SIAM J. Computing 7(1):291–301, 1998.

An $O(\log k)$ approximate min-cut max-flow theorem and approximation algorithm

Yonatan Aumann* Yuval Rabani†

Abstract

It is shown that the minimum cut ratio is within a factor of $O(\log k)$ of the maximum concurrent flow for k-commodity flow instances with arbitrary capacities and demands. This improves upon the previously best known bound of $O(\log^2 k)$ [27], and is existentially tight, up to a constant factor. An algorithm for finding a cut with ratio within a factor of $O(\log k)$ of the maximum concurrent flow, and thus of the optimal min cut ratio, is presented.

*This work was done while at MIT Laboratory for Computer Science, with support from the Wolfson postdoctoral fellowship, and DARPA contract N00014-92-J-1799. Present address: Department of Computer Science, Bar-Ilan University, Ramat-Gan 52900, Israel.

†This work was done while at MIT Lab. for Computer Science, with support from ARPA/Army contract DABT63-93-C-0038. Present address: Computer Science Department, The Technion, Haifa 32000, Israel.

1 Introduction

Multicommodity flow. Consider an undirected graph $G = (V, E)$ with an assignment of non-negative capacities to the edges, $c : E \to \mathbb{R}^+$. A *multicommodity flow* instance on G, is a set of ordered pairs of vertices $(s_1, t_1), (s_2, t_2), \ldots, (s_k, t_k)$. Each pair (s_i, t_i) represents a *commodity*, with *source* at s_i and *destination* or *target* at t_i. The s_i's and t_i's are also called *terminals*. The objective is maximize the amount of *flow* traveling from the sources to the corresponding destinations, subject to the capacity constraints. The problem comes in two flavors. In the first, called the *maximum throughput* problem, the total flow, summed over all commodities, is to be maximized. The second is called the *maximum concurrent flow* problem. Here, for each commodity (s_i, t_i) a non-negative demand D_i is specified. The objective is to maximize the *fraction* of the demand that can be shipped simultaneously for all commodities. A maximum concurrent flow instance is *uniform* if the set of commodities is the set of all ordered pairs of vertices, and all demands are equal. Both the maximum throughput problem and the maximum concurrent flow problem can be solved in polynomial time using linear programming.

Given a multicommodity flow instance (together with demands) the *minimum cut ratio* is defined. A *cut* $(S, \overline{S})$ is a partition of the vertices, with $S \cup \overline{S} = V$ and $S \cap \overline{S} = \emptyset$. The *capacity* of the cut $(S, \overline{S})$ is the sum of capacities of the edges with one endpoint in S and the other in $\overline{S}$. The *cut ratio* is this capacity divided by the sum of demands of commodities with one terminal in S and the other terminal in $\overline{S}$. Finally, the minimum cut ratio, R, is the minimum of cut ratios taken over all cuts $(S, \overline{S})$:

$$R = \min_{S \subset V} \frac{\sum_{e \in E \cap (S \times \overline{S})} c(e)}{\sum_{(s_i,t_i) \in (S \times \overline{S}) \cup (\overline{S} \times S)} D_i} .$$

In general, determining the minimum cut ratio is an NP-hard problem [9]. The maximum concurrent flow is a lower bound on the minimum cut ratio.

For the maximum throughput problem a different notion is useful, that of the *minimum multicut*. A *multicut* is a subset of edges $F \subseteq E$ whose removal disconnects all source-destination pairs. The *capacity* of a multicut F is the sum of capacities of the edges in F. Determining the value of the minimum multicut is an NP-hard and a MAX-SNP-hard problem [11]. The maximum throughput is a lower bound on the capacity of the minimum multicut.

Our results. In this paper, we study the problem of determining the worst case ratio between the minimum cut ratio and the maximum concurrent flow. We establish that for any k commodity instance, the ratio is $O(\log k)$. The bound holds for instances involving arbitrary capacities and demands, thus improving upon the best previous known bound of $O(\log^2 k)$ [27]. We also consider the related question of finding a cut whose ratio approximates the minimum cut ratio. Finding such a cut is a basic step in approximation algorithms for many NP-hard problems (see [20, 16]).

We give an algorithm that produces a cut whose ratio is within an $O(\log k)$ of the maximum concurrent flow, and thus within at most this factor of the minimum cut ratio.

We now give a brief overview of the proof method and the procedure for obtaining the cut. Starting off with the linear programming formulation of the maximum concurrent flow problem, a solution to the dual program is an assignment of points in $\mathbb{R}^k$ to the vertices of the graph. The *value* of the solution is the sum of L_∞ distances spanned by the edges, scaled by the capacities. The constraints of the dual program dictate that the sum of distances between source-destination pairs, scaled by the demands, equals one. Thus, we first solve the (dual) linear program, and obtain an embedding of the graph in $\mathbb{R}^k$. Next, we wish to produce the cut by placing a hyperplane in the space, where the two sides of the hyperplane determine the cut. Suppose that a random hyperplane would have the property that for each pair of vertices x, y, the hyperplane separates x from y with probability proportional to the distance between the two. Then, the expected amount of demands separated by a random hyperplane would be 1 (the sum of source-destination distances scaled by the demands). The expected capacity of the cut, induced by the random hyperplane, would be the sum of edge distances scaled by the capacities, which is exactly the *value* of the solution to the dual, which, in turn, equals the optimal value of the primal, which is the max-flow. Given these expectation, we could then try and find a hyperplane which maintains the expected ratio. Such a cut would have cut ratio equal to the max-flow. Clearly, this procedure fails. This is because in L_∞ norm the necessary distribution over hyperplanes does not necessarily exist. As it turns out, L_1 is the natural norm under which to produce such a distribution. In particular, a random hyperplane perpendicular to one of the axes has exactly the necessary property (i.e. it separates any two nodes with probability proportional to the distance between the two).

