Single Source Multiroute Flows and Cuts on Uniform Capacity Networks*

Henning Bruhn† Jakub Černý‡ Alexander Hall§
Petr Kolman¶

October 9, 2006

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 $t_1, \ldots, t_k \in V$; 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 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.

*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.

†Laboratoire LEIBNIZ, France. henning.bruhn@imag.fr

‡Department of Applied Mathematics, Faculty of Mathematics and Physics, Charles University, Malostranské nám. 25, 118 00 Praha 1, Czech Republic. kuba@kam.mff.cuni.cz

§Institut für Theoretische Informatik, ETH Zentrum, CH - 8092 Zürich, Switzerland. alex.hall@inf.ethz.ch

¶Department of Applied Mathematics, Faculty of Mathematics and Physics, Charles University, Malostranské nám. 25, 118 00 Praha 1, Czech Republic. kolman@kam.mff.cuni.cz

We study the relation between classical and multiroute single source flows on 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's is at most $2(1 - 1/h)$-times larger than the maximum h-route flow between s and the t_i's and this is the best possible bound for $h \geq 2$. This, as we show, is in contrast to the situation of general multicommodity multiroute flows that are up to $k(1 - 1/h)$-times smaller than their classical counterparts.

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 $s - t_i$ is zero. We show that the maximum h-route flow is within $2(h-1)$ of the mininimum h-disconnecting cut, independently of the number of commodities; we also describe a $2(h - 1)$-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. 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 [18]. 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 (non-obvious) claim is also true: every h-balanced (acyclic) flow is an h-route flow [1, 4, 18].

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 a 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 $t_1, \dots, t_k \in V$; we denote it $\mathcal{I} = (s; t_1, \dots, t_k)$. 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 t_i's is at most $2(1-1/h)$ times larger than the maximum h-route flow between s and t_i's; this bound is the best possible for $h \geq 2$. The result is in contrast with the situation of general 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 below). 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 $2(h-1)$ of the min-cut. As a corollary of this relation we get a $2(h-1)$-approximation algorithm for the h-disconnecting cut problem.

1.1 Related Results

Kishimoto and Takeuchi [19] and later Aggarwal and Orlin [1] studied single commodity multiroute flows (cf. [4, 11, 12]). 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., [3, 17, 9] 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 [6, 7, 8, 16, 20, 22, 23, 24, 31]; the h-splittable flow problem allows at most h, not necessarily disjoint, paths per commodity [5, 21, 28, 27]; particular attention has been given to single source unsplittable flow problems [10, 13, 22, 30]. 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 [29] 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 [14] 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 [25], Aumann and Rabani [2], Linial, London and Rabinovich [26], Garg, Vazirani and Yannakakis [15]).

1.2 Preliminaries

As indicated in the title, in this paper we deal with networks with uniform capacities. For simplicity, we assume throughout the paper, without loss of generality, that every edge has capacity one. The number of vertices is denoted n and the number of edges m; we allow multi-edges. The letter k denotes the number of commodities and the letter h the number of routes in the elementary multiroute 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 for the instance $\mathcal{I}$. 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., $cost(e) = 1,\ \forall e \in E$).

Consider a network $G = (V, E)$. Let $s_1, \dots, s_k$ be k sources and $t_1, \dots, t_k$ 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 $\mathcal{Q} = \bigcup_{i=1}^{k} \mathcal{Q}_i$. The maximum h-route flow problem can be stated as the following linear program (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):

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

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

$$\min\ h\cdot\sum_{e\in E} x(e) \tag{2}$$
$$\begin{aligned} \sum_{e\in q} x(e) &\geq 1 \quad \forall q\in\mathcal{Q} \\ x(e) &\geq 0 \quad \forall e\in E\ . \end{aligned}$$

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

2.1 A lower Bound

Theorem 2.1 *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 $i = 1, \ldots, k$.

