs of cuts.

Definition 3.6 (Compatible) Two restrictions are said to be compatible if their capacities differ in exactly one summand.

In the network in Fig. 12 the cuts $C_{1,1} = \{a_{1,1}, a_{2,1}, a_{3,1}\}$ and $C_{1,2} = \{a_{1,2}, a_{2,1}, a_{3,1}\}$ are compatible as their capacities differ only along the first path. Likewise, $U_P(C_{1,1}, C_{1,2}) = c_{2,1} + c_{3,1} + (c_{1,1} \cap c_{1,2})$ is compatible with $C_{1,3} = \{a_{1,3}, a_{2,1}, a_{3,1}\}$ since $U(C_{1,3}) = c_{2,1} + c_{3,1} + c_{1,3}$, which differs from $U_P(C_{1,1}, C_{1,2})$ in exactly one term ($c_{1,3}$ rather than $(c_{1,1} \cap c_{1,2})$). However, $C_{1,2}$ and $C_{2,2} = \{a_{1,1}, a_{2,2}, a_{3,1}\}$ are not compatible since they each have two unique arcs and hence two unique elements in their capacities.

The pairwise capacity of a pair of compatible restrictions is a generalization of the pairwise capacity calculation for a pair of cuts. Let $R_1 = \{a_1, \dots, a_n\}$ and $R_2 = \{a_1, \dots, a_{n-1}, b_n\}$. Then $U_P(R_1, R_2) = U(\{a_1, \dots, a_{n-1}\}) + (U(a_n) \cap U(b_n))$. For instance, we noted above that $U_P(C_{1,1}, C_{1,2})$ is compatible with $C_{1,3}$. Then $U_P(U_P(C_{1,1}, C_{1,2}), C_{1,3})$ is given by $c_{2,1} + c_{3,1}$, the common capacities, plus $c_{1,3} \cap (c_{1,1} \cap c_{1,2})$. In fact, we show that all valid flows must satisfy such mutual capacities.

Lemma 3.7 Any valid flow in a fully disjoint network satisfies mutual capacities of all compatible restrictions.

Proof Suppose f is a valid flow through a fully disjoint network. Then by Lemma 3.3, valid flows must abide by restrictions from pairs of compatible cuts. Now suppose that f must abide by two compatible restrictions R_1 and R_2 . Since they differ in exactly one summand, f must take the same value on the region of the network represented by common summands and thus the same value on the unique summand for each restriction. Thus the value of f on the unique summand in each restriction must lie in the intersection of the capacities of those summands. Hence the value of f is in the restriction formed by taking the sum of common summands in R_1 and R_2 plus the intersection of the value of their unique summands. \square

We now prove that the intersection over the values of all mutual capacities of cuts and all compatible restrictions gives the set of achievable flow values in the case of fully disjoint networks.

Theorem 3.8 *Let N be a fully disjoint network with p s - t paths. We assume that extra edges are introduced into each path except the longest one with capacity equal to that of the final arc in that path so that all paths have the same length ℓ . Let $a_{i,j}$ denote the j -th arc along the i -th path. Then the intersection over the values of all pairwise capacities of cuts and all compatible restrictions gives exactly the set of all achievable*

$$\text{flow values: } \sum_{i=1}^p \bigcap_{j=1}^{\ell} U(a_{i,j}).$$

Proof We first note that by Lemma 3.7, the set of achievable flows must be a subset of such an intersection. Hence it remains to show that such an intersection is a subset of the set of achievable flows. It will suffice to show that an intersection over some subset of the pairwise capacities equals the set of all achievable flow values.

We number the s - t paths from 1 to p . Since all paths are fully disjoint, each collection of arcs containing exactly one arc from each path forms a cut. All cuts of our interest have this form.

For any cut C_1 using the first arc in the first path, there is a set of compatible cuts C_j each identical to C_1 except that C_j uses arc j from path 1 rather than arc 1. This is a set of compatible cuts with capacity $C_{\text{cap}} + \bigcap_{j=1}^{\ell} c_j$, where C_{cap} is the sum of capacities from all arcs except the arc from path 1, and c_j is the capacity of arc j from path 1.

Then each cut has an associated, tighter restriction formed by replacing the capacity of an arc from path 1 with the intersection of capacities over that path. One can think of this step as replacing each cut with this tighter restriction, essentially reducing path 1 to a single arc with capacity equal to the intersection over capacities in the original path 1.

Since all cuts now take the same value on path 1, we can repeat this process over the remaining paths. At each step, we generate restrictions that replace the capacities from the current path with the intersection over capacities from the same path. Ultimately,

this gives us our tightest restriction with capacity $\sum_{i=1}^p \bigcap_{j=1}^{\ell} U(a_{i,j})$, i.e., for each path take the intersection of capacities of arcs over that path, then sum the values of those intersections. □

In our example above, we first partition cuts by equivalence relation into sets so that cuts in the one set pass through the same arcs in paths 2 and 3. For instance, one set is $\{\{a_{1,1}, a_{2,1}, a_{3,1}\}, \{a_{1,2}, a_{2,1}, a_{3,1}\}, \{a_{1,3}, a_{2,1}, a_{3,1}\}\}$, and another can be $\{\{a_{1,1}, a_{2,3}, a_{3,1}\}, \{a_{1,2}, a_{2,3}, a_{3,1}\}, \{a_{1,3}, a_{2,3}, a_{3,1}\}\}$. For each set, we generate a restriction with capacity $C_{\text{cap}} = \text{sum of capacities of arcs from path 2 and 3} + \bigcap c_{1,j}$, where $c_{1,j}$ is the capacity of arc $a_{1,j}$.

