THEORY OF COMPUTING, Volume 4 (2008), pp. 1–20
http://theoryofcomputing.org

Single Source Multiroute Flows and Cuts on Uniform Capacity Networks*

Henning Bruhn Jakub Černý† Alexander Hall Petr Kolman‡
Jiří Sgall§

Received: September 6, 2007; published: April 16, 2008.

Abstract: For an integer $h \geq 1$, an *elementary h-route flow* is a flow along h edge disjoint paths between a source and a sink, each path carrying a unit of flow, and a single commodity *h-route flow* is a non-negative linear combination of elementary h-route flows. An instance of a *single source multicommodity flow problem* for a graph $G = (V,E)$ consists of a source vertex $s \in V$ and k sinks $t_1,\ldots,t_k \in V$ corresponding to k commodities; we denote it $\mathcal{I} = (s;t_1,\ldots,t_k)$. In the *single source multicommodity multiroute flow problem*, we are given an instance $\mathcal{I} = (s;t_1,\ldots,t_k)$ and an integer $h \geq 1$, and the objective is to maximize the

*A preliminary version of this work appeared in Proceedings of the 18th ACM-SIAM Symposium on Discrete Algorithms [9].

†Supported by project 1M0021620808 (ITI) of Ministry of Education of the Czech Republic.

‡Supported by European Commission within the 6th Framework Programme under contract 001907 "Dynamically Evolving, Large Scale Information Systems" (DELIS) and by project 1M0021620808 (ITI) of Ministry of Education of the Czech Republic.

§Partially supported by Institutional Research Plan No. AV0Z10190503, grant 201/05/0124 of GA ČR, and grant IAA1019401 of GA AV ČR.

ACM Classification: G.2.2, F.2.2

AMS Classification: 68R10, 05C38, 68W25 ,90C27

Key words and phrases: Graphs, multi-route flow, paths, multicommodity flow, connectivity, approximation algorithms

Authors retain copyright to their work and grant Theory of Computing unlimited rights to publish the work electronically and in hard copy. Use of the work is permitted as long as the author(s) and the journal are properly acknowledged. For the detailed copyright statement, see http://theoryofcomputing.org/copyright.html.

© 2008 Henning Bruhn, Jakub Černý, Alexander Hall, Petr Kolman, and Jiří Sgall DOI: 10.4086/toc.2008.v004a001

total amount of flow that is transferred from the source to the sinks so that the capacity constraints are obeyed and, moreover, the flow of each commodity is an h-route flow.

We study the relation between classical and multiroute single source flows on undirected networks with uniform capacities and we provide a tight bound. In particular, we prove the following result. Given an instance $\mathcal{I} = (s; t_1, \ldots, t_k)$ such that each $s - t_i$ pair is h-connected, the maximum classical flow between s and the t_i is at most $(2 - 2/h)$-times larger than the maximum h-route flow between s and the t_i and this is the best possible bound for $h \geq 2$. This, as we show, is in contrast to the situation of general multicommodity (i. e., multiple sources or non-uniform capacities) multiroute flows that are up to $k(1 - 1/h)$-times smaller than their classical counterparts.

Furthermore, we introduce and investigate *duplex flows* defined so that the capacity constraints on edges are applied independently for each direction. We show that for networks with uniform capacities and for instances as above the maximum classical flow between s and the t_i is the same as the maximum h-route duplex flow between s and the t_i. Moreover, the total flow on each edge in the duplex flow can be restricted to $(2 - 2/h)C$, where C is the capacity of each edge.

As a corollary, we establish a max-flow min-cut theorem for the single source multicommodity multiroute flow and cut. An *h-disconnecting cut* for $\mathcal{I}$ is a set of edges $F \subseteq E$ such that for each i, the maximum h-route flow between s and t_i is zero. We show that the maximum h-route flow is within $2h-2$ of the minimum h-disconnecting cut, independently of the number of commodities; we also describe a $(2h-2)$-approximation algorithm for the minimum h-disconnecting cut problem.

1 Flows, multiroute flows and cuts

A classical flow is (roughly) a non-negative linear combination of unit flows along paths (cf. [2]). Classical flow theory is not much interested in the number of the paths or in interactions among them. It is plausible, for example, that there is an edge in the network that is used by every path of a given flow; a failure of this single edge results in a loss of the entire flow. This property of the classical flow is undesirable in some applications and motivated the definition of a multiroute flow. For a given integer $h \geq 1$, the *multiroute flow* (or an *h-route flow*) is a flow that is decomposable into a non-negative linear combination of elementary h-route flows where an *elementary h-route flow* is a flow along h edge disjoint paths between the source and the sink, each path carrying a unit of flow [21]. Closely related to this is the concept of h-balanced flows. A flow of size M between two vertices is *h-balanced* if the flow on every edge is at most M/h. Clearly, every h-route flow is an h-balanced flow; the opposite (less obvious) claim is also true: Every h-balanced (acyclic) flow is an h-route flow [1, 6, 21].

A necessary and sufficient condition for the existence of an h-route flow between two vertices is that the vertices are h-connected. A corollary of the equivalence of h-route flows and h-balanced flows is that on uniform capacity networks with an h-connected source s and sink t, every maximum s-t-flow is an h-route flow. However, for multicommodity flows and h-route flows, this relation is no longer valid. We investigate the relation between flows and h-route flows for a special case of multicommodity

problems, namely for single source problems. An instance of a *single source multicommodity flow problem* for a graph $G=(V,E)$ consists of a source vertex $s\in V$ and k sinks $t_1,\dots,t_k\in V$ corresponding to k commodities; we denote it $\mathcal{I}=(s;t_1,\dots,t_k)$. A (single-source) multicommodity flow is an h-route flow if the flow corresponding to each commodity is an h-route flow.

We show that for undirected networks with uniform capacities and for instances $\mathcal{I}=(s;t_1,\dots,t_k)$ such that s and t_i are h-connected, for each $i=1,\dots,k$, the maximum classical flow between s and the t_i is at most $2-2/h$ times larger than the maximum h-route flow between s and the t_i; this bound is the best possible for $h\geq 2$. In particular, for $h=2$ the ratio is 1, implying that by imposing the requirement that the flow be a 2-route flow, we do not lose anything with respect to the size of the flow. Moreover, if the uniform capacity of the edges is integral, then there always exists a half-integral h-route flow of size at least half of the maximum classical flow.

Furthermore, we introduce and investigate *duplex flows* defined so that the capacity constraints on edges are imposed independently for each direction, as if each undirected edge is replaced by two directed edges in the opposite direction. To give an example, an edge with capacity 1 is able to carry a flow of 1 in both direction simultaneously but it is not able to carry a flow of 1.5 in one direction even if the other direction is not used. This is a natural model for network flows and as far as we know no specific attention was given to it. For classical single commodity flow and single source multicommodity flow, the sizes of the maximum non-duplex and duplex flows coincide since any classical flow can be modified so that no edge is used in both directions. For h-route flows this simple transformation no longer works. Nevertheless, we show that for networks with *uniform* capacities and for instances $\mathcal{I}=(s;t_1,\dots,t_k)$ such that s and t_i are h-connected, for each $i=1,\dots,k$, the maximum classical flow between s and the t_i has the same size as the maximum h-route duplex flow between s and the t_i. Moreover, the total flow on each edge in the duplex flow can be restricted to $2-2/h$. Thus, our bound for duplex flows implies the results for non-duplex flows described in the previous paragraph (except for the half-integrality).

Our results for single source flows are in sharp contrast with the situation of general (i. e., multiple sources or non-uniform capacities) multiroute multicommodity flows: we describe an example with k commodities where the maximum classical flow is $k(1-1/h)$-times larger than the maximum h-route flow.

