Multicommodity Flows Handout

Roy Liu

Winter 2007

1 Background

In the traditional (s, t) flow setting, we want to push as many units of a single resource, or commodity, from the source s to the sink t. Upon generalizing to the case of multiple commodities, we introduce a set of pairs $(s_i, t_i), i \in [n]$, which act as the sources and sinks of n *simultaneous* flows. In other words, all flows use up some shared underlying capacity of the edges in the graph. The higher the demand D_i for commodity i, the higher proportion of it will routed, simultaneously with others, from s_i to t_i.

This problem setting was first explored by Leighton and Rao in [3], where the authors give new definitions to the concepts of **flow** and **cut**. To start off with, let a **flow** f be a function defined over paths in the underlying graph $G = (V, E, c)$. We introduce some definitions:

- Let $\mathcal{P}$ denote the set non-trivial of paths in G.
- Let $\mathcal{P}_{i,j}$ denote the set of paths starting from vertex i and ending at vertex j.
- In an abuse of notation, let $\mathcal{P}_e$ be the set of paths containing the edge e.
- Let $c(e)$ denote the capacity of edge e.
- Let $D_{i,j}$ denote the demand between vertex i and vertex j, for all pairs $i, j \in V$.

We use the notation of [6] to define a simultaneous flow f with *throughput* Z. Such a function satisfies:

$$\sum_{p \in \mathcal{P}_{i,j}} f(p) \geq Z * D_{i,j} \qquad \forall i, j \in V \tag{1}$$

$$\sum_{p \in \mathcal{P}_e} f(p) \leq c(e) \qquad \forall e \in E. \tag{2}$$

We say that f is a maximum flow when it's chosen so as to maximize Z. Intuitively, we can think of the quantity Z to be a guarantee that, for all i, a Z-fraction of commodity i's demand D_i can be simultaneously routed. Notice the inherent notion of fairness using the above measure – we want to make a guarantee for *all* commodities.

Where there's a notion for flow, there's one for **cut**. Each cut $(U, \overline{U})$ has a *ratio cost* R, which is defined to be

$$R = \frac{C(U, \overline{U})}{D(U, \overline{U})},$$

where

$$C(U, \overline{U}) := \sum_{e \in \langle U, \overline{U} \rangle} c(e)$$

$$D(U, \overline{U}) := \sum_{i \,\middle|\, \substack{(s_i \in U \wedge t_i \in \overline{U}) \vee \\ (s_i \in \overline{U} \wedge t_i \in U)}} D_i$$

$$\langle U, \overline{U} \rangle := \left\{ \{u, v\} \in E \,|\, u \in U \wedge v \in \overline{U} \right\}.$$

In other words, the cut ratio R is the capacity going over the cut divided by the demand going over the cut. Naturally, the minimum cut, or *sparsest cut*, minimizes such a ratio.

2 Relationship Between the Max-flow and Min-cut

In their now-famous result [1], Ford and Fulkerson proved that, in the case of a single commodity, the max-flow throughput Z is equal to the min-cut ratio R. Hu [2] then showed that, in the case of two commodities, equality also holds. With any greater number of commodities, however, one can only hope to give conditional statements of equality (true under special graph structure) or upper bound the gap R/Z (easy to see why $Z \leq R$ always). In a simple 5-node example, Seymour and Okamura [5] showed that $Z^* = 3/4$ and $R^* = 1$. See figure 1 for their construction.

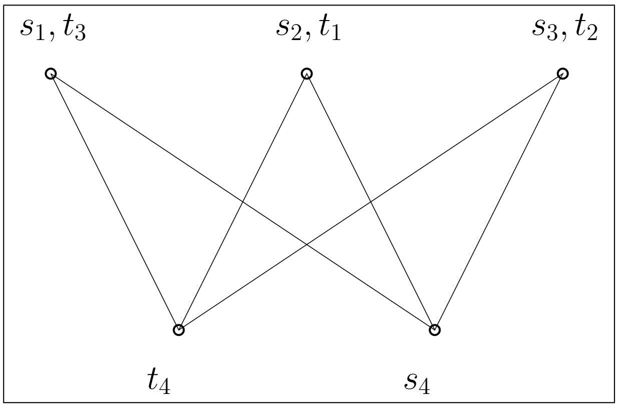

Figure 1: *An example where the max-flow is strictly less than the min-cut. All edges and demands are* 1. *The flows are each* $3/8$.

