

An LP Formulation and Approximation Algorithms for the Metric Labeling Problem*

C. Chekuri[†] S. Khanna[‡] J. Naor[§] L. Zosin[¶]

August 4, 2004

Abstract

We consider approximation algorithms for the metric labeling problem. This problem was introduced in a paper by Kleinberg and Tardos [26], and captures many classification problems that arise in computer vision and related fields. They gave an $O(\log k \log \log k)$ approximation for the general case where k is the number of labels, and a 2-approximation for the uniform metric case. (In fact, the bound for general metrics can be improved to $O(\log k)$ by the work of Fakcheroenphol, Rao, and Talwar [16].) Subsequently, Gupta and Tardos [18] gave a 4-approximation for the truncated linear metric, a metric motivated by practical applications to image restoration and visual correspondence. In this paper we introduce an integer programming formulation and show that the integrality gap of its linear relaxation either matches or improves the ratios known for several cases of the metric labeling problem studied until now, providing a unified approach to solving them. In particular, we show that the integrality gap of our LP is bounded by $O(\log k)$ for a general k -point metric and 2 for the uniform metric thus matching the known ratios. We also develop an algorithm based on our LP that achieves a ratio of $2 + \sqrt{2} \simeq 3.414$ for the truncated linear metric improving the earlier known ratio of 4. Our algorithm uses the fact that the integrality gap of the LP is 1 on a linear metric.

1 Introduction

Motivated by certain classification problems that arise in computer vision and related fields, Kleinberg and Tardos introduced the metric labeling problem [26]. In a typical classification problem, one wishes to assign labels to a set of objects so as to optimize some measure of the quality of the labeling. The metric labeling problem captures a broad range of classification problems where the quality of a labeling depends on the pairwise relations between the underlying set of objects. More precisely, the task is to classify a set V of n objects by assigning to each object a label from a set L of labels. The pairwise relationships between the objects are represented by a weighted graph $G = (V, E)$ where $w(u, v)$ represents the strength of the relationship between u and v . The objective is to find a labeling, a function $f : V \rightarrow L$, that maps objects to labels, where the cost of f , denoted by $Q(f)$, has two components.

*This version to appear in SIAM J. on Discrete Mathematics. A preliminary version appeared in the Proceedings of the 12th Annual ACM-SIAM Symposium on Discrete Algorithms, 2001, pp. 109-118.

[†]Bell Labs, Lucent Technologies, 600 Mountain Ave, Murray Hill, NJ 07974. E-mail: chekuri@research.bell-labs.com.

[‡]Department of Computer and Information Science, University of Pennsylvania, Philadelphia, PA 19104. E-mail: sanjeev@cis.upenn.edu. Supported in part by an Alfred P. Sloan Research Fellowship.

[§]Computer Science Dept., Technion, Haifa 32000, Israel. E-mail: naor@cs.technion.ac.il. Most of this work was done while the author was at Bell Labs, Lucent Technologies, 600 Mountain Ave, Murray Hill, NJ 07974. Research supported in part by US-Israel BSF Grant 2002276 and by EU contract IST-1999-14084 (APPOL II).

[¶]Bear Stearns & Co., One Federal St., 29th Floor, Boston, MA 02110. E-mail: LZosin@Bear.com. This work was done while the author was at NEC Research Institute, Princeton, NJ 08540.

- For each $u \in V$, there is a non-negative assignment cost $c(u, i)$ to label u with i . This cost reflects the relative likelihood of assigning labels to u .
- For each pair of objects u and v , the edge weight $w(u, v)$ measures the strength of their relationship. This models the assumption that strongly related objects should be assigned labels that are *close*. This is modeled in the objective function by the term $w(u, v) \cdot d(f(u), f(v))$ where $d(\cdot, \cdot)$ is a distance function on the labels L .

Thus

$$Q(f) = \sum_{u \in V} c(u, f(u)) + \sum_{(u,v) \in E} w(u, v) \cdot d(f(u), f(v))$$

and the goal is to find a labeling of minimum cost. In the metric labeling problem, the distance function d is assumed to be a metric. We remark that if the distance function d is allowed to be arbitrary, the graph coloring problem can be reduced to the labeling problem.

A prototypical application of the metric labeling problem is the *image restoration* problem in computer vision [4, 5, 6]. In the image restoration problem, the goal is to take an image corrupted by noise and restore it to its “true” version. The image consists of pixels and these are the objects in the classification problem. Each pixel has an integer intensity value associated with it that is possibly corrupted and we would like to restore it to its true value. Thus the labels correspond to intensity values and the goal is to assign a new intensity to each pixel. Further, neighboring pixels are assumed to be close to each other, since intensities typically change smoothly. Thus, neighboring pixels have an edge between them with a positive weight (assume a uniform value for concreteness). The original or observed intensities are assumed to be close to true values¹. The assignment cost is some positive cost associated with changing the intensity from its original value to a new value, the larger the change, the larger the cost. To see how the cost function models the restoration, consider a pixel corrupted by noise as a result of which its observed intensity is very different from its neighboring pixels. By changing the corrupted pixel’s intensity we incur a cost of relabeling but that can be offset by the edge cost saved by being closer to its neighbors. The assignment costs weigh the labeling in favor of the original values since most of the intensities are likely to be correct (see also the footnote).

The metric labeling problem naturally arises in other applications in image processing and computer vision. Researchers in these fields have developed a variety of good heuristics that use classical combinatorial optimization techniques such as network flow and local search [4, 21, 32, 17, 14, 12].

