Approximation Algorithms for the 0-Extension Problem

Gruia Calinescu * Howard Karloff † Yuval Rabani‡

July 7, 2000

Abstract

In the 0-*extension problem*, we are given a weighted graph with some nodes marked as *terminals* and a semimetric on the set of terminals. Our goal is to assign the rest of the nodes to terminals so as to minimize the sum, over all edges, of the product of the edge's weight and the distance between the terminals to which its endpoints are assigned. This problem generalizes the multiway cut problem of Dahlhaus, Johnson, Papadimitriou, Seymour, and Yannakakis and is closely related to the metric labeling problem introduced by Kleinberg and Tardos.

We present approximation algorithms for 0-EXTENSION. In arbitrary graphs, we present an $O(\log k)$-approximation algorithm, k being the number of terminals. We also give $O(1)$-approximation guarantees for weighted planar graphs. Our results are based on a natural *metric relaxation* of the problem, previously considered by Karzanov. It is similar in flavor to the linear programming relaxation of Garg, Vazirani, and Yannakakis for the multicut problem and similar to relaxations for other graph partitioning problems. We prove that the integrality ratio of the metric relaxation is at least $c\sqrt{\lg k}$ for a positive c for infinitely many k. Our results improve some of the results of Kleinberg and Tardos and they further our understanding on how to use metric relaxations.

*Department of Computer Science, Illinois Institute of Technology, Chicago, IL 60616. Research performed while at DIMACS and Georgia Institute of Technology, and supported in part by a DIMACS postdoctoral fellowship and by NSF grant CCR-9732746. E-mail: calinesc@cs.iit.edu.

†AT&T Labs–Research, Room C231, 180 Park Ave., Florham Park, NJ 07932, and College of Computing, Georgia Institute of Technology, 801 Atlantic Dr., Atlanta, GA 30332-0280. Research supported in part by NSF grant CCR-9732746. E-mail: howard@research.att.com.

‡Computer Science Department, Technion — IIT, Haifa 32000, Israel. Work at the Technion supported by BSF grant number 96-00402, by Ministry of Science contract number 9480198, by grant number 386/99-1 of the Israel Science Foundation founded by the Israeli Academy of Sciences and Humanities, through the CONSIST consortium by the Ministry of Trade and Industry's MAGNET program, and by the Fund for the Promotion of Research at the Technion. E-mail: rabani@cs.technion.ac.il.

1 Introduction

Let V be a finite set, let $T \subseteq V$, and let d be a semimetric on T.[1] Then a semimetric δ on V is an *extension of d to V* iff for every $i, j \in T$, $\delta(i,j) = d(i,j)$. If, in addition, for every $i \in V$ there exists $j \in T$ such that $\delta(i,j) = 0$, then δ is a *0-extension of d to V*.

We consider the following optimization problem (denoted 0-EXTENSION): Given a graph $G = (V, E)$ with nonnegative edge weights $c : E \to \mathbb{R}$, a subset T of the nodes, and a semimetric d on T, find a 0-extension δ of d to V that minimizes the total weighted length of the edges.

It helps to compare 0-EXTENSION to the multiway cut problem of Dahlhaus, Johnson, Papadimitriou, Seymour, and Yannakakis [5, 6]. MULTIWAY CUT is the following problem: Given a graph $G = (V, E)$ with nonnegative edge weights $c : E \to \mathbb{R}$, and a subset $T \subseteq V$ of terminals, find a mapping $f : V \to T$ such that $f(t) = t$ for all $t \in T$, so as to minimize

$$\sum_{uv \in E, f(u) \neq f(v)} c(u, v).$$

In other words, find a set of edges of minimum total weight whose removal disconnects all terminal pairs. If we define d to be the uniform metric on T, i.e., $d(i,j) = 1$ if $i \neq j$ and $d(i,i) = 0$, then MULTIWAY CUT is exactly this problem: Find $f : V \to T$ with $f(t) = t$ for all $t \in T$, so as to minimize

$$\sum_{uv \in E} c(u, v) \cdot d(f(u), f(v)),$$

as $d(f(u), f(v)) = 1$ if $f(u) \neq f(v)$ and $d(f(u), f(v)) = 0$ otherwise. Now 0-EXTENSION is the natural generalization of MULTIWAY CUT in which, instead of being the uniform metric, d is an arbitrary semimetric on T. In other words, we must find an $f : V \to T$ with $f(t) = t$ for all $t \in V$, so as to minimize

$$\sum_{uv \in E} c(u, v) \cdot d(f(u), f(v)).$$

Dahlhaus et al. [6] show that MULTIWAY CUT (and therefore 0-EXTENSION) is APX-hard. Thus there exists a constant $\alpha > 1$ such that no polynomial-time algorithm can find a solution within a factor of α of the optimum, unless P=NP.

In this paper we develop approximation algorithms for the 0-extension problem. We study what seems to us to be the most natural linear programming relaxation for this problem. Specifically, given the semimetric d on T, extend d to a semimetric δ on the larger set V so as to minimize $\sum_{uv \in E} c(u, v)\delta(u, v)$. (We call this the *metric relaxation*.) Obviously, the set of feasible extensions δ is defined by $O(|V|^3)$ linear constraints, and the objective function is linear. Thus, finding the best extension is a linear programming problem, so it can be solved in polynomial time. We derive approximation algorithms using the metric relaxation, thus bounding also the integrality ratio for the relaxation. For arbitrary graphs we give a randomized, $O(\log |T|)$-approximation algorithm. We show that the integrality ratio is at least a constant times $\sqrt{\log |T|}$ for infinitely many $|T|$. We improve the upper bounds to $O(1)$ for (weighted) planar graphs (or, in fact, for any family of graphs that excludes a $K_{r,r}$-minor for some fixed r).

Karzanov [11] considers the metric relaxation for the 0-extension problem, and characterizes some of the cases in which the relaxation gives the optimal solution. For the multiway cut problem, it was known that the metric relaxation yields a $(2 - 2/|T|)$-approximation algorithm (see [2]). Indeed, this observation uses the same idea underlying the combinatorial algorithm of Dahlhaus et al. [6] that has the same performance guarantee. For the general case, the quality of the metric relaxation was not known prior to our work. For multiway cut, a different relaxation gives better approximations, with an asymptotic ratio significantly below 2 (see Calinescu, Karloff, and Rabani [2] and

[1] A function $d : T \times T \to \mathbb{R}$ is a *semimetric* on T iff for every $i_1, i_2, i_3 \in T$, $d(i_1, i_1) = 0$, $d(i_1, i_2) \geq 0$, $d(i_1, i_2) = d(i_2, i_1)$, and $d(i_1, i_2) + d(i_2, i_3) \geq d(i_1, i_3)$. If, in addition, $d(i_1, i_2) = 0$ implies $i_1 = i_2$, then d is a *metric*.

the improved bounds of Cunningham and Tang [3] and of Karger, Klein, Stein, Thorup, and Young [10]). It is not clear if the Calinescu et al. relaxation can be extended to handle the general case of 0-EXTENSION.

