On a bidirected relaxation for the Multiway Cut problem

Chandra Chekuri* Anupam Gupta† Amit Kumar‡

May 10, 2005

Abstract

In the Multiway Cut problem, we are given an undirected edge-weighted graph $G = (V, E)$ with c_e denoting the cost (weight) of edge e. We are also given a subset S of V, of size k, called the *terminals*. The objective is to find a minimum cost set of edges whose removal ensures that the terminals are disconnected.

In this paper, we study a *bidirected* linear programming relaxation of Multiway Cut. We resolve an open problem posed by Vazirani [10], and show that the integrality gap of this relaxation is no better than that for a geometric linear programming relaxation given by Călinescu et al. [2], and may be strictly worse on some instances.

1 Introduction

The Multiway Cut problem is the following: we are given an undirected edge-weighted graph $G = (V, E)$ with c_e denoting the cost (weight) of edge e. We are also given a subset S of V, of size k, called the *terminals*. The objective is to find a minimum cost set of edges whose removal ensures that the terminals are disconnected. This fundamental cut problem has applications to parallel and distributed computing [9], VLSI chip design, and computer vision [1].

The Multiway Cut problem is NP-hard even for $k = 3$, and in fact is also Max-SNP hard [3]. Note that the case $k = 2$ is the classical s-t cut problem of Ford and Fulkerson [5]. In [3], a simple $2 - 2/k$ approximation is presented via *isolating cuts*; i.e., by considering, for each terminal s, the cheapest cut separating it from $S \backslash \{s\}$, and then taking the $k - 1$ cheapest cuts among these. The same bound can also be achieved via the following natural linear programming relaxation. For each edge e there is a variable $\mu_e \in [0, 1]$, and the LP minimizes $\sum_e c_e \mu_e$ subject to the constraint that for any two distinct terminals s_i and s_j, $\mu(s_i, s_j) \geq 1$. Here $\mu(s_i, s_j)$ is the distance from s_i to s_j in the graph with edge-lengths μ_e. The integrality gap of this relaxation is also $2 - 2/k$, and [7] showed that the LP has optimum solutions that are half-integral (even for the node weighted case). We will refer to this relaxation as DistLP.

More recently, Călinescu, Karloff and Rabani [2] gave an interesting and new geometric relaxation for the Multiway Cut problem and used it to obtain a $1.5 - 1/k$ approximation. The upper bound on the

*Lucent Bell Labs, 600 Mountain Avenue, Murray Hill NJ 07974. Email: `chekuri@research.bell-labs.com`.

†Department of Computer Science, Carnegie Mellon University. Email: `anupamg@cs.cmu.edu`. This work was done while the author was visiting Lucent Bell Labs and was partly supported by a DIMACS grant.

‡Dept. of Computer Science, Indian Institute of Technology, Hauz Khas, New Delhi, India - 110016. Email: `amitk@cse.iitd.ernet.in`. This work was done while the author was at Lucent Bell Labs.

integrality gap of their LP (henceforth referred to as the CKR relaxation) has been improved by Karger et al. [8] to $1.3438 - \epsilon_k$, where ϵ_k is a constant depending only on k and tends to zero as k tends to infinity. The results in [8] provide the current best approximation ratio for the problem.

In [10, Exercise 19.7], Vazirani proposes the following *bidirected* relaxation (henceforth referred to as the BiDir relaxation). From the given undirected graph $G = (V, E)$ a directed graph $H = (V, A)$ is obtained by replacing each undirected edge $e = \{u, v\}$ by two directed arcs (u, v) and (v, u) with the same cost as e. Given $s_1, s_2, \ldots, s_k$, some permutation of the terminals, the relaxation is the following: for each arc $a \in A$, there is a variable $x(a)$. Defining $x(u, v)$ as the distance from node u to node v in the bidirected graph with arc-lengths $x(a)$, the relaxation minimizes $\sum_a c_a x(a)$ subject to the constraint that $x(s_i, s_j) \geq 1$ for all $1 \leq i < j \leq k$.

It is simple to see that an optimal solution to BiDir is a lower bound on the value of the multiway cut, and that the integrality gap of BiDir is no more than that of DistLP. In fact, there was some hope that this LP relaxation would yield algorithms comparable to or even better than the CKR relaxation. In particular, one of the open problems posed by Vazirani in [10, Chapter 30] asks whether the CKR relaxation and the BiDir relaxation are related in some way. Implicit in this open problem is the question about the integrality gap of the BiDir relaxation. In this note we obtain several results about BiDir, in particular relating it to the CKR relaxation. Our results are summarized below.

- For any permutation of the terminals, the optimal value of the BiDir relaxation is no more than that of the CKR relaxation.
- The value of the BiDir relaxation depends on the permutation on the terminals, even when G is a tree.
- There exist instances on which, for every permutation of the terminals, the value of BiDir is *strictly* larger than the value of the DistLP relaxation.
- There are instances on which, for *every* permutation of the terminals, the value of BiDir is *strictly* smaller than the value of the CKR relaxation, thus showing that the two formulations are *not* equivalent.
- To further disambiguate the two relaxations, we can show that for $k = 4$, there exists an instance and a permutation for which the integrality gap of the BiDir relaxation is $6/5$, which is strictly larger than an upper bound on the integrality gap of the CKR relaxation for $k = 4$ (1.1539), shown in [8].