Kleinberg and Tardos [26] formalized the metric labeling problem and its connections to Markov random fields and other classification problems (see [26] for a thorough description of the various connections). Metric labeling also has rich connections to some well known problems in combinatorial optimization. It is related to the quadratic assignment problem, an extensively studied problem in Operations Research. A special case of metric labeling is the 0-extension problem, studied by Karzanov [24, 25]. There are no assignment costs in this problem, however, the graph contains a set of terminals, t_1, \dots, t_k , where the label of terminal t_i is fixed in advance to i and the non-terminals are free to be assigned to any of the labels. The 0-extension problem generalizes the well-studied multi-way cut problem [13, 7, 23]. Karzanov [24, 25] showed that certain special cases (special metrics) of the 0-extension problem can be solved optimally in polynomial time. Another special case of the metric labeling problem is the *task assignment problem* in distributed computing [28]. Here, tasks of a modular

¹This assumption is based on the connection of the labeling problem to the theory of Markov Random Fields (MRFs) [27]. In this theory, the observed data or labeling of the objects is assumed to be obtained from a true labeling by adding independent random noise. The idea is to decide the most probable labeling given the observed data. An MRF can be defined by a graph on the objects with edges indicating dependencies between objects. Under the assumption that the probability distribution of an object’s label depends only on its neighbors’ labels, and if the MRF satisfies two standard assumptions of homogeneity and pairwise interactions, the labeling problem can be restated as the problem of finding a labeling f that maximizes the *a posteriori* probability $\Pr[f|f']$ where f' is the observed labeling. We refer the reader to [26] for more details on the connection of metric labeling to MRFs.

program need to be assigned to the processors of a distributed system, while balancing between task execution cost and inter-task communication cost.

Kleinberg and Tardos [26] obtained an $O(\log k \log \log k)$ approximation for the general metric labeling problem, where k denotes the number of labels in L and a 2-approximation for the uniform metric. The approximation for general metrics improves to $O(\log k)$ by a recent result of Fakcheroenphol, Rao and Talwar [16]. Metric labeling is Max-SNP hard; this follows from the Max-SNP hardness of the multi-way cut problem which is a special case. Given the rich connections of this problem to other well-studied optimization problems and a variety of applications, a natural interesting question is to determine the approximability of the general problem as well as of the important special cases that arise in practice.

The Truncated Linear Metric Gupta and Tardos [18] considered the metric labeling problem for the truncated linear metric, a special case motivated by its direct applications to two problems in computer vision, namely *image restoration* and *visual correspondence*. We briefly describe the application of this metric to the image restoration problem discussed earlier, see [18] for more details. Consider the case of grey scale images where the intensities are integer values. Our earlier description assumed that intensities of neighboring pixels should be similar since the image is typically smooth. This motivates a linear-like metric on the labels where the distance between two intensities i and j is $|i - j|$. However, at object boundaries (here we are referring to objects in the image) sharp changes in intensities happen. Thus, for the metric to be robust, neighboring pixels that are actually at object boundaries (and hence naturally differ in their intensities by large amounts) should not be penalized by arbitrarily large quantities. This motivated Gupta and Tardos to consider a truncated linear metric, where the distance between two intensities i and j is given by $d(i, j) = \min(M, |i - j|)$. Thus, the penalty is truncated at M ; this is a natural (and non-uniform) metric for the problem at hand. A very similar reasoning applies to the visual correspondence problem, where the objective is to compare two images of the same scene for disparity. The labels here correspond to depth of the point in the image from the camera. For the truncated linear metric a 4-approximation algorithm was given in [18] using local search. The local moves in the algorithm make use of the flow network used in [4, 21] which gives an optimal solution to the linear metric case in polynomial time.

For the image restoration application, other distance functions have also been studied. In particular, the quadratic distance function $d(i, j) = |i - j|^2$ and its truncated version $d(i, j) = \min\{M, |i - j|^2\}$ have been considered (see [27] and [22]). Unfortunately, neither of these distance functions is a metric, and hence the algorithms for metric labeling problem cannot be used. However as we discuss shortly, we are able to provide non-trivial approximation ratios for them.

Results In this paper we address the problem of obtaining improved approximation guarantees for the metric labeling problem. Kleinberg and Tardos [26] point out the difficulty of the general case as having to do with the absence of a “natural” IP formulation for the problem. They give an IP formulation for tree metrics and use Bartal’s probabilistic tree approximations [2, 3] to reduce the problem with an arbitrary metric to that with a tree metric. In this work we give a natural IP formulation for the general problem. An advantage of our formulation is that it is applicable even to distance functions that are not metrics, for example, the quadratic distance function mentioned above. We substantiate the strength of this formulation by deriving both known results and new results using its linear relaxation. In particular, we show the following results on the integrality gap of our formulation.

- $O(\log k)$ for general metrics and a factor 2 for the uniform metric.
- 1 for the linear metric and distances on the line defined by convex functions (not necessarily metrics).
- $2 + \sqrt{2} \simeq 3.414$ for the truncated linear metric.
- $O(\sqrt{M})$ for the truncated quadratic distance function.

The integrality gaps we show for our formulation either match or improve upon the best previous approximation ratios for most of the cases known to us (in addition to the above, we can also show that if G is a tree the integrality gap is 1). Our formulation allows us to present these results in a unified fashion. In the process we improve the 4 approximation of Gupta and Tardos [18] for the important case of the truncated linear metric. We also show a reduction from the case with arbitrary assignment costs $c(u, i)$ to the case where $c(u, i) \in \{0, \infty\}$ for all u and i . The reduction preserves the graph G and the optimal solution, but increases the size of the label space from k labels to nk labels. Though this is undesirable in practice, it is useful from a theoretical point of view. In particular, this result implies that in order to find an $O(1)$ -approximation for the general case, it suffices to focus attention on restricted instances with no assignment costs, but where each vertex is required to be assigned a label from its own set of labels. We believe that our results and techniques are a positive step towards obtaining improved bounds for both the general metric labeling problem, as well as for interesting special cases.