With this intuition in mind, we use small distortion embeddings of finite metric spaces into ℓ_1 [3, 21]. Once the (solution) graph is embedded in ℓ_1, we show how to find (deterministically) a good cut. The factor of $O(\log k)$ is lost in the distortion of the embedding. We also show that similar ideas are applicable to bounding the min-multicut max-flow ratio for the maximum throughput problem, though the bounds derived are inferior to the best known bounds. Finally, as a consequence of known instances of graphs with a high min-cut max-flow ratio, we derive tight (up to a constant factor) lower bounds on the best possible distortion of embeddings into ℓ_1 and into ℓ_2, thus improving the lower bound of [3].

Related work. One of the most celebrated theorems in combinatorial optimization is the min-cut max-flow theorem of Ford and Fulkerson, and of Elias, Feinstein, and Shannon [8, 7], which states that for a single commodity, the maximum flow is equal to the minimum cut (here the two flavors of the problem coincide). Early on it has been noted that this is not the case for multicommodity flow.

Early work on multicommodity flow concentrated on characterizing instances where the maximum flow equals the minimum cut ratio. Equality has been established for instances where the *support* (i.e., the graph created by connecting source-destination pairs) is either the union of two stars, or the clique K_4, or the cycle C_5, through the consecutive works of Hu [14], Rothschild and Whinston [28], Dinits (see [1]), Seymour [29], and Lomonosov [23]. Seymour [30] shows that equality holds when the union of the network and the support is planar. Okamura and Seymour [26] establish equality for instances on planar graphs where all terminals reside on the boundary of a single face. For more on this vein of work, see [18].

In their ground-breaking work, Leighton and Rao [20] introduce the notion of an approximate min-cut max-flow theorem. They show that for uniform multicommodity flow, the min cut ratio is within a factor of $O(\log n)$ of the maximum concurrent flow, where n is the number of nodes in the network. Klein, Agrawal, Ravi, and Rao [16] extend their techniques to show an $O(\log C \log D)$ approximate min-cut max-flow theorem for general concurrent flow, where C is the sum of capacities, and D is the sum of demands. This ratio has been improved to $O(\log^2 k)$, where k is the number of commodities, through the works of Tragoudas [32], Garg, Vazirani, and Yannakakis [12], and Plotkin and Tardos [27]. The prevalent method in these papers is the use of graph partitions. The last paper in the sequence overcomes the dependency on the values of the demands by applying a scaling argument and combining flows. We note that while in [20, 16] the primal program considered is the problem of minimizing the capacity utilization (i.e., the factor by which edge capacities need to be multiplied in order to allow all the demand through, see also Shahrokhi and Matula [31]), in [12] the primal program considered is that of maximizing the throughput. This latter view is the one taken in this paper.

Multicuts have been considered by Garg, Vazirani, and Yannakakis [11, 12]. The first paper gives a 2-approximate min-multicut max-flow bound and approximation algorithm for maximum throughput instances on arbitrary capacitated trees. The second paper gives an $O(\log k)$-approximate min-multicut max-flow bound and approximation algorithm for arbitrary maximum throughput instances. Dahlhaus, Johnson, Papadimitriou, Seymour, and Yannakakis [5] consider the problem of *multiway cuts*; i.e., multicuts for maximum throughput instances defined by listing all pairs among a subset of k vertices. They show a $2 - \frac{2}{k}$ approximation algorithm.

Of related interest is the beautiful approximate max-cut algorithm given by Goemans and Williamson [13]. They use semidefinite programming to optimize distances among vertices of a graph. When embedded into the unit ball in $\mathbb{R}^n$ with distances determined by the L_2 norm, the original distances do not shrink too much (but they may have an unbounded increase). It is easy to cut the embedded graph by a hyperplane. Similar ideas have been exploited later by Karger, Motwani, and Sudan to give an improved approximation algorithm for vertex coloring [15]. We note that these techniques seem to fail when it is required to minimize a cut rather than to maximize it.

The relation between ℓ_1-embeddability and multicommodity flow has been noted by Avis and

Deza [2]. They show that Lomonosov's min-cut equals max-flow result (and its predecessors for two commodities etc.) is related directly to the ℓ_1-embeddability of certain fixed metric spaces.