In the spirit of approximation algorithms, we seek a tight upper bound on the ratio R/Z over an arbitrary demand structure. Leighton and Rao [3] solve a restricted version of the problem with their analysis of uniform multicommodity flows (all demands are 1 over all pairs of vertices). Our analysis will essentially follow the survey chapter of Shmoys [7], which uses the metric embeddings of Linial, London, and Rabinovich [4] as its big idea.

3 The Main Result

We state the main result, which we work towards proving:

Theorem 4
In any graph G with arbitrary edge capacities, the optimal cut ratio R^* and the optimal throughput Z^* satisfy

$$\frac{R^*}{O(\log n)} \leq Z^* \leq R^*,$$

where n is the number of vertices in G.

In what follows, we will prove this result in three parts. The first step is to show that the flow can be solved exactly using an LP formulation, and that the solution is a relaxation of the cut ratio. The next step will look upon the LP variables as a metric over the space of vertices, which will embed into an L_1 metric (with some distortion) by way of [4]. The last step is to show that the optimal IP solution (which describes the optimal sparsest cut), a kind of L_1 metric, is actually optimal over *all* L_1 metrics. Thus, the approximation ratio will be the distortion resulting from the embedding step. More specifically, we want to prove

$$\text{IP}^* \leq \text{LP}^* \cdot O(\log n) \leq \text{IP}^* \cdot O(\log n)$$

4.1 Solving for the Max-flow

The max-flow throughput Z can be solved by the system of linear equations

$$\sum_{p \in \mathcal{P}_{i,j}} f(p) \leq Z * D_{i,j} \qquad \forall i,j \in V$$

$$\sum_{p \in \mathcal{P}_e} f(p) \leq c(e) \qquad \forall e \in E$$

$$\max Z.$$

Unfortunately, there are a possibly exponential number of variables in each constraint, so we take the dual to get an exponential number of constraints instead, and define a separation oracle. The dual is a set of constraints

$$\sum_{\{i,j\} \in E(p)} \mu_{i,j} \geq \mu_{s_p,t_p} \qquad \forall p \in \mathcal{P} \tag{3}$$

$$\sum_{i,j \in V} \mu_{i,j} D_{i,j} = 1 \tag{4}$$

$$\mu_{i,j} \geq 0 \qquad \forall i,j \in V \tag{5}$$

$$\tag{6}$$

$$\min \sum_{\{i,j\} \in E} \mu_{i,j} \cdot c(\{i,j\}), \tag{7}$$

where s_p and t_p are the beginning and end points of the path p, respectively. The form of the dual is one that's much more palatable than the primal, since there's a natural interpretation: The variables can be thought of as distances, where $\mu_{i,j}$ is thought of as the shortest path distance from i to j. Also, observe the normalization in (4), and the fact that the $\mu_{i,j}$'s induce a metric.

There exists a separation oracle for the problem – given a setting for the $\mu_{i,j}$'s, find the all pairs shortest paths distances. If for some i, j of the calculated shortest path distances is less than $\mu_{i,j}$, then that's a violated constraint. Otherwise, all constraints must be satisfied (checking the last oddball constraint (4), of course). Thus, we can optimize over the LP dual.