The other subject of the paper is cuts for h-route flows. For the classical flow, a cut is a subset of edges whose removal disconnects the source and the sink (or every source-sink pair, in a case of the multicommodity flow). Analogously, we define cuts for h-route flows. A subset $F\subseteq E$ of edges is called an *h-disconnecting cut* for an instance of the multicommodity flow if no source-sink pair is h-connected in $(V,E\setminus F)$. The h-disconnecting cuts correspond to integral solutions of a dual of a natural linear programming formulation of the multiroute flow problem (see Section 1.2). We establish a max-flow min-cut theorem for the single source multiroute flow and the minimum disconnecting cut problems on networks with uniform capacities. In particular, we show that the max-flow for the problem is within $2h-2$ of the min-cut. As a corollary of this relation we get a $(2h-2)$-approximation algorithm for the h-disconnecting cut problem.

1.1 Related results

Kishimoto and Takeuchi [22] and later Aggarwal and Orlin [1] studied single commodity multiroute flows (cf. [6, 15, 14]). They provided the characterization of h-route flows as h-balanced flows and also

proved a duality of multiroute flows and multiroute cuts (for different cuts than those considered in this paper). Multiroute flows and integral variants of multiroute flows have applications in communication and routing problems (e. g., [4, 5, 20, 12] and references therein).

Another direction of research focuses on flows under the restriction that each commodity is allowed to use only a limited number of paths: the edge disjoint paths problem and the unsplittable flow problem allow one path per commodity [8, 10, 11, 19, 23, 25, 27, 26, 34]; the h-splittable flow problem allows at most h, not necessarily disjoint, paths per commodity [7, 24, 31, 30]; particular attention has been given to single source unsplittable flow problems [13, 16, 25, 33]. Though there is a certain similarity between the h-splittable flows and the h-route flows (in fact, they may even coincide for some instances), there is also a substantial difference. Whereas the h-splittable flows may split, the h-route flows have the obligation to split.

Relations between flows and cuts have been studied for over half a century. Menger [32] observed that the maximum number of edge disjoint paths between a pair of vertices is equal to the size of the minimum subset of edges whose removal disconnects the pair. Ford and Fulkerson [17] proved the celebrated theorem about the duality of (single commodity) flows and cuts in networks. Though an exact duality does not hold for multicommodity flows and cuts, there are several theorems establishing an approximate duality (with the gap of order $\log k$) for different variants of the problem (Leighton and Rao [28], Aumann and Rabani [3], Linial, London and Rabinovich [29], Garg, Vazirani and Yannakakis [18]).

1.2 Preliminaries

As indicated in the title, in this paper we deal with networks with uniform capacities. For simplicity and without loss of generality we assume throughout the paper that every edge has capacity one. The number of vertices is denoted n and the number of edges m; we allow parallel edges. The letter k denotes the number of commodities and the letter h the number of routes in the elementary multiroute flow. Several times we need the characterization of h-route flows as h-balanced flows that was first proved by Kishimoto and Takeuchi.

Theorem 1.1 ([1, 6, 21]). *A single commodity flow without cycles is h-balanced if and only if it is an h-route flow.*

For an instance $\mathcal{I}$ of the multicommodity flow problem, we use $\mathcal{F}^h(\mathcal{I})$ for the size of the maximum h-route flow and $\mathcal{F}^{h,\|}(\mathcal{I})$ for the size of the maximum h-route duplex flow for the instance $\mathcal{I}$. As mentioned in the introduction, for single source multicommodity flow, the sizes of the maximum non-duplex and duplex flows coincide. Thus we have $\mathcal{F}^h(\mathcal{I}) \leq \mathcal{F}^{h,\|}(\mathcal{I}) \leq \mathcal{F}^1(\mathcal{I})$.

For a given flow, an *empty* edge is an an edge unused by the flow. We will deal with *minimum cost flows* several times. In such cases we consider the uniform cost function (i. e., $cost(e) = 1$, $\forall e \in E$). Recall that a single source classical flow can be viewed as a single commodity flow problem and therefore there exists an integral maximum flow for every instance $\mathcal{I}$; there also exists a minimum cost maximum flow that is integral, and its cost is just the number of non-empty edges.

Consider a network $G = (V,E)$. Let $s_1,\ldots,s_k$ be k sources and $t_1,\ldots,t_k$ be k sinks of a multicommodity flow problem; we call the sources and sinks also *terminals*. Define $\mathcal{Q}_i$ as the set of all elementary h-route flows between s_i and t_i and let $\mathcal{Q} = \bigcup_{i=1}^{k} \mathcal{Q}_i$. As far as we know, no exact combinatorial algorithms for computing the maximum multicommodity flow are known (not even for the 1-route flow). Thus, for

completeness we provide a linear programming formulation of the maximum h-route flow problem (the variable $f(q)$ represents the size of the flow along the h-system q, that is, a flow of size $f(q)/h$ along each of the h paths of q):

$$\begin{aligned} \max \sum_{q\in\mathcal{Q}} f(q) & & (1.1)\\ \sum_{q\in\mathcal{Q}:e\in q} f(q)/h &\leq 1 \quad \forall e\in E \\ f(q) &\geq 0 \quad \forall q\in\mathcal{Q}\,. \end{aligned}$$

The dual program corresponds to the fractional relaxation of the the minimum h-disconnecting cut problem:

$$\begin{aligned} \min\ h\cdot\sum_{e\in E} x(e) & & (1.2)\\ \sum_{e\in q} x(e) &\geq 1 \quad \forall q\in\mathcal{Q} \\ x(e) &\geq 0 \quad \forall e\in E\,. \end{aligned}$$

By setting integrality constraints on the variables x, we get an integer linear programming formulation of the minimum h-disconnecting cut problem.

2 Relating flows and multiroute flows

In this section, we show that h-route flows are not much smaller than classical flows under certain assumptions: single source, uniform capacity, and connectivity. We prove the following theorems for non-duplex and duplex flows, respectively.

Theorem 2.1. *Let $G=(V,E)$ be an undirected graph and let $\mathcal{I}=(s;t_1,\ldots,t_k)$ be an instance of the single source multicommodity flow problem such that for each i, s and t_i are h-connected for a given $h\geq 2$. Then*

$$\mathcal{F}^1(\mathcal{I}) \leq (2-2/h)\cdot\mathcal{F}^h(\mathcal{I})\,. \tag{2.1}$$

There also exists a half-integral h-route flow of size at least $\mathcal{F}^1(\mathcal{I})/2$.

Theorem 2.2. *Let $G=(V,E)$ be an undirected graph and let $\mathcal{I}=(s;t_1,\ldots,t_k)$ be an instance of the single source multicommodity flow problem such that for each i, s and t_i are h-connected for a given $h\geq 2$. Then, for the duplex flows,*

$$\mathcal{F}^{h,\|}(\mathcal{I})=\mathcal{F}^1(\mathcal{I})\,,$$

and moreover the equality can be achieved by a duplex flow with a total flow on each edge of at most $2-2/h$.

Note that Theorem 2.2 implies Theorem 2.1, with the exception of the half-integrality: we can use the duplex flow scaled down so that we multiply a flow along each edge by a factor of $1/(2-2/h)$. Nevertheless, we give also a direct proof of Theorem 2.1. One consequence of this proof is that in case of $h=2$, even the sharper bound can be achieved by a half-integral flow. Since for $h=2$ the factor is $2-2/h=1$, and a trivial bound $\mathcal{F}^h(\mathcal{I}) \leq \mathcal{F}^1(\mathcal{I})$ holds for every h, we have the following corollary which shows that by imposing the requirement that the flow be a 2-route flow, we do not lose anything with respect to the size of the flow.