Călinescu, Karloff, and Rabani [8] considered approximation algorithms for the 0-extension problem. They consider a linear programming formulation (which they call the *metric relaxation*) originally studied by Karzanov [24], where they associate a length function with every edge of the graph and require that: (i) the distance between terminals t_i and t_j , for $1 \leq i, j \leq k$, is at least $d(i, j)$, and (ii) the length function is a semi-metric. We note that their formulation does not apply to the general metric labeling problem. They obtain an $O(\log k)$ -approximation algorithm for the 0-extension problem using this formulation and an $O(1)$ -approximation for planar graphs. Our linear programming formulation, when specialized to the 0-extension problem, induces a feasible solution for the metric relaxation formulation, by defining the length of an edge to be its transshipment cost (see Section 2). It is not hard to verify that this length function is a semi-metric. Călinescu, Karloff, and Rabani [8] also show a gap of $\Omega(\sqrt{\log k})$ on the integrality ratio of their formulation. Their lower bound proof does not seem to carry over in any straight forward way to our formulation. We note that the metric relaxation formulation optimizes over the set of all semi-metrics, while our formulation optimizes only over a subset of the semi-metrics. Whether our formulation is strictly stronger than the metric relaxation of [8] is an interesting open problem.

Outline Section 2 describes our linear programming formulation for the general metric labeling problem. In Section 3, we analyze our formulation for uniform and linear metrics. Building on our rounding scheme for the linear metric, we design and analyze a rounding procedure for the truncated linear metric in Section 4. In Section 5, we study the general metric labeling problem and show that the integrality gap of our formulation is bounded by $O(\log k)$. We also describe here a transformation that essentially eliminates the role of the label cost assignment function.

2 The Linear Programming Formulation

We present a new linear integer programming formulation of the metric labeling problem. Let $x(u, i)$ be a $\{0, 1\}$ -variable indicating that vertex u is labeled i . Let $x(u, i, v, j)$ be a $\{0, 1\}$ -variable indicating that for edge $(u, v) \in E$, vertex u is labeled i and vertex v is labeled j . See Fig. 1 for the formulation.

Constraints (1) simply express that each vertex must receive some label. Constraints (2) force consistency in the edge variables: if $x(u, i) = 1$ and $x(v, j) = 1$ they force $x(u, i, v, j)$ to be 1. Constraints (3) express the fact that (u, i, v, j) and (v, j, u, i) refer to the same edge; the redundancy helps in notation. We obtain a linear relaxation of the above program by allowing the variables $x(u, i)$ and $x(u, i, v, j)$ to take any non-negative value. We note that equality in (2) is important for the linear relaxation.

With each edge $(u, v) \in E$ we associate a complete bipartite graph $H(u, v)$. The vertices of $H(u, v)$ are $\{u_1, \dots, u_k\}$ and $\{v_1, \dots, v_k\}$, i.e., they represent all possible labelings of u and v . There is an edge (u_i, v_j) connecting the pair of vertices u_i and v_j , $1 \leq i, j \leq k$. In the sequel, we refer to the edges of $H(u, v)$ as *links* to distinguish them from the edges of G . Suppose that the value of the variables $x(u, i)$ for all u and i has been determined. For an edge $(u, v) \in E$, we can interpret the variables $x(u, i, v, j)$ from a flow perspective.

$$(I) \quad \min \sum_{u \in V} \sum_{i=1}^k c(u, i) \cdot x(u, i) + \sum_{(u, v) \in E} w(u, v) \sum_{i=1}^k \sum_{j=1}^k d(i, j) \cdot x(u, i, v, j)$$

subject to:

$$\sum_{i=1}^k x(u, i) = 1 \quad \forall v \in V \quad (1)$$

$$\sum_{j=1}^k x(u, i, v, j) - x(u, i) = 0 \quad \forall u \in V, (u, v) \in E, i \in 1, \dots, k \quad (2)$$

$$x(u, i, v, j) - x(v, j, u, i) = 0 \quad \forall u, v \in V, i, j \in 1, \dots, k \quad (3)$$

$$x(u, i) \in \{0, 1\} \quad \forall u \in V, i \in 1, \dots, k \quad (4)$$

$$x(u, i, v, j) \in \{0, 1\} \quad \forall (u, v) \in E, i, j \in 1, \dots, k \quad (5)$$

Figure 1: ILP Formulation

The contribution of edge $(u, v) \in E$ to the objective function of the linear program is the cost of the optimal transshipment of flow between $\{u_1, \dots, u_k\}$ and $\{v_1, \dots, v_k\}$, where: (i) the supply of u_i is $x(u, i)$ and the demand of v_j is $x(v, j)$ for $0 \leq i, j \leq k$, (ii) the cost of shipping a unit flow from u_i to v_j is $d(i, j)$. (The choice of the “supply” side and the “demand” side is arbitrary.)

For the rest of the paper, the quantity $d_{LP}(u, v)$ refers to the LP distance between u and v and is defined to be the transshipment cost $\sum_{i,j} d(i, j) \cdot x(u, i, v, j)$. The LP distance derived from an optimal (fractional) solution to (I) induces a metric on the graph, since for any $v_1, v_2, v_3 \in V$, the transshipment cost from v_1 to v_2 cannot be more than the sum of the transshipment costs from v_2 to v_3 and from v_3 to v_1 . The transshipment problem between two distributions was introduced by Monge [29] and is also referred to as the Monge-Kantorovich mass transference problem and has several applications [31]. In the image processing literature [33, 30] this metric has also been referred to as the *earth mover’s metric*.

A solution to the formulation has an interesting *geometric* interpretation. It defines an embedding of the graph into a k -dimensional simplex, where the distance between points in the simplex is defined by the earth mover’s metric on the labels. Our formulation specializes to that of Kleinberg and Tardos [26] for the uniform metric case which in turn specializes to that of Călinescu, Karloff and Rabani [7] for the multiway cut problem where the distance between points in the embedding is simply their ℓ_1 distance.

