

Rounding Algorithms for a Geometric Embedding of Minimum Multiway Cut

David R. Karger* Philip Klein† Cliff Stein‡ Mikkel Thorup§ Neal E. Young¶

Abstract

Given an undirected graph with edge costs and a subset of $k \geq 3$ nodes called *terminals*, a multiway, or k -way, cut is a subset of the edges whose removal disconnects each terminal from the others. The multiway cut problem is to find a minimum-cost multiway cut. This problem is Max-SNP hard. Recently Calinescu, Karloff, and Rabani (STOC'98) gave a novel geometric relaxation of the problem and a rounding scheme that produced a $(3/2 - 1/k)$ -approximation algorithm.

In this paper, we study their geometric relaxation. In particular, we study the worst-case ratio between the value of the relaxation and the value of the minimum multicut (the so-called integrality gap of the relaxation). For $k = 3$, we show the integrality gap is $12/11$, giving tight upper and lower bounds. That is, we exhibit a graph class with integrality gap $12/11 - o(1)$ and give an algorithm that finds a cut of value $12/11$ times the relaxation value. This is the best possible performance guarantee for any algorithm based purely on the value of the relaxation and improves on Calinescu et al.'s factor of $7/6$.

We also improve the upper bounds for all larger values of k . For $k = 4, 5$, our best upper bounds are based on computer constructed and analyzed rounding schemes, while for $k > 6$ we give an algorithm with performance ratio $1.3438 - \epsilon_k$.

Our results were discovered with the help of computa-

tional experiments that we also describe here.

1 Introduction

As the field of approximation algorithms matures, methodologies are emerging that apply broadly to many NP-hard optimization problems. One such approach (e.g., [7, 8, 1, 6, 5]) has been the use of metric and geometric embeddings in addressing graph optimization problems. Faced with a discrete graph optimization problem, one formulates a relaxation that maps each graph node into a metric or geometric space, which in turn induces lengths on the graph's edges. One solves this relaxation optimally, and then derives from the relaxed solution a near-optimal solution to the original problem.

This approach has been applied successfully [2] to the *min-cost multiway cut problem*, a natural generalization of the minimum (s, t) -cut problem to more than two terminals. An instance consists of a graph with edge-costs and a set of distinguished nodes (the *terminals*). The goal is to find a minimum-cost set of edges whose removal separates the terminals. If the number of terminals is k , we call such a set of edges a k -way cut.

The first approximation algorithm for the multiway cut problem in general graphs was given by Dahlhaus, Johnson, Papadimitriou, Seymour, and Yannakakis [4]. It used a traditional minimum (s, t) -cut algorithm as a subroutine and had a performance guarantee of $2 - 2/k$.

In the work that prompted ours, Calinescu, Karloff, and Rabani [2] used a novel geometric relaxation of k -way cut in a $(3/2 - 1/k)$ -approximation algorithm. Their relaxation uses the k -simplex $\Delta = \{x \in \mathbb{R}^k : x \geq 0, \sum_i x_i = 1\}$, which has k vertices; the i^{th} vertex is the point x in Δ with $x_i = 1$ (and all other coordinates 0). The relaxation is as follows: map the nodes of the graph to points in Δ such that terminal i is mapped to the i^{th} vertex of Δ . Each edge is mapped to the straight line between its endpoints. The goal is to minimize the *volume*

MT

*MIT Laboratory for Computer Science, Cambridge, MA 02138. karger@lcs.mit.edu. Research supported by NSF contract CCR-9624239, an Alfred P. Sloane Foundation Fellowship, and a David and Lucille Packard Foundation Fellowship.

†Brown University. klein@cs.brown.edu. Research supported by NSF Grant CCR-9700146.

‡Dartmouth College. cliff@cs.dartmouth.edu. Research supported by NSF Career award CCR-9624828.

§AT&T Labs-Research, Shannon Laboratory, 180 Park Avenue, Florham Park, NJ 07932. mthorup@research.att.com.

¶Dartmouth College. ney@cs.dartmouth.edu. Research supported by NSF Career award CCR-9720664.

of G ,

$$\text{vol}(G) \doteq \sum_{\text{edges } e} \text{cost}(e) \cdot |e|$$

where $|e|$ denotes the *length* of the embedded edge e , defined as half the L_1 distance between its endpoints, and $\text{cost}(e)$ is the cross-sectional area of edge e .

To see that the above is a relaxation of minimum k -way cut, consider any k -way cut and let S_i be the set of nodes reachable from terminal i in the graph with the cut-edges removed. Consider a geometric embedding in which all nodes in S_i are mapped to vertex i of Δ . For any edge, the distance between its edges is either 0, if the endpoints lie in the same S_i , or 1, if the endpoints lie in distinct sets S_i . Hence the volume of this embedding equals the cost of the k -way cut.

The algorithm of Calinescu et al. finds a minimum volume embedding by linear programming. It then uses a randomized rounding scheme to extract a cut from this embedding. Ignoring the graph, the scheme chooses (from a carefully selected distribution) a k -way cut of the simplex—a partition of the simplex into k subsets, each containing exactly one vertex of the simplex. The k -way cut of the simplex naturally induces a k -way cut in the embedded graph—namely, the set of edges with endpoints in different blocks of the partition. This cut has expected cost at most $3/2 - 1/k$ times the volume of the embedding.

Our results. Our goal is to further understand the geometric relaxation, with the hope of developing better approximation algorithms. We aim to determine the *integrality gap* of the relaxation and to find an algorithm whose approximation ratio matches the integrality gap. Note the the integrality gap is the best approximation ratio we can achieve for an algorithm that compares itself only to the embedding volume.

In this paper, we resolve this question for 3-cut and provide improved results for the general k -cut problem. For $k = 3$ we give a rounding algorithm with performance ratio $12/11$, improving Calinescu et al.’s bound of $3/2 - 1/3 = 7/6$. We also show that $12/11$ is the best possible bound, exhibiting a graph with a gap of $12/11$ between its embedded volume and minimum 3-way cut. Thus, for $k = 3$, we determine the exact integrality gap and give an optimal algorithm.

For larger k , we obtain results based on both computation and analysis. For $k = 4, 5$, we use LP-derived and -analyzed rounding schemes to give bounds of 1.1539 and 1.2161 respectively, improving the corresponding bounds of Calinescu et al. of 1.25 and 1.3. For larger k we give a single algorithm obtaining a (analytic) bound of $1.3438 - \epsilon_k$ where $\epsilon_k > 0$. The quantity ϵ_k can be evaluated computationally for any fixed k ; we use this to prove that $1.3438 - \epsilon_k < 3/2 - 1/k$ for all k .

Our efforts to find geometric cutting schemes that achieve good guarantees were guided by experiments: we formulated the problem of determining an optimal probability distribution on k -way cuts of the simplex as an infinite-dimensional linear program, and solved discrete approximations of this linear program and its dual. From these solutions we were able to deduce the lower bound and, using that, the upper bound for $k = 3$. These experiments also guided our search for cutting schemes that work for larger values of k .

The upper and lower bounds for $k = 3$ were discovered independently by Cunningham and Tang [3].

Presentation overview. In Section 2 we discuss the geometric ideas underlying the problem. In Section 3 we describe the computational experiments we undertook and the results it gave for small k . In Sections 4 and 5 we solve the 3-terminal case, giving matching upper and lower bounds. Finally, in Section 6, we present our improved algorithm for general k .