To summarize, we show that for any instance of Multiway Cut, the following holds true.

$$\mathsf{DistLP} \leq \mathsf{BiDir} \leq \mathsf{CKR} \leq \mathsf{OPT}$$

Furthermore, there are instances on which each of the above inequalities is strict. Thus we have essentially shown that the BiDir relaxation is unlikely to yield an improved approximation ratio for the Multiway Cut problem.

2 Notation and Preliminaries

Let $G = (V, E)$ be the undirected graph, and let A be the set of arcs formed by bidirecting the edges in E. Let $H = (V, A)$ be the resulting directed graph. For an arc $a = (u, v) \in A$, its reverse will be denoted by $\bar{a} = (v, u)$. We use $\{u, v\}$ to denote the edge joining u and v in G to distinguish it from the arcs (u, v) and (v, u) in H. Let $s_1, s_2, \ldots, s_k$ be the terminals ordered in some way. Then the BiDir relaxation with respect to this ordering is the following. For each arc $a \in A$ we have a variable $x(a)$. Let $\mathcal{P}$ denote the set of all directed paths in H from s_i to s_j, $i < j$.

$$\begin{array}{rll} & \min \sum_a c_a x(a) & \\ \text{s.t.} & \sum_{a \in p} x(a) \geq 1 & p \in \mathcal{P} \\ & x(a) \geq 0 & a \in A \end{array}$$

The dual of the above is a multicommodity flow problem. For each $p \in \mathcal{P}$ we have a variable $f(p)$ which is the flow on p.

$$\begin{array}{rll} & \max \sum_p f(p) & \\ \text{s.t.} & \sum_{p: p \ni a} f(p) \leq c_a & a \in A \\ & f(p) \geq 0 & p \in \mathcal{P} \end{array}$$

Let x^* and f^* denote optimal solutions to the primal and the dual, and $x^*(p) \doteq \sum_{a \in p} x^*(a)$ denote the length of a directed path p. Complementary slackness implies the following:

- if $x^*(a) > 0$, the arc a is used to capacity in f^*.
- if $f^*(p) > 0$, then $x^*(p) = 1$; in other words on a path with non-zero flow, the length is 1.

3 Properties of the Bidirected LP

We begin by showing that the value of the bidirected relaxation depends on the permutation on the terminals. This is contrary to the suggestion in [10].

Theorem 3.1 *The value of the bidirected relaxation for the* MULTIWAY CUT *problem is permutation dependent even when the input graph is a tree.*

Proof: Consider the instance in Fig. 1(a), where each edge has cost $c_e = 1$. The solution x^* assigns each of the arcs labeled a through e a value of $\frac{1}{2}$ and assigns zero to the rest of the arcs, giving us a total value of 2.5. To see that this is optimal, we set up a dual multiflow of value 2.5. The flow-paths are from s_1 to s_3, s_1 to s_2, s_2 to s_3, s_2 to s_4, and from s_3 to s_4, each carrying 1/2 unit of flow. The complementary slackness conditions can also be verified.

In fact, we can also show that this solution x^* is basic. Suppose x^* can be written as the convex combination of two solutions, and let x be one of these solutions that assigns values a through e to the arcs shown in Figure 1(a). (Note that the values given to the rest of the arcs must be 0.) Since x must also be optimal, it satisfies the complementary slackness conditions with respect to the optimal dual solution given above. Thus, we get a system of five linear equations, one for each of the five flow paths in the dual solution. (For example, the path from s_1 to s_3 gives us $a + c = 1$.) Solving this system yields $a = b = c = d = e = \frac{1}{2}$, and hence $x = x^*$.

If, however, we renumbered the terminals and changed the order of s_2 and s_3 in Fig. 1(a) to get the graph in Fig. 1(b), we would change the value of the linear program. A generic optimal feasible solution is shown;

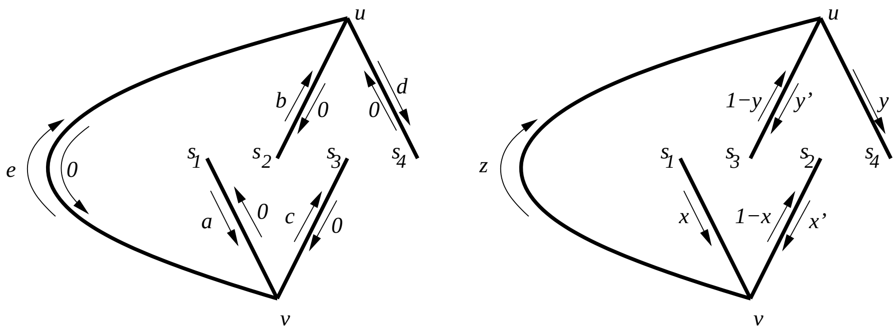

(a) Instance with $\mathsf{BiDir} < \mathsf{OPT}$ (b) Instance with $\mathsf{BiDir} = \mathsf{OPT}$

Figure 1: *Basic fractional solutions with different permutations.*

the arrows not shown are set to 0. The value of this optimal solution is $2 + x' + y' + z$. However, feasibility must ensure that $\min\{x, x'\} + z + \min\{y, y'\} \geq 1$, which implies that the LP value is at least 3, as opposed to 2.5 earlier. Thus the LP value *does* depend on the permutation used. ∎

Note that the optimal solution to this instance involves cutting at least 3 edges, and hence the integrality gap of the LP for the permutation in Fig. 1(a) is at least $\frac{6}{5}$. Note that this is worse than the integrality gap of the CKR relaxation for $k = 4$ which from [8], is no worse than 1.1539.