3 Uniform Metrics and Linear Metrics

We now analyze the performance of our linear programming formulation on two natural special cases, namely, uniform metrics and linear metrics. Kleinberg and Tardos [26] showed a 2-approximation for the uniform metric case. Their approach is based on rounding the solution of a linear program formulated specifically for uniform metrics. We will show that our linear programming formulation dominates the one in [26], and thus also has an integrality gap of at most 2. For the case of linear metrics, Boykov *et al* [4], and Ishikawa and Geiger [21], have obtained exact algorithms, by reducing the problem to a minimum $\{s, t\}$ -cut computation. We show here that, on linear metrics, our linear programming formulation gives an exact algorithm as well. Our analysis for the linear metric case plays an important role in the algorithm for the truncated linear metric case.

3.1 The Uniform Metric Case

Kleinberg and Tardos [26] formulated a linear program, denoted by (KT), for the uniform metric and gave the following iterative algorithm for rounding a solution to it. Initially, no vertex is labeled. Each iteration consists of the following steps: (i) choose a label uniformly at random from $1, \dots, k$ (say, a label i), (ii) choose a real threshold θ uniformly at random from $[0, 1]$, (iii) for all unlabeled vertices $u \in V$, u is labeled i if $\theta \leq x(u, i)$. The algorithm terminates when all vertices are labeled. Kleinberg and Tardos [26] showed that the expected cost of a labeling obtained by this algorithm is at most twice the cost of the LP solution.

$$(KT) \quad \min \sum_{v \in V} \sum_{i=1}^k c(v, i) \cdot x(v, i) + \sum_{(u,v) \in E} w(u, v) \cdot \frac{1}{2} \sum_{i=1}^k |x(u, i) - x(v, i)|$$

subject to:

$$\begin{aligned} \sum_{i=1}^k x(u, i) &= 1 && \forall u \in V \\ x(u, i) &\geq 0 && \forall u \in V \text{ and } i \in 1, \dots, k \end{aligned}$$

We show that applying the rounding algorithm to an optimal solution obtained from linear program (I), yields the same approximation factor. Let \bar{x} be a solution to (I). Note that for both (I) and (KT) the variables $x(u, i)$ completely determine the cost. We will show that cost of (KT) on \bar{x} is smaller than that of (I). Both linear programs (I) and (KT) coincide regarding the labeling cost. Consider edge $(u, v) \in E$. We show that the contribution of (u, v) to the objective function of (I) is at least as large as the contribution to the objective function of (KT).

$$\begin{aligned} \frac{1}{2} \cdot \sum_{i=1}^k |\bar{x}(u, i) - \bar{x}(v, i)| &= \frac{1}{2} \cdot \sum_{i=1}^k \left| \sum_{j=1}^k \bar{x}(u, i, v, j) - \sum_{j=1}^k \bar{x}(v, i, u, j) \right| \\ &\leq \frac{1}{2} \cdot \sum_{i=1}^k \left| \sum_{j=1, j \neq i}^k \bar{x}(u, i, v, j) + \sum_{j=1, j \neq i}^k \bar{x}(v, i, u, j) \right| \\ &= \frac{1}{2} \cdot \sum_{i=1}^k \sum_{j=1, j \neq i}^k 2 \cdot \bar{x}(u, i, v, j) \\ &= \sum_{i=1}^k \sum_{j=1}^k d(i, j) \cdot \bar{x}(u, i, v, j). \end{aligned}$$

The penultimate equality in the above set of equations is true since $\bar{x}(u, i, v, j) = \bar{x}(v, j, u, i)$. The final equality follows from the fact that $d(i, i) = 0$ and $d(i, j) = 1, i \neq j$. The last term is the contribution of (u, v) to the objective function of (I). We note that the example used in [26] to show that the integrality gap of (KT) is at least $2 - 1/k$, can be used to show the same gap for (I) as well.

3.2 The Line Metric Case

We now turn to the case of a linear metric and show that the value of an integral optimal solution is equal to the value of a fractional optimal solution of the LP (I). Without loss of generality, we can assume that the labels of the

metric are integers $1, 2, \dots, k$. If the label set contains non-consecutive integers, we can add all the “missing” intermediate integers to the label set and set the cost of assigning them to every vertex to be infinite. In fact the rounding can be generalized to the case where the labels are arbitrary points on the real line without any difficulty.

Rounding procedure: Let \bar{x} be an optimal fractional solution to the linear program. We round the fractional solution as follows. Let θ be a real threshold chosen uniformly at random from $[0, 1]$. For all i , $1 \leq i \leq k$, let

$$\alpha(u, i) = \sum_{j=1}^i \bar{x}(u, j).$$

Each vertex $u \in V$ is labeled by the unique label i that satisfies $\alpha(u, i-1) < \theta \leq \alpha(u, i)$. Clearly all vertices are labeled since $\alpha(u, k) = 1$.

Analysis: We analyze the expected cost of the assignment produced by the rounding procedure. For each vertex $u \in V$, let $L(u)$ be a random variable whose value is the label assigned to u by the rounding procedure. It can be readily verified that the probability that $L(u) = i$ is equal to $\bar{x}(u, i)$. This means that the expected labeling cost of vertex v is equal to $\sum_{i=1}^k \bar{x}(u, i) \cdot c(u, i)$ which is precisely the assignment cost of u in the linear program (with respect to solution \bar{x}). We now fix our attention on the expected cost of the edges.