2 The geometric problem

Finding the integrality gap of and a rounding scheme for the relaxation turns out to be expressible as a geometric question. That is, we can express integrality gaps and algorithmic performance purely in terms of the simplex, without considering particular graphs or embeddings.

As general notation, we let e_ℓ denote the interval $\{x_\ell | e \in e\}$. Then $\min e_\ell = \min_{x \in e} x_\ell$, $\max e_\ell = \max_{x \in e} x_\ell$, and $|e_\ell| = \max e_\ell - \min e_\ell$. Finally, we have defined the length $|e|$ of an edge e to be half its L_1 norm, so

$$|e| = \sum_{\ell=1}^k |e_\ell|/2.$$

MT

2.1 Density

Recall that a k -way cut of the simplex is a partition of the simplex into k subsets, each containing a unique vertex of the simplex, and that such a cut induces a k -way cut of any embedded graph. By a *cutting scheme*, we mean a probability distribution P on k -way cuts of the simplex. For any line segment e , the *density of P on segment e* , denoted $\tau_k(P, e)$, is the expected number of times a random cut from P cuts e , divided by the length¹ $|e|$ of e . Define the *maximum density of P* , $\tau_k(P)$ and the *minimal maximum density* τ_k^* as follows:

$$\tau_k(P) \doteq \sup_e \tau_k(P, e) \quad \text{and} \quad \tau_k^* = \inf_P \tau_k(P),$$

It is easy to see that the maximum density line segment will in fact be an edge of infinitesimal length, since any

¹By analogy to the length of an edge, the length of a segment is defined as half the L_1 distance between its endpoints.

segment can be divided into two edges, one of which has density no less than the original. Thus, in the remainder of this paper, we will focus discussion on such infinitesimal segments.

The relevance of τ_k^* is the following (this is implicit in the work of Calinescu et al.):

Lemma 2.1 *For any cutting scheme P and embedded graph G , the expected cost of the k -way cut of G induced by a random k -way cut from P is at most $\tau_k(P)$ times the cost of the embedding of G .*

Corollary 2.2 *Any cutting scheme P yields an approximation algorithm with approximation ratio at most $\tau_k(P)$.*

Proof Sketch: The endpoints of any edge e are embedded at two points in the simplex, so the edge corresponds to a segment connecting those two points. The expected number of times the edge is cut is $\tau_k(P, e) \cdot |e|$. By the Markov inequality this upper bounds the probability that the edge is cut. Thus, the expected cost of the k -way cut is at most $\sum_e (\tau_k(P, e) \cdot |e|) \text{cost}(e) \leq \tau_k(P) \sum |e| \cdot \text{cost}(e) = \tau_k(P) \text{vol}(G)$. \square

In fact, one can show that τ_k^* is both the integrality gap of the geometric relaxation and the best performance guarantee obtainable by any cutting scheme. That is, there is an embedded graph whose volume is arbitrarily close to τ_k^* times its minimum k -way cut and there is a cutting scheme with maximal density (and therefore performance guarantee) arbitrarily close to τ_k^* . This is a consequence of Yao's principle (i.e. von Neumann's min-max theorem, or equivalently strong linear programming duality, applied in the context of complexity theory). It also follows that a cutting scheme with optimum integrality gap can be defined obliviously, independent of the input graph.

Calinescu et al.'s algorithm gives a cutting scheme showing that $\tau_k^* \leq 3/2 - 1/k$. In this paper we show that $\tau_3^* = 12/11$, and that, for all k , $\tau_k^* \leq 1.3438$.

2.2 Alignment

We have just argued that the key question to study is the maximum density of line segments relative to a cutting scheme. Calinescu et al. showed that one can restrict attention to segments in certain orientations. We say a segment e in Δ is i, j -aligned if e is parallel to the edge connecting vertices i and j of Δ . We say it is aligned if it is i, j -aligned for some pair of vertices. Calinescu et al. observed that since length is proportional to the L_1 -norm, and since the aligned edges are the geodesics of the norm, the endpoints of any segment e can be connected by a piecewise linear path of total length $|e|$ whose segments are aligned. The segment e is cut iff some edge on this path is cut. Given any embedding of

a graph, Calinescu et al. apply this transformation separately to each segment connecting two embedded vertices, without changing the volume of the embedding. Thus, without loss of generality one may restrict attention to embeddings in which all edges are aligned.

Fact 2.3 *Segment e is i, j -aligned if and only if $|e| = |e_i| = |e_j|$ and $|e_\ell| = 0$ for $\ell \neq i, j$.*

MT

2.3 Side parallel cuts (SPARCS)

In this paper, we mainly restrict attention to a particular set of cutting schemes. Define $\Delta_{x_i=\rho} \doteq \{x \in \Delta : x_i = \rho\}$ and $\Delta_{x_i \geq \rho} \doteq \{x \in \Delta : x_i \geq \rho\}$. Note that $\Delta_{x_i=\rho}$ is a hyperplane that runs parallel to the face opposite terminal i and is at distance ρ from that face; it divides the simplex into two parts, of which $\Delta_{x_i \geq \rho}$ is the "corner" containing terminal i . An i, j -aligned segment (x, y) is cut by the hyperplane $\Delta_{x_i=\rho}$ iff $\ell \in \{i, j\}$ and ρ is between x_ℓ and y_ℓ .

We define a *side-parallel cut (sparc)* of the simplex:

1. Choose a permutation σ of the vertices;
2. For each vertex i in order by σ (except possibly the last), choose some $\rho_i \in [0, 1]$;
3. Assign to vertex i all points of $\Delta_{x_i \geq \rho_i}$ not already assigned to a previous terminal. We say terminal i *captures* all these points, and that terminal i *cuts* an edge e if it captures some but not all of e .

Thus we are slicing up the simplex using hyperplanes $\Delta_{x_i=\rho}$. In this context, we call each $\Delta_{x_i=\rho}$ a *slice*.

We consider algorithms that sample randomly from some probability distribution over sparcs. Our restriction to sparcs was motivated by several factors. The rounding algorithm of Calinescu et al. uses only sparcs. Furthermore, our computational study of the 3-terminal problem (discussed below) and some related analytic work gave some evidence that the optimal algorithm was a distribution over sparcs. Lastly, sparcs have concise descriptions (as a sequence of $k - 1$ slicing distances) that made them easy to work with computationally and analytically. It is conceivable, though, that one might do better with cuts that are not sparcs.

The key properties of SPARCS are expressed in the following fact. For segment e , let e_ℓ be the interval $\{x_\ell | x \in e\}$.

MT

Fact 2.4 *An i, j -aligned segment e is cut by a sparc if and only if it is cut by terminal i or j . Furthermore, for $\ell \in \{i, j\}$, the following conditions are all necessary for segment e to be cut by terminal ℓ :*

- (1) $\rho_\ell \in e_\ell$
- (2) For all terminals h preceding ℓ , $\rho_h > \min e_h$.

(3) Terminal ℓ is not last in the order

For probability distributions P on spars, one can obtain bounds on $\tau_k(P, e)$ by using Conditions 1–3 above. For example, we can restrict our attention to Condition 1: If ρ_i and ρ_j are uniformly distributed, Condition 1 holds for terminal i with probability $|e_i| = |e|$, and similarly for terminal j . Thus, by linearity of expectation, the expected number of times e is cut is at most $2|e|$.