A central tool in our proof is the small distortion embedding of a given metric space into ℓ_1. For this we use a result of Linial, London, and Rabinovich. In a fundamental work, they study algorithmic and other applications of embeddings of graphs into low dimensional normed spaces, introducing techniques from functional analysis. In particular, they consider a theorem of Bourgain [3], which asserts that any n-point metric space can be embedded into ℓ_1 (and also into ℓ_2) with logarithmic distortion. Bourgain's original proof is existential. Linial, London, and Rabinovich present an algorithmic version of the proof which also bounds the dimension by $O(\log^2 n)$ (see also Matoušek [25]). Following their initial work, Linial, London, and Rabinovich have obtained independently similar results to the ones reported here. A full account of their work, which also includes the earlier results, appears in [21]. We note that while the derivation in [21] is very similar to ours, we generate the cut by cutting with a hyperplane, whereas in [21] the cut is generated by a different procedure.

Subsequent to this work, Linial, London, and Rabinovich (in the journal version of their paper [22]), and independently Garg [10], have come up with modified algorithms that can be effectively derandomized.

2 Min-Cut-Ratio Max-Concurrent-Flow

Let $G = (V, E)$ be a graph. Let $c : E \to \mathbb{R}^+$ be an assignment of non-negative capacities to the edges of G. Consider a maximum concurrent flow problem on G, with demand D_i from source s_i to destination t_i, $i = 1, \ldots, k$. For each i, let $\{q_1^i, q_2^i, \ldots\}$ be an enumeration of the paths from s_i to t_i. Let $q_j^i(e)$ denote the characteristic function of the predicate $e \in q_j^i$. Let f_j^i be the amount of commodity i flowing along the path q_j^i, and let f be the minimum fraction of any of the demands which is transmitted in the system in total. We formulate the maximum concurrent flow problem as the following linear program:

$$\begin{array}{lll} \text{maximize} & f & \text{subject to} \\ & D_i f - \sum_j f_j^i \leq 0 & \forall i \in \{1, 2, \ldots, k\} \\ & \sum_{i,j} q_j^i(e) f_j^i \leq c(e) & \forall e \in E \\ & f_j^i \geq 0 & \forall i, j. \end{array} \qquad (P1)$$

The first set of constraints requires that the total flow of commodity i be at least an f fraction of the demand D_i. The second set of constraints prevents the edge capacities from being violated.

The dual is:

$$\begin{array}{llll} \text{minimize} & \sum_e c(e)d(e) & \text{subject to} & \\ & \sum_i D_i h_i = 1 & & \\ & \sum_e q_j^i(e)d(e) - h_i \geq 0 & \forall i,j & (D1) \\ & h_i \geq 0 & \forall i & \\ & d(e) \geq 0 & \forall e. & \end{array}$$

The dual has a pictorial physical interpretation. Imagine that the edges of the graph represent a network of pipes. The "cross-section" area of pipe e is $c(e)$, and $d(e)$ denotes the length of the pipe. Thus, the objective is to minimize the total volume of the network. The variable h_i corresponds to the distance between the terminals of commodity i (i.e. the minimum length of a path q_j^i). Now suppose that the network supports a flow such that for each commodity i, the rate at which this commodity is transmitted from source to destination is an f fraction the demand of the commodity. Then, at any given time, the total *volume* of flow in the network is at least $\sum_i fD_ih_i$. Clearly, the volume of flow cannot be greater than the total volume of the pipes, which is $\sum_e c(e)d(e)$. The first constraint in the dual scales $\sum_i D_ih_i = 1$. Thus we have have $f \leq \sum_e c(e)d(e)$, which is precisely the statement of weak duality in this case.

The program $(D1)$ may be of exponential size. Using the above interpretation, we convert it to an equivalent program of polynomial size. The purpose of this step is to allow the solution of the linear program in polynomial time. We consider the vertices of G as points in $\mathbb{R}^k$ with L_∞ norm. The distances between the vertices are the distances assigned to the corresponding edges. For each $i = 1, \ldots k$, the i-th coordinate of the location of the a vertex represents the distances of the vertex from the source s_i. Thus, we obtain the following equivalent, polynomial size, linear program.

$$\begin{array}{llll} \text{minimize} & \sum_e c(e)d(e) & \text{subject to} & \\ & d(e) \;\geq\; u_i^x - u_i^y & \multirow{2}{*}{$\forall e=(x,y),\ \forall i \in \{1,2,\ldots,k\}$} & \\ & \quad\;\;\;\; \geq\; u_i^y - u_i^x & & (D') \\ & \sum_i D_i h_i = 1 & & \\ & h_i \leq u_i^{t_i} - u_i^{s_i} & \forall i & \\ & u_i^{t_i} \geq 0,\ u_i^{s_i} \leq 0 & \forall i. & \end{array}$$

Here, vertex x is mapped to the point $u^x = (u_1^x, \ldots, u_k^x) \in \mathbb{R}^k$.