In a recent paper, Kleinberg and Tardos [14] give approximation algorithms for a similar problem of classification with pairwise relations, which they call METRIC LABELING. In their problem, the terminals are distinct from the vertices and are called *labels*. There is a semimetric on the labels, and for each node of the graph there is a vector of assignment costs to each of the labels. The goal is to minimize the total assignment cost plus the sum of weighted edge lengths. More formally, given a graph $G = (V, E)$ with nonnegative edge weights $c : E \to \mathbb{R}$, a set T of labels, a metric d on T, and a nonnegative assignment cost function $a : V \times T \to \mathbb{R}$, METRIC LABELING is the problem of finding a mapping $f : V \to T$ so as to minimize

$$\sum_{u \in V} a(u, f(u)) + \sum_{uv \in E} c(u, v) \cdot d(f(u), f(v)).$$

For the case that d is the uniform metric, Kleinberg and Tardos give a 2-approximation algorithm, based on a relaxation similar to the Calinescu et al. relaxation for multiway cut. Unlike the multiway cut problem, the integrality ratio for the relaxation here is at least $2 - 2/|T|$. Chuzhoy [4] improves their result for three and four terminals (achieving a tight $4/3$ bound for three terminals). Kleinberg and Tardos further give a constant approximation algorithm for a class of hierarchically well-separated tree metrics (HST). Following Bartal's small distortion embeddings of metrics into HST metrics [1], they use the result for trees to give an $O(\log |T| \log \log |T|)$-approximation algorithm for arbitrary metrics. Gupta and Tardos [9] later give a local search-based 4-approximation algorithm for the case that d is a truncated linear metric (i.e., $T = \{1, 2, \ldots, k\}$ and for some value m, $d(i, j) = \min\{|i - j|, m\}$).

Kleinberg and Tardos, and Gupta and Tardos motivate their work by several applications, mostly concerning computer vision, such as image restoration and visual correspondence. In these applications the nodes of the graph are pixels in a raster image and the edges model adjacency (so, in fact, the graph is a two-dimensional mesh). In image restoration applications, the labels model pixel intensities or colors. Assigning a label to a pixel amounts to determining the "true" intensity (or color) of the pixel from the observed values. The assignment cost penalizes for the difference between the observed and assigned intensity. In visual correspondence applications, the labels model possible shifts between two images. Assigning a label to a pixel amounts to determining the shift of that pixel between the two images. The assignment cost penalizes for the difference between the values of the supposedly matching pixels. In both types of applications, the structure of the graph arises from assuming that the *a priori* distribution of "true" labels is generated by a Markov random field (where the distribution of a pixel depends only on the distribution of its neighbors).

Notice that 0-EXTENSION is a special case of METRIC LABELING. (Set, for all $u \notin T$ and $i \in T$, $a(u, i) = \sum_{j \in T,\, j \neq i} c(u, j) \cdot d(i, j)$.) Thus, our results improve upon the results in [14] for this case. We note that the 0-extension formulation seems appropriate for many of the computer vision applications mentioned in [14, 9]. For example, if we connect each pixel by a weighted edge to the label corresponding to its observed intensity, we get an assignment cost proportional to the distance between the observed and assigned value. Our algorithm for weighted planar graphs, which actually only uses the fact that $V \setminus T$ induces a planar graph, is an alternative constant-ratio approximation algorithm for some of these computer vision problems.

Another problem related to ours is the multicut problem, first considered in the context of approximation algorithms in two papers by Garg, Vazirani, and Yannakakis [7, 8] (and implicitly in Klein, Agarwal, Ravi, and Rao [12]). In this problem, we are given a (weighted) graph and *k pairs* of terminals (nodes in the graph), and the goal is to find a set of edges of minimum weight whose removal disconnects every pair of terminals. This is a different generalization of multiway cut (the latter can be viewed as the multicut problem for all $\binom{k}{2}$ pairs of terminals). It is incomparable to the 0-extension problem, in the sense that neither problem is a special case of the other. In [8], Garg et al. give an $O(\log |T|)$ approximation algorithm for the multicut problem, based on a metric relaxation (assign lengths to edges so that the distance between every specified pair of terminals is at least 1). Their result is tight for the relaxation. The example achieving (asymptotically) the integrality ratio is an expander. For their upper bounds, they use a region-growing technique similar to that used by Leighton and Rao [15] for approximating the minimum flux (edge expansion) of a graph. Klein, Plotkin, and Rao [13] improve the Leighton-Rao technique for planar graphs (and more generally for graphs that exclude a $K_{r,r}$-minor), to get a constant factor approximation for the minimum flux. Using their technique, Tardos and Vazirani [17] exhibit a constant factor approximation algorithm for the multicut problem in planar graphs (and $K_{r,r}$-minor free graphs).

Our result can be seen as a counterpart to the Garg et al. and the Tardos-Vazirani results. The region-growing technique does not give a good approximation in the case of 0-extension. However, our results can be viewed as a form of (randomized) region growing after the application of a scaling function to the distances. This is implicit in the general case algorithm, and explicit in the planar graphs algorithm, where we use the Klein et al. technique on the scaled distances. It may be worth noting that expanders are *not* a particularly bad case for our relaxation (see Section 4).

The rest of the paper is organized as follows. In Section 2 we present the algorithm for the general case. Section 3 has the improved bounds for planar graphs. In Section 4 we discuss the quality of the linear programming relaxation underlying our approximation algorithms. Throughout the rest of the paper we use k to denote the number $|T|$ of terminals. We call the vertices in $V \setminus T$ *nonterminals*.

2 An $O(\log k)$-Approximation Algorithm

In this section we present the randomized algorithm which finds a 0-extension of weight at most $O(\log k)$ times the optimum. We begin by computing an optimal solution to the following natural linear programming relaxation, which we denote by (MET):

$$\text{Minimize} \sum_{uv \in E} c(u,v)\delta(u,v) \text{ subject to}$$

$$(V, \delta) \text{ is a semimetric} \tag{1}$$

$$\delta(i,j) = d(i,j) \qquad \forall i,j \in T. \tag{2}$$

If an assignment $f : V \to T$ defines an optimal solution to the 0-extension problem, then putting $\delta(u,v) = d(f(u), f(v))$ defines a feasible solution of (MET) of the same weight as the optimal solution. Therefore, the optimal value Z^* of (MET) is a lower bound on the minimum weight 0-extension. We exhibit a rounding procedure that takes any feasible solution δ of (MET) of value Z and constructs a 0-extension assignment $f : V \to T$ whose weight is $O(Z \log k)$.

Our rounding procedure works as follows. Pick uniformly at random a permutation σ of T and independently choose, uniformly at random in the interval $[1,2)$, a real number α. The rounding algorithm iteratively assigns some nodes to terminal σ_1, then some of the remaining nodes to terminal σ_2, and so on. For every $u \in V$, put $A_u = \min_{i \in T} \delta(u,i)$. The rounding procedure is given below.