Next, consider adding Condition 3. Suppose that the ordering of terminals is random, meaning that i is last with probability $1/k$. The probability that e is cut by i becomes $(1 - 1/k)|e|$, so $\tau_k(P, e) \leq (2 - 2/k)$. Thus, uniformly random ρ_ℓ 's and a random ordering gives a performance guarantee of $2 - 2/k$, matching the bound of Dahlhaus et al. [4].

To improve these bounds, one must use Condition 2. Calinescu et al. choose a sparc by selecting ρ uniformly at random in $[0, 1]$, setting $\rho_\ell = \rho$ for each terminal ℓ , and slicing off terminals in random order. Conditions 1 and 3 again derive a density bound of $1 - 1/k$ for terminal i and j . Calinescu et al. improve this analysis as follows. Suppose that the edge is farther from j than from i . We will argue that the density contribution from terminal j is only $1/2$. The point is that if ρ is such that j appears to cut e , and if i (which is closer to e) precedes j in the random slice ordering, then i will capture all of e and prevent j from cutting it. Formally, our assumption is that $\min e_i \geq \max e_j$. With probability $1/2$, i precedes j . If so, since $\rho_i = \rho_j$, Condition 1 - $\rho_j \in e_j$ - contradicts Condition 2 for i - $\rho_i > \min e_i$. Thus e can only be cut by terminal j if j precedes i , but then by Condition 1, the density contribution is $1/2$. Thus the density contribution from terminal j is $1/2$, as desired, leading to a total density of $3/2 - 1/k$. Finally note that any i, j aligned edge can be split in two with one part closer to i and one part closer to j , and our analysis then applies to each part separately.

To improve on the $3/2$ bound, we made stronger use of Condition 2. The analysis of Calinescu et al. only considers that a segment may be captured by the two terminals with which it is aligned. We derive stronger results by observing that other terminals may capture the edge as well. To do so, we had to change the cut distribution as well as the analysis. It can be shown that no distribution that holds all ρ_i equal can do better in the limit than $3/2$ of Calinescu et al. The $3/2$ can be improved somewhat in the limit by choosing the ρ_i independently with a non-uniform density distribution. However, the best cutting schemes we have found are based on different types of dependence between the ρ_i . One such scheme for 3-way cut gives us a bound of $12/11$, which is optimal over all schemes for 3-way cut. Another scheme gives us a bound of 1.3438 that

holds for any number k of terminals. This latter scheme is designed for large k .

2.4 Additional Observations

We now mention some additional observations whose full proofs must await the full paper.

What is the best embedding? Perhaps the first natural question to ask is whether the embedding chosen by Calinescu et al. is the best possible.

Lemma 2.5 *Among all embeddings in the simplex that minimize some norm (without adding other constraints) the L_1 norm has the smallest possible integrality gap.*

Proof: Very precisely, we will show that for any cutting distribution P , any other norm based distance measure d will have a maximal density at least as high as ours. First, for the embedding to be a relaxation, it is required that $d(\overline{t_i t_j}) \leq 1$ for any two terminals t_i and t_j . However, since any norm based distance measure is translation invariant and proportion preserving, this implies that for any i, j aligned edge e , that $d(e) \leq |e|$. This, in terms, implies that the density of e can only increase. With our L_1 based distance measure $|\cdot|$, we know that the maximal density is achieved over aligned edges, and hence a relaxation based on d cannot have a smaller maximal density. \square

Symmetry. A second observation is that there is no benefit in trying to identify a “good terminal order” in which to cut up the simplex.

Lemma 2.6 *There is an optimum sparc cutting scheme of the following form:*

1. choose slice distances (d_1, \dots, d_{k-1}) from some probability distribution
2. apply the slice distances (in order) to a uniform random permutation of the terminals

An analogous “order independence” statement holds for the best possible (possibly non-sparc) algorithm.

Proof: Consider a best sparc with integrality gap ρ . Consider any input embedding. We can “symmetrize” the embedding, without changing its volume, by averaging it over all permutations of the coordinates. Our sparc achieves integrality gap ρ on the symmetrized embedding. Since the embedding is symmetric, the order in which the sparc slices terminals is irrelevant. So we can assume it is some fixed order.

Note, however, that cut value achieved on the symmetrized graph but slicing in some fixed order is just the expected cut value achieved by applying the same slices

MT

MT

MT

MT

to the original embedding under a random ordering of the terminals. \square

The above lemma shows that there is no worst-case benefit to considering specific terminal ordering. The duality argument of Section 2.1 carries over to show that a sparc with optimum expected integrality gap can be specified simply as a distribution over slicing distances, without reference to an input graph embedding.

3 Our Computational Study

Neal said: The section on "LP results" disregards the initial work done with Maple and LP_SOLVE (not cplex). I don't know if we care (is crediting the software important?)

We've overlooked the point that the dual solutions were much more useful for the $k=3$ case than the primal solutions. Further, we do not describe the LOWER BOUND LP, which was actually the most useful for $k=3$, and was not the same as the upper bound LP that is described now.

This is probably a subtle point not worth pursuing, but there is a more important related issue. With $k=3$ we were able to extrapolate the general solution to the dual, but not to the primal, from the computational results. Then we figured out an analytic primal solution by reasoning about what it had to be like (according to complimentary slackness) if the analytic dual solution (lower bound) was to be tight. Recall we saw lots of optimal but very different primal solutions (from our computations), whereas the dual solutions tended to have a more canonical form. Maybe this is particular to $k=3$, but maybe not. I think using a lower bound LP (or at least the dual) might be important for higher k as well.

In fact I think we should pursue this lead ourselves by thinking carefully about the dual that we obtain when we restrict to sparc cuts.

Mikkel: I don't think we need to present this to the reader, so I wouldn't worry. Looking at the duals ourselves could be interesting, but first of all, I'd like this project to be completed and the journal version submitted.

In this section we describe some computational experiments we carried out to help us understand the behavior of the geometric embedding. One need read this section in order to understand the following ones.

As discussed above, our goal was to find a distribution over cuts of the k -simplex that minimized the den-

sity of any segment in the simplex. This problem can be formulated as an infinite dimensional linear program, with one variable per cut of the simplex, corresponding to the probability that that cut is chosen, and one constraint for every (aligned, infinitesimally small) line segment inside the simplex, which measures the expected number of times the chosen cut will cut that segment. Of course, it is not tractable to solve the infinite LP computationally, but we expected that discretized versions of it would be informative.

We applied this approach in two distinct ways. For the 3-terminal case, we devised an LP that exploited the planarity of the 3-terminal relaxation, and used it to home in on the optimal solution, which we then analytically proved to be optimal. For the general case, we devised an LP whose solutions are (provable) upper bounds on the performance of certain rounding algorithms. We solved this LP for small numbers of terminals (3–9), deriving algorithms with (computer aided) proofs of the best known performance ratios for these problems. The solution suggested certain properties that appear to hold in the "optimal" rounding scheme; we used these suggestions in our development of (analytic) solutions for arbitrary numbers of terminals.

3.1 The three-terminal case

For the 3-terminal problem we exploited planarity. The 3-simplex can be viewed as a triangle in the plane. We discretized the linear program by defining a triangular mesh over the simplex and considering only edges of the mesh instead of all line segments in the simplex.