Fact 1 *The optimal solutions to $D1$ and D' are equal.*

Proof. Consider a feasible solution to D'. The same $d(e)$'s and h_i's also constitute a feasible solution for $D1$. To see this we only have to show that $\sum_e q_j^i(e)d(e) \geq h_i$ for all i, j. Indeed

$$\sum_e q_j^i(e)d(i) \geq \sum_{e=(x,y)} q_j^i(e)\,|u_i^x - u_i^y| = u_i^{t_i} - u_i^{s_i} \geq h_i\ .$$

The equality holds because q_j^i is a path from s_i to t_i and thus the summation is a telescopic one. Thus, any feasible solution to D' is also a solution to $D1$. Conversely, given a feasible solution to $D1$, for all vertices x and $i = 1, \dots, k$, set $u_i^x = \text{dist}(s_i, x)$, where $\text{dist}(y, x)$ denotes the length of the shortest path from y to x, under the non-negative edge lengths $d(e)$. This assignment gives D' the same objective function value as for $D1$, while obeying all the constraints. ■

A feasible solution to D' induces a mapping of the vertices into the metric space ℓ_∞^k ($\mathbb{R}^k$ with L_∞ norm). The L_∞ distances among these points have the following properties:

1. $\forall i$, $\|u^{t_i} - u^{s_i}\|_\infty \geq h_i$; and
2. $\forall e = (x, y)$, $\|u^x - u^y\|_\infty \leq d(e)$.

Next we embed this n point subspace of ℓ_∞^k into ℓ_1^d ($\mathbb{R}^d$ with L_1 norm), with u^x mapped to $\tilde{u}^x$. Let $\tilde{h}_i$ denote $\|\tilde{u}^{s_i} - \tilde{u}^{t_i}\|_1$. The embedding has the following properties:

1. d is polynomial in the size of the input (in fact, the embedding will have $d = O(\log^2 k)$);
2. $\forall i$, $\tilde{h}_i \geq h_i / O(\log k)$;
3. $\forall e = (x, y)$, $\|\tilde{u}^x - \tilde{u}^y\|_1 \leq \|u^x - u^y\|_\infty$.

Such an embedding can be obtained by a slight variation of Bourgain's theorem [3], and the algorithmic version thereof by Linial, London, and Rabinovich. A sketch of the proof is provided in the appendix (Corollary 4, see also [21]).

For every pair of vertices x, y define the function $\delta^{x,y} : \{1, \dots, d\} \times \mathbb{R} \to \{0, 1\}$ as

$$\delta^{x,y}(j, \xi) = \begin{cases} 1 & \text{if } \min\{\tilde{u}_j^x, \tilde{u}_j^y\} \leq \xi < \max\{\tilde{u}_j^x, \tilde{u}_j^y\}; \\ 0 & \text{otherwise.} \end{cases}$$

The function $\delta^{x,y}$ is the characteristic function of the projection, on the j-th axis, of the straight line connecting $\tilde{u}^x$ and $\tilde{u}^y$.

Now, define two functions:

$$\begin{aligned} H &: \{1, 2, \dots, d\} \times \mathbb{R} \to \mathbb{R}^+; \\ C &: \{1, 2, \dots, d\} \times \mathbb{R} \to \mathbb{R}^+, \end{aligned}$$

as follows:

$$\begin{aligned} H(j, \xi) &= \sum_{i=1}^{k} D_i \delta^{s_i, t_i}(j, \xi); \\ C(j, \xi) &= \sum_{e=(x,y) \in E} c(e) \delta^{x,y}(j, \xi). \end{aligned}$$

Notice that $H(j,\xi)$ is the sum of demands cut by a hyperplane perpendicular to the jth axis at ξ, while $C(j,\xi)$ is the sum of capacities of edges cut by the same hyperplane.

We have that

$$\sum_{j=1}^{d}\int_{-\infty}^{\infty} H(j,\xi)d\xi = \sum_{j=1}^{d}\sum_{i=1}^{k}\int_{-\infty}^{\infty} D_i\delta^{s_i,t_i}(j,\xi)d\xi = \sum_{i=1}^{k} D_i\tilde{h}_i \geq \frac{1}{O(\log k)} \ . \tag{1}$$

Similarly,

$$\sum_{j=1}^{d}\int_{-\infty}^{\infty} C(j,\xi)d\xi \leq \sum_{e\in E} c(e)d(e). \tag{2}$$

Therefore, since both functions are non-negative, there exists a point (j_0,ξ_0) such that $H(j_0,\xi_0) > 0$ and

$$\frac{C(j_0,\xi_0)}{H(j_0,\xi_0)} \leq O(\log k)\sum_{e\in E} c(e)d(e).$$

To see this, consider the functions $H'(j,\xi) = O(\log k)H(j,\xi)$ and $C'(j,\xi) = C(j,\xi)/\sum_{e\in E} c(e)d(e)$. H' is not always 0 and C' is non-negative. If $H'(j,\xi) < C'(j,\xi)$ whenever $H'(j,\xi) > 0$, then we must have that $\sum_{j=1}^{d}\int_{-\infty}^{\infty} H'(j,\xi) < \sum_{j=1}^{d}\int_{-\infty}^{\infty} C'(j,\xi)$, in contradiction to inequalities (1) and (2).