We also note that the support of a basic solution to BiDir can have *digons*: i.e., pairs of anti-parallel edges. An instance that demonstrates this can be obtained from the one in Fig 1(a) by adding $(u, s_1), (u, s_3), (v, s_2)$, and (v, s_4) to the edge set and assigning a cost of 2 to each of these edges. A basic solution to this instance is the following: assign 1 to the arcs (s_1, v), (u, s_4) and 0.5 to arcs (s_1, u), (s_2, u), (s_2, v), (s_3, v), (u, s_2), (u, s_3), (u, v), (v, s_3), (v, s_4), and 0 to the rest. In this solution we have a digon (s_2, u) and (u, s_2).

A worse integrality gap The dependence of the value of the bidirected LP on the permutation can actually be made worse. We now show that the complete binary tree B_t with t levels, the 2^t leaves being terminals and each edge costing $c_e = 1$, has an integrality gap approaching $\frac{4}{3}$ for a certain permutation of the terminals.

Given a set of terminal names $N = s_{i_1}, s_{i_2}, \ldots, s_{i_k}$ and a complete binary tree with k leaves, the leaves are labeled inductively: the labels are divided into two sets with $L = s_{i_1}, s_{i_3}, \ldots, s_{i_{k-1}}$ being the names at the "odd" places in the set, and $R = N \setminus L$ being the rest of the labels. We then recursively label the left subtree with L and the right one with R. The base case is when we have a single vertex and a single label, in which case the labeling is trivial. E.g., given B_2 and $\{s_1, s_2, s_3, s_4\}$, the labeling will be s_1, s_3, s_2, s_4 from left to right. The reader will notice that this gives an instance isomorphic to the one in Theorem 3.1.

We now give a feasible fractional solution: consider any two leaf nodes s_i and s_j (with $i < j$) that have a common parent u. We assign $x(s_i, u) = x(u, s_j) = \frac{1}{2}$. For each internal node u with left child v and right child v', we assign $x(v, u) = x(u, v') = \frac{1}{4}$. It is easy to check that this is a feasible solution for the bidirected LP. However, the optimal integer solution to this instance is at least $2^t - 1$, since each edge deletion increases the number of components by 1. However, the fractional solution has value $\mathsf{BiDir} = 2^t \cdot \frac{1}{2} + (2^t - 1) \cdot \frac{1}{4}$, and thus the integrality gap is $\geq \frac{4}{3}$. ∎

Labeling the 2^t leaves from left to right gives an instance with an integrality gap of 1. Indeed, a dual multiflow solution with value $2^t - 1$ just sends one unit of flow from each leaf to the next one along the unique shortest path in the binary tree; this uses each arc exactly once. This shows that a poor choice of permutation can affect the LP value by as much as a factor of $\frac{4}{3}$.

Optimal solutions to the bidirected LP, irrespective of the permutation, satisfy some useful inequalities that are given below. We make the input graph complete by adding zero weight arcs. This does not change the value of either the integral solution or the fractional solution to BiDir, since we can set $x(u, v)$ for any zero-weight arc (u, v) to be the length of the shortest u-v path with respect to x. Whenever we refer to optimal solutions in the sequel, we will assume these solutions to be *minimal* in the sense that no $x(u, v)$ can be reduced without making the solution infeasible.

Lemma 3.2 *If x is an optimal solution, then $x(a) + x(\bar{a}) \leq 1$ for all $a = (u, v) \in A$.*

Proof: Since $x(a)$ cannot be reduced, there must be terminals s_p, s_q with $p < q$ such that $x(s_p, u) + x(a) + x(v, s_q) = 1$. Similarly, there must be s_r, s_t (with $r < t$) with $x(s_r, v) + x(\bar{a}) + x(u, s_t) = 1$. However, now at least one of $p < t$ or $r < q$ must hold true. Without loss of generality, it is the former (the other case is identical) – then $x(s_p, u) + x(u, s_t) \geq 1$, and hence $x(a) + x(\bar{a}) \leq 2 - x(s_p, u) - x(u, s_t) \leq 1$, proving the lemma. ∎

For any $v \in V$ and any $1 \leq i < k$, the constraints of the LP imply that the $x(s_i, v) + x(v, s_{i+1}) \geq 1$. The lemma below shows equality for an optimal solution.