Corollary 2.3. *Let $G=(V,E)$ be an undirected graph and let $\mathcal{I}=(s;t_1,\ldots,t_k)$ be an instance of the single source multicommodity flow problem such that for each i, s and t_i are 2-connected. Then*

$$\mathcal{F}^1(\mathcal{I}) = \mathcal{F}^2(\mathcal{I}).$$

In addition, the equality can be achieved by a half-integral 2-route flow.

Before proving the upper bounds, we first show, in Section 2.1, that for h-route flows with a single source, the factor of $2-2/h$ is the best possible and also that the assumptions of single source and unit capacity are essential. Then, in Section 2.2 we develop the common parts of the upper bound proofs and finally in Sections 2.3 and 2.4 we show the upper bounds for non-duplex and duplex flows, respectively.

2.1 Lower bounds

Theorem 2.4. *For every pair of integers $h,k \geq 2$ there exist an undirected graph G and an instance $\mathcal{I}=(s;t_1,\ldots,t_k)$ of the single source multicommodity flow problem such that for each i, s and t_i are h-edge-connected, and, at the same time,*

$$\mathcal{F}^1(\mathcal{I}) \geq \left(2-\frac{2}{h}\right)\cdot\mathcal{F}^h(\mathcal{I}).$$

Proof. The set of vertices of the graph G consists of $k+2$ distinct vertices $s,v,t_1,\ldots,t_k$. The set of edges contains $h-1$ parallel edges between s and t_i, and an edge between t_i and v, for $i=1,\ldots,k$ (Figure 1).

Consider the instance $\mathcal{I}=(s;t_1,\ldots,t_k)$. An elementary h-route flow between s and t_i, for $i=1,\ldots,k$,

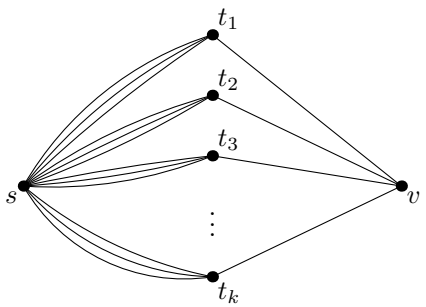

Figure 1: The graph G for the lower bound

has to use two edges from the set $F=\{\{t_jv\} : j=1,\ldots k\}$. Thus, the total h-route flow for the instance $\mathcal{I}$ is upper bounded by $h\cdot|F|/2$, that is, $\mathcal{F}^h(\mathcal{I}) \leq hk/2$. On the other hand, $\mathcal{F}^1(\mathcal{I}) = k(h-1)$. This yields the desired bound. □

The situation is completely different for general multicommodity h-route flows. Even though the maximum 2-route flow is as large as the maximum 1-route flow for single source multicommodity instances, for general instances the ratio between the sizes of a maximum 1-route flow and a maximum 2-route flow is as large as $k/2$. On the other hand, a trivial upper bound on the ratio is $\mathcal{F}^1(\mathcal{I}) \leq k\mathcal{F}^h(\mathcal{I})$.

Theorem 2.5. *For every pair of integers $h,k \geq 2$ there exists a graph $G = (V,E)$ and an instance $\mathcal{I} = (s_1,\ldots,s_k;t_1,\ldots,t_k)$ of the multicommodity flow problem such that for each i, the vertices s_i and t_i are h-connected, and, at the same time,*

$$\mathcal{F}^1(\mathcal{I}) \geq k\left(1-\frac{1}{h}\right)\mathcal{F}^h(\mathcal{I}).$$

Proof. Let G be a graph on $k+1$ distinct vertices $v_1,\ldots,v_{k+1}$ with v_i connected by $h-1$ parallel edges with v_{i+1}, for $i = 1,\ldots,k$, and v_{k+1} connected by an edge e with v_1 (Figure 2). Consider an instance $\mathcal{I}$ with $s_i = v_i$ and $t_i = v_{i+1}$, for $i = 1,\ldots,k$. Then, $\mathcal{F}^1(\mathcal{I}) = k(h-1)$. On the other hand, $\mathcal{F}^h(\mathcal{I}) \leq h$, since

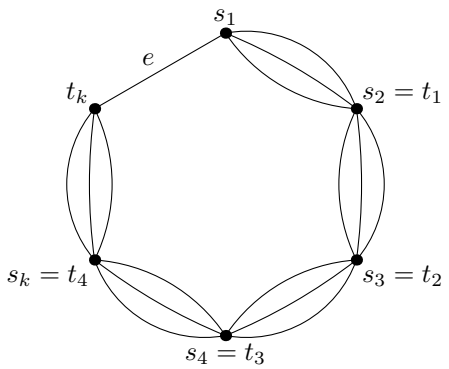

Figure 2: The graph G for $h = 4$ and $k = 5$

an elementary h-route flow between s_i and t_i has to use the edge $e = \{v_{k+1}, v_1\}$, for every $i = 1,\ldots,k$. This yields the desired bound. □

Theorem 2.1 relies on the assumption that the network has uniform edge capacities. The next theorem shows that without this assumption, the result does not hold.

Theorem 2.6. *For every $C \geq 1$ and every integer $h \geq 1$, there exists an undirected network $G = (V,E)$ with maximum edge capacity C and an instance $\mathcal{I} = (s;t_1,\ldots,t_k)$ of the single source multicommodity flow problem such that for each i, $\mathcal{F}^1(s,t_i) = \mathcal{F}^h(s,t_i)$, and, at the same time,*

$$\mathcal{F}^1(\mathcal{I}) \quad \geq \quad \left(C - \frac{C-1}{h}\right) \cdot \mathcal{F}^h(\mathcal{I}).$$

Proof. Choose $k = \lceil (C(h-1)+1)/h \rceil$ and consider a network G with $k+2$ vertices $V = \{s,u,t_1,t_2,\ldots t_k\}$ connected in the following way: s is connected with u by h edges, $h-1$ of them with capacity C and one with capacity 1, and for each $i \in \{1,\ldots,k\}$, u and t_i are connected by h edges, each of capacity 1 (Figure 3). Then, for an instance $\mathcal{I} = (s;t_1,\ldots,t_k)$ we have $\mathcal{F}^1(\mathcal{I}) = C(h-1)+1$ yet $\mathcal{F}^h(\mathcal{I}) = h$. □

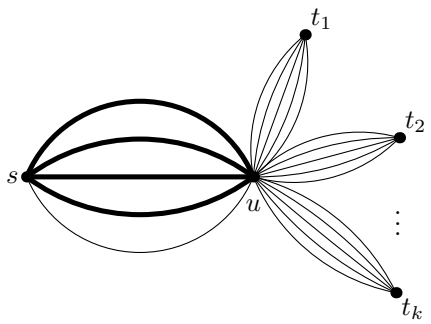

Figure 3: A bad network for nonuniform single source h-route flows (for $h = 5$).

2.2 Upper bounds: Preliminaries

In this section we cover three steps of the proof that are common for both non-duplex and duplex flows. The first step is essentially an induction: We restrict ourselves to instances with several useful properties, using the fact that any potential minimal counterexample has these properties. Second, on these instances, we choose some suitable maximal classical flow. Third, based on this flow, we define certain auxiliary structures on the empty edges. In the last step of the proof, which is done in the next subsections separately for non-duplex and duplex flows, we use these empty edges to reroute some flow and obtain an h-route flow of the appropriate size.

The following lemma shows that it is sufficient to prove Theorems 2.1 and 2.2 only for graphs $G = (V,E)$ and instances $\mathcal{I} = (s;t_1,\ldots,t_k)$ satisfying the following three properties:

A1 For each commodity i, the only minimum $s - t_i$ cut is the cut $\{t_i\}$ (we call it a *trivial* cut).