The point (j_0,ξ_0) determines a cut $(S,\overline{S})$ by setting $S = \left\{x \in V \mid \tilde{u}^x_{j_0} \leq \xi_0\right\}$. The amount of demand between S and $\overline{S}$ is exactly $H(j_0,\xi_0)$ and the capacity of the edges connecting the two sides of the cut is $C(j_0,\xi_0)$. The point (j_0,ξ_0) can be found in polynomial time. This is because there are at most $|V|$ points of interest to check in each dimension. A sweep in each dimension can be used to obtain the best cut.

We have obtained:

Theorem 2 *For every concurrent multicommodity flow instance involving k commodities, the minimum cut ratio is within a factor of $O(\log k)$ of the maximum concurrent flow. Finding such a cut can be done in random polynomial time.*

Proof. Obtain the optimal solution to D'; e.g., using Ye's interior point polynomial time algorithm [34]. (Alternatively, find a near-optimal solution using more efficient algorithms, e.g., Leighton et al [19].) By Fact 1 and linear programming duality, the value m of this solution is equal to the optimal value of $P1$. The above discussion shows how to find a cut with ratio within $O(\log k)$ of m. ■

3 Min-Multicut Max-Throughput

We can apply our method to derive bounds on the min-multicut to maximum-throughput ratio. However, the bounds we derive are inferior to the $O(\log k)$ bound presented in [12]. We get an

$O(\log^2 k)$ bound using a randomized algorithm to find an approximate cut. The algorithm in [12] is deterministic. A sketch the proof follows.

Using the notation of the previous section, we formulate the maximum throughput problem as the following linear program:

$$\begin{array}{llll} \text{maximize} & \sum_{i,j} f_j^i & \text{subject to} & \\ & \sum_{i,j} q_j^i(e) f_j^i \leq c(e) & \forall e \in E & (P2) \\ & f_j^i \geq 0 & \forall i, j. & \end{array}$$

The interpretation of this program is similar to that of $(P1)$. The variables f_j^i denote the flow due to commodity i along the path q_j^i. The constraints guarantee that none of the edge capacities are violated.

The dual is:

$$\begin{array}{llll} \text{minimize} & \sum_{e} c(e) d(e) & \text{subject to} & \\ & \sum_{e} q_j^i(e) d(e) \geq 1 & \forall i, j & (D2) \\ & d(e) \geq 0 & \forall e. & \end{array}$$

This program can be interpreted similarly to $D1$. We have a similar network of pipes. The constraints require that the distance between the pair of terminals of any commodity is at least 1. To accommodate a flow of rate $\sum_j f_j^i$ for every commodity i, the total volume of the pipes must be at least $\sum_{i,j} f_j^i$. This demonstrates weak duality.

As above, the dual (D2) can be converted into a polynomial size program where vertices are mapped to points in ℓ_∞^k. Note that in a feasible solution to $D2$ the distance between *each* source-destination pair is at least 1. We embed the metric space induced by the optimal solution into $\ell_1^{O(\log^2 k)}$, shrinking all distances and maintaining a distance of at least $1/O(\log k)$ among the terminal pairs. For the optimal solution, the embedded graph is contained in a cube of unit side length. Through the rest of this section we use the same notation as in the previous section.

Consider the following process. For each dimension j, independently choose a point $\xi_j \in [0, 1]$, uniformly at random. Place a hyperplane intersecting the j-th axis at ξ_j, and perpendicular to this axis. Consider a pair of vertices x, y. Let $a_j^{x,y}$ denote the length of the projection, onto the j-th axis, of the straight line connecting $\tilde{u}^x$ and $\tilde{u}^y$. Let $p^{x,y}$ be the probability that $\tilde{u}^x$ and $\tilde{u}^y$ are separated by at least one of the $d = O(\log^2 k)$ random hyperplanes. We have that

$$p^{x,y} \leq \sum_{j=1}^{d} a_j^{x,y} = \|\tilde{u}^x - \tilde{u}^y\|_1 \leq \|u^x - u^y\|_\infty,$$

and

$$p^{x,y} \geq 1 - \prod_{j=1}^{d} \left(1 - a_j^{x,y}\right) \geq 1 - \left(1 - \frac{\|\tilde{u}^x - \tilde{u}^y\|_1}{d}\right)^d.$$

For s_i, t_i, we have $\|u^{s_i} - u^{t_i}\|_1 \geq 1/O(\log k)$ for all i. Thus, if we repeat the placement of d random hyperplanes $c \log^2 k = O(\log^2 k)$ times, then the probability, p_i, that $\tilde{u}^{s_i}$ and $\tilde{u}^{t_i}$ are separated by any of these hyperplanes satisfies

$$p_i \geq 1 - \left(1 - \frac{1}{d \cdot O(\log k)}\right)^{cd \cdot O(\log^2 k)} \geq 1 - k^{-c \cdot O(1)}.$$

Thus, with c sufficiently large, with probability at least $1 - k^{-1} \geq 2/3$ (assuming $k \geq 3$, for $k = 1$ or $k = 2$ the min-multicut equals the max-throughput), all source-destination pairs are separated.