Lemma 3.3 *Let x be an optimal solution. Then for any $v \in V$ and any $1 \leq i < k$, $x(s_i, v) + x(v, s_{i+1}) = 1$.*

Proof: Let $x(s_i, v) = a$. Since this is an optimal solution, there must be some $l > i$ such that $x(v, s_l) = 1 - a$, else $x(s_i, v)$ can be decreased. Consider $i < j < l$: by feasibility, $x(s_j, v) \geq 1 - x(v, s_l) \geq a$, and $x(v, s_j) \geq 1 - x(s_i, v) \geq 1 - a$. However, by Lemma 3.2, $x(v, s_j) + x(s_j, v) \leq 1$, thus implying that $x(v, s_j) = 1 - a$. Now plugging in $j = i + 1$ proves the lemma. ∎

Let us define $y(\{u, v\}) = x(u, v) + x(v, u)$ for all edges $\{u, v\} \in E$. From Lemma 3.2, $0 \leq y(\{u, v\}) \leq 1$ for all $\{u, v\} \in E$. Thus y could be interpreted as assigning lengths to the edges of the undirected graph G. It is easy to check that y is feasible for DistLP and hence DistLP $\leq$ BiDir for all instances. The following corollary to Lemma 3.3 shows that y satisfies a property that is not necessarily satisfied by solutions to DistLP.

Corollary 3.4 *For any feasible solution x to* BiDir *and for any $v \in V$,*

$$\sum_{i=1}^{k} y(\{s_i, v\}) \geq k - 1 \tag{3.1}$$

with equality when x is an optimal solution.

Proof: We expand $\sum_{i=1}^{k} y(\{s_i, v\})$ as $x(v, s_1) + \sum_{i=1}^{k-1}(x(s_i, v) + x(v, s_{i+1})) + x(s_k, v)$. The two terms $x(v, s_1)$ and $x(s_k, v)$ are non-negative, and the other $k - 1$ terms are at least 1 each, since x is feasible. If x is optimal, then the terms $x(v, s_1)$ and $x(s_k, v)$ can be set to 0 since doing so does not violate any of the constraints of the LP. Furthermore, Lemma 3.3 ensures that each of the other terms in the sum is exactly 1. ∎

On some instances $\mathsf{DistLP} < \mathsf{BiDir}$. Consider a star rooted at r with unit weight edges connecting r to k leaves $s_1, s_2, \ldots, s_k$; the leaves are the terminals. The optimum integral solution to this instance is $k-1$. From Corollary 3.4 we see that for any optimum solution x to BiDir, $\sum_{i=1}^{k} y(\{r, s_i\}) = k-1$. Since this sum equals the integral optimum for the graph in question, the cost of the BiDir optimum is $k-1$. On the other hand, the optimum value of DistLP for this instance is $k/2$, obtained by assigning each edge a length of $1/2$.

The following lemma characterizes the requirements on a feasible solution y to DistLP that corresponds to some feasible solution x for BiDir. It will be useful later in comparing solutions to BiDir and CKR.

Lemma 3.5 *Let d be a semi-metric on $V(G)$ with $0 \le d(e) \le 1$ which satisfies $\sum_i d(u, s_i) = k-1$ for all $u \in V$. Then there is a feasible solution x to the bidirected LP (with the identity permutation) in which $y(\{u,v\}) \le d(u,v)$ for all $u, v \in V$ if and only if for all $u, v \in V$*

$$d(u,v) \ge \max_i \left\{ \sum_{j \le i} (d(v,s_j) - d(u,s_j)) \right\} + \max_i \left\{ \sum_{j \le i} (d(u,s_j) - d(v,s_j)) \right\}. \tag{3.2}$$

Proof: (IF) Suppose $d(u,v)$ satisfies conditions (3.2). Since $d(s_j,s_j) = 0$ and $\sum_i d(s_j,s_i) = k-1$, it must be the case that $d(s_j,s_i) = 1$ for all $i \neq j$. Let us now define $x(s_i,u)$ and $x(u,s_i)$ for all u and s_i.

For all $u \in V$ and i, $1 \le i \le k$, we define $x(s_i,u) = \sum_{j \le i} d(u,s_j) - (i-1)$. Since $d(u,s_j) \le 1$, the sum is at most i, and hence we get $x(s_i,u) \le 1$. Furthermore, $\sum_{j \le i} d(u,s_j) = (k-1) - \sum_{j > i} d(u,s_j)$, and $d(u,s_j) \le 1$ implies that the final sum is at most $k-i$. This ensures that $x(s_i,u) = \sum_{j \le i} d(u,s_j) - (i-1) \ge (k-1) - (k-i) - (i-1) \ge 0$.

Note that (1) $x(s_i,s_j) = 1$ if $i < j$ and 0 otherwise; (2) $x(s_i,u) \ge x(s_j,u)$ if $i < j$, and (3) $x(s_k,u) = 0$ for all u. Now define $x(u,s_i) = 1 - x(s_{i-1},u)$ for all u and $i > 1$, and $x(u,s_1) = 0$ for all u. Clearly, $x(u,s_i)$ also lies between 0 and 1. (Also note that this definition is consistent with the previous one when u is a terminal node.)