A2 In every integral maximum flow for the instance $\mathcal{I}$, each empty edge is incident to at least one of the sinks t_i, and, moreover, if an empty edge is incident to exactly one sink, then the degree of the sink is exactly h.

A3 Omitting any of the sinks from the instance $\mathcal{I}$ results in a decrease of the maximum flow (i. e., for every i, if we denote by $\mathcal{I}_{-i}$ the instance $\mathcal{I}$ without the sink t_i, $\mathcal{F}^1(\mathcal{I}_{-i}) < \mathcal{F}^1(\mathcal{I})$).

Lemma 2.7. *Let G and $\mathcal{I}$ be a graph and an instance that represent a conterexample to Theorem 2.1 or Theorem 2.2 with minimal $m+k$ (the number of edges plus the number of commodities). Then the Properties A1-A3 hold.*

Proof. Suppose we have a graph G and an instance $\mathcal{I}$ that do not satisfy the Properties A1-A3. We construct a smaller graph G' and an instance $\mathcal{I}'$ such that the classical maximum flow is the same in both cases and the maximal size of any type of h-route flow considered in the theorems can only decrease. Thus if G and $\mathcal{I}$ violate any claim in the theorems, then also G' and $\mathcal{I}'$ violate it and the proof is completed.

A1. Assume that there exists a commodity i and a minimum cut C for the commodity that is not trivial. Let δ_j denote the connectivity of s and t_j and let us denote by $\mathcal{F}$ an integral minimum cost maximum flow for $\mathcal{I}$. If the only commodity in the flow $\mathcal{F}$ that uses edges in the cut C is the commodity i, we perform the following modification of G: the t_i-side of G is merged into a single vertex t, that is,

keep every edge on the s-side, remove every edge on the t_i-side and for every edge $\{u,v\} \in C$ with v on the t_i-side, replace $\{u,v\}$ by a new edge $\{u,t\}$. We get a graph G' that is smaller than G and for an instance $\mathcal{I}' = (s;t_1,\dots t_{i-1},t,t_{i+1},\dots,t_k)$ on G', the connectivity of s and t_j is δ_j for $j \in \{1,\dots,k\}, j \neq i$, and the connectivity of s and t is δ_i, and the (classical) maximum flows for $\mathcal{I}$ in G and for $\mathcal{I}'$ in G' have the same size. The graph G' is smaller than G yet $\mathcal{F}^1(\mathcal{I}) = \mathcal{F}^1(\mathcal{I}')$ (note that multi-edges may occur). Any h-balanced flow for $\mathcal{I}'$ in G' can be easily extended into an h-balanced flow of the same size for the instance $\mathcal{I}$ in G.

If there are also some other commodities that use the cut C in the flow $\mathcal{F}$, we redirect the part of their flow going through C to t_i. This is possible since the cut is minimal and there will be no other commodities interfering. This way we maintain the same amount of the total flow and we argue as before.

A2. From now on we assume that for every commodity, every minimum cut is the trivial one. We denote by $\mathcal{F}$ an integral minimum cost maximum flow for $\mathcal{I}$ that does not satisfy the Property A2. Recall that the cost is uniform, i. e., the cost of an integral flow is just the number of edges used.

Assume first that there exists an edge e that is empty in $\mathcal{F}$ and that is not incident to any of the sinks t_i. Since e is not incident to any terminal node and since for every i each minimum $s-t_i$ cut is the trivial one, removing e from the graph G does not decrease the connectivity of any commodity and the maximum flow for the instance $\mathcal{I}$. As in the previous proof, an h-balanced flow for the smaller graph can be interpreted as a solution for G.

Similarly, if there exists an edge e that is empty in $\mathcal{F}$ and that is incident to exactly one sink and the degree of the sink is higher than h, deletion of e does not decrease the connectivity of any commodity below h and it does not decrease the maximum flow for the instance $\mathcal{I}$. Again, an h-balanced flow for the smaller graph can be interpreted as a solution for G.

A3. Suppose that the graph G and the instance $\mathcal{I}$ do not satisfy the Property $A3$, that is, there exists a commodity i such that $\mathcal{F}^1(\mathcal{I}_{-i}) = \mathcal{F}^1(\mathcal{I})$. We omit the commodity i to obtain the smaller instance $\mathcal{I}' = \mathcal{I}_{-i}$.

To finish the proof, note that all the reductions work also for half-integral and duplex h-route flows. □

Let G and $\mathcal{I}$ be a graph and an instance satisfying the three properties A1-A3 and consider an arbitrary integral minimum cost maximum flow for the instance $\mathcal{I}$. By the characterization of h-route flows as h-balanced flows (Theorem 1.1), the flow of every commodity with flow h or more is already an h-route flow. Our aim is, for every commodity with flow less than h, to find new edge disjoint paths between the source s and the relevant sink and to send some flow along each of them while not decreasing the flow of other commodities much. For this process we start with a particular minimum cost maximum flow that is described in Observation 2.8.

Given an integral flow for the instance $\mathcal{I}$, we denote, for a non-terminal vertex v, the number of empty edges incident to v by $p(v)$, and we denote the number of empty edges connecting v and the sink t_i by $m_i(v)$. By the Property A2, we have $\sum_{i=1}^{k} m_i(v) = p(v)$ for each non-terminal vertex v.

Observation 2.8. For the graph G and the instance $\mathcal{I}$ described above, there exists an integral minimum cost maximum flow such that for every non-terminal vertex v and for every i:

- $m_i(v) \leq \lceil p(v)/2 \rceil$.

Moreover, in every integral minimum cost maximum flow, for every non-terminal vertex v and for every i, the following holds:

- if $m_i(v) > p(v)/2$ then there exists at least one flow path of a commodity other than i going through v.

Proof. Consider an arbitrary integral minimum cost maximum flow and for a non-terminal vertex v denote by $r_{-i}(v)$ the number of flow paths of commodities other than i passing through v. Note that all empty edges incident to v are connected to a sink of degree exactly h (the Property A2). We are going to observe that $m_i(v) < p(v)/2 + r_{-i}(v)$, for every non-terminal vertex v and every commodity i.

Assume, for a contradiction, that $m_i(v) \geq p(v) - m_i(v) + 2r_{-i}(v)$ for some v and i, and consider the $s - t_i$ cut $\{v, t_i\}$. Due to our assumption, the size of this cut is smaller than or equal to the size of the trivial $s - t_i$ cut $\{t_i\}$ which is a contradiction with the Property A1. This completes the proof of the second part of Observation 2.8.

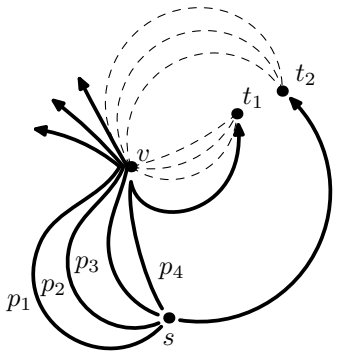

Figure 4: An example of a non-terminal vertex v satisfying the first property of Observation 2.8. Dashed lines represent empty edges and solid lines represent flow paths. We have $p(v) = 6$, $m_1(v) = 3$, $m_2(v) = 3$ and $r_{-1}(v) = 3$.

Now, if there is a non-terminal vertex v and a commodity i with $m_i(v) > \lceil p(v)/2 \rceil$, then there are $r_{-i}(v) > m_i(v) - p(v)/2$ flow paths of other commodities passing through v. Choose one of them, say a path p of a commodity j, and reroute it to t_i. To be more precise, the new path goes from the source s to the vertex v along the original path p, and then it continues to t_i along one of the empty edges connecting v and t_i. After the modification, $m_i(v)$ decreases by one and $m_j(v)$ increases by one; the cost and the size of the total flow are not affected. This way we continue until $m_i(v) \leq \lceil p(v)/2 \rceil$ for every i. Notice that the changes done in the flow around v will not destroy the desired property for any other vertex.