Lemma 3.1 Consider edge $(u, v) \in E$. Then,

$$\mathbf{E}[d(L(u), L(v))] = \sum_{i=1}^k |\alpha(u, i) - \alpha(v, i)|.$$

Proof. For the sake of the analysis we define auxiliary binary random variable Z_1, \dots, Z_{k-1} as follows. The variable Z_i is 1 if $\min\{L(u), L(v)\} \leq i$ and $\max\{L(u), L(v)\} > i$, and is 0 otherwise. In other words Z_i is 1 if i is in the interval defined by $L(u)$ and $L(v)$. It is easy to see that

$$d(L(u), L(v)) = \sum_{i=1}^{k-1} Z_i.$$

Therefore $\mathbf{E}[d(L(u), L(v))] = \sum_{i=1}^{k-1} \mathbf{E}[Z_i]$. We claim that

$$\mathbf{E}[Z_i] = \Pr[Z_i = 1] = |\alpha(u, i) - \alpha(v, i)|.$$

The lemma easily follows from this claim. To prove the claim, assume w.l.o.g. that $\alpha(u, i) \geq \alpha(v, i)$. If $\theta < \alpha(v, i)$ it is clear that $L(u), L(v) \leq i$ and if $\theta > \alpha(u, i)$ then $L(u), L(v) > i$: in both cases $Z_i = 0$. If $\theta \in (\alpha(v, i), \alpha(u, i)]$ then $L(u) \leq i$ and $L(v) > i$ which implies that $Z_i = 1$. Thus $\Pr[Z_i = 1]$ is exactly $|\alpha(u, i) - \alpha(v, i)|$. \square

We now estimate the contribution of an edge $(u, v) \in E$ to the objective function of the linear program. As indicated in Section 2, the contribution is equal to the cost of the optimal transshipment cost of the flow in the complete bipartite graph $H(u, v)$ between $\{u_1, \dots, u_k\}$ and $\{v_1, \dots, v_k\}$, where the supply of u_i is $\bar{x}(u, i)$ and the demand of v_j is $\bar{x}(v, j)$ for $1 \leq i, j \leq k$. Recall that $d_{LP}(u, v) = \sum_{i,j} d(i, j) \cdot \bar{x}(u, i, v, j)$.

Lemma 3.2 For the line metric

$$d_{LP}(u, v) \geq \sum_{i=1}^k |\alpha(u, i) - \alpha(v, i)|.$$

Proof. The crucial observation in the case of a linear metric is that flow can be *uncrossed*. Let $i \leq i'$ and $j \leq j'$. Suppose that ε amount of flow is sent from u_i to $v_{j'}$ and from $u_{i'}$ to v_j . Then, uncrossing the flow, i.e., sending ε amount of flow from u_i to v_j and from $u_{i'}$ to $v_{j'}$ will not increase the cost of the transshipment. This means that the amount of flow sent from $\{u_1, \dots, u_i\}$ to $\{v_1, \dots, v_i\}$, for all i , $1 \leq i \leq k$, is precisely $\min\{\alpha(u, i), \alpha(v, i)\}$. Therefore, $|\alpha(v, i) - \alpha(u, i)|$ amount of flow is sent “outside” the label set $1, 2, \dots, i$ and can be charged one unit of cost (with respect to i). Applying this argument to all values of i , $1 \leq i \leq k$, we get that the cost of the *optimal* transshipment of flow is precisely $\sum_{i=1}^k |\alpha(u, i) - \alpha(v, i)|$. The lemma follows. \square

Hence, together with Lemma 3.1, we get that the expected cost of edge $(u, v) \in E$ after the rounding is no more than its contribution to the objective function of the linear program.

The uncrossing of flow in the proof of Lemma 3.2 relies on the Monge-property of the distances induced by points on a line. Hoffman [20], in his classical paper, pointed out that the Monge-property can be exploited for transportation and related problems.

Theorem 1 *The integrality gap of LP (I) for the line metric is 1.*

3.3 Distance functions on the Line defined by Convex functions

We now consider distance functions on the labels $1, \dots, k$ on the integer line defined by strictly convex functions, that is $d(i, j) = f(|i - j|)$ where f is convex and non-decreasing. Notice that such distance functions do not satisfy the metric property, since d is a metric if and only if f is concave and increasing. Our motivation for studying these metrics comes from the quadratic function $d(i, j) = |i - j|^2$ which is of particular interest in the image restoration application [22, 27] described earlier. We note that for the special case where the assignment cost is also a convex function (of the label), efficient algorithms are given by [19]. We can show the following.

Theorem 2 *For any distance function on the line defined by a convex function, the integrality gap of LP (I) is 1.*

We sketch the proof, since it is similar to the linear case. Consider a feasible solution \bar{x} to the LP. For any edge (u, v) , if f is convex, the optimal cost transshipment flow in $H(u, v)$ is *non-crossing*. Further, for the rounding that we described for the linear case, if the flow is non-crossing, $\Pr[L(u) = i \wedge L(v) = j] = \bar{x}(u, i, v, j)$. The theorem follows from this last fact trivially. Ishikawa and Geiger [22] show that the flow graph construction for the line metric can be extended for convex functions to obtain an optimal solution. The advantage of our approach is that the solution is obtained from a general formulation. This allows us to extend the ideas to obtain the first non-trivial approximation for the truncated quadratic distance function.

4 Improved Approximation for the Truncated Line Metric

In this section we use LP (I) to give a $(2 + \sqrt{2})$ -approximation algorithm for the truncated line metric case. This improves the 4-approximation provided in [18]. We prefer to view the metric graph as a line with the truncation to M being implicit. This allows us to use, with appropriate modifications, ideas from Section 3.2. We round the fractional solution \bar{x} once again using only the values $\bar{x}(u, i)$. Let $M' \geq M$ be an integer parameter that we will fix later. We repeat the following iteration until all vertices are assigned a label.

- Pick an integer ℓ uniformly at random in $[-M' + 2, k]$. Let I_ℓ be the interval $[\ell, \ell + M' - 1]$.
- Pick a real threshold θ uniformly at random from $[0, 1]$.
- Let u be an unassigned vertex. If there is a label $i \in I_\ell$ such that

$$\sum_{j=\ell}^{i-1} \bar{x}(u, j) < \theta \leq \sum_{j=\ell}^i \bar{x}(u, j),$$

we assign i to u . Otherwise u is unassigned in this iteration.