Finally, we define for all $u, v \in V$, $x(u,v) = \max_i\{x(s_i,v) - x(s_i,u)\}$. It is easy to see that this is an identity if we replace u or v by a terminal node s_j, and hence this definition is consistent with the previous one as well. Also note that $x(u,v) \ge 0$ because $x(u,v) \ge x(s_k,v) - x(u,s_k) = 0$. To finish the proof of the feasibility of x we need to establish that x satisfies the triangle inequality: $x(u,v) + x(v,w) \ge x(u,w)$ for all triples (u,v,w). From the definition of x it follows that for $1 \le i \le k$, $x(u,v) + x(v,w) \ge x(s_i,v) - x(s_i,u) + x(s_i,w) - x(s_i,v) = x(s_i,w) - x(s_i,u)$. Therefore $x(u,v) + x(v,w) \ge \max_i\{x(s_i,w) - x(s_i,u)\} = x(u,w)$. Now, defining $y(\{u,v\})$ as $x(u,v) + x(v,u)$, we see that $y(\{u,v\})$ is exactly same as the right hand side of (3.2), and hence at most $d(u,v)$.

(Only IF) We know that $x(s_i,u) + x(u,s_{i+1}) \ge 1$ for all i, $1 \le i < k$. Adding these inequalities and using the fact $x(u,s_1) = x(s_k,u) = 0$, we get $\sum_{i=1}^{k} (x(s_i,u) + x(u,s_i)) \ge k-1$. But, $x(s_i,u) + x(u,s_i) = y(\{u,s_i\}) \le d(u,s_i)$ and we know that $\sum_i d(u,s_i) = k-1$. This implies that $x(s_i,u) + x(u,s_{i+1}) = 1$ for all i, $1 \le i < k$, and $x(s_i,u) + x(u,s_i) = d(u,s_i)$ for all i.

Solving these equations, we get $x(s_i,u) = \sum_{j \le i} d(u,s_j) - (i-1)$ and $x(u,s_i) = 1 - x(s_{i-1},u)$. Also, since x can be assumed to satisfy the triangle inequality, we must have $x(u,v) \ge \max_i\{x(s_i,v) - x(s_i,u)\}$. Now using the fact that $x(u,v) + x(v,u) = y(\{u,v\}) \le d(u,v)$, we get inequality (3.2). ∎

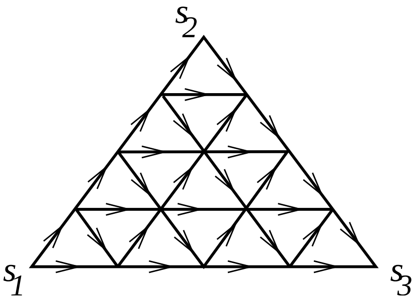

Figure 2: *How to orient aligned edges: the 3-terminal case.*

4 Relationship to the CKR relaxation

Călinescu, Karloff and Rabani [2] gave the following simplex-based relaxation for the MULTIWAY CUT problem: each vertex u is mapped to a point in the k-simplex with the terminals mapped to the vertices of the simplex. Each edge $e = \{u, v\}$ gets length μ_e that is half the ℓ_1 distance between the points corresponding to u and v. In the integral setting the mapping can be thought of as assigning for each vertex u, k non-negative values $z(u, i)$ such that $\sum_i z(u, i) = 1$. For each terminal s, exactly one i is chosen and $z(s, i)$ is set to 1. Note that the relaxation does not depend on how the terminals are mapped to the vertices of the simplex. Then for an edge $e = \{u, v\}$, $\mu_e = 1/2 \sum_i |z(u, i) - z(v, i)|$. The relaxation finds the mapping that minimizes $\sum_{e \in E} c_e \mu_e$.

The main result of this section is the following:

Theorem 4.1 *For any ordering π on the terminals, the value of the* CKR *relaxation is at least as large as that of the* BiDir *relaxation. Hence, the integrality gap of the* BiDir *relaxation is at least as large as that of the* CKR *relaxation.*

We present two proofs for this theorem, one more visual, the other more algebraic.

4.1 Proof I: By picture

Proof: Let $s_1, s_2, \ldots, s_k$ be the ordering of the terminals for the bidirected relaxation. We will set $z(s_i, i) = 1$ in the CKR relaxation. From the symmetry argument it suffices to argue that the value of the bidirected relaxation with this ordering is no more than the value of the CKR relaxation with the corresponding assignment.

Let (z, μ) be some feasible solution to the CKR relaxation. As shown in [2], we can assume without loss of generality that the solution z is *aligned* with respect to the simplex. This means that for each edge $e = \{u, v\}$, the coordinates of u and v differ in at most two coordinates. The edge e is ij aligned, $i < j$, if u and v differ in the i^{th} and j^{th} coordinates. Given z we construct a feasible solution x to the bidirected relaxation as follows. If u and v are mapped to the same point in the simplex, then they do not differ in any coordinate and we set $x(a) = 0$ for both the arcs (u, v) and (v, u). Otherwise let $e = \{u, v\}$ be ij aligned, $i < j$. For the arc $a = (u, v)$, we set $x(a) = \max\{0, z(u, i) - z(v, i)\}$. In other words $x(a) = \mu_e$ if $z(u, i) - z(v, i) \geq 0$ and 0 otherwise. Similarly for the arc $\bar{a} = (v, u)$ we set $x(\bar{a}) = \max\{0, z(v, i) - z(u, i)\}$. Note that $z(u, i) - z(v, i) = z(v, j) - z(u, j)$ since e is ij aligned. Hence for each edge $\{u, v\}$, exactly one of the arcs (u,v) and (v,u) gets length μ_e the other gets length 0. From this, it is obvious that the value of the constructed solution to the bidirected relaxation is exactly the same as that of the CKR relaxation. In other words each edge is oriented from the lower indexed terminal to the higher indexed terminal. See Figure 2 for an example.