Using these restrictions, we repeat the process along path 2. So now, with $c_1 = \bigcap_j U(a_{1,j})$ our partitioned sets include $\{\{a_{2,1}, a_{3,1}\}, \{a_{2,2}, a_{3,1}\}, \{a_{2,3}, a_{3,1}\}\}$, for instance. Note that the restrictions from our first iteration are uniquely identified by two arcs, since the path 1 component in all restrictions is the same.

Applying the same restriction-generating process to each set, we get a restriction with capacity $\text{Cap} = \text{capacity of arcs from path 3} + \bigcap_j c_{1,j} + \bigcap_j c_{2,j}$.

This set of restrictions is already compatible, as they differ only in the arc each of them include from path 3. Applying our restriction-generating process again, we now get a single restriction with capacity $\text{Cap} = \bigcap_j c_{1,j} + \bigcap_j c_{2,j} + \bigcap_j c_{3,j}$. This is the true capacity, as we can verify independently.

4 General mutual capacity and duality result

Our approach to prove Theorem 3.8 does not generalize to all networks. For instance, Example 3.4 presents a network for which two-cut mutual capacities reduce but do not eliminate the gap. We need a further generalization to calculate the true feasible region in all cases.

To this end, we use the enriched Minkowski sum and enriched intersection (Definitions 2.26 and 2.27) instead of using Minkowski sums and intersections. Then for a single cut, the capacity is the enriched Minkowski sum. For a set of cuts \mathcal{C} , the mutual capacity $U_M(\mathcal{C})$ should capture the collection of flows which can fit through all cuts. The general duality result should then give that the mutual capacity of a set of cuts is the set of all local flows over the collection of cuts. We present Lemmas 4.1, 4.2, and 4.3 in preparation for the main result in Theorem 4.4.

Lemma 4.1 *Given a flow f , net flow across any cut is equal to the flow value of f .*

Proof Decompose f into separate flows in each commodity in single commodity networks. Then for each commodity, net flow is the same over each cut in the network. Hence net flow of f is the same over each cut, as net flow of f over a cut is exactly a vector returning the net flow in each commodity. \square

Lemma 4.2 *A set of compatible local flows admits a local flow over the union of their arc sets and cut sets.*

Proof For each arc in the union we find a contributing local flow which assigns that arc a value, then assign that value to the arc. Compatibility guarantees that all local flows which assign that arc a value assign it the same one. The flow so admitted, called the gluing of the local flows, is the unique local flow over the union of the arc sets and cut sets which, when restricted to the arc set and cut set of one of the contributing flows, returns that local flow.

Suppose $\{L_i\}_{i \in I}$ is a set of compatible local flows, and suppose L is the product of the above gluing operation. We must show that L is a local flow over the union of arc sets and that L reduces to L_i when restricted to its arc set and cut set.

Because the local flows are compatible, each cut in the union of cut sets has the same value, the value each assigns to e . Furthermore, compatibility guarantees that L does not assign multiple values to any arc. Because it assigns a value to an arc if and only if at least one local flow assigns a value, it assigns values to arcs if and only if they are in the union of the local flows' arc sets. Thus L is a local flow over the

union of the cut sets of the contributing local flows. Furthermore, this assignment is uniquely determined.

Likewise, the gluing condition guarantees that L agrees with each L_i over all arcs to which L_i assigns values, so L restricted to the arc and cut sets of L_i gives exactly the assignment L_i . □

Lemma 4.3 (Mutual Capacity) *The mutual capacity U_M of a set of cuts is the set of local flows which can be constructed as the gluing of compatible local flows from each cut in the set of cuts.*

Proof Suppose f is in the mutual capacity of a set of arcs. Then, by definition, f is a local flow over the set of cuts. Let f_i be the local flow obtained by restricting f to cut i . Note that $\{f_i\}$ is a compatible set of local flows which glue to f .

Thus each element of U_M is a local flow in the gluing of local flows over each cut.

Suppose f is the gluing of a set of local flows over each cut in a collection of cuts. Then f is a local flow over the collection of cuts, and hence in U_M .

Thus a local flow is in U_M if and only if it is the gluing of a set of compatible local flows from each cut in the set of cuts. □

Theorem 4.4 *The set of feasible flows through a network is precisely the mutual capacity of all cuts in the network.*

Proof We show that a flow is feasible if and only if it is in the mutual capacity of all cuts in the network.

Suppose a flow f is feasible. The restriction $f|_c$ of f to each cut $c \in \mathcal{C}$, the set of cuts in the network, gives a local flow over that cut. Then f is the gluing of all $\{f|_c\}_{c \in \mathcal{C}}$ (by Lemma 4.2), and hence in the mutual capacity of all cuts.

Suppose f is in the mutual capacity of all cuts. Then f respects all capacity constraints, so all that remains to show is that the flow respects conservation. For each node v and each commodity κ , take C_1 to be a cut with cut set including the in-arcs of v for commodity κ , and take C_2 to have the same cut set but including the out-arcs of v for commodity κ rather than the in-arcs, i.e., the same cut but with v moved from the t -set to the s -set.

Since the net flow across each cut must be equal by Lemma 4.1, we have that

$$\sum_{a \in C_1} \phi(a) = \sum_{b \in C_2} \phi(b),$$

and hence

$$\sum_{a \in \text{in-arcs of } v} \phi(a) + \sum_{c \in C_1 \cap C_2} \phi(c) = \sum_{b \in \text{out-arcs of } v} \phi(b) + \sum_{c \in C_1 \cap C_2} \phi(c),$$

giving

$$\sum_{a \in \text{in-arcs of } v} \phi(a) = \sum_{b \in \text{out-arcs of } v} \phi(b).$$

This result holds for each commodity $\kappa \in \{1, \dots, k\}$ and at each node $v \in V$. Hence flow conservation is respected for each commodity at all nodes, and hence over all commodities. Therefore f is a feasible flow.