The above algorithm generalizes the rounding algorithms for the uniform and line metrics in a natural way. Once the index ℓ is picked, the rounding is similar to that of the line metric in the interval I_ℓ . The difference is that a vertex might not get a label in an iteration. If two vertices u and v get separated in an iteration, that is only one of them gets assigned, our analysis will assume that their distance is M . This is similar to the analysis in [26] for the uniform metric case. Our improvement comes from a careful analysis of the algorithm that treats links differently, based on whether their distance is linear or truncated. The analysis guides the choice of M' to obtain the best guarantee.

Let $L(u)$ and $L(v)$ be random variables that indicate the labels that get assigned to u and v by the algorithm.

Lemma 4.1 *In any given iteration, the probability of an unassigned vertex u getting a label i in that iteration is exactly $\bar{x}(u, i) \cdot M' / (k + M' - 1)$. The probability of u getting assigned in the iteration is $M' / (k + M' - 1)$. Therefore $\Pr[L(u) = i] = \bar{x}(u, i)$.*

Proof. If ℓ is picked in the first step of an iteration, the probability of assigning i to u is exactly $\bar{x}(u, i)$, if $i \in I_\ell$, and zero otherwise. The number of intervals that contain i is M' , hence the probability of u getting i in an iteration is simply $\bar{x}(u, i) \cdot M' / (k + M' - 1)$. \square

It also follows from Lemma 4.1 that with high probability all vertices are assigned in $O(k \log n)$ iterations. The following lemma bounds the expected distance between $L(u)$ and $L(v)$ as a function of M' . Recall $d_{LP}(u, v) = \sum_{i,j} d(i, j) \cdot \bar{x}(u, i, v, j)$.

Lemma 4.2 *The expected distance between $L(u)$ and $L(v)$ satisfies the following inequality.*

$$\mathbf{E}[d(L(u), L(v))] \leq \left(2 + \max \left\{ \frac{2M}{M'}, \frac{M'}{M} \right\}\right) d_{LP}(u, v).$$

Theorem 3 *There is a $(2 + \sqrt{2})$ -approximation algorithm for the metric labeling problem when the metric is truncated linear.*

Proof. We note that the algorithm and the analysis easily generalizes to the case when M' is a real number. We choose $M' = \sqrt{2}M$ and the theorem follows from Lemmas 4.1 and 4.2. \square

For the rest of the section, we restrict our attention to one particular edge $(u, v) \in E(G)$. We analyse the effect of the rounding on the expected distance with the goal of proving Lemma 4.2. To analyze the process, we consider an iteration before which neither u and v are assigned. If, in the current iteration, only one of u or v is assigned a label, that is they are *separated*, we will pay a distance of M . If both of them are assigned, then we pay a certain distance based on their labels in the interval. The main quantity of interest is the expected distance between the labels of u and v in a single iteration conditioned under the event that both of them are assigned in this iteration. Recall that a link refers to the edges in the complete bipartite graph $H(u, v)$.

Given an interval $I_\ell = [\ell, \ell + M' - 1]$, we partition the interesting links for I_ℓ into three categories, *internal*, *left crossing*, and *right crossing*. The internal links denoted by $\text{INT}(I_\ell)$ are all the links (u, i, v, j) with $i, j \in I_\ell$, the left crossing links denoted by $\text{LCROSS}(I_\ell)$ are those with $\min\{i, j\} < \ell$ and $\max\{i, j\} \in I_\ell$, and right crossing link denoted by $\text{RCROSS}(I_\ell)$ are those with $\min\{i, j\} \in I_\ell$ and $\max\{i, j\} > \ell + M' - 1$. It is clear that no link is both left and right crossing. Let $\text{CROSS}(I_\ell)$ denote $\text{LCROSS}(I_\ell) \cup \text{RCROSS}(I_\ell)$. See Figure 2 for an example. It is easy to see that e is not relevant to I_ℓ if $i, j \notin I_\ell$.

We set up some notation that we use for the rest of this section. For a link $e = (u, i, v, j)$ let $d_{\text{lin}}(e) = |i - j|$ and $d(e) = d(i, j) = \min(M, d_{\text{lin}}(e))$ be the linear distance and truncated linear distance respectively, between

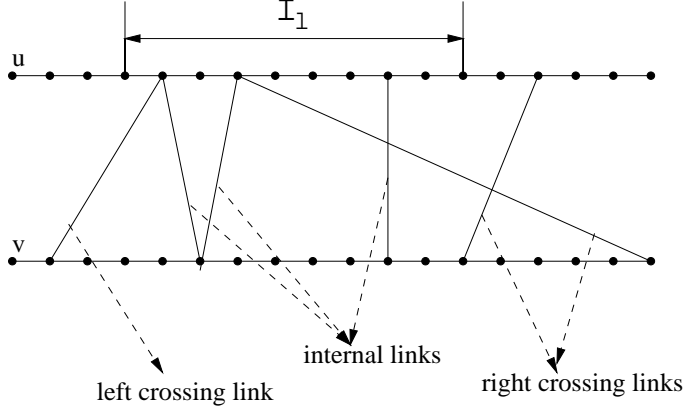


Figure 2: Illustration of interesting links for I_ℓ .

i and j . The quantity \bar{x}_e is compact notation for $\bar{x}(u, i, v, j)$. Consider a link e that crosses the interval I_ℓ and let i be the label of e that is internal to I_ℓ . We denote by $d_\ell(e)$ the quantity $(\ell + M' - 1 - i)$, the linear distance from i to the right end of interval. The quantity $\bar{x}(u, I_\ell)$ refers to $\sum_{i \in I_\ell} \bar{x}(u, i)$, the flow of u assigned by the LP to labels in I_ℓ .