For any edge $e = (x, y)$, $\|u^x - u^y\|_\infty \leq d(e)$. Thus, the probability, p_e, that e is cut by any of the hyperplanes satisfies $p_e \leq O(\log^2 k) d(e)$. Therefore, the expected size of the multicut is $O(\log^2 k) \sum_e c(e) d(e)$. By Markov's inequality, the probability that we get a multicut whose size is more than twice the expectation is at most $1/2$. Therefore, with probability at least $1/6$ all commodities are cut and the size of the cut is no more than twice the expectation. This guarantees that in random polynomial time we get a multicut that separates all source-destination pairs, and whose size is $O(\log^2 k) \sum_e c(e) d(e)$.

References

[1] G.M. Adelson-Welsky, E.A. Dinits, and A.V. Karzanov. *Flow Algorithms* (in Russian). Nauka, Moscow, 1975.

[2] D. Avis and M. Deza. The cut cone, L^1-embeddability, complexity and multicommodity flows. *Networks*, 21:595–617, 1991.

[3] J. Bourgain. On Lipschitz embedding of finite metric spaces in Hilbert space. *Israel Journal of Mathematics*, 52:46–52, 1985.

[4] J. Bretagnolle, D. Dacunha-Castelle, and J.L. Krivine. Lois stables et espaces L^p. *Annales de l'Institut Henri Poincaré*, 2:231–259, 1966.

[5] E. Dahlhaus, D.S. Johnson, C.H. Papadimitriou, P.D. Seymour, and M. Yannakakis. The complexity of multiway cuts. In *Proceedings of the 24th Annual ACM Symposium on Theory of Computing*, pages 241–251, 1992.

[6] M. Deza and M. Laurent. Applications of cut polyhedra. Technical report BS-R9221, CWI, Amsterdam, 1992.

[7] P. Elias, A. Feinstein, and C.E. Shannon. A note on the maximum flow through a network. *IRS Transactions on Information Theory*, 2:117–119, 1956.

[8] L.R. Ford and D.R. Fulkerson. Maximal flow through a network. *Canadian Journal of Mathematics*, 8:399–404, 1956.

[9] M.R. Garey and D.S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*, Freeman, San Francisco, 1979.

[10] N. Garg. A deterministic $O(\log k)$-approximation algorithm for sparsest-cut. Preprint, 1995.

[11] N. Garg, V.V. Vazirani, and M. Yannakakis. Primal-dual approximation algorithms for integral flow and multicut in trees, with applications to matching and set cover. In *Proceedings of ICALP '93*, pages 64–75.

[12] N. Garg, V.V. Vazirani, and M. Yannakakis. Approximate max-flow min-(multi)cut theorems and their applications. In *Proceedings of the 25th Annual ACM Symposium on Theory of Computing*, pages 698–707, 1993.

[13] M.X. Goemans and D.P. Williamson. .878-approximation algorithms for MAX CUT and MAX 2SAT. In *Proceedings of the 26th Annual ACM Symposium on Theory of Computing*, pages 422–431, 1994.

[14] T.C. Hu. Multicommodity network flows. *Operations Research*, 11:344-360, 1963.

[15] D. Karger, R. Motwani, and M. Sudan. Approximate graph coloring by semidefinite programming. In *Proceedings of the 35th Annual IEEE Symposium on Foundations of Computer Science*, pages 2–13, 1994.

[16] P. Klein, A. Agrawal, R. Ravi, and S. Rao. Approximation through multicommodity flow. In *Proceedings of the 31st Annual IEEE Symposium on Foundations of Computer Science*, pages 726–737, 1990.

[17] P. Klein, S. Plotkin, and S. Rao. Excluded minors, network decomposition, and multicommodity flow. In *Proceedings of the 25th Annual ACM Symposium on Theory of Computing*, pages 682–690, 1993.

[18] B. Korte, L. Lovász, H.J. Promel, and A. Schrijver, editors. *Paths, flows, and VLSI-layout*. Springer-Verlag. 1990.

[19] F.T. Leighton, F. Makedon, S. Plotkin, C. Stein, É. Tardos, and S. Tragoudas. Fast approximation algorithms for multicommodity flow problems. *Journal of Computer and Systems Sciences*, to appear. A preliminary version appeared in *Proceedings of the 23rd Annual ACM Symposium on Theory of Computing*, pages 101–111, 1991.

[20] F.T. Leighton and S. Rao. An approximate max-flow min-cut theorem for uniform multicommodity flow problems with applications to approximation algorithms. In *Proceedings of the 29th Annual IEEE Symposium on Foundations of Computer Science*, pages 422–431, 1988.

[21] N. Linial, E. London, and Y. Rabinovich. The geometry of graphs and some of its algorithmic applications. In *Proceedings of the 35th Annual IEEE Symposium on Foundations of Computer Science*, pages 577-591, 1994.

[22] N. Linial, E. London, and Y. Rabinovich. The geometry of graphs and some of its algorithmic applications. *Combinatorica*, 15(2):215–245, 1995.

[23] M. Lomonosov. Combinatorial approaches to multiflow problems. *Discrete Applied Mathematics*, 11:1–93, 1985.