By Lemma 4.3, this gives that the set of feasible flows is exactly the mutual capacity over all cuts. \square

4.1 Min-cut algorithm

We consider computational implications of the tight duality result in Theorem 4.4. While identifying the minimum cut in our multicommodity setting is expected to be expensive, we explore situations where computing the mutual capacity may turn out to be more efficient than exhaustive enumeration.

Let X_i be the capacity semimodule assigned to arc i , with x_i an arbitrary element of X_i . We consider the enriched Minkowski sum in Definition 2.26 as a more informative version of the Minkowski sum over a cut-set $\mathcal{C}_j = \{a_\kappa\}$ the function $f_j(x_1, x_2, \dots, x_m) = \sum_\kappa x_\kappa$. In the worst case, this is no more computationally expensive than the Minkowski sum, since one must already calculate all element-wise sums. We note that there are some special cases where the Minkowski sum is less computationally expensive, most notably in the case of convex polyhedra, but in most cases the expense is comparable.

The intersection of these objects is their intersection as graphs of functions over $\prod_{i \in I} X_i$, or, equivalently, the set of all (x_1, x_2, \dots, x_m) such that all $f_j((x_1, x_2, \dots, x_m))$ are equal.

Note that this algorithm cannot run in sub-exponential time since a fully connected graph will have 2^{n-2} edge-minimal $s-t$ cuts, and we must intersect over all of them. However, note that this is the case for Krishnan's special cases (Krishnan 2021) as well. On the other hand, these computations can be done on sparse networks with less expense.

To find a realization of a particular flow value, it will suffice to find an element in the levelset of that value in the intersection of the graphs, or, again equivalently, an equalizer for all cuts. We illustrate this computation on a practical example. We finish this section by presenting a class of networks on which our algorithm can run exponentially faster (in a capacity parameter) than brute force search.

Example 4.5 Consider a network with capacities given as sets of integer vectors in Fig. 13.

First we find the local flows over each cut by listing all combinations of flow values that can be assigned to arcs in that cut. Then we calculate the flow value over the cut for each combination by adding flow values on forward arcs and subtract flow values on backward arcs in the cut. This gives the net flow value for each local flow over the cut.

To glue together the local flows over two cuts, we find all combinations of local flows that agree on all arcs the two cuts share (including the net flow as assigned to the edge e from t to s). We then construct the local flows over the collection of cuts by creating for each pair of compatible local flows discovered a local flow that assigns to each arc the value which at least one of the contributing cuts assigns. For instance,

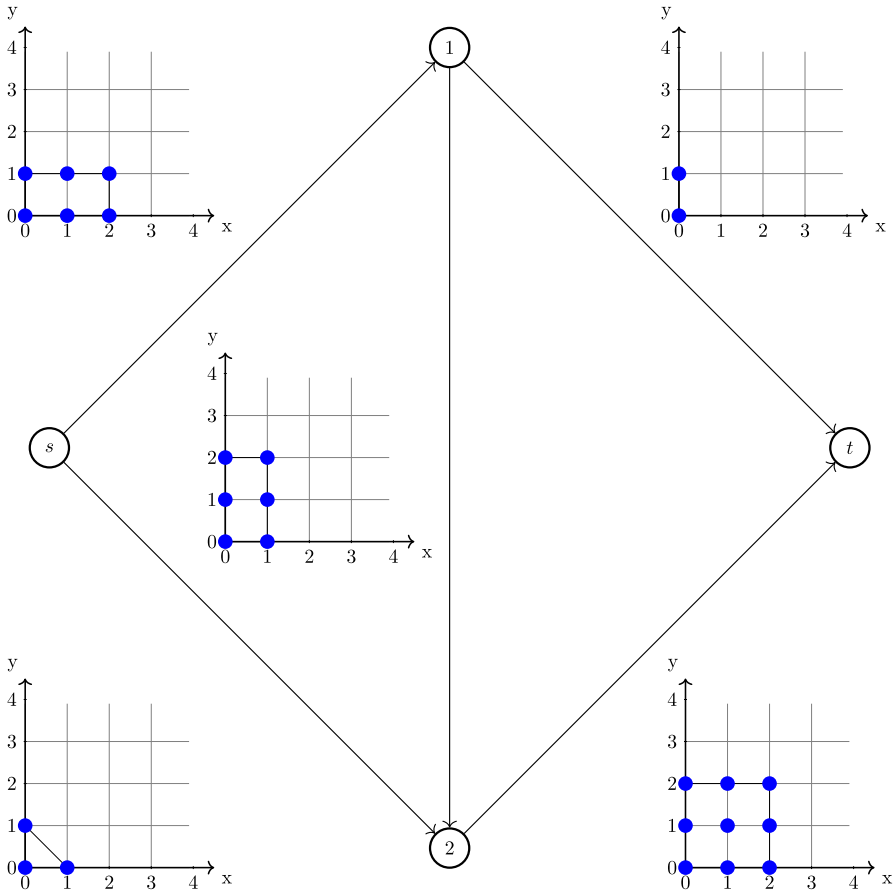


Fig. 13 Network considered in Example 4.5

the local flow that assigns $(0, 1)$ to $s-2$, $(1, 0)$ to $1-2$, and $(0, 1)$ to $1-t$ can be glued to the local flow that assigns $(0, 1)$ to $s-2$ and $(1, 1)$ to $s-1$ to give a local flow over both cuts assigning $(0, 1)$ to $s-2$, $(1, 1)$ to $s-1$, $(1, 0)$ to $1-2$, and $(0, 1)$ to $1-t$.