The Rounding Procedure

Set $f(t) = t$ for all terminals t.
Pick a random permutation $\sigma = \langle \sigma_1, \sigma_2, \ldots, \sigma_k \rangle$ of the terminals.
Pick α uniformly at random in the interval $[1,2)$.
for $j = 1$ **to** k **do**
  for all unassigned nonterminals u such that $\delta(u, \sigma_j) \le \alpha A_u$, **do**
    set $f(u) = \sigma_j$ (i.e., assign u to σ_j).
  endfor
endfor

We first show that the rounding procedure produces a 0-extension:

Claim 1. The rounding procedure assigns every nonterminal to a terminal.

Proof. Consider a nonterminal v and let $t \in T$ be a terminal with $\delta(v,t) = A_u$. Choose j such that $t = \sigma_j$. Then, if v is not assigned to a terminal in iterations $1, 2, \ldots, j-1$, it must be assigned to t in iteration j, because $\alpha \geq 1$. ∎

For any pair of nodes $u, v \in V$, define a random variable $s(u,v) = d(f(u), f(v))$. We say that $u, v \in V$ are *separated* if $f(u) \neq f(v)$. Notice that if u, v are not separated, then $s(u,v) = 0$. Our goal is to bound the expectation $E[s(u,v)]$. We first state a bound on the probability that u, v are separated.

Lemma 2. Fix $u, v \in V$ and let $\delta = \delta(u,v)$. If $0 < \delta \leq \frac{1}{4}\min\{A_u, A_v\}$, then

$$\Pr[u, v \text{ are separated}] \leq 4\mathcal{H}_k \left(\frac{\delta}{A_u} + \frac{\delta}{A_v} \right),$$

where $\mathcal{H}_k = 1 + \frac{1}{2} + \cdots + \frac{1}{k}$ is the kth harmonic number.

Before we prove this bound, we state and prove its consequence, a bound for $E[s(u,v)]$ which is the core of the analysis of our algorithm.

Lemma 3. For any pair of distinct vertices $u, v \in V$, $E[s(u,v)] \leq 38\mathcal{H}_k \delta(u,v)$. [2]

Proof. Fix $u \neq v$ and put $\delta = \delta(u,v)$. By the triangle inequality, $A_v \leq A_u + \delta$ and $A_u \leq A_v + \delta$. As $\alpha \in [1,2)$, we have

$$s(u,v) = d(f(u), f(v)) = \delta(f(u), f(v)) \leq \delta(f(u), u) + \delta(u,v) + \delta(v, f(v)) \leq 2A_u + \delta + 2A_v. \tag{3}$$

If u is a terminal, then $A_u = 0$ and $A_v \leq \delta$. Therefore, by Inequality (3), $s(u,v) \leq 3\delta$ always. If both u and v are nonterminals and $\delta = 0$, then by the triangle inequality, for any terminal $j \in T$, $\delta(u,j) = \delta(v,j)$. Therefore, u and v are assigned to the same terminal, regardless of the choice of σ and α, so $s(u,v) = 0 = \delta$ always.

Thus, we may assume that both u and v are nonterminals, and that $\delta > 0$. We consider two cases, depending on whether δ is small compared to A_u or A_v, or not. If $A_u < 4\delta$, then $A_v < 5\delta$, and by Inequality (3), $s(u,v) < 2(4+5)\delta + \delta = 19\delta$. Similarly, if $A_v < 4\delta$, then $s(u,v) < 19\delta$. Therefore, if $A_u < 4\delta$ or $A_v < 4\delta$, the lemma holds.

Assume, therefore, that $\delta \leq \frac{1}{4}\min\{A_u, A_v\}$. (Recall that we also assume that $\delta > 0$.) We have

$$\begin{aligned}
E[s(u,v)] &\leq 4\mathcal{H}_k \left(\frac{\delta}{A_u} + \frac{\delta}{A_v} \right) (2A_u + 2A_v + \delta) \\
&\leq 4\mathcal{H}_k \delta \left(\frac{4A_u + 3\delta}{A_u} + \frac{4A_v + 3\delta}{A_v} \right) \\
&\leq 4\mathcal{H}_k \delta \left(4 + \frac{3}{4} + 4 + \frac{3}{4} \right) \\
&= 38\mathcal{H}_k \delta,
\end{aligned}$$

where the first inequality follows from Lemma 2 and Inequality (3). ∎

We conclude with

Theorem 4. There is a randomized, polynomial-time, $O(\log k)$-approximation algorithm for 0-EXTENSION.

Proof. Let δ^* be an optimal solution of (MET) of cost Z^*. By Lemma 3, the expected weight of the 0-extension obtained by the rounding procedure is $O(Z^* \log k)$. Therefore, there exists a choice of σ and of α that produces a solution of weight $O(Z^* \log k)$. To obtain a polynomial-time algorithm, notice that for a given permutation σ, two different values of α, $\alpha_1 > \alpha_2$, produce combinatorially distinct solutions only if there is a terminal j and a node u such that $\delta^*(u,j) \leq \alpha_1 A_u$ but $\delta^*(u,j) > \alpha_2 A_u$. Thus we can enumerate over at most $k|V|$ "interesting" values of α. We can determine these values by sorting the fractions $\delta^*(u,j)/A_u$, over all nodes u with $A_u > 0$ and over all $j \in T$. ∎

[2] The constant 38 is somewhat arbitrary, and definitely could be improved.

Proof of Lemma 2. Let $\mathcal{E}(u,v)$ denote the event that there is a terminal j such that when j is processed u is assigned to j whereas v remains unassigned; define $\mathcal{E}(v,u)$ similarly. We show that

$$\Pr[\mathcal{E}(u,v)] \le 4\mathcal{H}_k\delta/A_u. \tag{4}$$

By symmetry, $\Pr[\mathcal{E}(v,u)] \le 4\mathcal{H}_k\delta/A_v$. Therefore, the lemma follows from Inequality (4).

For $j = 1,2,\dots,k$, let $l_j = \delta(u,j)/A_u$ and let $r_j = \delta(v,j)/A_v$. For $\alpha \ge 1$, let

$$M(\alpha) = \{j \in T |\ l_j \le \alpha \text{ and } r_j > \alpha\},$$

and let

$$S(\alpha) = \{j \in T |\ r_j \le \alpha\}.$$

Notice that $M(\alpha)$ and $S(\alpha)$ are disjoint subsets of terminals. Put $m(\alpha) = |M(\alpha)|$ and $s(\alpha) = |S(\alpha)|$.

