A Simple Combinatorial Proof of Duality of Multiroute Flows and Cuts

Amitabha Bagchi, Amitabh Chaudhary
Department of Information and Computer Science
University of California
Irvine, CA 92697, USA
{bagchi,amic}@ics.uci.edu

Petr Kolman*
Institute of Theoretical Computer Science
Charles University
Malostranské nám. 25
118 00 Praha 1, Czech Republic
kolman@kam.mff.cuni.cz

Jiří Sgall
Mathematical Institute, AS CR, Žitná 25
115 67 Praha 1, Czech Republic
sgall@math.cas.cz

March 2, 2004

*Research done while visiting University of California at Riverside

Abstract

A classical flow is a nonnegative linear combination of unit flows along simple paths. A multiroute flow, first considered by Kishimoto and Takeuchi, generalizes this concept. The basic building blocks are not single paths with unit flows but rather tuples consisting of k edge disjoint paths, each path with a unit flow. A multiroute flow is a nonnegative linear combination of such tuples.

We present a simple combinatorial proof of the duality theorem for multiroute flows and cuts and its corollary which characterizes multiroute flows in terms of classical flows. Specifically, we show that a (classical) flow of size F is a k-flow if and only if the flow through every edge is at most F/k. This duality then immediately yields an efficient algorithm.

1 Introduction

A classical flow is a nonnegative linear combination of unit flows along simple paths. A multiroute flow, first considered by Kishimoto and Takeuchi [9], generalizes this concept. The basic building blocks are not single paths with unit flows but rather tuples consisting of k edge disjoint paths, each path with a unit flow. Such a structure of k edge disjoint paths (from s to t), each with a unit flow, is called an elementary k-flow. A multiroute flow (or a k-flow) is a nonnegative linear combination of elementary k flows. We postpone for later discussion what a *size* of a cut means for the multiroute flow.

Whereas classical flow has the *flexibility* to split (and merge), k-flow has the *obligation* to split. The motivation for enforced splitting comes from communication applications: a multiroute channel is more tolerant against link failures (cf. [3, 5, 7, 10]).

The celebrated Max-Flow Min-Cut theorem [6] occupies a central position in classical flow theory. Is there an analogue for the multiroute flow? Menger's theorem immediately provides a necessary and sufficient condition for a non-zero multiroute flow: it exists if and only if the source and target vertices are k-edge-connected. However, the question about the maximal multiroute flow is more subtle and the answer is not obvious.

A 2-flow was first introduced by Kishimoto and Takeuchi in the early 90's. Kishimoto and Takeuchi [9] and Kishimoto [8], and Aggarwal and Orlin [1] gave two efficient algorithms for finding a maximal k-flow. Given the maximal k-flow, they both find a cut of the same size [9, 1], which implies the duality of the general multiroute flows and cuts. A similar algorithm for computing the minimal cut of a multiroute flow (and thus also an algorithm for the maximal k-flow) is given by Aneja et al. [2] in the context of parametric analysis of minimal cuts (edge capacities depend to some extend on a varying parameter). A discrete multicommodity variant of the multiroute flow problem, namely the k Disjoint Flow Problem, was studied by Bagchi et al. [4].

In this note, we present a simple combinatorial proof of the duality of multiroute flows and cuts. As a by product of the duality theorem we get a characterization of multiroute flows in terms of classical flows: A (classical) flow of size F is a k-flow if and only if the flow through every edge is at most F/k (see Theorem 3.2). This duality also immediately yields an efficient algorithm for finding the maximum k-flow. Although these results are not new, a direct proof of the duality was not known before. Overall, our proofs are significantly simpler than previously known. We believe that they provide a valuable new insight into the structure of k-flows.

2 Preliminaries

A *network* is an undirected graph $G = (V, E)$ together with a *capacity* function $c : E \to R$ and two distinguished vertices s and t called the *source* and the *sink* respectively.

An *elementary $s - t$ flow* is a flow of one unit from s to t along a simple path. A famous result in classical flow theory says that every $s - t$ flow can be decomposed into (a finite number of) $s - t$ paths and cycles, or in other words, that every acyclic flow is a non-negative linear combination of elementary flows. We take this approach to define to k-flows.

An *elementary k-flow* is an $s - t$ flow along k edge disjoint paths from s to t, each path carrying a unit of flow. A *k-flow* is any flow that is a non-negative linear combination of elementary k-flows [8, 1].