To approximate the best cutting scheme, we computed the best distribution over 3-way cuts of the mesh. We used the planarity of the 3-simplex to simplify our LP formulation. Any 3-way cut of the mesh corresponds to a collection of paths (representing the boundary of the cut) through the planar dual of the mesh. Thus the distribution of cuts corresponds to a packing of these paths, which can be seen as a kind of flow. So instead of enumerating all possible cuts, we could define a linear program that assigned a (multicommodity) flow to each edge of the dual mesh. This gave us a tractable representation of the linear program.

We also found it helpful to solve the dual of our flow-based linear program, which assigns weights to the mesh edges to minimize the total weight such that every 3-cut has value at least 1. Since each 3-cut corresponds to a set of two or three paths in the planar dual of the mesh, the latter constraint can be represented efficiently by constraining shortest-path lengths (as a function of the variable edge lengths) in the planar dual. A solution to the dual can be interpreted as an embedded graph demonstrating the integrality gap. The dual showed us the important idea of "ball cuts" versus "corner cuts"

which we will discuss in the following sections, and thus led to the discovery of the optimum cutting scheme for three terminals.

3.2 The general case

In the general case, the lack of a planar embedding prevented us from exploiting nice properties of its cuts; we were faced with the problem of enumerating cuts as well as edges. Based on the work of Calinescu et al. and our own results for the optimal 3-terminal solution, we decided to limit our exploration to spars as discussed above.

There is still an infinite space of possible spars, so we discretized our problem. Fix an integer *grid size* N . A *discrete sparc* is described by a vector (q_1, \dots, q_{k-1}) where each q_i is an integer in the range $[0, N-1]$. Given such a vector, we choose a random sparc by setting d_i uniformly in the range $[q_i/N, (q_i + 1)/N]$. This defines a probability distribution on (continuous) spars. We now define a linear program to search for a probability distribution over discrete spars (which induces a probability distribution over continuous spars). We define a variable for each discrete sparc, which reflects the probability of choosing that discrete sparc, and provide constraints that aim to minimize the density of any segment under the probability distribution.

There still appear to be infinitely many constraints (segments) but we reduce this to a finite number as follows. The slices at distances q/N for each terminal that determine our sparc distribution partition the simplex into cells. For a given distribution on the discrete spars, we can compute a (linear) upper bound on the density induced on *any* segment with a given alignment within a cell, and specify one constraint saying that this upper bound should be small. Since the cells are small, we expect all segments with a given alignment to have roughly the same density under our cutting scheme, so we hope that the upper bound is reasonable tight. With this simplification, the number of constraints is bounded by the number of cells times the number of segment alignments per cell, which is at most $k^2 N^k$.

We determine the upper bound for a cell as follows. For any discrete cut, the slices generated from it will fall into one of three categories. If the i^{th} coordinate of the discrete cut is different from that of the cell, then the i^{th} slice will not pass through that cell: depending on whether the coordinate is larger or smaller it will either capture the entire cell or none of the cell. If the i^{th} coordinates are the same, then the slice might pass through the cell; we can use that the slice is uniformly distributed over a range to determine its density contribution.

If we consider an i, j -aligned segment, it can only be cut if the slices for terminal i or j go through its cell. If only one of the two slices goes through the terminal then

its contribution to a segment’s density is at most $1/N$. If both slices go through the cell, their contribution is at most $2/N$. We ignore the fact that different slices within the cell might capture the segment before it can be cut, thus introducing some slack in our upper bound.

We can exploit symmetry to further reduce the number of constraints we consider. Since by assumption our sparc slices terminals in random order, two segments that are identical under permutation of coordinates will have the same densities, so we need consider only one of them. Thus, we restrict our constraints to 1, 2-aligned segments in which the remaining coordinates are in non-decreasing order.

3.3 LP Results

Exploiting symmetry as discussed above, we were able to solve relatively fine discretizations of the problem. We wrote a simple program to generate the linear programs automatically, and used CPLEX to solve them. While it is difficult to “prove” programs correct, our computations did converge to the correct 12/11 approximation ratio for the 3-terminal case.

We give our results below in tabular form. We derived improved bounds for 4–9 terminals. Note that (under the assumption that the programs were correct) these are provable upper bounds. In fact, since the programs output a particular distribution over discrete cuts, their performance ratio could be proven analytically via a tedious case analysis (which we have not performed).

k	Grid	LP Gap	$3/2 - 1/k$	corner cut probability
3	90	1.0941	1.16666	.2849
4	36	1.1539	1.25	.2891
5	18	1.2161	1.3	.3144
6	12	1.2714	1.33333	.3760
7	9	1.320	1.357	.3973
8	6	1.3322	1.375	.4146

Our experiments also revealed one interesting fact: in all cases, the optimum cut distribution made use of “corner cuts.” That is, the output distribution had the following form: with some probability, place each slice at a distance chosen uniformly between 0 and $1/3$ from its terminal; otherwise, use a (joint) distribution that places every slice at distance greater than $1/3$ from its terminal.

Adding constraints that forced the corner cuts to operate over a range other than $1/3$ of the way from the terminals worsened the computed performance ratio, hinting that perhaps the optimal algorithm uses corners of size exactly $1/3$. This result is consistent with the optimal 3-terminal algorithm, however it could be a misleading artifact of working with a discretized problem.




Figure 1: This figure illustrates the cuts used for the case $k = 3$. The leftmost diagram shows how r might be chosen for the ball cut. The middle diagram shows one possible resulting ball cut (in bold). The rightmost diagram shows a corner cut (in bold).

4 Upper Bound for $k = 3$

Our analytic upper bound of $12/11$ for $k = 3$ comes from a new cutting scheme that we call the ball/corner scheme. Though for simplicity we present a non-sparse scheme, there is a similar scheme using sparses that achieves the same bound.

For $k = 3$, the simplex Δ can be viewed as a triangle in the plane, which simplifies our pictures. However, we continue to use the original three-dimensional coordinate system to locate points in the simplex. Our cut of the simplex is determined by some lines and rays drawn through the triangle; we refer to them as *boundaries*. We will show that no segment has high density with respect to our random choice of boundaries.

As illustrated in Figure 1, denote the vertices of the simplex 1, 2, 3. Let points a, b, \dots, f divide the edges in thirds, so that $a-b-f-d-c-e-a$ is the hexagon in Δ with side length $1/3$. Note that the hexagon is (a scaled version of) the unit ball for our distance metric. The points on the boundary of the hexagon are each at distance $1/3$ from the hexagon's center.² Outside the hex, we have a corner for each terminal i consisting of the points x with $x_i > 2/3$

The ball/corner scheme chooses a *ball cut* with probability $8/11$, otherwise it chooses a *corner cut*. These two types of cuts are defined next. The scheme is illustrated in Figure 1.

Ball cut: Choose a point r uniformly at random from either line $a-c$ or line $b-d$. Consider the three lines $\Delta_{x_i=r_i}$ ($i = 1, 2, 3$) parallel to the triangle's sides and passing through the point r . Each such line is divided at the point r into two rays. Thus we get six rays. Each side of the triangle intersects two of these rays. For each side, randomly choose of the two rays that hit it. This gives three rays; they form the boundary of the 3-way cut.