Now we argue that the above assignment is feasible for the bidirected relaxation. Let P be a directed path from terminal s_i to s_j in the bidirected graph, where $i < j$. Let $x(P) = \sum_{a \in P} x(a)$ denote the length of P according to x. For a vertex u, let $\alpha(u,h) = \sum_{l \le h} z(u,l)$. We have the following lemma.

Lemma 4.2 *For any directed path P from u to v and for $1 \le i \le k$, $x(P) \ge \max\{0, \alpha(u,i) - \alpha(v,i)\}$.*

Proof: We fix an index ℓ, $1 \le \ell \le k$ and show that $x(P) \ge \max\{0, \alpha(u,i) - \alpha(v,i)\}$. The proof is by induction on the number of arcs in P. The base case is when P consists of a single arc (u,v). The edge $e = \{u,v\}$ is ij aligned for some $i < j$. Then it follows that $\alpha(u,h) - \alpha(v,h) = 0$ for all $h < i$ and $h \ge j$ and $\alpha(u,h) - \alpha(v,h) = z(u,i) - z(v,i)$ for $i \le h < j$. Further $\mu_e = 1/2(|z(u,i) - z(v,i)| + |z(u,j) - z(v,j)|) = |z(u,i) - z(v,i)|$. From the construction of x, $x(u,v) = \max\{0, z(u,i) - z(v,i)\}$. If $\ell < i$ or $\ell \ge j$, $\alpha(u,\ell) - \alpha(v,\ell) = 0$ and hence trivially $x(P) \ge \max\{0, \alpha(u,\ell) - \alpha(v,\ell)\}$. If $i \le \ell < j$, then $x(P) = x(u,v) = \max\{0, z(u,i) - z(v,i)\} = \max\{0, \alpha(u,\ell) - \alpha(v,\ell)\}$.

Suppose the induction hypothesis is true for all paths of length up to $t-1$ and let $P = P_{uv}$ be a path of length t. Let (u,w) be the first arc in P and let P_{wv} be the restriction of P from w to v. Since the number of edges in P_{wv} is $t-1$ we can apply the induction hypothesis to P_{wv} and hence $x(P_{wv}) \ge \max\{0, \alpha(w,\ell) - \alpha(v,\ell)\}$. If $\alpha(w,\ell) \ge \alpha(u,\ell)$ we are done. Hence we can restrict ourselves to the case in which $\alpha(u,\ell) > \alpha(w,\ell)$. In this case it must be that the edge $\{u,w\}$ is ij aligned such that $i \le \ell < j$ and $z(u,i) - z(w,i) = \alpha(u,\ell) - \alpha(w,\ell)$. It follows that $x(u,w) = \alpha(u,\ell) - \alpha(w,\ell)$. Therefore

$$\begin{aligned} x(P) &= x(u,w) + x(P_{wv}) \\ &= \alpha(u,\ell) - \alpha(w,\ell) + x(P_{wv}) \\ &\ge \alpha(u,\ell) - \alpha(w,\ell) + \max\{0, \alpha(w,\ell) - \alpha(v,\ell)\} \\ &\ge \max\{0, \alpha(u,\ell) - \alpha(v,\ell)\} \qquad (\text{since } \alpha(u,\ell) - \alpha(w,\ell) \ge 0) \end{aligned}$$

which finishes the induction step. ■

An immediate corollary to the above is the following.

Corollary 4.3 *For terminals s_i and s_j such that $i < j$, the shortest x-distance in the bidirected graph is at least 1.*

This shows that x is a valid solution for the bidirected LP, and hence proves the theorem. ■

A concise proof of Corollary 4.3 was suggested by an anonymous referee and we present it below.

Lemma 4.4 *Let P be a directed path from terminal s_i to s_j, $i < j$. Then $x(P) \ge 1$.*

Proof: For a vertex u, define $d_F(u) = \sum_{t=1}^{j-1} z(u,t)$ (the half ℓ_1 distance between u and F where F is the simplex facet containing $s_j, s_{j+1}, \ldots, s_k$). Consider how d_F changes as we traverse P from s_i to s_j. Initially we have $d_F(s_i) = 1$ and at the end $d_F(s_j) = 0$. For an arc $a = (u,v)$, let $\Delta_a = d_F(u) - d_F(v)$ denote the change in d_F value as we traverse the arc a. Let a be hl aligned for $h < l$. Recall that $z(u,h) - z(v,h) = z(v,l) - z(u,l)$. Therefore $d_F(u) - d_F(v) = 0$ if $h \ge j$ or $l < j$. If $h < j \le l$ then