We construct the set of global flows over the network by starting from the first cut, gluing the second cut to it, gluing the third cut to the local flows over the first two cuts, and iterating through all cuts of the graph. The result, a local flow over all cuts, is precisely the set of feasible flows on the graph.

Example 4.6 Our algorithm can greatly outperform brute force approaches. Consider the network in Fig. 14. Note that (s, v_1) has six elements in its capacity. Numbering the cuts 1 to 3 from left to right, the gluing of cut 2 to cut 1 compares each of 6 objects to each of $4(U + 1)$ objects (local flows over cut 2, giving $24(U + 1)$ operations), resulting in 6 local flows over cuts 1 and 2. Gluing cut 3 on requires another comparison of 6 objects to $4(U + 1)$ objects, for a total of $48(U + 1)$ operations.

The brute force method, on the other hand, requires examining $6 \times 4(U + 1) \times 4(U + 1) = 96(U + 1)^2$ coordinates and checking that inflow equals outflow at all

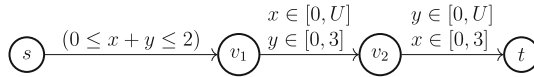


Fig. 14 A multicommodity network. Capacities are restricted to integer values of each commodity within the intervals listed on the edges. The parameter U is a large number

nodes. This results in $384(U + 1)^2$ operations, $8U$ times that required for our gluing algorithm.

5 Framework based on the work of Krishnan

We briefly discuss a theorem presented by (Krishnan 2021) and discuss a way to frame the concept of mutual capacity using the structures through which he studied this problem.

Krishnan (2021) provides two cases where the duality gap disappears and the maximum flow and minimum cut coincide. First, if each s - t path passes through a local minimum, i.e., an arc whose capacity equals the intersection of all capacities for arcs along that path, we can avoid the duality gap. Second, the duality gap vanishes if we can model the capacities as sufficiently nice objects in the sense of the following theorem.

Theorem 5.1 (Krishnan 2021, Corollary 7.12) *Consider the following data.*

- (1) Lattice-ordered flat \mathbb{N} -semimodule M .
- (2) Finite M -weighted digraph $(X; \omega)$ with distinguished edge e .

$$\text{Then } \sup_{\text{flow } \phi} \phi(e) = \inf_{e\text{-cut } C} \sum_{c \in C} \omega_c.$$

We wish to build such a lattice-ordered flat \mathbb{N} -semimodule which will model our version of the multicommodity flow problem. To this end, we first introduce definitions on and certain properties that such a set must satisfy. All definitions in this section are taken from Krishnan’s paper (Krishnan 2021).

Definition 5.2 (Lattice) A lattice is a poset S such that each pair of elements s_1, s_2 admits a least upper bound $s_1 \wedge s_2$ and a greatest lower bound $s_1 \vee s_2$. A lattice is complete if there exist a minimum 0_S and a maximum 1_S .

Definition 5.3 (Partial monoid) A partial monoid is a set \mathcal{M} with operation $+_{\mathcal{M}} : \mathcal{M} \times \mathcal{M} \rightarrow \mathcal{M}$ and element $0_{\mathcal{M}}$ such that in the following equations both sides are defined and coincide if one side is defined:

$$0_{\mathcal{M}} +_{\mathcal{M}} x = x \text{ and } x +_{\mathcal{M}} (y +_{\mathcal{M}} z) = (x +_{\mathcal{M}} y) +_{\mathcal{M}} z.$$

Definition 5.4 (S-semimodule) A partial commutative monoid \mathcal{M} is an S -semimodule provided there exists a form of scalar multiplication $\times_{\mathcal{M}} : S \times \mathcal{M} \rightarrow \mathcal{M}$ such that

for the following equations both sides are defined and coincide if one side exists:

$$\begin{aligned}(\lambda_1 \times_S \lambda_2) \times_{\mathcal{M}} x &= \lambda_1 \times_{\mathcal{M}} (\lambda_2 \times_{\mathcal{M}} x); \\(\lambda_1 +_S \lambda_2) \times_{\mathcal{M}} x &= (\lambda_1 \times_{\mathcal{M}} x) +_{\mathcal{M}} (\lambda_2 \times_{\mathcal{M}} x); \\ \lambda \times_{\mathcal{M}} (a +_{\mathcal{M}} b) &= (\lambda \times_{\mathcal{M}} a) +_{\mathcal{M}} (\lambda \times_{\mathcal{M}} b); \\ 1_S \times_{\mathcal{M}} x &= x; \\ 0_S \times_{\mathcal{M}} x &= 0_{\mathcal{M}}; \text{ and} \\ \lambda \times_{\mathcal{M}} 0_{\mathcal{M}} &= 0_{\mathcal{M}}.\end{aligned}$$

Definition 5.5 (*Partially Ordered Commutative Monoid*) A partially ordered commutative monoid \mathcal{M} is a commutative monoid with a reflexive, transitive, antisymmetric partial order $\leq_{\mathcal{M}}$ such that $a +_{\mathcal{M}} b \leq_{\mathcal{M}} a +_{\mathcal{M}} c$ whenever $b \leq_{\mathcal{M}} c$.

Definition 5.6 (*lattice ordered commutative monoid*) A lattice ordered commutative monoid \mathcal{M} is a complete (i.e. the underlying poset is a complete lattice) partially ordered partial monoid such that $a +_{\mathcal{M}} b \vee a +_{\mathcal{M}} c = a +_{\mathcal{M}} (b \vee c)$.