Corner Cut: Choose two terminals in $\{1, 2, 3\}$, and a value $\rho \in (2/3, 1]$, uniformly at random. For each of the two chosen terminals i , let $l_i = \Delta_{x_i=\rho}$. The two lines l_i form the boundaries of the 3-way cut.

²Remember that we measure length as half the L_1 norm.

Analysis. We first state two simple properties of the ball cut that we need to analyze the performance of the cutting scheme:

Fact 4.1 *Each of the 3 coordinates of the random point r is uniformly distributed in $[0, 2/3]$.*

Fact 4.2 *Once r is chosen, each one of the six candidate rays connecting r to one side of the triangle is chosen with probability $1/2$.*

Theorem 4.3 *The maximum density of the ball/corner scheme is $12/11$, so $\tau_3^* \leq 12/11$.*

Proof: We show that the expected density of any segment e is at most $|e| \cdot 12/11$. For the ball cuts, we use only the two facts claimed above. Since these two facts, as well as the corner cut scheme, are symmetric with respect to the three coordinates, it suffices to prove the claim only for a 1,2-aligned segment e . Further, we may assume that e is entirely contained either in a corner, or in the hex; for otherwise, as discussed in Section 2.1, we can just split e into corresponding pieces, calculating the density for each piece separately. We will consider several cases, depending on where e is located.

First, assume e is located entirely in the hex. Such a segment cannot be cut by a corner cut, so we need only consider the density when a ball cut is made and multiply by the probability of choosing a ball cut, namely $8/11$. Assume a ball cut is made. Then e can only be cut by rays of in $\Delta_{x_i=r_i}$ for $i = 1, 2$. By Fact 4.1, r_i is uniformly distributed in $[0, 2/3]$. Hence, the probability that $\Delta_{x_i=r_i}$ goes through e is $|e|/(2/3)$ since e is 1,2-aligned. If $\Delta_{x_i=r_i}$ touches e , it is at a single point. By Fact 4.2, the ray of $\Delta_{x_i=r_i}$ containing this point is picked for the cut with probability $1/2$. Thus the expected number of times e is cut is $\frac{8}{11} \cdot 2 \cdot \frac{|e|}{2/3} \cdot \frac{1}{2} = \frac{12}{11}|e|$.

Exactly the same argument applies if the edge is in the corner closest to terminal 3. The ball cut contributes the same $12/11$ density, while the corner cut contributes

MT

MT

MT




Figure 2: The lower bound for $k = 3$ (here $N = 7$). The paths from 2 to 3 are on the left. The entire graph is on the right. On the border, overlapping paths are drawn side-by-side for clarity, so line width represents edge cost.

nothing (note that a 1, 2-aligned edge is parallel to the line $\Delta_{x_3=r_i}$, so cannot be cut by it).

MT

Finally, suppose segment e is in the corner closest to terminal 1 (a symmetric argument applies if e is in the corner closest to terminal 2). In this case, if a ball cut is made, the above analysis applies except that only the line $\Delta_{x_2=r_2}$ can cut e (the line $\Delta_{x_1=r_1}$ never enters the corner), so the density contribution of the ball cut is halved to $|e|\frac{6}{11}$. But the edge can also be cut by a corner cut. A corner cut is chosen with probability $3/11$. When it is, two of the three terminals are chosen, so terminal 1 is chosen with probability $2/3$. If terminal 1 is chosen, then, since the cutting line near terminal 1 is of the form $\Delta_{x_1=1-p}$, where p is chosen uniformly in $[0, 1/3]$, the probability that the line cuts e is $|e|/(1/3)$. Thus, the expected number of times that the edge e is cut (by a ball cut or corner cut) is $|e|\frac{6}{11} + \frac{3}{11} \cdot \frac{2}{3} \cdot \frac{|e|}{1/3} = |e|\frac{12}{11}$.

Finally, if e spans several regions (e.g. it lies in a corner and in the hex), e can be partitioned into sub-segments each contained entirely in one region, and the previous analysis applied to the sub-segments. \square

MT

5 Lower Bound for $k = 3$

Theorem 5.1 *For $k = 3$, the minimum maximum density $\tau_3^* \geq 12/11$. Hence, the integrality gap for the geometric relaxation is $12/11$.*

Note that this theorem applies to all cutting schemes, not just spars.

Proof: Fix N to be any positive integer. We construct an embedded weighted graph G_N with no 3-way cut of cost less than $12N - 3$, but with an embedding of cost $11N + 3$. This implies that no cutting scheme has maximum density less than $(12N - 3)/(11N + 3)$, because by Lemma 2.1 such a cutting scheme applied to G_N would yield a 3-way cut with expected cost less than $12N - 3$, a contradiction. Since N is arbitrary, the result follows. Our construction (for $N = 7$) is shown in Figure 4.

For any pair of distinct terminals i, j and number $d \in [0, 1]$, define embedded path $p(i, j, d)$ as follows. Let ℓ be the terminal in $\{1, 2, 3\} - \{i, j\}$; let a be the

point on segment $i\ell$ at distance d from i ; let b be the point on segment $j\ell$ at distance d from j . Then $p(i, j, d)$ is the union of the three segments ia , ab , and bj .

We form the graph from $9N$ paths $p(i, j, d)$ for $0 \leq d \leq 2/3$; where d is an integer multiple of $1/(3N)$. Although we describe the graph as a set of paths, technically it is a planar graph consisting of nodes and edges as follows: for every point in Δ whose coordinates are integer multiples of $1/(3N)$, there is a node in the graph embedded at that point; for every pair of nodes embedded $1/(3N)$ units apart, G has an edge with cost equal to the number of paths that pass through both nodes.

With this understanding, we now specify the graph. For each of the 3 distinct pair of terminals i, j , there are $3N$ paths. Of these paths, N run directly between the terminals; that is, there are N copies of $p(i, j, 0)$. The remaining $2N$ paths are the paths $p(i, j, m/(3N))$ where $m = 1, 2, \dots, 2N$.

The total cost of the embedding is the total length of the paths. Since a path $p(i, j, m/(3N))$ has length $1 + m/(3N)$, a direct calculation shows that the total length of the paths is $11N/3 + 1$.

Next we lower bound the cost of any 3-way cut. Since the graph is planar, any minimal 3-way cut corresponds either to a disconnected cut (meaning that the cut is the union of two disjoint 2-way cuts, each separating some terminal from both other terminals), like our upper bound's corner cut, or a connected cut (meaning that the cut edges give, in the planar dual, three paths connected at some central node and going to the three sides of the triangle), like our upper bound's ball cut.

Any 3-way cut must cut all of the $9N$ paths at least once. To finish the proof, we will argue that for either type of 3-way cut (connected or not), at least $3N - 3$ paths are cut twice, so that the edges cut by the 3-way cut cost at least $12N - 3$. This is easy to verify for a disconnected cut: a disconnected cut is the union of two 2-way cuts, so the $3N$ paths running between the two terminals that are cut off must be cut twice.