As usual, an *$s - t$ cut* is a partition of vertices V into two groups S and $\bar{S} = V \setminus S$ such that $s \in S$ and $t \in \bar{S}$. Often we will just talk about a cut and we will also often view the cut as a set of edges $\{(u, v) \mid u \in S \wedge v \in \bar{S}\} = \{e_1, \cdots, e_l\}$. The *size* $C(S, V \setminus S)$ of the cut is then equal to $\sum_{i=1}^{l} c(e_i)$. An alternative definition is that the size of the cut $\{e_1, \cdots, e_l\}$ is the size of the maximal flow in a simple network G' where the source s' and sink t' are connected by l edges with respective capacities $c(e_1), \cdots, c(e_l)$. The meaning is that the size of a cut is the maximal amount of flow that can get through if there are no other restrictions.

For multiroute flows, we are going to define the k-size of a cut in an analogous way. The *k-size* of a cut $\{e_1, \cdots, e_l\}$ is the maximal amount of k-flow in a network where the source and destination are connected by l edges with respective capacities $c(e_1), \cdots, c(e_l)$. Note that the k-size of a cut depends only on the capacities of the l edges. In the following, $\gamma_k(c_1, \cdots, c_l)$ will denote the k-size of a cut $\{e_1, \cdots, e_l\}$, $c_i = c(e_i)$. Sometimes we will talk about a k-cut instead of a cut to stress that we are interested in the k-size of the cut.

To give an example, the size of a cut consisting of two edges with capacities 1 and 100 is 101 but the maximal 2-flow that can pass through this cut is only 2. Thus, the 2-size of this cut is two.

Our definition deviates from the original definition of Kishimoto [8] and also from the definition given by Aggarwal and Orlin [1]. However, we think that our definition provides a better insight in the meaning of the k-size of a cut. Moreover, for $k = 1$ the k-size of a cut clearly coincides with the size of the cut. The next lemma shows how to calculate the k-size of a cut and also shows that our definition is equivalent to the previous one, as the previous definitions essentially define the k-size as the value calculated in the lemma.

Lemma 2.1 *Given a cut C with l edges, let $c_1, \cdots, c_l$ denote the edge capacities in nonincreasing order (i.e., $c_i \geq c_{i+1}$) and let $M_j = \frac{k}{k-j} \sum_{i=j+1}^{l} c_i$. Then the k-size of the cut is*

$$\gamma_k(c_1, \cdots, c_l) = \min_{j=0,\ldots,k-1} M_j.$$

Furthermore, if the minimum is achieved at M_j, then $c_i \geq M_j/k \geq c_{i'}$, for any $i \leq j$ and $i' > j$, and

$$\gamma_k(c_1, \cdots, c_l) = \gamma_k(M_j/k, \cdots, M_j/k, c_{j+1}, \ldots, c_l)$$

Proof. Suppose that the minimum is achieved at M_j. Obviously, no k-flow is larger then M_j: in each elementary k-flow, there are at least $k - j$ edges from $c_{j+1}, \ldots, c_l$ and their total capacity allows only a total flow of M_j.

By the minimality of M_j, for any $i \leq j$ and $i' > j$, we have $c_i \geq M_j/k \geq c_{i'}$, as otherwise M_{i-1} or $M_{i'}$ would be smaller than M_j. The last part of the lemma follows by the definition of M_j, as decreasing all c_i, $i < j$, to M_j/k does not change M_j and does not decrease any other $M_{j'}$ below M_j.

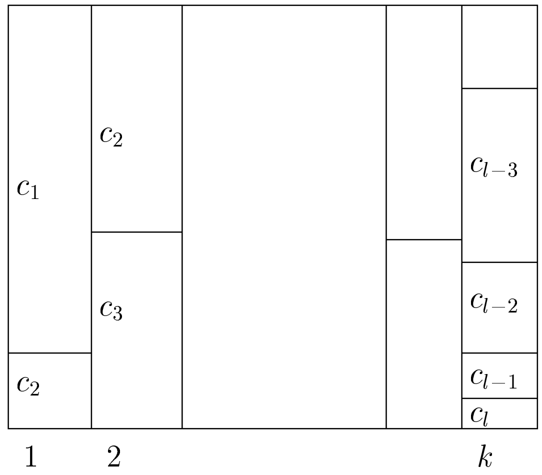

Figure 1: Size of a cut $\{e_1, \cdots, e_l\}$, if $c(e_1) \leq \frac{1}{k} \sum_{i=1}^{l} c(e_i)$.