Assume that the rounding procedure uses $\alpha \in [1,2)$. Then, v must be assigned to a terminal in $S(\alpha)$, because $\delta(v,j) \le \alpha A_v$ is equivalent to $r_j \le \alpha$. Similarly, u cannot be assigned to a terminal which is not in $M(\alpha) \cup S(\alpha)$. Indeed, u can only be assigned to a terminal j with $l_j \le \alpha$, and based on whether $r_j \le \alpha$ or not, j is either in $S(\alpha)$ or in $M(\alpha)$. Moreover, $\mathcal{E}(u,v)$ happens if and only if the first terminal in $M(\alpha) \cup S(\alpha)$ to be processed is in $M(\alpha)$. Indeed, if the first such terminal is $j \in S(\alpha)$, then v (and possibly also u) will be assigned to j. Thus,

$$\Pr[\mathcal{E}(u,v) \mid \alpha] = \frac{m(\alpha)}{m(\alpha)+s(\alpha)}.$$

(Notice that there exists a j for which $r_j = 1$, so $S(\alpha)$ is never empty.) As α is distributed uniformly in $[1,2)$, we get

$$\Pr[\mathcal{E}(u,v)] = \int_1^2 \frac{m(\alpha)}{m(\alpha)+s(\alpha)} d\alpha. \tag{5}$$

We need the following claim:

Claim 5. Fix a positive integer k and let β be a nonnegative real. Let $\langle 1 = l_1 \le l_2 \le \cdots \le l_k\rangle$ be a nondecreasing sequence of reals. Let $\langle r_1, r_2, \dots, r_k\rangle$ be another sequence of reals such that $r_j \ge 1$ for all j, and for all $j \in \{1,2,\dots,k\}$, either $r_j - l_j \le \beta$ or $l_j > 2$. Define functions $m : [1,\infty) \to \mathbb{Z}$ and $s : [1,\infty) \to \mathbb{Z}$ as follows: $m(\alpha) = |\{j |\ l_j \le \alpha \text{ and } r_j > \alpha\}|$ and $s(\alpha) = |\{j |\ r_j \le \alpha\}|$. Then

$$\int_1^2 \frac{m(\alpha)}{m(\alpha)+s(\alpha)} d\alpha \le \mathcal{H}_k\beta.$$

Proof. Notice that $m(\alpha) \ge 1$ for $\alpha \in [1, r_1)$ and $s(\alpha) \ge 1$ for $\alpha \in [r_1, \infty)$, so the function we are integrating is well-defined. Let t be the largest index for which $l_t \le 2$ (so if $t < k$, $l_{t+1} > 2$). We prove by induction on t that the value of the integral is at most $\mathcal{H}_t\beta \le \mathcal{H}_k\beta$. For $t = 1$ the claim holds, because for $\alpha \in [1, r_1)$, $m(\alpha) = 1$ and $s(\alpha) = 0$, whereas for $\alpha \in [r_1, 2]$, $m(\alpha) = 0$ and $s(\alpha) = 1$. As $r_1 - 1 = r_1 - l_1 \le \beta$, we get

$$\int_1^2 \frac{m(\alpha)}{m(\alpha)+s(\alpha)} d\alpha \le \int_1^{r_1} 1 d\alpha \le \beta = \mathcal{H}_1\beta.$$

So, assume that the claim is true for the case that $t-1$ values l_j are at most 2, and consider sequences $\langle l_1, l_2, \dots, l_k\rangle$ and $\langle r_1, r_2, \dots, r_k\rangle$ where exactly t j's are such that $l_j \le 2$. We compare the value I of the integral in this case to the value I' of the integral for the sequences $\langle l_1, l_2, \dots, l_{t-1}, l_{t+1}, l_{t+1}, l_{t+2}, \dots, l_k\rangle$ and $\langle r_1, r_2, \dots, r_k\rangle$, except that if $t = k$, the first sequence becomes $\langle l_1, l_2, \dots, l_{k-1}, 3\rangle$. Notice that the latter sequences satisfy the conditions of the claim, and furthermore, only $t-1$ j's are such that $l_j \le 2$. Therefore, by the inductive hypothesis, $I' \le \mathcal{H}_{t-1}\beta$.

If $l_t \ge r_t$, then $I = I'$, so the claim holds in this case. Thus, we may assume that $l_t < r_t$. The only range where the integrands differ is $\alpha \in [l_t, \min\{r_t, 2\})$. In this range, for every $j < t$, $l_j \le \alpha$. Therefore, for the second pair of sequences, $m(\alpha) + s(\alpha) = t - 1$. For the original sequences, for $\alpha \in [l_t, \min\{r_t, 2\})$, $m(\alpha)$ is greater by 1, and $s(\alpha)$ is the same. The difference in the integrands is therefore

$$\frac{m(\alpha)+1}{m(\alpha)+1+s(\alpha)} - \frac{m(\alpha)}{m(\alpha)+s(\alpha)} = \frac{m(\alpha)+1}{t} - \frac{m(\alpha)}{t-1} \leq \frac{1}{t}.$$

Therefore, $I - I' \leq (\min\{r_t, 2\} - l_t)/t \leq (r_t - l_t)/t$. Because $l_t \leq 2$, $r_t - l_t \leq \beta$. Hence, $I \leq I' + \beta/t \leq \mathcal{H}_{t-1}\beta + \beta/t = \mathcal{H}_t\beta$. This completes the proof of Claim 5. ∎

We now proceed with the proof of Lemma 2. Notice that if $l_j \leq 2$, then $\delta(u,j) \leq 2A_u$. Using the assumption that $\delta \leq \frac{1}{4}A_u$, we have for such j,

$$r_j - l_j \leq \frac{\delta(u,j)+\delta}{A_v} - \frac{\delta(u,j)}{A_u} \leq \frac{\delta(u,j)+\delta}{A_u-\delta} - \frac{\delta(u,j)}{A_u} = \frac{\delta(\delta(u,j)+A_u)}{A_u(A_u-\delta)} \leq \frac{3\delta}{A_u-\delta} \leq \frac{4\delta}{A_u}. \tag{6}$$

Hence, using Claim 5 with $\beta = 4\delta/A_u$, we have

$$\int_1^2 \frac{m(\alpha)}{m(\alpha)+s(\alpha)} d\alpha \leq 4\mathcal{H}_k\delta/A_u. \tag{7}$$

Combining Equation (5) and Inequality (7), we get $\Pr[\mathcal{E}(u,v)] \leq 4\mathcal{H}_k\delta/A_u$, which proves Inequality (4) and thus the lemma. ∎

3 An $O(1)$-Approximation Algorithm for Planar Graphs

In this section we use the linear programming relaxation (MET) to get improved bounds for planar graphs. To achieve the improved bounds, we present a different rounding procedure. We show that if the input graph $G = (V, E)$ does not have a $K_{r,r}$-minor, then the rounding procedure presented in this section guarantees an $O(r^3)$ approximation ratio. As planar graphs are $K_{3,3}$-minor free, this gives a polynomial-time $O(1)$-approximation algorithm for planar graphs (and, more generally, for $K_{r,r}$-minor free graphs, for every fixed r).

The main tool that we use is the following theorem of Klein, Plotkin, and Rao [13] (the extension to the weighted case was stated by Tardos and Vazirani [17]).