Now consider any connected cut. In the planar dual of G_N , the connected cut corresponds to a central node and three paths from the node to each side of the trian-

gle. Let $x = (x_1, x_2, x_3)$ be any point inside the face of G_N corresponding to the central node. Consider a path $p(i, j, d)$ such that $d \geq x_\ell$, where $\ell \neq i, j$. That is, X is inside the cycle formed by the union of $p(i, j, d)$ and $p(i, j, 0)$. Then the path $p(i, j, d)$ is cut twice by the connected cut. For fixed i and j , the number of such paths (with $d \geq x_\ell$) is at least $(2/3 - x_\ell)N/3 - 1$. Thus, the total number of such paths is at least $(2/3 - x_1 + 2/3 - x_2 + 2/3 - x_3)3N - 3 = 3N - 3$. \square

6 Improvement for general k

Theorem 6.1 *For all k , $\tau_k^* \leq 1.3438$. Moreover, there is a k -way cut approximation algorithm with an approximation guarantee of 1.3438.*

Our bound improves on the Calinescu et. al. bound of $1.5 - 2/k$ for all $k \geq 14$. For $3 < k < 14$, we can also obtain improvements by taking advantage of k being small (see Section 6.1).

As discussed in Section 2.3, the essential observation in this analysis is that many slices can capture an edge before it has a chance to be cut.

We will use a (sparc) cutting scheme called ICUT: we choose k slicing thresholds ρ_i , and apply the slices $\Delta_{x_i=\rho_i}$ to a random permutation σ of the terminals.

To bound the cutting density of our scheme, we will bound the density of every segment. As justified in Section 2.2, we consider a segment of length $\epsilon > 0$, and let ϵ approach zero. As in the ball/corner scheme, by symmetry we can assume without loss of generality that the segment is 1, 2-aligned.

Define $d_k(x_1, \dots, x_k)$ to be the density which which ICUT cuts a 1, 2-aligned segment of infinitesimal length located at x_1, x_2, \dots, x_k . We will show:

Theorem 6.2

$$d_k(x_1, \dots, x_k) \leq \begin{cases} 2.012096 & \text{if } x_1, x_2 \leq 6/11 \\ 11/12 & \text{otherwise.} \end{cases}$$

The final cutting scheme chooses to ICUT with probability $\alpha = 0.66719$ and otherwise chooses a corner cut. The corner cut is chosen by the natural generalization of the scheme for $k = 3$: choose a value $\rho \in [6/11, 1]$. The k -cut consists of the hyperplanes $l_i = \Delta_{x_i=\rho}$, for each i . Note that the last corner cut need not technically be made but it simplifies the analysis.

This combined scheme gives a maximum density of $\max\{(2.012096)\alpha, (11/12)\alpha + (11/5)(1 - \alpha)\} \leq 1.3438$, proving Theorem 6.1. It remains to prove Theorem 6.2.

The cumulative probability distribution function for any ρ_i is $F(z) = \min\{(11/6)z, 1\}$. The correspond-

ing probability density function is

$$F'(z) = \begin{cases} 11/6 & \text{if } z \in [0, 6/11] \\ 0 & \text{otherwise.} \end{cases}$$

Consider a 1, 2-aligned segment of length ϵ with one endpoint fixed at x_1, x_2, \dots, x_k . As ϵ goes to zero, the density of this segment goes to

$$\begin{aligned} & d_k(x_1, \dots, x_k) \\ &= \frac{1}{k!} \sum_{\sigma} \left(F'(x_1) \prod_{i:\sigma(i)<\sigma(1)} [1 - F(x_i)] \right. \\ & \quad \left. + F'(x_2) \prod_{i:\sigma(i)<\sigma(2)} [1 - F(x_i)] \right) \quad (1) \end{aligned}$$

where the sum is over all $k!$ orderings of the terminals. This formula follows from Fact 2.41-2. Particularly, it accounts for the probability of the 1, 2-aligned edge being captured by the terminals going before 1 or 2, and this saving is crucial for improving the factor of $3/2$ for large k . The formula assumes that F is continuous around each x_i and that F' is continuous in an open region around x_1 and x_2 . The latter is not the case around $6/11$. However, as discussed Section 2.1, we may assume that all segments e considered have been subdivided so that for each $i = 1, 2$, either $\max x_i \leq 6/11$ or $\min x_i > 6/11$. MT

Note that $d_k(x_1, \dots, x_i, 0, \dots, 0) = d_i(x_1, \dots, x_i)$ (provided $i \geq 2$), because $x_j = 0$ implies terminal j cannot save the edge. Note also that d_k is symmetric with respect to the variables x_i for $i > 2$. Define

$$\begin{aligned} D_k(x_1, x_2) &\doteq \max_{x_3, \dots, x_k} d_k(x_1, x_2, \dots, x_k) \\ C_k(x_1, x_2) &\doteq d_k(x_1, x_2, c, \dots, c) \\ &\quad \text{where } c = (1 - x_1 - x_2)/(k - 2), \\ D_\infty(x_1, x_2) &\doteq \lim_{k \rightarrow \infty} D_k(x_1, x_2), \\ C_\infty(x_1, x_2) &\doteq \lim_{k \rightarrow \infty} C_k(x_1, x_2). \end{aligned}$$

In these definitions, (x_1, x_2, \dots, x_k) is required to lie in the k -simplex.

D_k is the maximum density of any 1, 2-aligned infinitesimal segment with an endpoint whose first two coordinates are x_1, x_2 . Note that the maximum is well-defined and achieved by some x_3, \dots, x_k because the simplex is closed under limit.

To understand ICUT, our first goal is to characterize D_k . We consider C_k as it is one candidate for D_k .

Lemma 6.3 $D_k(x_1, x_2) \leq D_{k+1}(x_1, x_2)$ for all k .

Proof: $d_k(x_1, \dots, x_k) = d_{k+1}(x_1, \dots, x_k, 0)$, \square

Thus the D_k are a nondecreasing sequence bounded from above (by 2). This implies that D_∞ is well-defined. We will see later that C_∞ is also well-defined.

Next we show that for fixed x_1 and x_2 , the worst case occurs at either the ‘‘central point’’ x_1, x_2, c, c, \dots, c or the ‘‘three-terminal’’ point $x_1, x_2, 1 - x_1 - x_2, 0, \dots, 0$. (Analogous results hold for *any* convex or concave F .)

Lemma 6.4

$$d_k(x_1, \dots, x_k) \leq \begin{cases} C_k(x_1, x_2) & \text{if } \forall i > 2 : x_i \leq 6/11 \\ C_3(x_1, x_2) & \text{if } \exists i > 2 : x_i \geq 6/11. \end{cases}$$

Proof: Fix x_1 and x_2 . Let $c = (1 - x_1 - x_2)/(k - 2)$.

Claim 1: Among all x_3, \dots, x_k such that $0 \leq x_i \leq 6/11$ for all $i > 2$ (and x_1, x_2, \dots, x_k is in the simplex), the unique maximizer of $d_k(x_1, x_2, x_3, \dots, x_k)$ satisfies $x_3 = x_4 = \dots = x_k$. Suppose for contradiction that some other such x_3, x_4, \dots, x_k maximizes d_k . Then $x_i < x_j$ for some $i, j > 2$. Considered just as a function of x_i and x_j (holding the other coordinates fixed)

$$d_k(x_1, \dots, x_k) = p + q[1 - F(x_i)] + r[1 - F(x_j)] + s[1 - F(x_i)][1 - F(x_j)] \quad (2)$$

where p, q, r and s are nonnegative and independent of x_i and x_j . Furthermore $q = r$ because d_k is symmetric in x_i and x_j . Consider increasing x_i and decreasing x_j at equal rates. This maintains $0 \leq x_i, x_j \leq 6/11$ but increases d_k at a rate proportional to

$$q[F'(x_j) - F'(x_i)] + s\left(F'(x_j)[1 - F(x_i)] - F'(x_i)[1 - F(x_j)]\right).$$

This is positive because $F'(z) = 11/6$ for $z \leq 6/11$ and $F(x_j) > F(x_i)$ (recall that $x_i < x_j \leq 6/11$). This contradicts the choice of x_3, \dots, x_k .