It remains to construct a k-flow of size M_j. We organize the capacity from the smallest $l-j$ edges in a rectangle $(k-j) \times (M_j/k)$ ($k-j$ columns of height M_j/k). We "fill" the rectangle by the available capacity $c_{j+1} + c_{j+2} + \cdots + c_l$ by columns from left to right, in each column from top to bottom, starting with c_{j+1} in the leftmost column and ending with c_l in the rightmost. The rectangle can now be divided into some number of horizontal slices such that in each slice the portion of capacity in each column comes from a single edge. Moreover, since the edge capacities $c_{j+1}, \cdots, c_l$ are at most M_j/k, in each slice the $k-j$ portions of capacities belong to distinct $k-j$ edges. Thus, each slice corresponds to (a multiple of) an elementary $(k-j)$-flow, of size proportional to the height of the slice. Together, we have a $(k-j)$-flow of size $\sum_{i=j+1}^{l} c_i$ that does not use the j heaviest edges. Adding a flow of M_j/k along the j heaviest edges, we get a k-flow of size M_j. $\square$

3 Duality

We say that an (s,t)-flow of size F is *k-balanced*, if every edge carries at most F/k units of flow. An edge e is *critical* if it carries F/k units of flow. A *k-system* is a set of k edge disjoint path between two vertices s and t; we view a k-system as a set of edges. A technical observation crucially useful in proving the duality theorem is the following lemma.

Lemma 3.1 *For every k-balanced flow f without cycles there exists a k-system that uses all critical edges in f.*

Proof. Let F denote the size of the flow f. Let $G' = (V', E', c')$ be a directed network obtained from the flow f as follows. $V' = V \cup \{s', t'\}$, where s' and t' are new source and target vertices. The set of edges E' contains (i) all edges $e \in E$ with non-zero flow, their capacity $c'(e) = f(e)$, i.e., the value of the flow, and (ii) k edges from s' to s and k edges from t to t', the capacity of each edge e incident with s' or t' is $c'(e) = F/k$.

By the definition of s', t' and edges incident to them, and the assumption that f is acyclic, a set of edges $Q \subseteq E'$ is a k-system between s' and t' in G' if and only if Q contains all k edges from s' to s as well as all k edges from t to t', and at each node $v \in V' - \{s', t'\}$, Q has the same in-degree and out-degree (i.e., satisfies the flow conservation, if all the edges carry the same flow).

Let Q be the k-system in G' that uses the largest number of critical edges (among all possible k-systems). Towards contradiction, assume that there exists a critical edge $(u,v) \notin Q$. Let S be the set of all vertices reachable from v by an "augmenting" path of a special kind. The paths are allowed to use two types of edges: (i) in forward direction, edges not used by Q, and (ii) in backward direction, edges used by Q but only if they are not critical. Formally, the set S consists of all vertices reachable from v in a graph $\bar{G} = (V', E_+ \cup E_-)$ where $E_+ = E' - Q$ and $E_- = \{(y,x) | (x,y) \in Q \text{ and } (x,y) \text{ is not a critical edge}\}$. Obviously, s' and t' are not in S as no edges incident to them are in $E_+ \cup E_-$.

We claim that u is not in S. Assume, towards contradiction, that $u \in S$. Since $(u,v) \in E_+$, by definition of S there is a cycle C in $\bar{G}$ containing (u,v). Add to Q all edges in $C \cap E_+$ and remove all edges $(x,y) \in Q$ such that $(y,x) \in C$. The flow conservation is preserved, thus thus the new set is a k-system. However, the number of critical edges increased: none was removed and at least one, namely (u,v) was added. This is a contradiction with the choice of Q and we conclude that $u \notin S$.

Finally, we get a contradiction by observing that the amount of flow f flowing into the set S is larger that the amount of flow f flowing out of S. Since $s', t' \notin S$, Q enters S exactly the same number of times as it leaves S. By definition of E_-, Q enters S along edges with flow F/k. Therefore the flow into S along edges from Q is at least as large as the flow from S along edges from Q. There is also a flow of size F/k going into S along the edge $(u, v) \notin Q$. By definition of E_+, there is no outgoing flow from S along edges not in Q. Thus, the flow going into S is strictly larger (by at least F/k) than the flow out of S, a contradiction. We conclude that Q contains all critical edges. $\square$

From Lemma 3.1 we almost immediately get the characterization of multiroute flows in terms of classical flows.

Theorem 3.2 *A flow without cycles is k-balanced if and only if it is a k-flow.*

Proof. Every k-flow is k-balanced by definition.