Theorem 6 (Klein, Plotkin, and Rao). There are constants κ and λ and a polynomial-time algorithm KPR$(H, \delta, c, \gamma, r)$ which takes as input a graph $H = (V_H, E_H)$ with nonnegative integral edge lengths $\delta : E_H \to \mathbb{Z}$ and nonnegative edge costs $c : E_H \to \mathbb{Q}$, a positive rational γ, and a positive integer r, and finds either (1) a $K_{r,r}$-minor in H or (2) a set of edges of total c-cost at most $\kappa\frac{r}{\gamma}\sum_{e\in E_H}\delta(e)c(e)$ whose removal decomposes H into connected components such that the shortest path (in H, using edge lengths δ) between any two nodes in the same component is at most $\lambda r^2\gamma$. ∎

Let r be a positive integer. Let $\delta : V \times V \to \mathbb{R}$ be a feasible solution of (MET) of weight Z. Using Theorem 6, we exhibit a deterministic rounding procedure that obtains a 0-extension of weight $O(Z)$, assuming that the input graph G is $K_{r,r}$-minor free.

The main idea of the rounding procedure is to partition the nonterminals into clusters such that, for any two nodes u and v in the same cluster, A_u is at most twice A_v. We then assign all the nodes in a cluster to a terminal closest to one of the nodes in the cluster. More formally, the rounding procedure computes a 0-extension $f : V \to T$ as follows.

The Second Rounding Procedure

Set $f(t) = t$ for every terminal t.

for every nonterminal $u \in V$ such that $A_u = 0$, **do**
  Set $f(u) \leftarrow i$ for some $i \in T$ with $\delta(u,i) = 0$.
endfor
Let $\bar{G} = (\bar{V}, \bar{E})$ be the subgraph of G induced by the remaining nonterminals.
$\tilde{\delta}_{\min} \leftarrow \min\{\delta(u,v)/\max\{A_u, A_v\} \mid uv \in \bar{E},\ \delta(u,v) > 0\}$.
for every edge $uv \in \bar{E}$, **do**
  $\tilde{\delta}(uv) \leftarrow \lceil \delta(u,v)/(\tilde{\delta}_{\min} \cdot \max\{A_u, A_v\}) \rceil$
  $\tilde{c}(uv) \leftarrow c(u,v) \cdot \max\{A_u, A_v\}$.
endfor
Execute KPR$(\bar{G}, \tilde{\delta}, \tilde{c}, 1/(2\lambda r^2 \tilde{\delta}_{\min}), r)$.
for each resulting cluster $C \subseteq \bar{V}$, **do**
  Choose $x \in C$ to minimize A_x.
  Choose $i \in T$ such that $\delta(x,i) = A_x$.
  Set $f(u) \leftarrow i$ for all $u \in C$.
endfor

We first establish a few simple facts about this rounding procedure. Let $\tilde{Z} = \sum_{uv \in \bar{E}} \tilde{\delta}(uv)\tilde{c}(uv)$.

Claim 7. $\tilde{Z} \leq 2Z/\tilde{\delta}_{\min}$.

Proof. Notice that if $\delta(u,v) > 0$, then $\delta(u,v)/(\tilde{\delta}_{\min} \cdot \max\{A_u, A_v\}) \geq 1$; hence $\tilde{\delta}(uv) = \lceil \delta(u,v)/(\tilde{\delta}_{\min} \cdot \max\{A_u, A_v\}) \rceil \leq 2\delta(u,v)/(\tilde{\delta}_{\min} \cdot \max\{A_u, A_v\})$. We have

$$\begin{aligned}
\sum_{uv \in \bar{E}} \tilde{\delta}(uv)\tilde{c}(uv) &= \sum_{uv \in \bar{E}} \left\lceil \frac{\delta(u,v)}{\tilde{\delta}_{\min} \cdot \max\{A_u, A_v\}} \right\rceil \cdot (c(u,v) \cdot \max\{A_u, A_v\}) \\
&\leq \frac{2}{\tilde{\delta}_{\min}} \sum_{uv \in \bar{E}} \delta(u,v)c(u,v) \\
&\leq \frac{2}{\tilde{\delta}_{\min}} \sum_{uv \in E} \delta(u,v)c(u,v) \\
&= 2Z/\tilde{\delta}_{\min}. \qquad \blacksquare
\end{aligned}$$

Claim 8. The total $\tilde{c}$-cost of the edges removed by KPR$(\bar{G}, \tilde{\delta}, \tilde{c}, 1/(2\lambda r^2 \tilde{\delta}_{\min}), r)$ is $O(r^3 Z)$, and each of the resulting clusters C has $\tilde{\delta}$-diameter at most $1/(2\tilde{\delta}_{\min})$.

Proof. By Theorem 6, the sum of $\tilde{c}(uv)$ over edges uv with u,v in different clusters is at most $(\kappa r/\gamma)\tilde{Z}$. By Claim 7, this is at most $4\kappa\lambda r^3 Z$, which is $O(r^3 Z)$. Also, by Theorem 6, the $\tilde{\delta}$-diameter of each resulting cluster C is at most $1/(2\tilde{\delta}_{\min})$. $\blacksquare$

We now relate the $\tilde{\delta}$-distances to the original δ-distances.

Lemma 9. Let $u, v \in \bar{V}$. If the length of a shortest path in $\bar{G}$ between u and v with respect to edge lengths $\tilde{\delta}$ is at most $1/(2\tilde{\delta}_{\min})$, then $\delta(u,v) \leq A_u$.

Proof. Let $\langle u = x_0, x_1, \ldots, x_j = v \rangle$ be a shortest path in $\bar{G}$ between u and v with respect to the edge lengths $\tilde{\delta}$. For $1 \leq t \leq j$, let $s_t = \sum_{i=1}^{t} \delta(x_{i-1}, x_i)$. By the triangle inequality, $A_{x_t} \leq A_u + \sum_{i=1}^{t} \delta(x_{i-1}, x_i) = A_u + s_t$. Notice that $s_1 \leq s_2 \leq \cdots \leq s_j$. Therefore, for $i \leq j$, $\frac{\delta(x_{i-1},x_i)}{\max\{A_{x_{i-1}}, A_{x_i}\}} \geq \frac{\delta(x_{i-1},x_i)}{A_u + s_i} \geq \frac{\delta(x_{i-1},x_i)}{A_u + s_j}$. Also, $\tilde{\delta}(x_{i-1}x_i) \geq \frac{\delta(x_{i-1},x_i)}{\max\{A_{x_{i-1}}, A_{x_i}\}\tilde{\delta}_{\min}}$, and therefore $\frac{\delta(x_{i-1},x_i)}{\max\{A_{x_{i-1}}, A_{x_i}\}} \leq \tilde{\delta}_{\min}\tilde{\delta}(x_{i-1}x_i)$. Using this, we have $\frac{s_j}{A_u + s_j} = \sum_{i=1}^{j} \frac{\delta(x_{i-1},x_i)}{A_u + s_j} \leq \tilde{\delta}_{\min} \sum_{i=1}^{j} \tilde{\delta}(x_{i-1}x_i) \leq \tilde{\delta}_{\min} \frac{1}{2\tilde{\delta}_{\min}} = \frac{1}{2}$, where the last inequality follows from the hypothesis that the length of $\langle x_0, x_1, \ldots, x_j \rangle$ is at most $1/(2\tilde{\delta}_{\min})$. We conclude that $s_j \leq A_u$. Finally, notice that by the triangle inequality $\delta(u,v) \leq s_j$. $\blacksquare$

We are ready to analyze the performance of the rounding procedure.