5.1 Construction

Now let \mathcal{M} be constructed as follows: c_1, c_2, \dots, c_m be the capacities of the m arcs in a network (these will be regions in \mathbb{R}^k in the network with k commodities), and let $\mathcal{C} = c_1 \times c_2 \times \dots \times c_m$. Let $\{\pi_1, \pi_2, \dots, \pi_m\} \in \mathcal{M}$, with $\pi_i =$ the canonical projection onto the i th component of \mathcal{C} . This captures the capacity restrictions on each arc. The image of \mathcal{C} under π_i gives the set of feasible flow values that can flow through arc i , and the level set for any value gives the set of flows which achieve that value on that arc.

For $f_1, f_2 \in \mathcal{M}$, let $f_1 +_{\mathcal{M}} f_2 = f_3$ be standard function addition, and let $(f_1 \cdot_{\mathcal{M}} f_2)(c) = f_1(c)$ if $f_1(c) = f_2(c)$, and $-\infty$ otherwise. Let $f_1 \leq_{\mathcal{M}} f_2$ if and only if $f_1(c) = f_2(c)$ or $-\infty$ for all $c \in \mathcal{C}$, or if $f_2 = \infty \forall c \in \mathcal{C}$.

Addition is essentially the same idea as taking Minkowski sums of the images of the functions, except it requires the vectors to be associated with at least one of the same flows. Multiplication, which will be the greatest lower bound in our lattice, captures the values and flows which satisfy both functions at the same time. A function f is less than or equal to another function provided they agree wherever f is defined. We also include an ‘‘infinity’’ function which is greater than or equal to all functions.

Let $g(c) = \infty$ for all c , and let $0_{\mathcal{M}} = s$ where $s(c) = 0$ for all $c \in \mathcal{C}$. We need a universal upper and lower bound, which we get in the form of g and $-g$, respectively. We also need the 0 flow, which we get in the form of $0_{\mathcal{M}}$.

Let \mathcal{M}^- be the set generated by $\{\pi_1, \pi_2, \dots, \pi_m\}$ under $+_{\mathcal{M}}, \times_{\mathcal{M}}$, and function scalar multiplication by -1 . Let $\mathcal{M} = \mathcal{M}^- \cup g \cup 0_{\mathcal{M}}$.

Now we will show that \mathcal{M} has all the required properties to satisfy Theorem 5.1.

Lemma 5.7 *The underlying poset of \mathcal{M} is a complete lattice.*

Proof We claim that $f_1 \wedge f_2 = f_1$ if $f_2 \leq_{\mathcal{M}} f_1$, vice-versa if $f_1 \leq_{\mathcal{M}} f_2$, and if neither of those holds $f_1 \wedge f_2 = g$. If one function is $\leq_{\mathcal{M}}$ the other, e.g., suppose

$f_1 \leq_{\mathcal{M}} f_2$ without loss of generality, then by definition any function greater than both must be greater than f_2 , and hence f_2 is the lowest upper bound on f_1, f_2 . Otherwise, there exists a $c \in \mathcal{C}$ where $-\infty \neq f_1(c) \neq f_2(c) \neq -\infty$. Then there is no function which agrees with both on c , and hence only $g \geq_{\mathcal{M}} f_1, f_2$. \square

Lemma 5.8 $f_1 \vee f_2 = f_1 \cdot_{\mathcal{M}} f_2$.

Proof Suppose $f_3 \leq_{\mathcal{M}} f_1$ and $f_3 \leq_{\mathcal{M}} f_2$. Then for all $c \in \mathcal{C}$ f_3 agrees with f_1 or maps to $-\infty$ and f_3 agrees with f_2 or maps to $-\infty$. Then f_3 maps all values on which f_1 and f_2 don't agree to $-\infty$, so $f_3 \leq_{\mathcal{M}} f_1 \cdot_{\mathcal{M}} f_2$. \square

Lemma 5.9 \mathcal{M} is a partial monoid.

Proof We observe that $0_{\mathcal{M}} +_{\mathcal{M}} f = f$ for all f , and function addition is associative. \square

Lemma 5.10 \mathcal{M} is a commutative \mathbb{N} -semimodule.

Proof Since the functions over real numbers are a semimodule over \mathbb{N} , \mathcal{M} inherits the required properties to be a commutative \mathbb{N} -semimodule. \square

Lemma 5.11 \mathcal{M} is a partially ordered commutative monoid.

Proof Clearly $\leq_{\mathcal{M}}$ is reflexive, transitive, and antisymmetric. Furthermore, suppose $f_1 \leq_{\mathcal{M}} f_2$. Then for all $c \in \mathcal{C}$, $f_1(c) = f_2(c)$ or $f_1(c) = -\infty$. Then $(f_1 +_{\mathcal{M}} f_3)(c) = f_1(c) + f_3(c)$ whenever $f_1(c) = f_2(c)$, and $(f_1 + f_3)(c) = -\infty + f_3(c) = -\infty$ otherwise. Thus \mathcal{M} is a partially ordered commutative monoid. \square

Lemma 5.12 \mathcal{M} is a lattice ordered commutative monoid.