Consider the instance $\mathcal{I} = (s; t_1, \ldots, t_k)$. An elementary h-route flow between s and t_i, for $i = 1, \ldots, k$, has to use two edges from the set $F = \big\{\{t_i v\} : i = 1, \ldots k\big\}$. Thus, the total h-route flow for the instance $\mathcal{I}$ is

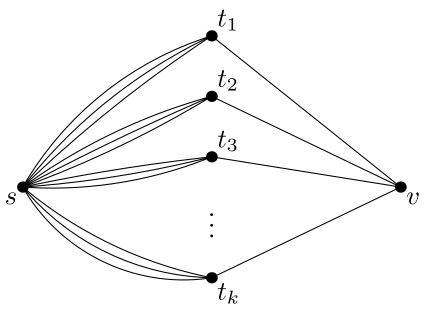

Figure 1: The graph G for the lower bound

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. □

In the next subsection we show that for h-route flows with a single source, this is the worst that can happen.

2.2 A tight upper bound

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*

$$\mathcal{F}^1(\mathcal{I}) \quad \leq \quad (2 - 2/h) \cdot \mathcal{F}^h(\mathcal{I}) \ . \tag{3}$$

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

Proof. 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) $\mathcal{F}^1(\mathcal{I})/(2-2/h)$.

The proof is by induction on the sum $n + m + k$ where the base case is proved for graphs $G = (V, E)$ and instances $\mathcal{I} = (s; t_1, \ldots, t_k)$ satisfying the following three assumptions:

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 adjacent to at least one of the sinks t_i, and, moreover, if an empty edge is adjacent 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, $\mathcal{F}(\mathcal{I}_{-i}) < \mathcal{F}(\mathcal{I})$).

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.

Base case Let G and $\mathcal{I}$ be a graph and an instance as above 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 described in Introduction, 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 a half unit of 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 2.3.

Given an integral flow for the instance $\mathcal{I}$, we denote, for a non-terminal vertex v, the number of empty edges adjacent 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 Assumption A2, we have $\sum_{i=1}^{k} m_i(v) = p(v)$ for each non-terminal vertex v.

Observation 2.3 *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 adjacent to v are connected to a sink of degree exactly h (Assumption 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 assumption A1. This completes the proof of the second part of Observation 2.3.

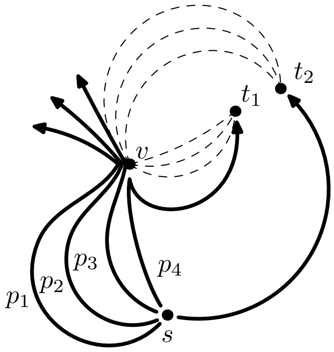

Figure 2: An example of a non-terminal vertex v satisfying the first property of Observation 2.3. Dashed lines represent empty edges and solid lines represent flow paths. We have $p(v) = 6$, $m_1(v) = 3$, $m_2(v) = 3$ and $r_{-1}(v) = 3$.

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 $m_i'(v') > \lceil p(v')/2 \rceil$.
$\square$

From now on we denote by $\mathcal{F}$ the minimum cost maximum flow from Observation 2.3. By the choice of $\mathcal{F}$ and by the Assumption A2, each empty edge is adjacent either to two different sinks or to a sink and to a vertex adjacent to another sink. The idea of the proof of the base case 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.

Octopuses 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 3: 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 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 adjacent 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. Assumption 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 $\tau_1, \ldots, \tau_{p(v)}$ such that for each $l \in \{1, \ldots, p(v)\}$, the tentacles τ_l and $\pi(\tau_l)$ belong to different octopuses (for example order the tentacles according to the number of octopus they belong to and set $\pi(\tau_l)$ to $\tau_{l+\lceil p(v)/2\rceil \bmod p(v)}$; since $m_i(v) \leq p(v)/2$ for all i, tentacles τ_l and $\pi(\tau_l)$ belong to different octopuses). We lengthen the tentacle τ_l through the edge used by the tentacle $\pi(\tau_l)$, so that τ_l now terminates in the vertex 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 2.3, 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 for each of them, the tentacles τ and $\pi(\tau)$ belong to different octopuses. 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 2.3 there exists a flow path passing through v that does not belong to the commodity i, and the minimum cost of the flow $\mathcal{F}$ implies that 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 2.4 *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 Assumption A3. □

Rerouting For each tentacle of the octopus O_i touching the commodity $j \neq i$, we reroute a half unit of the flow of commodity j to t_i along the edges that the corresponding tentacle is stretched through. Observation 2.4 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.