Lemma 10. Let $r > 0$ be an integer. Then for every input graph G which is $K_{r,r}$-minor free, for every feasible solution to (MET) of weight Z, the above rounding procedure produces a 0-extension of weight $O(r^3 Z)$.

Proof. Let $uv \in E$ be an edge of G. If both endpoints $u, v \notin \bar{V}$, then each endpoint is either a terminal or a node at distance 0 from some terminal; hence $d(f(u), f(v)) = \delta(u, v)$.

If $u \notin \bar{V}$ and $v \in \bar{V}$, then $\delta(f(u), u) = 0$, and v, together with the cluster C that contains it, is assigned to some terminal i. By the definition of the rounding procedure, there is a node $x \in C$ such that $\delta(x, i) = A_x \le A_v$. Combining Claim 8 and Lemma 9, we have $\delta(v, x) \le A_v$. Therefore, using the triangle inequality, $\delta(v, i) \le 2A_v$. Using the triangle inequality again, $d(f(u), f(v)) = \delta(f(u), f(v)) \le \delta(f(u), v) + \delta(v, i) \le \delta(f(u), v) + 2A_v$. However, $\delta(u, v) = \delta(f(u), v) \ge A_v$. Therefore, $d(f(u), f(v)) \le 3\delta(u, v)$.

We are left with the edges $uv \in \bar{E}$. For $u \in \bar{V}$, let $C(u)$ denote the cluster containing u. Then

$$\sum_{uv \in \bar{E}} d(f(u), f(v)) c(u, v) = \sum_{uv \in \bar{E}\colon\, C(u) \neq C(v)} d(f(u), f(v)) c(u, v).$$

However, $d(f(u), f(v)) = \delta(f(u), f(v)) \le \delta(f(u), u) + \delta(u, v) + \delta(v, f(v)) \le \delta(u, v) + 2A_u + 2A_v \le \delta(u, v) + 4\max\{A_u, A_v\}$. The second inequality follows from the fact that by the definition of the algorithm, for every nonterminal u, $f(u)$ is a terminal closest to some $x \in C(u)$ with $\delta(f(u), x) = A_x \le A_u$. As argued in the previous case, $\delta(x, u) \le A_u$. The inequality follows because by the triangle inequality, $\delta(f(u), u) \le \delta(f(u), x) + \delta(x, u)$. Therefore,

$$\begin{aligned} \sum_{uv \in \bar{E}\colon\, C(u) \neq C(v)} d(f(u), f(v)) c(u, v) &\le \sum_{uv \in \bar{E}\colon\, C(u) \neq C(v)} \delta(u, v) c(u, v) + 4 \sum_{uv \in \bar{E}\colon\, C(u) \neq C(v)} \tilde{c}(uv) \\ &\le Z + 4 \sum_{uv \in \bar{E}\colon\, C(u) \neq C(v)} \tilde{c}(uv), \end{aligned}$$

Now Claim 8 states that $\sum_{uv \in \bar{E}\colon\, C(u) \neq C(v)} \tilde{c}(uv)$ is $O(r^3 Z)$. ■

We conclude with the main result of this section:

Theorem 11. Let $r > 0$ be a fixed integer. There is a deterministic polynomial-time $O(1)$-approximation algorithm for 0-EXTENSION in $K_{r,r}$-minor free weighted graphs.

Proof. Solve (MET) optimally and then use the rounding procedure from this section, which clearly can be implemented in polynomial time. Lemma 10 establishes the performance guarantee of this algorithm. ■

4 The Integrality Ratio

In this section we use the max flow-min cut theorem to prove the following lower bound on the integrality ratio of the natural relaxation.

Theorem 12. There are $c > 0$ and infinitely many positive integers k such that there is an instance of 0-EXTENSION with k terminals for which the optimal value of the objective function is at least $c\sqrt{\lg k}$ times the optimal value of the relaxation.

Proof. There are reals $\alpha > 0$ and Δ such that there are infinitely many graphs G (known as expanders) such that the maximum degree in G is at most Δ and such that for any subset S of at most $|V(G)|/2$ nodes, there are at least $\alpha|S|$ nodes not in S which are adjacent to at least one node of S. In fact, given a prime $q \ge 13$ which is congruent to 1 mod 4 (of which there are infinitely many), Lubotzky, Phillips, and Sarnak construct a 6-regular expander with either $q(q^2 - 1)$ or $q(q^2 - 1)/2$ nodes. (Take $p = 5$ in [16, pp. 262-263].) Given sufficiently large q, we use an expander $G = (V, E)$ having $n = |V| \in \{q(q^2 - 1)/2, q(q^2 - 1)\}$ nodes; we need $n \ge \max\{n_0, 16\}$ (n_0 to be determined later). Define $l = \lceil \sqrt{\lceil \lg n \rceil} \rceil \le n$ and $k = \lceil n/l \rceil$. Choose any k distinct nodes $h_1, h_2, ..., h_k$ in V. For $i = 1$ to k, add l new nodes and l new edges to the current graph, so as to get a new path P_i starting at h_i and ending at some new node; label that new node i. When finished, let the new graph be $G' = (V', E')$. Now $n' = |V'| = n + kl \le n + (1 + n/l) \cdot l \le n + (n + l) \le n + 2n = 3n$ vertices. $|E'| = |E| + kl \le n(\Delta/2 + 2)$.

Now define an instance I of 0-EXTENSION as follows. The vertex set is V'. The set T of terminals is $\{1, 2, 3, ..., k\}$. Define $d(i, j)$, for terminals $1 \le i, j \le k$, to be the distance in G' between i and j. Define $c(u, v)$ to be 1 if $uv \in E'$ and $c(u, v) = 0$ otherwise.

We now show that the integrality ratio for this instance I is large. First, we study the relaxation. Define $\delta(u, v)$ to be the distance in G' between u and v. It is clear that $\delta(i, j) = d(i, j)$ if i, j in T. It is also clear that δ is a semimetric on G'. It follows that $\sum_{u<v} c(u, v)\delta(u, v) = |E'| \le (\Delta/2 + 2)n$ (since adjacent vertices in G' are at distance 1).