Proof Since $(f_1 +_{\mathcal{M}} f_2) \vee (f_1 +_{\mathcal{M}} f_3) = (f_1 +_{\mathcal{M}} f_2)(c)$ when $(f_1 +_{\mathcal{M}} f_2)(c) = (f_1 +_{\mathcal{M}} f_3)(c)$ and $-\infty$ otherwise, and since $(f_1 +_{\mathcal{M}} f_2)(c) = (f_1 +_{\mathcal{M}} f_3)(c)$ if and only if $f_2(c) = f_3(c)$ or $f_1(c) = -\infty$, then for $c \in \mathcal{C}$ either $f_1(c) = -\infty$, and hence $(f_1 + (f_2 \times_{\mathcal{M}} f_3))(c) = -\infty = ((f_1 +_{\mathcal{M}} f_2) \times_{\mathcal{M}} (f_1 +_{\mathcal{M}} f_3))(c)$; or $f_2(c) = f_3(c)$, and thus $((f_1 +_{\mathcal{M}} f_2) \times_{\mathcal{M}} (f_1 +_{\mathcal{M}} f_3))(c) = f_1(c) + (f_2 \times_{\mathcal{M}} f_3)(c)$. Thus \mathcal{M} is a lattice ordered commutative monoid. \square

Note that \mathbb{N} is flat over itself and direct sums of flat modules are flat (nLab authors 2021). Hence we can build flat \mathbb{N} semimodules representing the flows over any collection of edges, and subsemimodules of any such construction remain flat (Krishnan 2021).

Thus Corollary 7.12 from Krishnan (2021) gives us that, provided the semimodule for which \mathcal{M} is a subsemimodule is flat over \mathbb{N} (such as, for instance, any set of capacities modeling rational flows), $\sup_{\text{flow } \phi} \phi(e) = \inf_{e\text{-cut } \mathcal{C}} \sum_{c \in \mathcal{C}} \omega_c$. In other words, the feasible region of flows is given by the infimum over e-cuts of the sum of the projection functions associated with that cut. The region of feasible flow values must then be the image of \mathcal{C} under that function, while all possible flows achieving that value is determined by the level sets of the function.

6 Ratio maximum flow

While Theorem 4.4 presents a tight duality result for a general multicommodity max flow problem, the resulting algorithms presented in Sect. 4.1 could still be computationally expensive. Hence we consider problems closely related to the MULT-COM MAX-FLOW problem for which we could develop computationally tractable approaches. Motivated by the work of Leighton and Rao (1999), we study RATIO MAX-FLOW and INT RATIO MAX-FLOW as multicommodity maximum flow problems that include an explicit optimization component rather than searching for the full feasible region. Our approach is motivated by augmenting path algorithms for single commodity max flow problem (Ahuja et al. 1993). However, we search for an augmenting *cycle set* instead of augmenting paths.

Definition 6.1 (*Augmenting Cycle Set*) An *augmenting cycle set* in a network with flow f having coefficients in \mathbb{R}^k is a set of cycles such that if we modify the flow of every arc in a cycle by that cycle's coefficients, we obtain a new flow with a higher value than f .

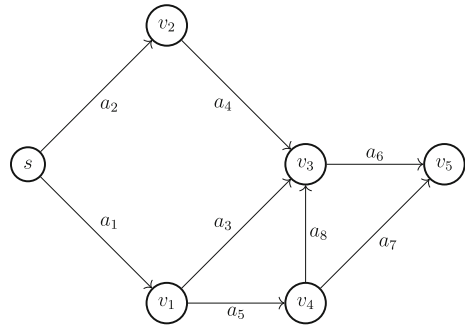
Lemma 6.2 Any two flows with the same flow value in a network differ by an element of the cycle space.

Proof We first observe that this result is elementary in the case of single commodity flows. we show that it also holds for our general multicommodity setting. Let flows f_1 and f_2 in the enhanced network have the same flow value. Since both flows are cycles in the enhanced network, $f_1 - f_2$ must also be in the cycle space. Furthermore, since both flows have the same flow value, and hence the same value on the distinguished edge e , $f_1 - f_2$ has a value 0 on edge e , and thus lies in the cycle space of the base network. Thus it follows that any two flows with the same flow value differ by a sum of cycles. \square

We first address the following decision problem: given a flow value f , is there a flow on the network N with value f ? We solve this problem by constructing a pseudoflow F' with value f and search for a set of augmenting cycles that transform F' into a flow. To this end, we transform the problem into one of finding an area of cycle space within particular bounds. We show that the decision problem returns true if and only if the bounds in cycle space are non-empty, and that augmenting by a set of cycles in that region gives a flow.

Since we know that the flows on a network are cycles in the enhanced network, and since we know two potential flows share a flow value if and only if they differ by a cycle in the network (Lemma 6.2), we can approach the problem of finding a valid realization of a particular flow value by considering it in terms of the cycle space of the network. Since all arcs in our network are intrinsically bi-directional, the cycle space of our network is in fact the cycle space of the underlying graph. A generating set for a graph's cycle space can be found in polynomial time by first finding a spanning tree of the graph rooted at s Ahuja et al. (1993). Each edge not on the spanning tree determines a unique cycle with the tree: take the unique path to the head and tail of the edge and include in our cycle each edge unique to one of the paths, then add the

Fig. 15 Example network to illustrate cycle constraints in Definition 6.3



edge in question not on the tree. There are then $n - 1$ edges on the tree and $m - n + 1$ edges determining cycles.

Next we establish constraints on what combinations of cycles will give a flow. We start by partitioning the edge set by the cycles to which the edge belongs. The $m - n + 1$ edges determining unique cycles belong to exactly one cycle. The remaining $n - 1$ edges may belong to one or more cycles, and we can find this information directly from the cycle-membership matrix. Finally, we determine the capacity constraints for each set by intersecting the orientation corrected capacity of each arc involved in that set.

Definition 6.3 (Cycle Constraints) For a network with a set of generating cycles $\{C_i\}_{i \in I}$, the cycle constraints $\{c_i\}_{i \in I}$ are specified on k -dimensional vectors. For each arc a in the network, we have a constraint R_a that the sum of all c_i that use a as a forward arc minus the sum of all c_i that use a as a backward arc must lie within the capacity of a minus the flow assigned to a .