Claim 2: Among all x_3, \dots, x_k such that $x_i \geq 6/11$ for some $i > 2$ (and x_1, \dots, x_k is in the simplex), the unique maximizer of $d_k(x_1, x_2, x_3, \dots, x_k)$ satisfies $x_i = 1 - x_1 - x_2$ and $x_j = 0$ for $j \neq i$. Suppose for contradiction that some other such x_3, x_4, \dots, x_k maximizes d_k . Fix some $j > 2$ such that $0 < x_j < 6/11 \leq x_i$. Since $F(x_i) = 1$, the expression (2) reduces to $p + q(1 - F(x_j))$. If we increase x_i and decrease x_j at the same rate, the rate of increase in d_k is $qF'(x_j) > 0$, contradicting the choice of x_3, \dots, x_k .

The two claims together prove the lemma. \square

Lemma 6.5 For $k \geq 4$, $C_k(x_1, x_2) \leq C_{k+1}(x_1, x_2)$.

Proof:

$$\begin{aligned} C_k(x_1, x_2) &= d_k(x_1, x_2, c, \dots, c) \\ &= d_{k+1}(x_1, x_2, c, \dots, c, 0) \\ &\leq C_{k+1}(x_1, x_2). \end{aligned}$$

Here $c = (1 - x_1 - x_2)/(k - 2)$. The last inequality follows from Lemma 6.4 (using $c \leq 1/2 < 6/11$). \square

An immediate corollary is that $C_\infty(x_1, x_2)$ is well-defined and $C_k(x_1, x_2) \leq C_\infty(x_1, x_2)$ for all k . Using this and Lemma 6.4, to bound D_∞ it suffices to bound C_3 and C_∞ . We begin with C_∞ .

Lemma 6.6

$$C_\infty(x_1, x_2) \leq \begin{cases} 2.012096 & \text{if } x_1, x_2 \leq 6/11 \\ 11/12 & \text{otherwise.} \end{cases}$$

Proof: Fix x_1 and x_2 . Our first goal is to derive a closed-form expression for $C_k(x_1, x_2)$ for any k . Fix k for now and let $x_i = c = (1 - x_1 - x_2)/(k - 2)$ for $i > 2$.

For $j = 1, 2$, let S_j denote the probability that the segment at (x_1, x_2, \dots, x_k) is not captured by a terminal other than j before the j th cut is made:

$$S_j \doteq \frac{1}{k!} \sum_{\sigma} \prod_{i: \sigma(i) < \sigma(j)} 1 - F(x_i).$$

Then $C_k(x_1, x_2) = S_1 F'(x_1) + S_2 F'(x_2)$.

We will derive a closed-form expression for S_1 (and by symmetry for S_2). Recall that $x_i = c$ for $i > 2$. We thus rewrite

$$S_1 = \frac{1}{k} \sum_{q=0}^{k-1} \frac{q}{k-1} (1 - F(c))^{q-1} [1 - F(x_2)] + \left(1 - \frac{q}{k-1}\right) (1 - F(c))^q.$$

Here we condition on q , the number of j such that $\sigma(j) < \sigma(1)$. Note that q is uniform in $[0, k - 1]$ while $\frac{q}{k-1}$ is the probability that $\sigma(2) < \sigma(1)$, given q .

A change of variables and rewriting give

$$S_1 = \left(1 + \frac{1 - F(x_2)}{k-1}\right) \sum_{q=0}^{k-2} \frac{(1 - F(c))^q}{k} - F(x_2) \sum_{q=0}^{k-2} \frac{q(1 - F(c))^q}{k^2 - 2k}.$$

Now we let $k \rightarrow \infty$. The two sums above have standard closed forms that tend respectively to

$$[1 - e^{-a}]a^{-1} \text{ and } [1 - (1 + a)e^{-a}]a^{-2},$$

where $a \doteq \lim_{k \rightarrow \infty} k F(c) = (1 - x_1 - x_2)F'(0)$. Thus,

$$S_1 \rightarrow [1 - e^{-a}]a^{-1} - F(x_2)[1 - (1 + a)e^{-a}]a^{-2}.$$

Of course S_2 is the above with x_1 replacing x_2 . This gives us our closed-form expression for $C_\infty(x_1, x_2)$:

$$\begin{aligned} C_\infty(x_1, x_2) &= [F'(x_1) + F'(x_2)] \times \frac{1 - e^{-a}}{a} \\ &\quad - [F'(x_1)F(x_2) + F'(x_2)F(x_1)] \times \frac{1 - (1 + a)e^{-a}}{a^2}. \end{aligned} \quad (3)$$

where $a = (1 - x_1 - x_2)F'(0)$.

The above equality holds for any F . Using this closed form and our particular choice of F , we now show the two desired bounds on C_∞ .

Case 1: $x_1, x_2 \leq 6/11$. In this case $a = 11/6(1 - x_1 - x_2)$, $F'(x_1) = F'(x_2) = 11/6$, and $F(x_1) + F(x_2) = 11/6(x_1 + x_2) = 11/6 - a$. So (3) gives

$$C_\infty(x_1, x_2) = 11/3 \frac{1 - e^{-a}}{a} - \frac{121}{36} \left(1 - \frac{6}{11}a\right) \frac{1 - (1+a)e^{-a}}{a^2}$$

where $a = 11/6(1 - x_1 - x_2)$ so $a \in [0, 11/6]$. Let $C(a) = C_\infty(x_1, x_2)$. In the rest of this case (Case 1), we will prove that $C(a) \leq 2.012096$ for $a \in (0, 11/6)$. The cases $a = 0$ and $a = 11/6$ follow by the continuity of C . The claim is ‘‘obvious’’ from a plot but the somewhat technical proof appears below.

We show that $C(a)$ is strictly concave for $a \in (0, 11/6)$. It therefore has a unique maximum at some a_0 , where $C'(a_0) = 0$. By substitution, $C'(.294) \geq 0.00045 > 0$ and $C'(.295) \leq -0.00009 < 0$, so $a_0 \in (.294, .295)$. Hence

$$\max_{a \in [0, 11/6]} C(a) \leq C'(.295) - 0.001 \cdot C'(.295) \leq 2.012096$$

To show $C(a)$ is strictly concave, we show that $C''(a)$ is strictly negative. Now, $C'(a) = \frac{11}{36} \frac{7e^{-a}a^2 - 18a - 4e^{-a}a}{a^3} + \frac{11}{36} \frac{6e^{-a}a^3 + 22 - 22e^{-a}}{a^3}$ and $C''(a) = -\frac{11}{36a^4} (7e^{-a}a^3 + 3e^{-a}a^2 - 36a - 30e^{-a}a + 6e^{-a}a^4 + 66 - 66e^{-a})$.