$\Delta_a = d_F(u) - d_F(v) = z(u,h) - z(v,h)$ and hence $\Delta_a \le x(a) = \max\{0, z(u,h) - z(v,h)\}$. From this we obtain the the following:

$$d_F(s_i) - d_F(s_j) = 1 = \sum_{a \in P} \Delta_a \le \sum_{a \in P} x(a) = x(P).$$

■

4.2 Proof II: By algebra

Proof: Since the CKR relaxation does not depend on any ordering on the vertices, we can assume that π is the identity permutation, else we can permute the names of the vertices to get this. Our proof shows that the optimal metric μ given by the CKR LP can be massaged into a feasible solution to the bidirected LP with the same cost. The crucial fact we will use is that solutions to the CKR relaxation satisfy two useful inequalities, as pointed out in [2]:

$$\textstyle\sum_i \mu(u, s_i) = k - 1 \qquad \text{for all } u \in V \tag{4.3}$$

$$\mu(u,v) \ge \textstyle\sum_i \max\{\mu(u,s_i) - \mu(v,s_i), 0\}. \qquad \text{for all } u, v \in V \tag{4.4}$$

In fact, these two inequalities, along with the constraint that μ is restricted to being a metric, are *equivalent* to the CKR relaxation [2, Proposition 1].

Since this metric μ satisfies (4.3), it satisfies the conditions of Lemma 3.5, and we can get a solution with $y(e) \le \mu(e)$ (and hence with total cost at most that of the CKR relaxation) if μ satisfies (3.2); we now show this. Consider the two maxima in (3.2), and let them be attained at i_1 and i_2 respectively, with $i_1 \ge i_2$ (the other case is symmetric). Then

$$\begin{aligned}
&\textstyle\sum_{j \le i_1} (\mu(u,s_j) - \mu(v,s_j)) + \sum_{j \le i_2} (\mu(v,s_j) - \mu(u,s_j)) \\
&\quad = \textstyle\sum_{j \le i_1} (\max\{\mu(u,s_j) - \mu(v,s_j), 0\} - \max\{\mu(v,s_j) - \mu(u,s_j), 0\}) \\
&\qquad\quad + \textstyle\sum_{j \le i_2} (\max\{\mu(v,s_j) - \mu(u,s_j), 0\} - \max\{\mu(u,s_j) - \mu(v,s_j), 0\}) && (4.5) \\
&\quad \le \textstyle\sum_{j \le i_1} \max\{\mu(u,s_j) - \mu(v,s_j), 0\} \\
&\qquad\quad - \textstyle\sum_{j \le i_1} \max\{\mu(v,s_j) - \mu(u,s_j), 0\} + \sum_{j \le i_2} \max\{\mu(v,s_j) - \mu(u,s_j), 0\} && (4.6) \\
&\quad \le \textstyle\sum_{j \le i_1} \max\{\mu(u,s_j) - \mu(v,s_j), 0\} \le \sum_{j \le k} \max\{\mu(u,s_j) - \mu(v,s_j), 0\} \le \mu(u,v).
\end{aligned}$$

The first manipulations above ensure that all terms within the summations are positive. Then, in (4.5), we drop the final summation that was being subtracted off; in (4.6), we drop two summations, since one of them subtracts more than the other one adds. Finally, we use (4.4) for the last inequality, which finishes the proof of the theorem. ■

We can also show that the bidirected LP has a value that is *strictly lower* than the CKR relaxation for *any* permutation.

Theorem 4.5 *There exists a graph for which the optimal value of* BiDir *is* 8.5 *for any permutation* π *on the terminals, whereas the value of the* CKR *relaxation (as well as the optimal cut) is* 9.

Proof: Consider the graph in Figure 3. It is constructed by taking three copies of the tree in Figure 1(a), and identifying the terminals. (The different trees are indicated by different patterns on the edges.) This has been done in such a way that for any labeling of the resulting terminals, at least one of the three trees is isomorphic to the instance in Figure 1(a), i.e., for any assignment of labels $\{s_1, s_2, s_3, s_4\}$ to the terminal nodes (the figure shows one such assignment), at least one of the trees will have a common parent for s_1 and s_2.

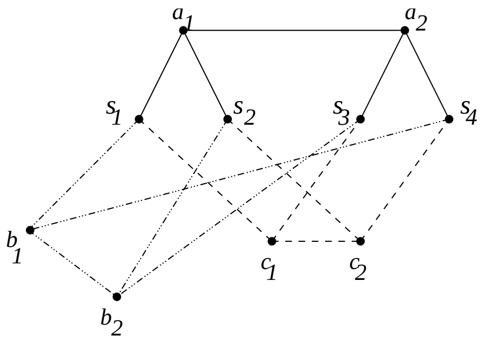

Figure 3: *Three trees glued together at the terminals.*

We first observe that for any setting of the arc lengths x, the solution is feasible for the bidirected relaxation on G if and only if it is feasible for each of the three trees. It is clearly feasible for each of the trees if it is feasible for G. Conversely, suppose x is feasible for the trees but not for G. Then there exist terminals s_i and s_j, $i < j$, and a directed path P from s_i to s_j such that $x(P) < 1$. Clearly, we can break this path into sub-paths such that each sub-path connects two terminals and stays inside one of the three trees. Since $i < j$, there must exist such a sub-path which connects two terminals s_l and s_p, where $l < p$. But then, the length of this path is at least 1, a contradiction.