We illustrate cycle constraints on an example network in Fig. 15. Suppose we have cycles $c_1 = \{a_1, a_3, -a_4, -a_2\}$, $c_2 = \{a_5, a_8, -a_3\}$, and $c_3 = \{a_5, a_7, -a_6, -a_3\}$ in the network.

Then to find the cycle constraints, we observe a_1, a_2, a_4 are unique to c_1 , a_3 is common to all three, a_5 is unique to c_2 and c_3 , a_6, a_7 are unique to c_3 , and a_8 is unique to c_2 . Then c_1 has a capacity of $v(a_1) \cap -v(a_4) \cap -v(a_2)$, c_2 has a capacity of $v(a_8)$, and c_3 has a capacity of $-v(a_6) \cap v(a_7)$. Additionally, $c_2 + c_3$ has capacity $v(a_5)$, and $c_1 - c_2 - c_3$ has capacity $v(a_3)$.

Then any feasible solution to c_1, c_2, c_3 such that:

- $c_1 \in v(a_1) \cap -v(a_4) \cap -v(a_2)$
- $c_2 \in v(a_8)$
- $c_3 \in -v(a_6) \cap v(a_7)$
- $c_2 + c_3 \in v(a_5)$
- $c_1 - c_2 - c_3 \in v(a_3)$

is a set of cycle augmentations that give a feasible flow of the current infeasible flow value.

In general, each arc contributes a constraint that the sum of all generating cycles through that arc in the positive orientation minus the sum of generating cycles through

the arc in the negative orientation must lie within the adjusted capacity of that arc. Hence we get an $(m - n + 1)k$ -dimensional optimization problem with m constraints.

Theorem 6.4 *Augmenting along a cycle set results in a flow if and only if the cycle set abides by the cycle constraints of the network.*

Proof Suppose we have a cycle set in the solution set to the problem. Then the constraints imposed by arcs unique to a single cycle C restrict the possible values of C to ones that result in no arc unique to that cycle overflowing when augmented along. Likewise, the constraints imposed by the arcs common to two cycles restrict the values those cycles can take to pairs of values which will result in none of the common arcs overflowing after augmentation, and so on. Thus since all arcs belong to one set in the partition, and since all constraints are satisfied, augmenting along the cycle set must result in no arcs overflowing. Since we started with a pseudoflow, which respects flow conservation, and augmented along cycles (which preserve flow conservation) we get a flow. Since augmenting along cycles does not change the value of the flow, we maintain a flow value of f .

Suppose a set of cycles gives us a flow after augmentation. Then for each arc, the cycle set must fall within the constraints imposed by that arc. Thus the cycle set must fall in the solution set to the above problem. Again, augmenting along cycles guarantees that we maintain the same flow value. \square

We use the solution to the decision problem in algorithms that find an ϵ -approximate solution to RATIO MAX-FLOW (Algorithm 1) and an ϵ -approximate solution to INT RATIO MAX-FLOW using a binary search (Algorithm 2).

We show that these algorithms terminate and that they are accurate.

Lemma 6.5 *Algorithms 1 and 2 terminate.*

Proof Because the capacities of the network are bounded, we get a finite initial value for B^+ . As a binary search, the algorithm reduces the difference of B^+ and B^- by a factor of $\frac{1}{2}$ at each step. Thus $B^+ - B^- < \epsilon$ when $\frac{B^+}{2^s} < \epsilon$ after s steps, which happens in $s = \lceil \log(\frac{B^+}{\epsilon}) \rceil$ steps.

Theorem 6.6 *Algorithm 1 gives a feasible max flow in the desired ratio which is within ϵ of the true max flow.*

Proof Since all capacities are reducible, the $\mathbf{0}$ flow is feasible. Since B^- initializes at 0 and is replaced only with feasible multiples, B^- is guaranteed to be feasible at all steps.

Now suppose there exists $P' \in \mathbb{R}^+$ such that $P'R$ has a realization, but $P' - B^- > \epsilon$. We have $B^+ - B^- < \epsilon$ by the termination condition, so $P' > B^+$.

Since P' has a realization, multiplying the values at each arc by $\frac{B^+}{P'}$ gives a realization for B^+ , which is feasible because each arc has a reducible capacity. Thus B^+ has a realization. But the algorithm only assigns values for which no realization exists to B^+ , so this is a contradiction.

Hence the max flow is within ϵ of B^- . \square

Data: Network N with compact (i.e., closed and bounded) reducible capacities $\{C_i\}$;
 Desired Ratio: R , the ratio of commodities we wish to transport;
 Upper bound: B^+ , an overestimate on the multiples of R that can be sent;
 Lower bound: $B^- = 0$; and
 Error allowed: $\epsilon > 0$ (returned multiple is within ϵ of true max flow multiple).
Result: P^* , an approximation of the true max flow multiple P such that $P^*R \leq PR \leq (P^* + \epsilon)R$

Initialization:
 $D = B^+ - B^-$;
while $D > \epsilon$ **do**
 $T = \frac{B^+ + B^-}{2}$;
 Test(TR, B^-, B^+);
 $D = B^+ - B^-$;
end
return B^-

Function Test (V, B^-, B^+):
 $T = \frac{B^+ + B^-}{2}$;
 Test decision problem on V ;
 if decision problem returns *TRUE* **then**
 $B^+ = T$
 else
 $B^- = T$
 end

Algorithm 1: ϵ -error approximation for RATIO MAX-FLOW.

Data: Same as in Algorithm 1, except we require $0 < \epsilon < 0.5$
 (so that the returned value can be within ϵ of at most one integer).