To show that $C''(a)$ is negative, it suffices to prove that

$$D(a) = -7e^{-a}a^3 - 3e^{-a}a^2 + 36a + 30e^{-a}a - 6e^{-a}a^4 - 66 + 66e^{-a}$$

is negative. By substitution, $D(0) = 0$ and $D(11/6) = 0$, so it suffices to show that D' has only one zero a_1 , $D'(a) < 0$ for $a < a_1$ and $D'(a) > 0$ for $a > a_1$. Here

$$D'(a) = -17e^{-a}a^3 - 18e^{-a}a^2 - 36e^{-a}a + 36 - 36e^{-a} + 6e^{-a}a^4$$

and $D''(a) = e^{-a}a^2(-6a^2 + 41a - 33)$. For $a \in (0, 11/6]$, D'' has only one zero $a_2 = \frac{41 - \sqrt{889}}{12} \approx 0.93$ and $D''(a) < 0$ for $a < a_2$ and $D''(a) > 0$ for $a > a_2$. That is, D' is first decreasing and then increasing. Since $D'(0) = 0$ and $D'(11/6) \geq 4.108 > 0$ it follows that D' has only one zero a_1 for $a \in (0, 11/6]$.

Case 2: x_1 or $x_2 \geq 6/11$. Assume $x_1 \geq 6/11$ (the case $x_2 \geq 6/11$ is symmetric). In this case, $F'(x_1) = 0$ and $F(x_1) = 1$, so we get

$$C_\infty(x_1, x_2) = \frac{11}{6} \frac{1 - e^{-a}}{a} - \frac{11}{6} \frac{1 - (1+a)e^{-a}}{a^2}.$$

As before, let $C(a) = C_\infty(x_1, x_2)$. We will prove that $C(a) \leq 11/12$ for $a \in [0, 11/6]$. First, $\lim_{a \rightarrow 0} C(a) = 11/12$, so $C(a) \leq 11/12$ follows if we can show that $C'(a) \leq 0$ for $a \in (0, 11/6]$. We have

$$C'(a) = \frac{11}{6a^3} (-a - e^{-a}a + 2 - 2e^{-a}).$$

Define $E(a) = -a - e^{-a}a + 2 - 2e^{-a}$. Since $\frac{11}{6a^3} > 0$ for $a > 0$, $C(a) \leq 0$ if and only if $E(a) \leq 0$. Since $E(0) = 0$, we can infer $E(a) \leq 0$ if $E'(a) \leq 0$ for all $a \in (0, 11/6]$. We have $E'(a) = -1 + e^{-a}(a+1)$. Note that $E'(0) = 0$, so $E'(a) \leq 0$ follows if $E''(a) \leq 0$ for $a \in (0, 11/6]$. We have $E''(a) = -e^{-a}a$, so $E''(a) \leq 0$. We conclude that $C_\infty(x_1, x_2) \leq 11/6$ if $x_1 > 6/11$. \square

Lemmas 6.4 through 6.6 prove that, for x such that $x_i \leq 6/11$ for all $i > 2$,

$$d_k(x_1, \dots, x_k) \leq \begin{cases} C_\infty(x_1, x_2) & \\ \leq \begin{cases} 2.012096 & \text{if } x_1, x_2 \leq 6/11 \\ 11/12 & \text{otherwise.} \end{cases} \end{cases}$$

The remaining case is when $x_i \geq 6/11$ for some $i > 2$. In this case by Lemma 6.4,

$$d_k(x_1, \dots, x_k) \leq C_3(x_1, x_2) = d_3(x_1, x_2, 1 - x_1 - x_2)$$

and $x_1 + x_2 \leq 5/11$. Thus, to finish the proof of the theorem, it suffices to show the following lemma.

Lemma 6.7 *If $x_1 + x_2 \leq 5/11$, $C_3(x_1, x_2) \leq 11/6$.*

Proof: Let $x_3 = 1 - x_1 - x_2 \geq 6/11$.

Then $F(x_3) = 1$ while $F(x_1) = 11/6 x_1$, $F(x_2) = 11/6 x_2$, and $F'(x_1) = F'(x_2) = 11/6$.

By inspection of (1), $C_3(x_1, x_2) = d_3(x_1, x_2, x_3) = (1/6)(11/6)(6 - 11/6(x_1 + x_2)) \leq 11/6$. \square

This proves Theorem 6.2.

6.1 Improvements for small values of k

For particular values of k it is possible to refine the analysis in the proof of Theorem 6.1 to get improved bounds. In this case it is useful to modify the algorithm so that it only uses $k - 1$ cuts instead of k . In particular, we do not use the cut for the terminal j with $\sigma(j) = k$. The analysis for this case goes similarly, with our definitions appropriately modified to reflect that we are using $k - 1$ instead of k cuts.

Then, instead of passing to the limit, $C_k(x_1, x_2)$ can be evaluated directly. Following this approach we obtained the following performance guarantees for particular k :

k	corner	p	ratio
3	0.641	0.675	1.131
4	0.607	0.663	1.189
5	0.588	0.659	1.223
6	0.576	0.659	1.244
7	0.565	0.657	1.258
8	0.557	0.656	1.269
9	0.557	0.659	1.277
10	0.557	0.661	1.284
12	0.554	0.661	1.293
20	0.554	0.666	1.314
35	0.550	0.666	1.327

Here, “corner” is the placement of the corner (analogous to 6/11), p is the probability of choosing ICUT, and “ratio” is an upper bound on the resulting ratio. The corner sizes and p 's are approximate and only close to optimal and the ratios were evaluated numerically without formal verification.

Foundations of Computer Science, pages 577–591, Santa Fe, New Mexico, 20–22 November 1994. IEEE.

References

- [1] Yonatan Aumann and Yuval Rabani. An $O(\log k)$ approximate min-cut max-flow theorem and approximation algorithm. *SIAM Journal on Computing*, 27(1):291–301, February 1998.
- [2] G. Calinescu, H. Karloff, and Y. Rabani, “An improved approximation algorithm for MULTIWAY CUT”, *Proceedings of the Thirtieth Annual ACM Symposium on Theory of Computing*, pages 48–52, Dallas, Texas 230026 May, 1998. ACM Press.
- [3] W. H. Cunningham and L. Tang, “Optimal 3-terminal cuts and linear programming,” *Proceedings of the Seventh Conference on Integer Programming and Combinatorial Optimization (IPCO)*, to be published as *Lecture Notes in Computer Science*, Springer-Verlag, to appear, 1999.
- [4] E. Dahlhaus, D. S. Johnson, C. H. Papadimitriou, P. D. Seymour, and M. Yannakakis. The complexity of multiterminal cuts. *SIAM Journal on Computing*, 23(4):864–894, August 1994.
- [5] Guy Even, Joseph (Seffi) Naor, Satish Rao, and Baruch Schieber. Divide-and-conquer approximation algorithms via spreading metrics (extended abstract). In *Proceedings of the 36th Annual Symposium on Foundations of Computer Science*, pages 62–71, Milwaukee, Wisconsin, 23–25 October 1995. IEEE.
- [6] Michel X. Goemans and David P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM*, 42(6):1115–1145, November 1995.
- [7] 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 Symposium on Foundations of Computer Science*, pages 422–431, 1988.
- [8] Nathan Linial, Eran London, and Yuri Rabinovich. The geometry of graphs and some of its algorithmic applications. In *Proceedings of the 35th Annual Symposium on*