If we restrict attention to one of the three trees, there is an integral solution of value 3 (cut the edges incident with three of the four terminals). So the BiDir relaxation must have value at most 3 on any of the three trees. Further, we know that at least one of the three trees has an instance isomorphic to the instance in Figure 1(a) – and in this case, the BiDir relaxation has value at most 2.5. Thus, the optimal BiDir value is at most 8.5.

On the other hand, a similar argument shows that μ is a solution to the CKR relaxation on G if and only if μ restricted to each of the trees is a solution to the CKR LP on the tree. Now we know that there exists a permutation for which BiDir is 3 on the trees, and hence, by Theorem 4.1, the value of the CKR formulation is at least 3. Hence the solution value of the CKR relaxation on G is at most 9. Since there is an integral solution of value 9 as well, it must be the case that the optimal solution to the CKR relaxation has value 9. This proves the theorem. ■

5 Discussion

In a paper by Erdős, Frank and Székely [4], another lower bound for MULTIWAY CUT was proposed. Let $\vec{G}$ be some orientation of the edges of G, and let $\lambda_{\vec{G}}(s)$ be the min-cut in this directed graph between a terminal s and $S - s$. Let us define $\nu = \max_{\vec{G}}\{\sum_{s \in S} \lambda_{\vec{G}}(s)\}$. The following theorem is proved in that paper:

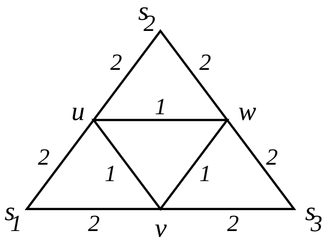

Figure 4: *The example for $k = 3$ from [4] and [6]. The numbers indicate edge costs.*

Theorem 5.1 (Erdős, Frank & Székely) *The quantity ν is a lower bound on the* MULTIWAY CUT*, and is exact when $G - S$ is a tree.*

This implies that example in Figure 1(a) must have $\nu = 3$, and since $\mathsf{BiDir} = 2.5$, this shows that $\nu > \mathsf{BiDir}$; conversely, it can be shown for that the graph in Fig. 4, $\nu = 7$ (see [4]) is less than $\mathsf{BiDir} = 7.5$ (see [10, Exercise 19.7]), showing that the two lower bounds are incomparable.

We can show that the CKR relaxation is also exact when $G - S$ is a tree. The proof is technical and we do not include it in this paper. Even though we have shown that the BiDir relaxation is weaker than the CKR relaxation, there are interesting questions about the relaxation that we have not answered. Is the integrality gap of the BiDir relaxation strictly smaller than 2 for all permutations? If so, what is the exact integrality gap? Can the permutation that yields the largest value for the BiDir relaxation be computed in polynomial time?

Acknowledgments: We are grateful to the anonymous reviewers for their extensive comments and also to one of them for suggesting the proof of Lemma 4.4.

References

[1] Y. Boykov, O. Veksler, and R. Zabih. Fast approximate energy minimization via graph cuts. In *Proceedings of the 7th IEEE International Conference on Computer Vision*, volume 1, pages 377–384, 1999.

[2] G. Călinescu, H. Karloff, and Y. Rabani. An improved approximation algorithm for MULTIWAY CUT. *J. Comput. System Sci.*, 60(3):564–574, 2000. A preliminary version appeared in *STOC*, 1998.

[3] E. Dahlhaus, D. S. Johnson, C. H. Papadimitriou, P. D. Seymour, and M. Yannakakis. The Complexity of Multiterminal Cuts. *SIAM J. Comput.*, 23(4): 864-894, 1994.

[4] P. L. Erdős, A. Frank, and L. Székely. Minimum multiway cuts in trees. *Discrete Appl. Math.*, 87(1-3):67–75, 1998.

[5] L. R. Ford, Jr. and D. R. Fulkerson. *Flows in Networks.* Princeton University Press, Princeton, NJ, 1973.

[6] A. Freund and H. Karloff. A lower bound of $8/(7+1/(k-1))$ on the integrality ratio of the Călinescu-Karloff-Rabani relaxation for multiway cut. *Information Processing Letters*, 75(1-2): 43-50, 2000.

[7] N. Garg, V. Vazirani, and M. Yannakakis. Multiway Cuts in Directed and Node Weighted Graphs. In *Proceedings of ICALP*, 487-498, 1994.

[8] D. Karger, P. Klein, C. Stein, M. Thorup, and N. Young. Rounding Algorithms for a Geometric Embedding of Minimum Multiway Cut. *Math. of Operations Research*, 29(3): 436-461, 2004. Preliminary version appeared in *Proc. of STOC*, 1999.

[9] H. S. Stone. Multiprocessor Scheduling with the Aid of Network Flow Algorithms. *IEEE Transactions on Software Engineering*, 3(1): 85-93, 1977.

[10] Vijay Vazirani. **Approximation Algorithms**. First Edition, Springer-Verlag, Berlin Heidelberg, 2001.