Inductive step Let G and $\mathcal{I}$ be a graph and an instance that do not satisfy the Assumptions A1-A3. We distinguish several cases in the inductive step.

1. If there exists a commodity i with a non-trivial minimum $s - t_i$ cut, we do the inductive step on the number of vertices.

2. If for each commodity every minimum cut is trivial *and* there exists an integral maximum flow with an empty edge that is not adjacent to a sink or is adjacent to exactly one sink and the degree of the sink is higher than h, we do the inductive step on the number of edges.

3. Otherwise we do the inductive step on the number of commodities.

1. 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 that uses C in the flow $\mathcal{F}$ 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). The theorem holds for G' and an 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 way we maintain the same amount of the total flow and we argue as before.

2. 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 second assumption.

Assume first that there exists an empty edge e that is not adjacent to any of the sinks t_i. Since e is not adjacent 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 empty edge e that is adjacent 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.

3. From now on we further assume that every edge unused by the integral minimum cost maximum flow $\mathcal{F}$ (if there are several such flows, choose arbitrarily one of them) is adjacent to one of the terminals. We proceed as in the proof of the base case and at the point where the Assumption A3 is used we observe that it is possible to decrease the number of commodities in the instance $\mathcal{I}$ without decreasing the total flow.

Sharper bound 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 $1/(2(h-1))$ units of flow for each path that it provided to other commodity. Commodities are able to pay these taxes since every commodity had initial flow by at least one greater than the number of tentacles touching it (Observation 2.4) 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 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 the commodity j (and not directly touching the sink t_j) (i.e., the tentacle τ reaching a non-terminal vertex) the flow flows from s to t_i along an original path of 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 τ 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 $l \geq h$ will have a final h-balanced flow at least $l(h/(2(h-1)))$,

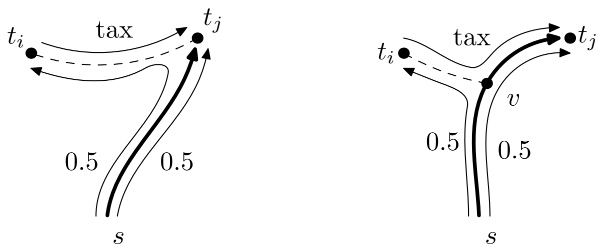

Figure 4: Taxation: on the left side is depicted the case when a tentacle touches a terminal vertex, and on the right side is depicted the case when a tentacle touches a path of other commodity.

which corresponds to an h-route flow of the same size. $\square$

Remark 1 Notice that for every h, a trivial bound $\mathcal{F}^h(\mathcal{I}) \leq \mathcal{F}^1(\mathcal{I})$ holds. Therefore for $h = 2$ the inequality (3) simplifies to

$$\mathcal{F}^1(\mathcal{I}) \quad = \quad \mathcal{F}^2(\mathcal{I}) \ . \tag{4}$$

The equality (4) tells us 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.

Remark 2 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 $k/2$.