We apply the same rerouting procedure for every other non-terminal vertex v' for which there exists a commodity i' such that $m_{i'}(v') > \lceil p(v')/2 \rceil$. □

From now on we fix some minimum cost maximum flow satisfying Observation 2.8 and denote it $\mathcal{F}$. By the choice of $\mathcal{F}$ and by the Property A2, each empty edge is incident either to two different sinks or to a sink and to a vertex adjacent to another sink; the last assertion holds since otherwise there exists a smaller cost flow of the same size. The idea of the proof is to exploit these empty edges to reroute some

flow from other commodities to each sink with flow less than h. If we succeed to provide a non-zero flow along at least h edges to each sink, we get a non-zero h-balanced flow for each commodity.

We define an auxiliary structure, called octopus, which will help us to organize the rerouting. Formally, an *octopus* is a (multi)graph that is a union of edge disjoint paths of length one and two that start in the same vertex; the paths are called *tentacles*. If an octopus O is a subgraph of the graph G and the initial vertex of the paths (i. e., of the tentacles) is a vertex v, we say that the octopus *is sitting in* v.

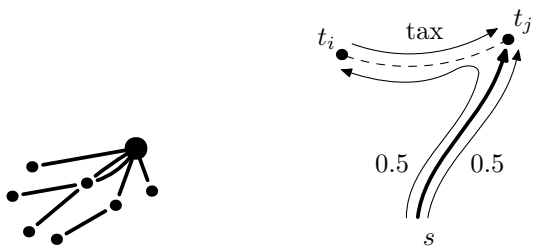

Figure 5: An octopus

For every commodity i with flow smaller than h, we define an octopus O_i. The octopus O_i is sitting in the terminal t_i and has $h - f_i$ tentacles, where f_i denotes the amount of flow of a commodity i in $\mathcal{F}$, and the tentacles reach through different empty edges to neighboring vertices (if there are more than $h - f_i$ empty edges incident to t_i, we choose any $h - f_i$ of them). Later we will amend the octopuses, namely we will lengthen some of the tentacles.

Consider a non-terminal vertex v. The Property A2 implies that the number of tentacles reaching v is $p(v)$ and we denote them by $\tau_1, \ldots, \tau_{p(v)}$. If none of the octopuses reaches v by more than $p(v)/2$ tentacles, there exists a permutation π of the tentacles $\tau_1, \ldots, \tau_{p(v)}$ which consists only of 2-cycles and possibly one 3-cycle such that tentacles τ_l and $\pi(\tau_l)$ belong to different octopuses. For example, always greedily form a 2-cycle between tentacles of two distinct octopuses with the maximal number of remaining tentacles ending in v. Do this until 2 or 3 tentacles remain (depending on the parity of $p(v)$), and then form the last cycle (only this last cycle can be a 3-cycle). We lengthen the tentacle τ_l through the edge used by the tentacle $\pi(\tau_l)$ so that τ_l now terminates in the sink in which the octopus with tentacle $\pi(\tau_l)$ is sitting.

If there exists an octopus O_i that reaches the non-terminal vertex v by more than $p(v)/2$ tentacles, then such an octopus is exactly one. For such an octopus, by Observation 2.8, the number of its tentacles reaching v is exactly $\lceil p(v)/2 \rceil$. There exists a permutation π of $p(v) - 1$ tentacles reaching v such that it consists of 2-cycles of tentacles belonging to different octopuses, namely a matching of all but one tentacles of O_i to all the others. In a similar way as before, each tentacle τ involved in the permutation is lengthened to the sink in which the octopus with the tentacle $\pi(\tau)$ is sitting. Recall that by Observation 2.8 there exists a flow path passing through v that does not belong to the commodity i, and by the minimality of the cost of the flow $\mathcal{F}$, the terminal vertex of the path is adjacent to v.

At this point, each tentacle of an octopus reaches either another terminal vertex (we say that the tentacle *touches* the corresponding commodity), or a flow path of another commodity that no other tentacle reaches (again we say that the tentacle *touches* the corresponding commodity). Moreover, each tentacle τ is stretched only through empty edges and at most one tentacle is stretched through each

empty edge in each direction; if there are two tentacles stretched through the same edge (in opposite direction) they belong to different octopuses.

Observation 2.9. For each i, the number of tentacles that touch the commodity i is strictly less than f_i.

Proof. Were it not the case, it would be possible to redirect the complete flow of the commodity i, through the tentacles touching it, to other terminals without decreasing the total flow, contradicting the Property A3. □

2.3 Upper bounds: Non-duplex flows

Proof of Theorem 2.1. We start by constructing the half-integral h-route flow of size (at least) $\mathcal{F}^1(\mathcal{I})/2$. Then we explain how to increase the size of the flow to (at least) $\mathcal{F}^1(\mathcal{I})/(2-2/h)$.

For each tentacle of the octopus O_i touching the commodity $j \neq i$, we reroute a half unit of the flow of a commodity j to t_i along the edges that the corresponding tentacle is stretched through. Observation 2.9 guarantees that every commodity j has enough flow to provide a half unit for each tentacle touching it and yet to keep more than $f_j/2$ units for itself. We decrease the flow of every unaffected path to one half.

At this point, the amount of flow of a commodity i with $f_i < h$ is $h/2$ and the amount of flow of a commodity i with $f_i \geq h$ is $f_i/2$. Moreover, since the initial flow was integral (flow paths from source to terminals were disjoint), the new flow paths of each individual commodity will be edge disjoint. Thus, we have an h-balanced flow of size at least $\mathcal{F}(\mathcal{I})/2$, for the instance $\mathcal{I}$, and by construction, the flow is half integral.

To prove the sharper bound (not necessarily with half-integral flows) we observe that for every commodity with flow at most $h-1$ in the initial flow, its h-balanced flow at the end is at least $h/2$ which corresponds to the ratio $2-2/h$. The only problem is with commodities with original flow h or more. Thus, if we manage to slightly increase the final flow of these commodities, the proof is completed. Recall that no octopus is sitting in a terminal vertex of a commodity with flow h or more.

We proceed as follows: every commodity t_j with initial flow h or more will demand a *tax* of $1/(2h-2)$ units of flow for each path that it provided to another commodity. Commodities are able to pay these taxes since every commodity had initial flow that was at least one greater than the number of tentacles touching it (Observation 2.9) and every commodity requires help from at most $h-1$ other commodities (more precisely, needs at most $h-1$ new edge disjoint paths). In the worst case, it keeps (only) a half unit of flow for itself and spends the other half on taxes for the $h-1$ helpers.

The flow corresponding to a tax of a commodity t_i paid to a commodity t_j flows from s to t_i along an original path of a commodity i and then from t_i to t_j along the tentacle of the octopus sitting in t_i; in the case of an octopus O_i touching a path of the commodity j (and not directly touching the sink t_j) the flow flows from s to t_i along an original path of the commodity i, then along the tentacle of the octopus O_i and finally along an edge of the flow path of the commodity j that the tentacle of O_i touches. In addition to this, we set the flow along each path that was unaffected by the rerouting process to $1/(2-2/h)$ (and not to $1/2$ as we did for the half-integral flow). In this way, a commodity with an initial flow $f_i \geq h$ will have a final h-balanced flow at least $f_i(h/(2h-2))$, which corresponds to an h-route flow of the same size. □

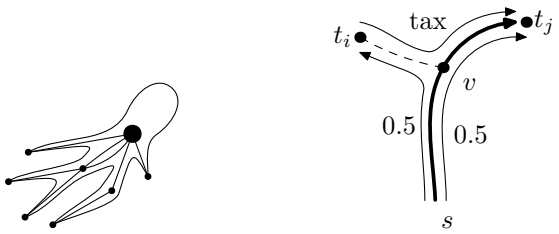

Figure 6: Taxation: on the left side is depicted the case when a tentacle touches a terminal vertex, and on the right side is depicted the case when a tentacle touches a path of other commodity.