Beyond being the *optimal* solution for the flow, the LP dual is a *relaxation* for the IP

$$\sum_{\{i,j\} \in E(p)} \mu_{i,j} \geq \mu_{s_p,t_p} \qquad \forall p \in \mathcal{P} \tag{8}$$

$$\mu_{i,j} \in \{0,1\} \qquad \forall i,j \in V \tag{9}$$

$$\min \frac{\sum_{\{i,j\} \in E} \mu_{i,j} \cdot c(\{i,j\})}{\sum_{i,j \in V} \mu_{i,j} D(i,j)}. \tag{10}$$

Not surprisingly, the objective function (10) is the sparsest cut ratio. Once restricted to being either 0 or 1, the variables $\mu_{i,j}$ now induce what is called a *cut metric*. In other words,

$$\mu_{i,j} = 1 \iff |\{i,j\} \cap S| = 1$$

for some underlying cut $(S, \overline{S})$.

4.2 Cut Metrics and L_1 Metrics

Cut metrics are a specific kind of L_1 metric. Let $\mathcal{L}$ be the space of L_1 metrics and $\mathcal{C}$ be the space of cut metrics over V. Clearly, $\mathcal{C} \subseteq \mathcal{L}$, and so it must be the case that

$$\min_{\mu \in \mathcal{L}} \frac{\sum_{\{i,j\} \in E} \mu_{i,j} \cdot c(\{i,j\})}{\sum_{i,j \in V} \mu_{i,j} D_{i,j}} \leq \min_{\mu \in \mathcal{C}} \frac{\sum_{\{i,j\} \in E} \mu_{i,j} \cdot c(\{i,j\})}{\sum_{i,j \in V} \mu_{i,j} D_{i,j}}. \tag{11}$$

We show that the relation is actually much stronger, in that equality holds. The idea of the proof will be to represent every L_1 metric as a linear combination of cut metrics. More formally, we state and prove the following lemma from [7]:

Lemma 5
If d is an L_1 metric over $\mathbb{R}^m$ and gives an IP objective function value of α, then one can find, in polynomial time, a cut with sparsity at most α.

Proof of 5 By assumption, there is some embedding $f : V \to \mathbb{R}^m$ for which

$$d(u,v) = ||f(u) - f(v)||_1$$

for all $u, v \in V$. For further notational convenience, let $d_k(u,v) := |f_k(u) - f_k(v)|$. Let μ_S be the cut metric arising from the cut $(S, \overline{S})$. In other words,

$$\mu_S(u,v) := \begin{cases} 1 & |\{u,v\} \cap S| = 1 \\ 0 & - \end{cases}.$$

In trying to represent d as a linear combination of cut metrics, we show that there exists an assignment of coefficients $\lambda_S \geq 0$ for all $S \subseteq V$ for which

$$d(u,v) = \sum_{S \subseteq V} \lambda_S \cdot \mu_S(u,v) \qquad \forall u, v \in V.$$

To show this, we concern ourselves with the first coordinate only. Without loss, assume the vertices are labeled so that

$$f_1(v_1) \leq f_1(v_2) \leq \ldots \leq f_1(v_n).$$

Let $v_j, v_{j'}$ be given, and consider the case where $j > j'$. Write the contribution of the first coordinate to be a telescoping sum:

$$d_1(v_j, v_{j'}) = f_1(v_j) - f_1(v_{j'}) = \sum_{l=j'}^{j-1} f_1(v_{l+1}) - f_1(v_l).$$

Define $v_{\leq k} := \{v_1, \ldots, v_k\}$. Instead of letting the range of the summation be determined by j and j', we introduce cut metrics into the sum:

$$d_1(v_j, v_{j'}) = f_1(v_j) - f_1(v_{j'}) = \sum_{l=1}^{n-1} (f_1(v_{l+1}) - f_1(v_l)) \cdot \mu_{v_{\leq l}}(v_j, v_{j'}).$$

We have made a seemingly cosmetic change here by using cut metrics *as a filter* to reject out of range indices. See figure 2 for an illustration. Notice that even in the case where $j < j'$, the above is true. Thus,