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 5). Consider an instance $\mathcal{I}$ with

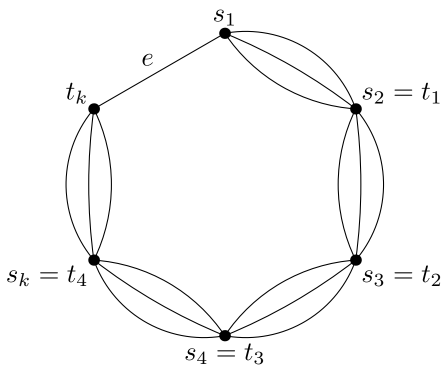

Figure 5: The graph G for $h = 4$ and $k = 5$

$s_i = v_i$ 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 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. □

On the other hand, $\mathcal{F}^1(\mathcal{I}) \leq k\mathcal{F}^h(\mathcal{I})$.

Remark 3 Theorem 2.2 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,*

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

Proof. Choose $k = \lceil \frac{C(h-1)+1}{h} \rceil$ 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 6). Then, for an instance $\mathcal{I} = (s; t_1, \ldots, t_k)$ we have $\mathcal{F}^1(\mathcal{I}) = C(h-1) + 1$ yet $\mathcal{F}^h(\mathcal{I}) = h$. □

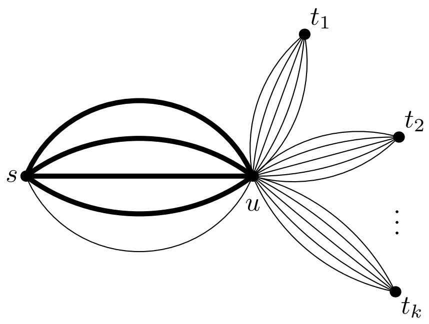

Figure 6: A bad network for nonuniform single source h-route flows (for $h = 5$).

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,*

$$\frac{\mathcal{F}^h(\mathcal{I})}{h} \leq \mathcal{C}^h(\mathcal{I}) \leq (2 - 2/h) \cdot \mathcal{F}^h(\mathcal{I}) \ . \tag{5}$$

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

$$(1 - \epsilon) \cdot \mathcal{F}^h(\mathcal{I}) \leq \mathcal{C}^h(\mathcal{I}) \ , \tag{6}$$

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

$$\frac{\mathcal{F}^h(\mathcal{I})}{h} = \mathcal{C}^h(\mathcal{I}) \ . \tag{7}$$

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 2.2 (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 $\epsilon > 0$ there exists an integer m such that $(m-h+1)/m \geq 1-\epsilon$, and thus, there exists an instance $\mathcal{I} = \{s;t\}$ satisfying the inequality (6). 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 the instance and the network described at the end of the previous section (Figure 6) with every edge capacity set to one. Then, $\mathcal{F}^h(\mathcal{I}) = hk$ and $\mathcal{C}^h(\mathcal{I}) = k$. $\square$

Corollary 3.2 *For every $h \geq 2$, there exists a polynomial time $2(h-1)$-approximation algorithm for the h-disconnecting problem with a single source.*

Remark 4 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 (5) can be slightly improved to

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

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 multi-route flows. The approximation ratio of the algorithm for disconnecting cuts for single source flow problems described in the last section is $2(h-1)$;

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.

Acknowledgement The last author thanks Neal Young, Jiří Sgall and Yefim Dinitz for stimulating discussions.

References

[1] C. C. Aggarwal and J. B. Orlin. On multiroute maximum flows in networks. *Networks*, 39:43–52, 2002.

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

[3] A. Bagchi, A. Chaudhary, and P. Kolman. Short length Menger's theorem and reliable optical routing. *Theoretical Computer Science*, 339:315–332, 2005.

[4] A. Bagchi, A. Chaudhary, P. Kolman, and J. Sgall. A simple combinatorial proof for the duality of multiroute flows and cuts. Technical Report 2004-662, Charles University, Prague, 2004.

[5] G. Baier, E. Koehler, and M. Skutella. The k-splittable flow problem. *Algorithmica*, 42(3-4):231–248, 2005.

[6] A. Baveja and A. Srinivasan. Approximation algorithms for disjoint paths and related routing and packing problems. *Mathematics of Operations Research*, 25(2):255–280, 2000.