[24] G.A. Margulis. Explicit constructions of concentrators. *Problems of Information Transformation*, 9:325–332, 1973.

[25] J. Matoušek. Note on bi-Lipschitz embeddings into normed spaces. *Commentationes Mathematicae Univ. Carolinae*, 33(1):51–55, 1992.

[26] H. Okamura and P.D. Seymour. Multicommodity flows in planar graphs. *Journal of Combinatorial Theory (B)*, 31:75–81, 1989.

[27] S. Plotkin and É. Tardos. Improved bounds on the max-flow min-cut ratio for multicommodity flows. In *Proceedings of the 25th Annual ACM Symposium on Theory of Computing*, pages 691–697, 1993.

[28] B. Rothschild and A. Whinston. On two-commodity network flows. *Operations Research*, 14:377–387, 1966.

[29] P.D. Seymour. Four terminous flows. *Networks*, 10:79–86, 1980.

[30] P.D. Seymour. Matroids and multicommodity flows. *European Journal of Combinatorics*, 2:257–290, 1981.

[31] F. Shahrokhi and D.W. Matula. The maximum concurrent flow problem. *Journal of the ACM*, 37:318–334, 1990.

[32] S. Tragoudas. *VLSI partitioning approximation algorithms based on multicommodity flow and other techniques*. Ph.D. thesis, University of Texas at Dallas, 1991.

[33] J.H. Wells and L.R. Williams. *Embeddings and Extentions in Analysis*. Springer-Verlag, 1975.

[34] Y. Ye. An $O(n^3L)$ potential reduction algorithm for linear programming. *Mathematical Programming*, 50:239–258, 1991.

Appendix: Embeddings of Finite Metric Spaces

An *embedding* of a finite metric space $\mathcal{M} = (X, d)$ into a (larger) metric space $\mathcal{M}' = (X', d')$ is a 1—1 mapping $\varphi : X \to X'$. An embedding φ is a *contraction* if $\forall x, y \in X$, $d'(\varphi(x), \varphi(y)) \leq d(x, y)$ (see [33]). Any embedding into a normed space can be converted into a contraction without changing the ratios of distances among the points, by scaling all distances. For a contraction φ, the *distortion* is $\max_{x,y\in X} \{d(x, y)/d'(\varphi(x), \varphi(y))\}$. Obviously, the distortion is at least 1. A contraction is *isometric* if it has distortion exactly 1. For any p, a finite metric space is ℓ_p-*embeddable* if it can be isometrically embedded into the r dimensional normed space ℓ_p^r, for some r. A finite metric space is *Euclidean* if it is ℓ_2-embeddable. We say that φ is into ℓ_p, if $\mathcal{M}'$ is ℓ_p^r, for some r. For $x \in X$, $Y \subset X$, denote by $d(x, Y)$ the distance from x to Y, defined as $\min_{y\in Y} d(x, y)$. For $x \in X$, $\rho \geq 0$, denote $B(x, \rho) = \{z \mid d(x, z) \leq \rho\}$ (the closed ball around x of radius ρ).

The following lemma is due to Linial, London, and Rabinovich [21, 22] (see also [25]). It gives an algorithmic version of a theorem of Bourgain [3]. For completeness, we give a sketch of the proof.

Lemma 3 ([22]) *Let $\mathcal{M} = (X, d)$ be a finite metric space with $|X| = n$. There exists a contraction, φ, of $\mathcal{M}$ into $\ell_1^{O(\log^2 n)}$ that has distortion $O(\log n)$. The contraction can be constructed in random polynomial time.*

Proof. For every $t = 1, \ldots, \log n - 1$, choose $L = O(\log n)$ subsets $Q_{t,j} \subset X$, $j = 1, \ldots, L$, such that $|Q_{t,j}| = \frac{n}{2^t}$. Each subset is chosen uniformly at random among all subsets of the specified size. The choices are mutually independent. For each $x \in X$, define a vector $\phi(x)$ with entries $\phi_{t,j}(x) = d(x, Q_{t,j})$.

Consider two points $x, y \in X$. For $t = 0, 1, \ldots$ define $\rho_t(x) = \min \rho$ such that $|B(x\rho)| \geq 2^t$, and similarly define $\rho_t(y)$. Let $\rho_t = \max\{\rho_t(x), \rho_t(y)\}$, and let $\hat{t}$ be the largest t such that $\rho_t \leq d(x, y)/3$. Set $\rho_{\hat{t}+1} = d(x, y)/3$. (Notice that $\rho_0 = 0$.) Consider a particular $t > 0$, and one of the corresponding sets $Q_{t,j}$. W.l.o.g. $\rho_t(x) = \rho_t$ (i.e. ρ_t is obtained around x). Then, by assumption $|B(x, \rho_t)| \leq 2^t$, and by definition $|B(y, \rho_{t-1})| \geq 2^{t-1}$. Thus, a random set $Q_{t,j}$ of size $n/2^t$, has a constant probability of intersecting $B(y, \rho_{t-1})$ and not intersecting $B(x, \rho_t)$. (Notice that these balls are disjoint.) Thus, with constant probability $|\phi_{t,j}(x) - \phi_{t,j}(y)| \geq \rho_t - \rho_{t-1}$. Thus, with L sufficiently large, with high probability, $\sum_{j=1}^{L} |\phi_{t,j}(x) - \phi_{t,j}(y)| \geq cL(\rho_t - \rho_{t-1})$, for some constant c. Thus,