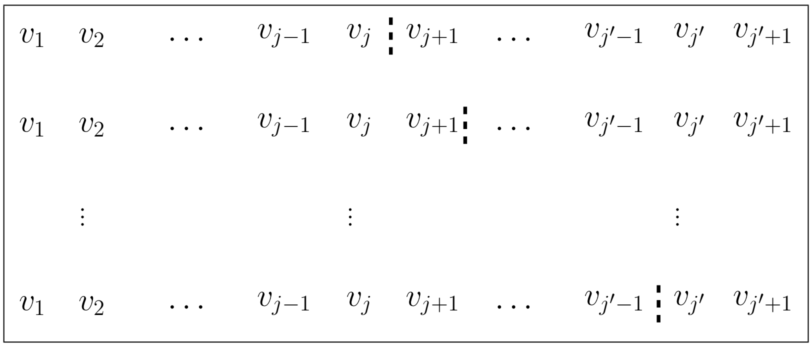

Figure 2: *The telescoping terms that get summed up are precisely those that straddle the $v_{\leq l}$ cuts ($l \in [n]$), denoted by dashed lines.*

if the contribution of the first coordinate is a linear combination of n cut metrics, then the total contribution (sum of all the individual coordinate contributions) is a linear combination of nm cut metrics, and so the set $\mathcal{Z} := \{S \subseteq V \mid \lambda_S > 0\}$ is polynomial in the size of n and m. Summing up the contributions of each coordinate yields

$$d(u, v) = \sum_{S \in \mathcal{Z}} \lambda_S \cdot \mu_S(u, v).$$

To finish, we apply some simple algebraic manipulations to upper bound the value of the best cut by α:

$$\alpha = \frac{\sum_{\{i,j\} \in E} d(i,j) \cdot c(\{i,j\})}{\sum_{i,j \in V} d(i,j) \cdot D(i,j)} = \frac{\sum_{S \in \mathcal{Z}} \lambda_S \sum_{\{i,j\} \in E} c(\{i,j\}) \cdot \mu_S(i,j)}{\sum_{S \in \mathcal{Z}} \lambda_S \sum_{i,j \in V} D(i,j) \cdot \mu_S(i,j)} \tag{12}$$

$$\geq \min_{S \in \mathcal{Z}} \frac{\sum_{\{i,j\} \in E} \mu_S(i,j) \cdot c(\{i,j\})}{\sum_{i,j \in V} \mu_S(i,j) \cdot D(i,j)} \tag{13}$$

The nm vectors we derive in the μ_S's tell us how to find a cut with sparsity at most α – just choose the one that minimizes the sparsity. □

5.1 Embedding into L_1

The last step applies the main theorem of [4] due to Bourgain, which states:

Theorem 6

For any metric space (X, d), one can find a function $f : X \to \mathbb{R}^{O(\log^2 n)}$ that satisfies

$$d(x, y) \leq ||f(x) - f(y)||_1 \leq O(\log n) \cdot d(x, y) \qquad \forall x, y \in X$$

It's not hard to see that the theorem, along with everything developed above, implies the main result.

Let d be the distance variables assigned by the LP, and remember that they induce a metric space (V, d). Embed (V, d) into $\mathbb{R}^{O(\log^2 n)}$ under the L_1 norm with distortion $O(\log n)$ using Theorem 6. Call the resulting metric space $(\widehat{V}, \widehat{d})$. Applying Lemma 5, we get

$$\begin{aligned} \text{IP}^* &\leq \frac{\sum_{\{i,j\} \in E} \widehat{d}(i,j) \cdot c(\{i,j\})}{\sum_{i,j \in \widehat{V}} \widehat{d}(i,j) \cdot D_{i,j}} \leq \frac{\sum_{\{i,j\} \in E} O(\log n) \cdot d(i,j) \cdot c(\{i,j\})}{\sum_{i,j \in V} O(\log n) \cdot d(i,j) \cdot D_{i,j}} \\ &\leq O(\log n) \frac{\sum_{\{i,j\} \in E} d(i,j) \cdot c(\{i,j\})}{\sum_{i,j \in V} d(i,j) \cdot D_{i,j}} \leq O(\log n) \cdot \text{LP}^* \leq O(\log n) \cdot \text{IP}^*, \end{aligned}$$

hence proving Theorem 4:

$$\frac{\text{IP}^*}{O(\log n)} \leq \text{LP}^* \leq \text{IP}^*.$$

6.1 Extra: Bourgain's Embedding is Tight

As it turns out, Bourgain-style embeddings are tight on metric spaces induced by expander graphs. We show this with lemma 7, adapted from [4].