Concerning the proof of Theorem 2.3, namely the half-integrality, notice that for $h = 2$ the taxes in the previous proof are equal to $1/2$. Thus the resulting flow is half-integral.

2.4 Upper bound: Duplex flows

Proof of Theorem 2.2. We now construct an h-balanced duplex flow of the same size as the original flow $\mathcal{F}^1(\mathcal{I})$. To do this, for each octopus O_i we reroute some flow from other sinks to the sink t_i. More exactly, for a certain amount $z_i \in [0,1]$, we reroute z_i units of the flow along each of the tentacles of O_i to t_i. At the same time, we guarantee that from the original flow f_i to t_i, exactly $f_i(1-z_i)$ units are rerouted to other sinks using their octopuses. If this rerouted flow is taken evenly from all f_i incoming paths, then the resulting flow of the commodity i has size at least hz_i and thus it is an h-route flow. If a tentacle of O_i touches a flow path of a commodity j at v which is not the sink t_j, then we take z_i/f_j units of the rerouted flow from this flow path and the remaining flow is routed from t_j back one edge along to v. (Cf. the last paragraph of the proof.)

The choice of z_i guarantees that the commodities with original flow less than h can be rerouted. On the other hand, each commodity with flow h or more is touched by less then h tentacles due to Observation 2.9. Consequently there is enough original flow for the rerouting and if taken evenly from all paths, the flow of this commodity is also an h-route flow afterwards. After the rerouting, each edge has in each direction a flow of at most 1, either at most 1 from the original flow or at most $z_i \leq 1$ from the rerouting along one tentacle.

It remains to guarantee the existence of numbers z_i described above. For simplicity, renumber the commodities so that the first k' sinks t_i are exactly those with the initial flow $f_i < h$, that is, exactly those with an octopus. Let a_{ij} be the number of tentacles of O_j touching the commodity i. We need the values z_i to satisfy, for each $i \leq k'$,

$$\sum_{j=1}^{k'} a_{ij} z_j = f_i(1-z_i)\,. \tag{2.2}$$

Define a function $F : \mathbf{R}^{k'} \to \mathbf{R}^{k'}$ so that its ith coordinate is

$$(F(\vec{z}))_i = 1 - \frac{1}{f_i} \sum_{j=1}^{k'} a_{ij} z_j\,.$$

Then the system of equations (2.2) is equivalent to the equation $\vec{z} = F(\vec{z})$. Due to Observation 2.9 we know that for each $i \leq k'$

$$\sum_{j=1}^{k'} a_{ij} < f_i \,. \tag{2.3}$$

This implies that F maps the unit cube $[0,1]^{k'}$ to itself. Obviously, F is continuous as it is a linear function. Using Brouwer's fixed-point theorem (which asserts that any continuous mapping from a ball to itself has a fixed point) and the fact that a ball and a cube are homeomorphic, we conclude that F has a fixed point, that is, the equation $\vec{z} = F(\vec{z})$ has a solution. This solution also satisfies the system (2.2), which concludes the proof. Note that the proof is still constructive, since F is linear and we can find the solution efficiently.

In fact, inequalities (2.3) imply that F is a strong contraction under the l_∞ norm, which means that there exists a constant $\alpha < 1$ such that for any $\vec{y}, \vec{z} \in [0,1]^{k'}$, it holds that

$$\|F(\vec{z}) - F(\vec{y})\|_\infty \leq \alpha \cdot \|\vec{z} - \vec{y}\|_\infty \,,$$

where $\|\vec{z}\|_\infty = \max_{i=1,\ldots,k'} |z_i|$. For the constant $\alpha = \max_{i=1,\ldots,k'} (\sum_{j=1}^{k'} a_{ij}/f_i)$ the strong contraction property follows easily for each coordinate of $F(\vec{z}) - F(\vec{y})$ by subtracting the defining equations for $(F(\vec{z}))_i$ and $(F(\vec{y}))_i$. The strong contraction property implies that in the unit cube, there exists a (unique) solution of the equation $\vec{z} = F(\vec{z})$, namely a limit of points obtained by repeated applications of F, starting at any point x the unit cube. (Note that the distances between subsequent points in the sequence x, $F(x)$, $F(F(x))$, etc., decrease in a geometric sequence due to the contraction property.) This gives an alternative elementary proof without use of Brouwer's theorem.

To prove the second part of Theorem 2.2, we slightly modify the octopuses. In the case of a non-terminal vertex v and a 3-cycle (ijl) in the permutation π used for extending the tentacles, we do not extend the tentacles along the cycle but instead we split each tentacle into two halves and extend them to the remaining two vertices. Thus this 3-cycle will contribute $1/2$ to each a_{ij}, a_{jl}, a_{li}, a_{ji}, a_{lj}, a_{il}. Then, for example, the edge vt_i will have flow z_i to t_i and $(z_j + z_l)/2$ from t_i.

The same argument as above guarantees that the system of equations (2.2) has again a solution in the unit cube, with flow at most 1 in each direction along any edge. It remains to verify that on any edge, the sum of flows in both directions is at most $2 - 2/h$. There are three types of such edges.

1. An edge used by two tentacles, one tentacle of O_i touching the commodity j and one tentacle of O_j touching the commodity i. This may be an edge $t_i t_j$, or an edge incident to a non-terminal v and involved in a 2-cycle (ij) of the permutation π. The total flow is $z_i + z_j$. We have $f_i, f_j \leq h - 1$ as there are octopuses at t_i and t_j, and the corresponding equations in (2.2) imply (after removing the left-hand side terms for other tentacles) that $z_j \leq (h-1)(1-z_i)$ and $z_i \leq (h-1)(1-z_j)$. Adding these inequalities and dividing by h we obtain the desired bound $z_i + z_j \leq 2 - 2/h$.

2. An edge vt_i with a non-terminal vertex v and involved in a 3-cycle (ijl) of π. This edge has total flow $z_i + (z_j + z_l)/2$. Again $f_i, f_j, f_l \leq h - 1$ and the corresponding equations in (2.2) imply

$$\begin{aligned} 2(h-1)z_i + z_j + z_l &\leq 2(h-1)\,, \\ z_i + 2(h-1)z_j + z_l &\leq 2(h-1)\,, \\ z_i + z_j + 2(h-1)z_l &\leq 2(h-1)\,. \end{aligned}$$

Adding the first inequality multiplied by $2(h-1)$ and the remaining two inequalities multiplied by $h-2$ yields the desired bound.

3. An edge vt_j for a tentacle of O_i touching the commodity j at a non-terminal vertex v. The flow is at most $1-z_i/f_j$ of the original flow to t_j, and $z_i - z_i/f_j$ of the flow that is rerouted from t_j. Since $z_i \leq 1$, it suffies to show that $f_j \leq h$ to conclude that the total flow on the edge vt_j is at most $1+z_i(1-2/f_j) \leq 2-2/h$. We argue as follows. If we redirect the unit of the original flow going through the edge vt_j to the sink v_i, we get another maximal flow with an empty edge adjacent to t_j. By the Property A2 the degree of t_j is at most h which in turn implies that $f_j \leq h$. □

3 Disconnecting cuts

We will denote the size of a minimum h-disconnecting cut for an instance $\mathcal{I}$ by $\mathcal{C}^h(\mathcal{I})$.