$$\|\phi(x) - \phi(y)\|_1 = \sum_{t,j} |\phi_{t,j}(x) - \phi_{t,j}(y)| \geq \sum_{t=1}^{\hat{t}+1} cL(\rho_t - \rho_{t-1}) = cL\rho_{\hat{t}+1} \geq cL\frac{d(x, y)}{3}.$$

On the other hand, by the triangle inequality, for any $x, y \in X$, and any set $Q_{t,j}$, $|d(x, Q_{t,j}) -$

$d(y, Q_{t,j})| \leq d(x, y)$. Thus,

$$\|\phi(x) - \phi(y)\|_1 = \sum_{t,j} |\phi_{t,j}(x) - \phi_{t,j}(y)| \leq \log n \cdot L \cdot d(x, y) \ .$$

Finally, define $\varphi(x) = \frac{1}{L \log n}\phi(x)$. ■

We get the following corollary:

Corollary 4 *Let $\mathcal{M} = (X, d)$ be a finite metric space. Let $\mathcal{N} = (Y, d)$ be a subspace of $\mathcal{M}$ induced by a subset of points of cardinality k. There exists a contraction, φ, of $\mathcal{M}$ into $\ell_1^{O(\log^2 k)}$ such that the restriction of φ to $\mathcal{N}$ has distortion $O(\log k)$. The contraction can be constructed in random polynomial time.*

Proof. Using the construction for $\mathcal{N}$ given by the proof of Lemma 3, we get the subsets $Q_{t,j} \subset Y$, $t = 1, \dots, \log k - 1$, $j = 1, \dots, L$, $L = O(\log k)$. For each $x \in X$, set $\phi_{t,j}(x) = d(x, Q_{t,j})$, and define $\varphi(x) = \frac{1}{L \log k}\phi(x)$. Now, for any $x, y \in X$, $\|\varphi(x) - \varphi(y)\|_1 \leq d(x, y)$, and for all $x, y \in Y$, $\|\varphi(x) - \varphi(y)\|_1 \geq d(x, y)/c \log k$, for some constant c. ■

Theorem 5 *There are (explicitly constructible) metric spaces $\mathcal{M}_n$ over n points, such that any contraction of $\mathcal{M}_n$ into ℓ_1 has distortion $\Omega(\log n)$.*

Proof. By contradiction. Consider n node bounded degree expander graphs G_n. (The expansion property we need is that any subset of $s \leq n/2$ vertices is connected to at least αs other vertices, for some absolute constant $\alpha > 0$. Such graphs can be constructed, e.g., via [24].) There are $\Theta(n)$ edges in G_n, and there exist $k_n = \Theta(n^2)$ pairs of vertices at distance $\Theta(\log n)$ apart. For each n, consider a multicommodity flow instance where all the edges of G_n have unit capacity, and there is a unit demand between each of the k_n pairs. It is not difficult to see that any cut in G_n has cut ratio at least c/n, for some constant c. Consider a feasible solution to the dual problem, where $d(e)$ is set to $d_n = \Theta(1/k_n \log n)$ for all e, with an appropriate choice of constant. The value of the dual for this solution is at most $c'/n \log n$, for some constant c'. Now, consider the path metric of G_n (i.e., the n-point metric space over vertices of G_n, where the distance between two vertices is the minimum length path connecting them). Suppose this metric space can be embedded into ℓ_1 with distortion less than $c \log k/c'$. Using the arguments from Section 2, this implies that there exists a cut in G_n with cut ratio less than c/n, a contradiction. ■

Remark. An alternative proof can be derived by noticing that for multicuts one can disconnect a constant fraction of the commodities by a multicut whose size is within an order of the distortion factor of the optimal fractional solution. The result then follows since there are instances of multicommodity flow on expander graphs where cutting a constant fraction of the commodities requires a cut whose size is $\Omega(\log k)$ times the value of the maximum throughput (see [12]).

Bourgain [3] gives a lower bound of $\Omega(\log n/ \log\log n)$ for the worst case distortion of any contraction of n-point metric spaces into ℓ_2. From Theorem 5 we get the following improvement over Bourgain's lower bound:

Corollary 6 *Any contraction of* $\mathcal{M}_n$ *into* ℓ_2 *has distortion* $\Omega(\log n)$.

Proof. Any finite Euclidean space is ℓ_1-embeddable. (It is well known that L_2 isometrically embeds in L_1, see e.g., [4]. For finite metric spaces, L_1 embeddability implies ℓ_1 embeddability, see [6].) ■

Remark. The question of improving Bourgain's lower bound was raised by Linial, London, and Rabinovich. A similar theorem, derived independently, is shown in [21, 22].