[7] A. Chakrabarti, C. Chekuri, A. Gupta, and A. Kumar. Approximation algorithms for the unsplittable flow problem. In *Proceedings of the 5th International Workshop on Approximation Algorithms for Combinatorial Optimization*, pages 51–66, 2002.

[8] C. Chekuri and S. Khanna. Edge disjoint paths revisited. In *Proceedings of the 14th ACM-SIAM Symposium on Discrete Algorithms*, pages 628–637, 2003.

[9] I. Cidon, R. Rom, and Y. Shavitt. Analysis of multi-path routing. *IEEE/ACM Transactions on Networking*, 7(6):885–896, 1999.

[10] Y. Dinitz, N. Garg, and M. Goemans. On the single source unsplittable flow problem. *Combinatorica*, 19:17–41, 1999.

[11] D. Du. *Multiroute Flow Problem*. Ph.D. thesis, The University of Texas at Dallas, 2003.

[12] D. Du and R. Chandrasekaran. The multiroute maximum flow problem revisited. *Networks*, 47(2):81–92, 2006.

[13] T. Erlebach and A. Hall. NP - Hardness of broadcast scheduling and inapproximability of single-source unsplittable min-cost flow. In *Proceedings of the 13th Annual ACM-SIAM Symposium On Discrete Algorithms*, pages 194–202, 2002.

[14] L. R. Ford and D. R. Fulkerson. Maximum flow through a network. *Canad. J. Math.*, 8:399–404, 1956.

[15] N. Garg, V. V. Vazirani, and M. Yannakakis. Approximate max-flow min-cut theorems and their applications. *SIAM Journal on Computing*, 25(2):235–251, 1996.

[16] 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.

[17] K. Kar, M. Kodialam, and T. V. Lakshman. Routing restorable bandwidth guaranteed connections using maximum 2-route flows. *IEEE/ACM Transactions on Networking*, 11:772–781, 2003.

[18] W. Kishimoto. A method for obtaining the maximum multiroute flows in a network. *Networks*, 27:279–291, 1996.

[19] W. Kishimoto and M. Takeuchi. On m-route flows in a network. *IEICE Transactions*, J-76-A(8):1185–1200, 1993. (in Japanese).

[20] J. M. Kleinberg. Single-source unsplittable flow. In *Proceedings of the 37th Annual Symposium on Foundations of Computer Science*, pages 68–77, 1996.

[21] R. Koch, M. Skutella, and I. 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*, pages 244–257, 2005.

[22] S. G. Kolliopoulos and C. Stein. Approximation algorithms for single-source unsplittable flow. *SIAM Journal on Computing*, 31(3):919–946, 2002.

[23] P. Kolman and C. Scheideler. Improved bounds for the unsplittable flow problem. In *Proceedings of the 13th ACM-SIAM Symposium on Discrete Algorithms*, pages 184–193, 2002. To appear in Journal of Algorithms.

[24] P. Kolman and C. Scheideler. Simple on-line algorithms for the maximum disjoint paths problem. *Algorithmica*, 39(3):209–233, 2004.

[25] T. Leighton and S. Rao. Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *Journal of the ACM*, 46(6):787–832, Nov. 1999.

[26] N. Linial, E. London, and Y. Rabinovich. The geometry of graphs and some of its algorithmic applications. *Combinatorica*, 15:215–245, 1995.

[27] M. Martens and M. Skutella. Flows on few paths: Algorithms and lower bounds. *Networks*, 48(2):68–76, 2006.

[28] M. Martens and M. Skutella. Length-bounded and dynamic k-splittable flows. In *Operations Research Proceedings 2005*, pages 297–302, 2006.

[29] K. Menger. Zur allgemeinen Kurventheorie. *Fundam. Math.*, 10:96–115, 1927.

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

[31] A. Srinivasan. Improved approximations for edge-disjoint paths, unsplittable flow, and related routing problems. In *Proceedings of the 38th Annual Symposium on Foundations of Computer Science*, pages 416–425, 20–22 1997.