Now we prove that there is a universal $c > 0$ such that any feasible solution to I, i.e., function $f : V' \to T$ with $f(t) = t$ for all $t \in T$, satisfies $\sum_{u<v} c(u, v)d(f(u), f(v)) \ge cn\sqrt{\lg n}$. Note first of all that the minimum distance between two distinct terminals i, j is at least $2l \ge 2\sqrt{\lg n}$. We will see below, however, that there are at least $k/2$ terminals for which the distance to terminal i^* is at least $\epsilon \lg n$, ϵ a fixed positive constant, not just $\sqrt{\lg n}$, for any $i^* \in T$.

We need a lemma.

Lemma 13. For any $i^* \in T$, there are at least $k/2$ vertices h_i in G whose distance from h_{i^*} exceeds $a = \lceil \frac{\lg k}{2 \lg \Delta} \rceil$.

Proof. For a contradiction, assume that there are more than $k/2$ vertices of G at distance at most a from h_{i^*}. By the degree bound, the number of vertices in G at distance at most a from h_{i^*} is at most $1 + \Delta^1 + \Delta^2 + \cdots + \Delta^a < \Delta^{a+1}$. Hence $k/2 < \Delta^{a+1}$. Hence $-1 + \frac{-1+\lg k}{\lg \Delta} < a$, $\lg k < 2 + 4 \lg \Delta$, and $k < 4\Delta^4$. Now choose n_0 large enough that k, which is a function of n, must be at least $4\Delta^4$, giving us a contradiction. ■

Now we show that without loss of generality, every vertex in $V' - V$ in P_i maps to i.

Lemma 14. Given $f : V \to T$ with $f(i) = i$ for all $i \in T$, set $f' : V \to T$ by $f'(v) = f(v)$ for all v, except that for all $i \in T$, all the new vertices in path P_i map under f' to i. Then $f'(i) = i$ for all $i \in T$ and $\sum_{u<v} c(u, v)d(f(u), f(v)) \ge \sum_{u<v} c(u, v)d(f'(u), f'(v))$.

Proof. Name the vertices on the $i - h_i$ path $v_0(= i), v_1, v_2, ..., v_l(= h_i)$, in order. Then

$$\sum_{u<v} c(u, v)d(f'(u), f'(v)) - \sum_{u<v} c(u, v)d(f(u), f(v))$$

$$= \sum_{i=1}^{k} d(i, f(h_i)) - \sum_{i=1}^{k} \sum_{j=0}^{l-1} d(f(v_j), f(v_{j+1})).$$

By the triangle inequality,

$$\sum_{j=0}^{l-1} d(f(v_j), f(v_{j+1})) \ge d(i, f(h_i)),$$

so the difference is nonpositive. ■

Let us assume from now on that $f : V' \to T$ itself maps the new vertices in P_i to terminal i. Let $R_i = \{v \in V \subseteq V' | f(v) = i\}$. We have two cases.

- Case 1: $|R_i| \le n/2$ for all i.

  If $uv \in E$, $u \in R_i$, $v \in R_j$, $i \ne j$, then $d(f(u), f(v)) = d(i, j) = \delta(i, j) \ge 2\sqrt{\lg n}$. Because G is an expander and $|R_i| \le n/2$ for all i, for each i the number of edges uv, $u \in R_i$, $u \notin R_i$, is at least $\alpha|R_i|$. Hence $\sum_{u<v, u,v \in V} c(u, v)d(f(u), f(v)) \ge (1/2) \sum_{i=1}^{k} \alpha|R_i|(2\sqrt{\lg n}) = \alpha \cdot n\sqrt{\lg n}$. (Each "cross edge" is counted twice.)

- Case 2: Some R_i, say R_1, has size exceeding $n/2$. We will use expansion and the max flow-min cut theorem to prove our theorem.

  Choose any $\lceil k/2 \rceil$ terminals i such that the distance in G between h_i and h_1 is at least $a = \lceil \frac{\lg k}{2 \lg \Delta} \rceil$; by Lemma 13, they exist. Let F be the set of chosen terminals. Because $n \geq 16$, $a \geq \frac{\lg n}{4 \lg \Delta}$.

  For each $i \in T$, let x_i be the unique neighbor of h_i in $V' - V$. Let $X = \{x_i | i \in F\}$ and $V^* = V \cup X$. Build a (directed) network N on $V^* \cup \{s, t\}$, s, t being new nodes, as follows. Start with the subgraph of G' induced by V^* and replace each edge by a pair of antiparallel arcs, each of capacity one. Add arcs (u, t) for all $u \in R_1$, each of capacity ∞. Add arcs (s, x) for all $x \in X$, each of capacity ∞. Because $R_1 \subseteq V$ and $x_i \notin V$, N has a finite capacity $s \to t$ cut defined by $\{s\} \cup X$.

  Now choose any finite capacity $s \to t$ cut $C^* = (S^*, \bar{S^*})$ in N, $s \in S^*$, $t \notin S^*$. $S^* \supseteq X$ and $S^* \cap R_1 = \emptyset$. Let $S = S^* \cap V$. (Possibly $S = \emptyset$.) Because $|R_1| \geq n/2$, $|S| \leq n/2$. Let C be the set of arcs (u, v) with $u \in S, v \in V - S$. By the expansion of G, $|C| \geq \alpha |S|$. Let $M = \{x_i | i \in F, h_i \notin S\}$. Corresponding to each $i \in F$ such that $h_i \notin S$ is an arc (x_i, h_i) in $C^* - C$, $|M|$ in total. Thus the total number of arcs in C^* is at least $\alpha |S| + |M| \geq \alpha(|S| + |M|) \geq \alpha |F| \geq \alpha k/2$, the penultimate inequality following from the fact that either h_i or x_i is in $S \cup M$ for all $i \in F$. It follows that the minimum capacity of an $s \to t$ cut in N is at least $\alpha k/2$.

  By the max flow-min cut theorem, there is a maximum $s \to t$ flow which has value 0 or 1 on each arc. If the flow on (u, v) and on (v, u) is 1 for some $u \neq v$, change the flow on both arcs to 0. The result is a maximum $s \to t$ flow that "uses" at most one of $(u, v), (v, u)$ for any $u \neq v$. This flow can be decomposed into at least $\alpha k/2$ arc-disjoint $s \to t$ paths in N, such that for all $u \neq v$, at most one of $(u, v), (v, u)$ appears in any path. It follows that there are at least $\alpha k/2$ arc-disjoint paths from an x_i with $i \in F$ to some vertex in R_1. If $Q =< x_i = v_{i0}, v_{i1}, v_{i2}, v_{i3}, ..., v_{is} \in R_1 >$ is such a path, then, using $f(x_i) = i$, $f(v_{is}) = 1$, we have $\sum_{j=0}^{s-1} d(f(v_{ij}), f(v_{i,j+1})) \geq d(f(x_i), f(v_{is})) = d(i, 1) = \delta(i, 1) \geq a$, a being at least $\frac{\lg n}{4 \lg \Delta}$. Since the paths are arc-disjoint and there are at least $\alpha k/2$ of them, we infer that $\sum_{u<v:u,v \in V'} c(u, v) d(f(u), f(v)) \geq \frac{\alpha k}{2} \frac{\lg n}{4 \lg \Delta}$.