Lemma 4.3 *The probability of u and v being separated given that ℓ was chosen in the first step of the iteration is at most $\sum_{e \in \text{CROSS}(I_\ell)} \bar{x}_e$.*

Proof. The probability of separation is exactly $|\bar{x}(u, I_\ell) - \bar{x}(v, I_\ell)|$ which can easily be seen to be upper bounded by $\sum_{e \in \text{CROSS}(I_\ell)} \bar{x}_e$. \square

Lemma 4.4 *For two vertices u and v , unlabeled before an iteration, let p_ℓ be the expected distance between them, conditioned on the event that ℓ was chosen in the first step of the iteration and both were assigned a label in I_ℓ . Then*

$$p_\ell \leq \sum_{e \in \text{CROSS}(I_\ell)} d_\ell(e) \bar{x}_e + \sum_{e \in \text{INT}(I_\ell)} d_{in}(e) \bar{x}_e.$$

We provide some intuition before giving the formal proof. Once ℓ is fixed, the rounding is exactly the same as that for the line metric when restricted to the interval I_ℓ . From Lemma 3.1 we know the exact expected cost of the rounding. However, we do not have an equivalent of Lemma 3.2 to bound the LP cost, because I_ℓ is only a subinterval of the full line and also because of the truncation. The main difficulty is with links that belong to $\text{CROSS}(I_\ell)$. By charging each of the crossing links e , an amount equal to $d_\ell(e)$ (instead of $d(e)$ that LP paid), we are able to pay for the expected cost of rounding. In other words, we charge the interesting links of I_ℓ to pay for the optimal linear metric transshipment flow induced by the fractional values $\bar{x}(u, i)$ and $\bar{x}(v, i)$, $i \in I_\ell$ when restricted to I_ℓ .

Proof. Fix an ℓ and w.l.o.g. assume that $\bar{x}(u, I_\ell) \geq \bar{x}(v, I_\ell)$. With probability $q = \bar{x}(v, I_\ell)$, both u and v get labels from I_ℓ . We analyze the expected distance conditioned on this event. Once the interval I_ℓ is fixed the rounding is very similar to that of the linear metric case. For $0 \leq i < M'$ let $\alpha(u, i) = \sum_{j=\ell}^{j+i} \bar{x}(u, j)$. The quantity $\alpha(u, i)$ sums the amount of flow of u in the first $i + 1$ labels of the interval I_ℓ . In the following analysis we assume that distances within I_ℓ are linear and ignore truncation. This can only hurt us. Following the reasoning in Lemma 3.1 the expected distance between u and v is equal to $\sum_{i=0}^{M'-1} |\min\{q, \alpha(u, i)\} - \alpha(v, i)|$

which we upper bound by $\sum_{i=0}^{M'-1} |\alpha(u, i) - \alpha(v, i)|$. We claim that

$$\sum_{i=0}^{M'-1} |\alpha(u, i) - \alpha(v, i)| \leq \sum_{e \in \text{CROSS}(I_\ell)} d_\ell(e) \bar{x}_e + \sum_{e \in \text{INT}(I_\ell)} d_{\text{lin}}(e) \bar{x}_e.$$

To prove this claim, we consider each link $e \in \text{CROSS}(I_\ell) \cup \text{INT}(I_\ell)$ and sum its contribution to the terms $q_i = |\alpha(u, i) - \alpha(v, i)|$, $0 \leq i < M'$. Let $e = (u, a, v, b) \in \text{INT}(I_\ell)$. It is clear that e contributes exactly \bar{x}_e to q_i if $a \leq i$ and $b > i$ or if $a > i$ and $b \leq i$. Otherwise its contribution is 0. Therefore, the overall contribution of e to $\sum q_i$ is $\bar{x}_e |a - b| = \bar{x}_e d_{\text{lin}}(e)$.

Now suppose $e = (u, a, v, b) \in \text{LCROSS}(I_\ell)$. Assume w.l.o.g. that $a \geq \ell$ and $b < \ell$, the other case is similar. Link e will contribute \bar{x}_e to $\alpha(u, i)$ for $a - \ell \leq i < M'$ and contributes 0 to $\alpha(v, i)$ for $0 \leq i < M'$ since b is outside the interval I_ℓ . Therefore, the contribution of e to q_i is \bar{x}_e for $a - \ell \leq i < M'$ and 0 otherwise. The overall contribution of e to $\sum q_i$ is $\bar{x}_e |\ell + M' - 1 - a| = d_\ell(e) \bar{x}_e$. A similar argument holds for the case when $e \in \text{RCROSS}(I_\ell)$. \square

Proof of Lemma 4.2. For a given iteration before which neither u nor v has a label, let $\Pr[u \wedge v]$, $\Pr[u \oplus v]$, and $\Pr[u \vee v]$ denote the probabilities that u and v are both assigned, exactly one of them is assigned, and at least one of them is assigned respectively. We upper bound the quantity $\mathbf{E}[d(L(u), L(v))]$ as follows. If u and v are separated in some iteration we upper bound their resulting distance by M . Using this simplification $\mathbf{E}[d(L(u), L(v))]$ is bounded by the quantity below.

$$\frac{\Pr[u \oplus v] \cdot M + \Pr[u \wedge v] \cdot \mathbf{E}[d(L(u), L(v)) | u \wedge v]}{\Pr[u \vee v]}$$

We upper bound the above expression as follows.