Theorem 3.1. *For every $h \geq 2$ and every instance $\mathcal{I}$ of the single source flow problem,*

$$\frac{\mathcal{F}^h(\mathcal{I})}{h} \;\leq\; \mathcal{C}^h(\mathcal{I}) \;\leq\; \left(2-\frac{2}{h}\right)\cdot \mathcal{F}^h(\mathcal{I})\,. \tag{3.1}$$

Moreover, for every $h \geq 2$ and every $\varepsilon > 0$ there exists an instance $\mathcal{I} = \{s;t\}$ of the problem such that

$$(1-\varepsilon)\cdot \mathcal{F}^h(\mathcal{I}) \;\leq\; \mathcal{C}^h(\mathcal{I})\,, \tag{3.2}$$

and for every $k \geq 1$ and every $h \geq 2$ there exists an instance $\mathcal{I}$ such that

$$\frac{\mathcal{F}^h(\mathcal{I})}{h} \;=\; \mathcal{C}^h(\mathcal{I})\,. \tag{3.3}$$

Proof. Given a decomposition of an h-route flow into a linear combination of elementary h-route flows, we have to cut at least one of the h paths of every h-system in the decomposition. Altogether we have to cut edges of total capacity at least $\mathcal{F}^h(\mathcal{I})/h$ which proves the first inequality.

To prove the inequality $\mathcal{C}^h(\mathcal{I}) \leq (2-2/h)\cdot\mathcal{F}^h(\mathcal{I})$ we observe that a minimum classical cut is also an h-cut, and from the duality of flows and cuts we know that the size of this cut is equal to $\mathcal{F}^1(\mathcal{I})$. We apply the bound $\mathcal{F}^1(\mathcal{I}) \leq (2-2/h)\cdot\mathcal{F}^h(\mathcal{I})$ of Theorem 2.1 (without loss of generality we assume that all sinks in the instance $\mathcal{I}$ are h-connected with the source) and the proof is completed.

Concerning the second part of the theorem, consider a graph consisting of two vertices s and t connected by m parallel edges, with $m \geq h$. The maximum h-route flow has size m and the minimum h-disconnecting cut has size $m-(h-1)$. We conclude that for every $\varepsilon > 0$ there exists an integer m such that $(m-h+1)/m \geq 1-\varepsilon$ and, thus, there exists an instance $\mathcal{I} = \{s;t\}$ satisfying the inequality (3.2). Note that a fractional disconnecting cut is in this case (almost) h-times better: take a fraction $1/h$ of each edge in the cut.

For the last part of the theorem, consider the instance and the network in Figure 3 with every edge capacity set to one. Then, $\mathcal{F}^h(\mathcal{I}) = hk$ and $\mathcal{C}^h(\mathcal{I}) = k$. □

The proof above implies that the minimal classical cut is a good approximation for the h-disconnecting problem. It can be found efficiently, and thus we have:

Corollary 3.2. *For every $h \geq 2$, there exists a polynomial time $(2h-2)$-approximation algorithm for the h-disconnecting problem with a single source.*

Remark. The bound on the performance of the algorithm is not far from what happens for "bad" instances. Think about a simple graph consisting of two vertices u, v connected by h parallel edges and an instance with one commodity with source in u and sink in v: the minimum disconnecting cut has size 1 while the disconnecting cut obtained by the algorithm has size h.

We also note that the bound (3.1) can be slightly improved to

$$\frac{\mathcal{F}^h(\mathcal{I})}{h} \leq \mathcal{C}^h(\mathcal{I}) \leq \left(2-\frac{2}{h}\right)\cdot\mathcal{F}^h(\mathcal{I})-(h-1)$$

by deleting all but $h-1$ edges from the minimum classical cut (instead of deleting all the edges). If there is only one commodity, this slightly modified procedure computes an optimal h-disconnecting cut.

4 Open problems

We conclude with two open problems about disconnecting cuts for multiroute flows. The approximation ratio of the algorithm for disconnecting cuts for single source flow problems described in the last section is $2h-2$; design a better algorithm. Similarly, design an approximation algorithm for the disconnecting cut problem for the more general multiroute multicommodity flow problems (e. g., single source and non-uniform capacities, multiple sources and uniform capacities). As the close relation between classical flows and multiroute flows is lost in these cases, a novel approach will be needed.

Acknowledgment The authors thank Neal Young and Yefim Dinitz for stimulating discussions.

References

[1] * CHARU C. AGGARWAL AND JAMES B. ORLIN: On multiroute maximum flows in networks. *Networks*, 39:43–52, 2002. [Wiley:10.1002/net.10008]. 1, 1.1, 1.1

[2] * R. K. AHUJA, T. L. MAGNANTI, AND J. B. ORLIN: *Network Flows: Theory, Algorithms and Applications*. Prentice-Hall, 1993. 1

[3] * YONATAN AUMANN AND YUVAL RABANI: An $O(\log k)$ approximate min-cut max-flow theorem and approximation algorithm. *SIAM Journal on Computing*, 27(1):291–301, 1998. [SICOMP:10.1137/S0097539794285983]. 1.1

[4] * AMITABHA BAGCHI, AMITABH CHAUDHARY, AND PETR KOLMAN: Short length Menger's theorem and reliable optical routing. *Theoretical Computer Science*, 339:315–332, 2005. [TCS:10.1016/j.tcs.2005.03.009]. 1.1

[5] * AMITABHA BAGCHI, AMITABH CHAUDHARY, PETR KOLMAN, AND CHRISTIAN SCHEIDELER: Algorithms for fault-tolerant routing in circuit switched networks. *SIAM Journal on Discrete Mathematics*, 21:141–157, 2007. [SIDMA:10.1137/S0895480102419743]. 1.1

[6] * AMITABHA BAGCHI, AMITABH CHAUDHARY, PETR KOLMAN, AND JIŘÍ SGALL: A simple combinatorial proof for the duality of multiroute flows and cuts. Technical Report 2004-662, Charles University, Prague, 2004. 1, 1.1, 1.1

[7] * GEORG BAIER, EKKEHARD KOEHLER, AND MARTIN SKUTELLA: The k-splittable flow problem. *Algorithmica*, 42(3–4):231–248, 2005. [Algorithmica:jtl3r3k633523486]. 1.1

[8] * ALOK BAVEJA AND ARAVIND SRINIVASAN: Approximation algorithms for disjoint paths and related routing and packing problems. *Mathematics of Operations Research*, 25(2):255–280, 2000. 1.1

[9] * HENNING BRUHN, JAKUB ČERNÝ, ALEXANDER HALL, AND PETR KOLMAN: Single source multiroute flows and cuts on uniform capacity networks. In *Proc. 18th Ann. ACM-SIAM Symposium on Discrete Algorithms (SODA'07)*, pp. 855–863. SIAM, 2007. [SODA:1283383.1283475]. *

[10] * AMIT CHAKRABARTI, CHANDRA CHEKURI, ANUPAM GUPTA, AND AMIT KUMAR: Approximation algorithms for the unsplittable flow problem. *Algorithmica*, 47(1):53–78, 2007. [Algorithmica:m5710531h2514815]. 1.1

[11] * CHANDRA CHEKURI AND SANJEEV KHANNA: Edge-disjoint paths revisited. *ACM Transactions on Algorithms*, 3, 2007. [ACM:1290672.1290683]. 1.1

[12] * ISRAEL CIDON, RAPHAEL ROM, AND YUVAL SHAVITT: Analysis of multi-path routing. *IEEE/ACM Transactions on Networking*, 7(6):885–896, 1999. [ACM:323983.323992]. 1.1

[13] * Y. DINITZ, N. GARG, AND M. GOEMANS: On the single source unsplittable flow problem. *Combinatorica*, 19(1):17–41, 1999. [Springer:pg0crc9nqlfdmajj]. 1.1