  Using the definition of k, this last sum is at least $(\alpha/(16 \lg \Delta)) \cdot n \sqrt{\lg n}$. Since we have a feasible solution to (MET) of value at most $|E'| \leq (\Delta/2 + 2)n$, the ratio between the two is at least

$$\frac{\alpha}{(8 \lg \Delta)(\Delta + 4)} \sqrt{\lg n} \geq \frac{\alpha}{(8 \lg \Delta)(\Delta + 4)} \sqrt{\lg k}.$$

  Choose $c = \alpha/((8 \lg \Delta)(\Delta + 4))$ and the proof of Theorem 12 is complete. ■

The following theorem shows that the above analysis is asymptotically tight. It also suggests an alternative rounding procedure that for some instances performs better than the results in Section 2 (though in general it is far worse).

Theorem 15. There is a polynomial-time algorithm that takes as input a connected graph $G = (V, E)$ and a subset $T \subseteq V$ of terminals and computes a function $f : V \to T$ with $f(i) = i$ for all $i \in T$ such that

$$\sum_{uv \in E} d(f(u), f(v)) \leq 3\sqrt{d_{\max}} |E|,$$

where $d(u, v)$ is the minimum number of edges in a path between u and v and $d_{\max} = \max_{uv} d(u, v)$.

Proof. Add a new vertex s to G and connect s to all the terminals. Run a breadth-first search starting at s, computing for every node $v \in V$ its distance $l(v)$ from s. Notice that for every $v \in V$, $1 \leq l(v) \leq d_{\max} + 1$. Partition E into classes C_i, for $i = 1, 2, \ldots, d_{\max} + 1$: Place edge $uv \in E$ in C_i for $i = \min\{l(u), l(v)\}$. Let r be the smallest positive integer such that $|C_r| \leq |E|/\sqrt{d_{\max}}$. As $\sum_i |C_i| = |E|$, $r \leq \sqrt{d_{\max}}$. We now define f. Let t be an arbitrary terminal. For every $v \in V$ with $l(v) > r$, set $f(v) = t$. For every $v \in V$ with $l(v) \leq r$, set $f(v) = t_v$, where t_v is a terminal closest to v in G.

If $uv \in C_i$, then $d(u, t_u) \le i$ and $d(v, t_v) \le i$, and at least one of the inequalities is strict. Therefore $d(t_u, t_v) \le d(t_u, u) + 1 + d(v, t_v) \le 2i$. Consider an edge $uv \in C_i$. If $i \le r-1$, then both $l(u), l(v) \le r$, so $d(f(u), f(v)) = d(t_u, t_v) \le 2i \le 2r-2$. If $i > r$, then both $l(u), l(v) > r$, so $d(f(u), f(v)) = d(t, t) = 0$. For the remaining case of $i = r$ we use the trivial bound $d(f(u), f(v)) \le d_{\max}$. Using the bounds on r and on $|C_r|$, we have

$$\begin{aligned}\sum_{uv \in E} d(f(u), f(v)) &\le \sum_{i<r} \sum_{uv \in C_i} (2r-2) + \sum_{uv \in C_r} d_{\max} + \sum_{i>r} \sum_{uv \in C_i} 0 \\ &< 3\sqrt{d_{\max}}|E|. \qquad \blacksquare\end{aligned}$$

It is interesting to note that for bounded degree expanders the bounds are much better. Using arguments similar to those of the proofs of Theorems 12 and 15, we can prove

Theorem 16. For every sufficiently large positive integer Δ, there is a constant $\kappa > 0$ such that for infinitely many k there is an expander $G = (V, E)$ with maximum degree at most Δ, $|V|$ being $O(k \log k)$, and a set $T \subseteq V$ of size k, such that the integrality ratio of (MET) on the 0-extension instance defined by G, T, and the G-path metric on T is at least $\kappa \log \log k$. On the other hand, for every Δ there is a λ such that for every n-node expander G of maximum degree at most Δ and choice of terminals T, the integrality ratio of (MET) on the 0-extension instance defined by G, T and the G-path metric on T is at most $\lambda \log \log n$ (and there is a polynomial-time algorithm that achieves this bound).

References

[1] Y. Bartal. On approximating arbitrary metrics by tree metrics. In *Proc. STOC '98*.

[2] G. Calinescu, H. Karloff, and Y. Rabani. An improved approximation algorithm for multiway cut. To appear in *J. Comput. Syst. Sci.*. Preliminary version appeared in *STOC '98*.

[3] W. H. Cunningham and L. Tang. Optimal 3-terminal cuts and linear programming. In *Proc. IPCO '99*.

[4] J. Chuzhoy. Private communication, 1999.

[5] E. Dahlhaus, D. S. Johnson, C. H. Papadimitriou, P. D. Seymour, and M. Yannakakis. The complexity of multiway cuts. Extended abstract, 1983.

[6] E. Dahlhaus, D. S. Johnson, C. H. Papadimitriou, P. D. Seymour, and M. Yannakakis. The complexity of multiterminal cuts. *SIAM J. Comput.* 23:864-894, 1994. Preliminary version appeared in *STOC '92*.

[7] 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 *Proc. ICALP '93*.

[8] N. Garg, V. V. Vazirani, and M. Yannakakis. Approximate max-flow min-(multi)cut theorems and their applications. *SIAM J. Comput.* 25:235-251, 1996. Preliminary version appeared in *STOC '93*.

[9] A. Gupta and É. Tardos. A constant factor approximation algorithm for a class of classification problems. To appear in *Proc. STOC 2000*.

[10] D. R. Karger, P. Klein, C. Stein, M. Thorup, and N. E. Young. Rounding algorithms for a geometric embedding of minimum multiway cut. In *Proc. STOC '99*, pages 668–678.

[11] A. V. Karzanov. Minimum 0-extension of graph metrics. *Europ. J. Combinat.*, 19:71–101, 1998.

[12] P. Klein, A. Agarwal, R. Ravi, and S. Rao. Approximation through multicommodity flow. In *Proc. FOCS '90*, pages 726–737.

[13] P. Klein, S. Plotkin, and S. Rao. Planar graphs, multicommodity flow, and network decomposition. In *Proc. STOC '91*, pages 682–690.

[14] J. Kleinberg and É. Tardos. Approximation algorithms for classification problems with pairwise relationships: Metric labeling and Markov random fields. In *Proc. FOCS '99*, pages 14–23.

[15] F. T. Leighton and S. Rao. An approximate max-flow min-cut theorem for uniform multicommodity flow problems with applications to approximation algorithms. In *Proc. FOCS '88*, pages 422–431.

[16] A. Lubotzky, R. Phillips, and P. Sarnak. Ramanujan graphs. *Combinatorica* 8:261–277, 1988.

[17] É. Tardos and V. V. Vazirani. Improved bounds for the max-flow min-multicut ratio for planar and $K_{r,r}$-free graphs. *Inform. Proc. Lett.*, 47:77–80, 1993.