Given a k-balanced flow f, let Q be the k-system that uses all critical edges, obtained by Lemma 3.1. Let c_Q be the minimal flow on edges from Q and let c'_Q be the maximal flow on edges in $E - Q$. In our flow f, we subtract the value $\min\{c_Q, F/k - c'_Q\}$ from each edge in Q. We observe that the remaining flow is balanced again. Since either the number of edges used by the remaining flow decreases or the number of critical edges increases, eventually we distribute all the original flow into k-systems. More precisely, let I be the number of all edges in f plus the number of non-critical edges in f. In each iteration, I decreases by at least one, so in $O(m)$ iterations we find a complete decomposition in k-systems. $\square$

Theorem 3.3 (Duality of k-flows and k-cuts) *The size of the maximal k-flow in G is equal the size of the minimal k-cut in G.*

Proof. The nontrivial part of the theorem is to find in G a cut with k-size equal the maximal k-flow. Let f be the maximal k-flow in G, F be its size and let G' be the network where all capacities larger than F/k are decreased to F/k. Observe that f fits in G'. Moreover, the maximal flow in G' has size F (larger flow would imply there is also a larger k-flow in G). Let $C = \{e_1, \cdots, e_l\}$ be edges in a minimal cut in G', ordered nonincreasingly by their original capacities in G, and let j be the maximal index such that $c_j > F/k$. Then $\frac{k}{k-j} \sum_{i=j+1}^{l} c_i = F$. Considering Lemma 2.1 we conclude that C is a cut of k-size F, and thus a minimal k-cut in G. $\square$

Finally, let us comment on the algorithmic issue. The duality in Theorem 3.3 implies that a regular flow algorithm can be used to find a k-flow of a given size if it exists: knowing there is a k-flow of size F, it is sufficient to restrict the capacity on each edge e to $\min\{c(e), F/k\}$ and then run the regular flow algorithm. Moreover, the Lemma 2.1 implies that for networks with integral capacities, the value of a cut is a rational number $x/y \leq kU$ with $y \leq k$, where U is the maximal capacity of an edge. This implies that the maximal k-flow can be found by binary search, using $O(\log kU)$ runs of maximum flow algorithms, as Kishimoto [8] and Aggarwal and Orlin [1] did.

Open problems

A lot of attention has been paid to max-flow min-cut theorems for multicommodity flows in the last ten years (cf. an excellent survey by Shmoys [11]). This immediately raises the question: Is there an analogous max k-flow min k-cut theorem for multicommodity multiroute flows? Though there is a clear relation between single commodity (classical) flows and multiroute flows, the above question is unanswered so far and the methods used to prove the min-cut max-flow theorems for classical multicommodity flows do not seem to apply easily to multiroute multicommodity flows.

Acknowledgments

J. Sgall was partially supported by Institute for Theoretical Computer Science, Prague (project LN00A056 of MŠMT ČR) and grant A1019401 of GA AV ČR. P. Kolman was partially supported by Institute for Theoretical Computer Science, Prague (project LN00A056 of MŠMT ČR), by the NSF grant CCR-0208856 and by the EU within the 6th Framework Programme under contract 001907 (DELIS).

References

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

[2] Y. P. Aneja, R. Chandrasekaran, and K. P. K. Nair. Parametric analysis of overall min-cuts and applications in undirected networks. *Information Processing Letters*, 85:105–109, 2003.

[3] A. Bagchi, A. Chaudhary, M. T. Goodrich, and S. Xu. Constructing disjoint paths for secure communication. In *Proc. of 17th Intl. Conference on Distributed Computing (DISC '03)*, volume 2848 of *Lecture Notes in Computer Science*, Sorrento, Italy, 2003. Springer.

[4] A. Bagchi, A. Chaudhary, P. Kolman, and C. Scheideler. Algorithms for fault-tolerant routing in circuit switched networks. In *Proceedings of 14th Annual ACM Symposium on Parallel Algorithms and Architectures*, pages 265–274, 2002.

[5] D. Dolev, C. Dwork, O. Waarts, and M. Yung. Perfectly secure message transmission. *JACM*, 40(1):17–47, 1993.

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

[7] 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.

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

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

[10] S.-J. Lee and M. Gerla. Split multipath routing with maximally disjoint paths in ad hoc networks. In *Proc. of IEEE ICC 2001*, pages 3201–3205, Helsinki, Finland, June 2001.

[11] D. B. Shmoys. Cut problems and their application to divide-and-conquer. In D. S. Hochbaum, editor, *Approximation Algorithms for NP-hard Problems*, pages 192–235. PWS Publishing Company, 1997.