[14] * D. DU AND R. CHANDRASEKARAN: The multiroute maximum flow problem revisited. *Networks*, 47(2):81–92, 2006. [Wiley:10.1002/net.20099]. 1.1

[15] * DONGLEI DU: *Multiroute Flow Problem*. Ph.D. thesis, The University of Texas at Dallas, 2003. 1.1

[16] * THOMAS ERLEBACH AND ALEXANDER HALL: NP-hardness of broadcast scheduling and inapproximability of single-source unsplittable min-cost flow. *Journal of Scheduling*, 7(3):223–241, 2004. [Springer:h57n764rq0124578]. 1.1

[17] * L. R. FORD AND D. R. FULKERSON: Maximum flow through a network. *Canad. J. Math.*, 8:399–404, 1956. 1.1

[18] * NAVEEN GARG, VIJAY V. VAZIRANI, AND MIHALIS YANNAKAKIS: Approximate max-flow min-cut theorems and their applications. *SIAM Journal on Computing*, 25(2):235–251, 1996. [SICOMP:10.1137/S0097539793243016]. 1.1

[19] * V. GURUSWAMI, S. KHANNA, R. RAJARAMAN, B. SHEPHERD, AND M. YANNAKAKIS: Near-optimal hardness results and approximation algorithms for edge-disjoint paths and related problems. *Journal of Computer and System Sciences*, 67(3):473–496, 2003. [JCSS:10.1016/S0022-0000(03)00066-7]. 1.1

[20] * KOUSHIK KAR, MURALI KODIALAM, AND T. V. LAKSHMAN: Routing restorable bandwidth guaranteed connections using maximum 2-route flows. *IEEE/ACM Transactions on Networking*, 11:772–781, 2003. [ACM:948928.948935]. 1.1

[21] * WATARU KISHIMOTO: A method for obtaining the maximum multiroute flows in a network. *Networks*, 27(4):279–291, 1996. 1, 1.1

[22] * WATARU KISHIMOTO AND M. TAKEUCHI: On m-route flows in a network. *IEICE Transactions*, J-76-A(8):1185–1200, 1993. (in Japanese). 1.1

[23] * J. M. KLEINBERG: Single-source unsplittable flow. In *Proc. 37th FOCS*, pp. 68–77. IEEE Computer Society Press, 1996. [FOCS:10.1109/SFCS.1996.548465]. 1.1

[24] * RONALD KOCH, MARTIN SKUTELLA, AND INES SPENKE: Approximation and complexity of k-splittable flows. In *Proceedings of the Third Workshop on Approximation and Online Algorithms*, volume 3879 of *Lecture Notes in Computer Science*, pp. 244–257. Springer, 2005. [WADS:l871111265475j11]. 1.1

[25] * STAVROS G. KOLLIOPOULOS AND CLIFFORD STEIN: Approximation algorithms for single-source unsplittable flow. *SIAM Journal on Computing*, 31(3):919–946, 2002. [SICOMP:10.1137/S0097539799355314]. 1.1

[26] * PETR KOLMAN AND CHRISTIAN SCHEIDELER: Simple on-line algorithms for the maximum disjoint paths problem. *Algorithmica*, 39(3):209–233, 2004. [Algorithmica:ppc4pcxagm69d4jv]. 1.1

[27] * PETR KOLMAN AND CHRISTIAN SCHEIDELER: Improved bounds for the unsplittable flow problem. *Journal of Algorithms*, 61(1):20–44, 2006. [Elsevier:10.1016/j.jalgor.2004.07.006]. 1.1

[28] * TOM LEIGHTON AND SATISH RAO: Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *Journal of the ACM*, 46(6):787–832, November 1999. [JACM:331524.331526]. 1.1

[29] * N. LINIAL, E. LONDON, AND YU. RABINOVICH: The geometry of graphs and some of its algorithmic applications. *Combinatorica*, 15(2):215–245, 1995. 1.1

[30] * MAREN MARTENS AND MARTIN SKUTELLA: Flows on few paths: Algorithms and lower bounds. *Networks*, 48(2):68–76, 2006. [Wiley:10.1002/net.20121]. 1.1

[31] * MAREN MARTENS AND MARTIN SKUTELLA: Length-bounded and dynamic k-splittable flows. In *Operations Research Proceedings 2005*, pp. 297–302, 2006. [Springer:w2u3032j1804481t]. 1.1

[32] * K. MENGER: Zur allgemeinen Kurventheorie. *Fundam. Math.*, 10:96–115, 1927. 1.1

[33] * M. SKUTELLA: Approximating the single source unsplittable min-cost flow problem. *Mathematical Programming, Ser. B*, 91(3):493–514, 2002. [Springer:txfq9gvdtw0cac0g]. 1.1

[34] * ARAVIND SRINIVASAN: Improved approximations for edge-disjoint paths, unsplittable flow, and related routing problems. In *Proc. 38th FOCS*, pp. 416–425. IEEE Computer Society Press, 20–22 1997. [FOCS:10.1109/SFCS.1997.646130]. 1.1

AUTHORS

Henning Bruhn [About the author]
Mathematisches Seminar, Universität Hamburg, Germany
bruhns@math.uni-hamburg.de
http://www.math.uni-hamburg.de/home/bruhn/

Jakub Černý [About the author]
Department of Applied Mathematics, Faculty of Mathematics and Physics
Charles University, Czech Republic
kuba@kam.mff.cuni.cz
http://kam.mff.cuni.cz/~kuba/

Alexander Hall [About the author]
Google Switzerland GmbH, Zurich, Switzerland
alex.hall@gmail.com

Petr Kolman [About the author]
Department of Applied Mathematics, Faculty of Mathematics and Physics
Charles University, Czech Republic
kolman@kam.mff.cuni.cz
http://kam.mff.cuni.cz/~kolman/

Jiří Sgall [About the author]
Institute of Mathematics, AS ČR, Czech Republic
sgall@math.cas.cz
http://www.math.cas.cz/~sgall/

ABOUT THE AUTHORS

HENNING BRUHN has been a postdoc at Universität Hamburg. since summer 2006. In 2005 he obtained his Ph. D. under the supervision of Reinhard Diestel. He spent the following year in Grenoble learning combinatorial optimisation and French. He likes rock climbing and tries to learn Japanese.

ALEXANDER HALL received a Master's degree ("Diplom") in Computer Science at the Technische Universität München in 1998. In December 2003 he completed his doctoral studies in the group of Thomas Erlebach at the ETH Zürich and received a Ph. D. for his thesis "Scheduling and Flow-Related Problems in Networks." After being a post-doc at ETH Zürich and UC Berkeley, he is now with Google in Zurich.

JAKUB ČERNÝ is finishing his Ph. D. at the Department of Applied Mathematics of Charles University in Prague. His advisor is Pavel Valtr. He is interested in computational and discrete geometry and efficient algorithms in general. His hobbies are Aikido, improvisational comedy theatre, outdoor activities.

PETR KOLMAN is an Assistant Professor at Charles University in Prague, Czech Republic. After obtaining his Ph. D., he spent a year at the Heinz Nixdorf Institut at the University of Paderborn, Germany and a year at the University of California, Riverside.

JIŘÍ SGALL grew up in Prague, Czechoslovakia, where he received his RNDr. degree (equivalent of Master's) at Charles University under supervision of Antonín Sochor. Then he went to Carnegie Mellon University and received his Ph. D. under the supervision of Steven Rudich. After that he went back to Prague, Czech Republic, where he is now a senior researcher at the Institute of Mathematics of the Academy of Sciences of the Czech Republic. His main research interests are online and approximation algorithms for scheduling and other combinatorial problems. He also worked in communication complexity and proof complexity.