Result: P^* , an approximation of the true max flow multiple P such that $P^*R \leq PR \leq (P^* + \epsilon)R$

Initialization:
 $B^- = 0$;
 $D = B^+ - B^-$;
while $D > \epsilon$ **do**
 $T = \frac{B^+ + B^-}{2}$;
 Test(TR, B^-, B^+);
 $D = B^+ - B^-$;
end
 EndTest($\lceil B^- \rceil R, B^-$)
return B^-

Function EndTest (V, B^-):
 Test decision problem on V ;
 if decision problem returns *TRUE* **then**
 $B^- = \lceil B^- \rceil$
 else
 $B^- = \lfloor B^- \rfloor$
 end

Algorithm 2: ϵ -error approximation for INT RATIO MAX-FLOW. The function Test is defined in Algorithm 1.

Lemma 6.7 Algorithm 2 gives a maxflow in the desired ratio which the greatest integer multiple of the desired ratio which can pass through the network.

Proof By Theorem 6.6, B^- is within ϵ of the true max flow. Since $\epsilon < 0.5$ here, B^- can be within ϵ of at most a single integer. Since B^- has a realization, the maximum integer value with a realization must be either $\lceil B^- \rceil$ or $\lfloor B^- \rfloor$, and can be $\lceil B^- \rceil$ only if $\lceil B^- \rceil - B^- < \epsilon$. Since Algorithm 2 checks $\lceil B^- \rceil$ in such a case, it must return the maximum integer value with a feasible realization. \square

7 Discussion

We addressed two problems using topological techniques and with graphs to gain insight into data and complex systems. First, we discussed our work on the MULT-COM MAX-FLOW problem, a generalization of the traditional Max Flow problem from the study of network flows. The generalization removes some key pieces from the traditional Max-Flow Min-Cut theorem and the algorithms to which it gives rise—in particular, the idea of a flow saturating a cut does not translate nicely into the multicommodity setting. Our work developed a way to connect the feasible region for a network flow problem to the flows that can fit through all s - t cuts in the network. We have also given some methods to calculate those values for the general problem and some useful sub-problems. It would be interesting to extend the results in Sect. 3.2 on fully disjoint networks to the case of multicommodity flows in *series-parallel digraphs*, which are obtained by composing in an edge-disjoint manner (i.e., by merging the source and sink nodes) the copies of smaller digraphs in series or in parallel fashion.

Our main result is Theorem 4.4, in which we showed that the set of local flows over individual cuts that can be “glued together” into true flows is precisely the set of feasible flows over the network. By generalizing max flow to be the feasible region of flow through the network and the min cut to be the set of consistent collections of local flows over cuts, we achieve a generalization of the duality between flows and cuts which the traditional Max-Flow Min-Cut theorem provides.

We also studied a special case of MULT-COM MAX-FLOW that considers maximizing the multiple of a specific ratio of flows transported through the network. We show how to model the possible rearrangements of flow as a problem on the cycle space of the network. We show that this method of assigning coefficients exactly captures the set of feasible rearrangements of flows of particular flow values, meaning that a pseudoflow has a feasible realization if and only if it has a feasible set of cycle coefficients. This gives us a method to check the feasibility of a particular flow value which, in turn, allows us to use a binary search of possible multiples of the desired ratio to determine an arbitrarily precise approximation of the max flow.

There are several avenues for future research in this area. Primarily, more efficient methods for finding the mutual capacities for a collection of cuts is necessary for any large-scale implementations of this work. Even efficient methods for calculating better approximations, e.g., mutual capacities for all collections of three cuts, could give useful heuristics for solving max-flow problems.

Industry applications of multicommodity flows may be solvable by more computationally tractable subproblems. Finding additional assumptions about network structure (e.g., acyclic networks or limits on degrees of nodes) or capacities (e.g.,

polygonal regions or convexity) that give more computationally efficient solutions is another potentially fruitful area of research.

Acknowledgements We thank the National Science Foundation for support through Grants 1661348 and 1819229.

Author Contributions The two authors contributed equally to all aspects of the research and manuscript preparation.

Funding This work was partly supported by the US National Science Foundation (NSF) through grants 1661348 and 1819229.

Data Availability There are no datasets used. All enquiries about data could still be directed to the authors.

Declarations

Conflict of interest The authors have no relevant financial or non-financial competing interests to disclose.

References

- Ahuja RK, Magnanti TL, Orlin JB (1993) Network flows: theory, algorithms, and applications. Pearson, London
- Delos V, Teissandier D (2014) Minkowski sum of polytopes defined by their vertices. [arXiv:1412.2564](https://arxiv.org/abs/1412.2564)
- Even S, Itai A, Shamir A (1976) On the complexity of timetable and multicommodity flow problems. *SIAM J Comput* 5(4):691–703. <https://doi.org/10.1137/0205048>
- Ford LR Jr, Fulkerson DR (1956) Maximal flow through a network. *Can J Math* 8:399–404. <https://doi.org/10.4153/CJM-1956-045-5>
- Ghrist R, Krishnan S (2013) A topological max-flow-min-cut theorem. In: 2013 IEEE global conference on signal and information processing (GlobalSIP), pp 815–818. <https://doi.org/10.1109/GlobalSIP.2013.6737016>
- Krishnan S (2021) Flow-cut dualities for sheaves on graphs. <https://people.math.osu.edu/krishnan.118/papers/mfmc.pdf>, last Accessed 18 Sep 2021
- Leighton T, Rao S (1999) Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *J ACM* 46(6):787–832. <https://doi.org/10.1145/331524.331526>
- nLab authors. Flat Module. <http://ncatlab.org/nlab/show/flat+module>, March 2021. Revision 29

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

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.