Lemma 7

Any *embedding from the metric space induced by a constant-degree expander graph with unit edge weights into L_1 has distortion $\Omega(\log n)$.*

Proof of 7 Let G be a constant degree expander, and consider the expression

$$\Phi = \frac{\sum_{\{i,j\} \in E(G)} d(i,j)}{\sum_{i,j \in V} d(i,j)}.$$

In the case of G, we have

$$\Phi = O\left(\frac{1}{n \log n}\right).$$

To see why, if G has constant degree k, notice that

$$\frac{\sum_{\{i,j\} \in E(G)} d(i,j)}{\sum_{i,j \in V} d(i,j)} \leq \frac{(k/2)|V|}{\Theta(n^2) \log n} = O\left(\frac{1}{n \log n}\right)$$

Notice that if the expansion constant is upper bounded by k, then each node can reach no more than k^i nodes in i hops. This means that given a vertex v, there exist $\Theta(n)$ nodes with distance at least $\log n$ from it, which means that the total distance between all pairs of nodes is $\Theta(n^2 \log n)$.

The minimum Φ achievable over a cut metric d (defining some cut $(S, \overline{S})$), however, is

$$\Phi = \Omega\left(\frac{1}{n}\right) = \frac{|S|}{|S| \cdot |V \setminus S|} \geq \frac{1}{n}.$$

Since every l_1 metric is a linear combination of cut metrics, then by lemma 5 we have that

$$\begin{aligned}
\Phi &= \frac{\sum_{\{i,j\} \in E(G)} d(i,j)}{\sum_{i,j \in V} d(i,j)} \\
\Phi \sum_{i,j \in V} d(i,j) &= \sum_{\{i,j\} \in E(G)} d(i,j) \\
\Phi \sum_{S \in \mathcal{Z}} \lambda_S \sum_{i,j \in V} \mu_S(i,j) &= \sum_{S \in \mathcal{Z}} \lambda_S \sum_{\{i,j\} \in E(G)} \mu_S(i,j) \\
\Phi \cdot \min_{S \in \mathcal{Z}} \sum_{i,j \in V} \mu_S(i,j) &= \min_{S \in \mathcal{Z}} \sum_{\{i,j\} \in E(G)} \mu_S(i,j) \\
\Phi &\geq \min_{S \in \mathcal{Z}} \frac{\sum_{\{i,j\} \in E(G)} \mu_S(i,j)}{\sum_{i,j \in V} \mu_S(i,j)} \\
\min_{d \in \mathcal{L}} \Phi &= \min_{d \in \mathcal{C}} \Phi,
\end{aligned}$$

which implies that any embedding of a unit edge weight, constant-degree expander graph metric into l_1 will result in distortion $\Omega\left(1/\log n\right)$.

□

References

[1] L.R. Ford and D.R. Fulkerson. *Flows in Networks*. Princeton University Press, Princeton, NJ, 1962.

[2] T.C. Hu. Multicommodity network flows. *Operations Research*, 11(3):344–360, 1963.

[3] 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, 1999.

[4] Nathan Linial, Eran London, and Yuri Rabinovich. The geometry of graphs and some of its algorithmic applications. *Combinatorica*, 15:215–245, 1995.

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

[6] Farhad Shahrokhi and D. W. Matula. The maximum concurrent flow problem. *Journal of the ACM*, 37(2):318–334, 1990.

[7] David Shmoys. Cut problems and their applications to divide-and-conquer, 1996.