- We lower bound $\Pr[u \vee v]$ by $\Pr[u]$ which by Lemma 4.1 is equal to $M'/(k + M' - 1)$.
- We upper bound $\Pr[u \oplus v]$ by $\frac{1}{k + M' - 1} \sum_\ell \sum_{e \in \text{CROSS}(I_\ell)} \bar{x}_e$ using Lemma 4.3.
- By the definition of p_ℓ in Lemma 4.4 we get the following.

$$\Pr[u \wedge v] \mathbf{E}[d(L(u), L(v)) | u \wedge v] = \frac{1}{k + M' - 1} \sum_\ell p_\ell.$$

Putting all these together and using Lemma 4.4 to bound p_ℓ ,

$$\begin{aligned} \mathbf{E}[d(L(u), L(v))] &\leq \frac{1}{M'} \sum_\ell \left(\sum_{e \in \text{CROSS}(I_\ell)} (M + d_\ell(e)) \bar{x}_e + \sum_{e \in \text{INT}(I_\ell)} d_{\text{lin}}(e) \bar{x}_e \right) \\ &\leq \frac{1}{M'} \sum_e \bar{x}_e \left(\sum_{\text{CROSS}(I_\ell) \ni e} (M + d_\ell(e)) + \sum_{\text{INT}(I_\ell) \ni e} d_{\text{lin}}(e) \right). \end{aligned}$$

Lemma 4.5 below shows that

$$\sum_{\text{CROSS}(I_\ell) \ni e} (M + d_\ell(e)) + \sum_{\text{INT}(I_\ell) \ni e} d_{\text{lin}}(e) \leq (2M' + \max\{2M, (M')^2/M\}) d(e).$$

It follows that

$$\begin{aligned} \mathbf{E}[d(L(u), L(v))] &\leq \frac{1}{M'} (2M' + \max\{2M, (M')^2/M\}) \sum_e \bar{x}_e d(e) \\ &\leq (2 + \max\{2M/M', M'/M\}) \cdot d_{LP}(u, v). \end{aligned}$$

This finishes the proof. \square

Lemma 4.5 *Let $e = (u, i, v, j)$ be a link. Then*

$$\sum_{\text{CROSS}(I_\ell) \ni e} (M + d_\ell(e)) + \sum_{\text{INT}(I_\ell) \ni e} d_{\text{lin}}(e) \leq (2M' + \max\{2M, (M')^2/M\}) d(e).$$

Proof. Let $\bar{d}(e) = \sum_{\text{CROSS}(I_\ell) \ni e} (M + d_\ell(e)) + \sum_{\text{INT}(I_\ell) \ni e} d_{\text{lin}}(e)$. We evaluate $\bar{d}(e)$ separately for three different types of links based on their lengths. Recall that $M' \geq M$ and hence $d(e) \leq M \leq M'$ for all links e . Let e correspond to the link (u, i, v, j) in $H(u, v)$ and w.l.o.g. assume that $i \leq j$, the other case is similar. Also recall that $d_{\text{lin}}(e) = |i - j|$.

- $d_{\text{lin}}(e) \geq M'$: In this case it is clear that e is not an internal edge for any I_ℓ , hence $\sum_{\text{INT}(I_\ell) \ni e} d_{\text{lin}}(e) = 0$. Also $d(e) = M$. Therefore

$$\begin{aligned} \bar{d}(e) &= \sum_{\text{CROSS}(I_\ell) \ni e} (M + d_\ell(e)) \\ &= \sum_{\text{LCROSS}(I_\ell) \ni e} (M + d_\ell(e)) + \sum_{\text{RCROSS}(I_\ell) \ni e} (M + d_\ell(e)) \\ &= \sum_{\ell=j-M'+1}^j (M + M' + \ell - 1 - i) + \sum_{\ell=i-M'+1}^i (M + M' + \ell - 1 - i) \\ &\leq M'(2M + M') \\ &= M'(2 + M'/M)M = M'(2 + M'/M)d(e). \end{aligned}$$

- $d_{\text{lin}}(e) < M$: In this case $d(e) = d_{\text{lin}}(e)$.

$$\begin{aligned} \bar{d}(e) &= \sum_{\text{CROSS}(I_\ell) \ni e} (M + d_\ell(e)) + \sum_{\text{INT}(I_\ell) \ni e} d_{\text{lin}}(e) \\ &= \sum_{\text{LCROSS}(I_\ell) \ni e} (M + d_\ell(e)) + \sum_{\text{RCROSS}(I_\ell) \ni e} (M + d_\ell(e)) \\ &\quad + \sum_{\text{INT}(I_\ell) \ni e} d_{\text{lin}}(e) \\ &= \sum_{\ell=i+1}^j (M + M' + \ell - 1 - i) + \sum_{\ell=i-M'+1}^{j-M'} (M + M' + \ell - 1 - i) \\ &\quad + \sum_{\ell=j-M'+1}^i d_{\text{lin}}(e) \\ &\leq (2M' + 2M - d_{\text{lin}}(e))d_{\text{lin}}(e) \\ &\leq (2M' + 2M)d_{\text{lin}}(e) \\ &= M'(2 + 2M/M')d(e). \end{aligned}$$

- $M \leq d_{\text{lin}}(e) < M'$: In this case $d(e) = M$.

$$\begin{aligned}
\bar{d}(e) &= \sum_{\text{LCROSS}(I_\ell) \ni e} (M + d_\ell(e)) + \sum_{\text{RCROSS}(I_\ell) \ni e} (M + d_\ell(e)) + \sum_{\text{INT}(I_\ell) \ni e} d_{\text{lin}}(e) \\
&= \sum_{\ell=i+1}^j (M + M' + \ell - 1 - i) + \sum_{\ell=i-M'+1}^{j-M'} (M + M' + \ell - 1 - i) \\
&\quad + \sum_{\ell=j-M'+1}^i d_{\text{lin}}(e) \\
&= (2M' + 2M - d_{\text{lin}}(e))d_{\text{lin}}(e) \\
&\leq M'(2M + M') \\
&= M'(2 + M'/M)d(e).
\end{aligned}$$

□

4.1 The Truncated Quadratic Distance on the Line

Consider the label space $1, 2, \dots, k$ on the line where the distance function $d(i, j) = \min\{M, |i - j|^2\}$. This is the truncated version of the quadratic distance. We note that the quadratic distance is not a metric. However, as mentioned earlier, this distance function arises in image processing applications. In Subsection 3.3 we showed that our LP formulation gives an optimal solution for the quadratic distance on the line. For the truncated version of this distance we can use the algorithm from Section 4. By choosing $M' = \sqrt{M}$ we can show the following theorem.

Theorem 4 *The integrality gap of LP (I) for the truncated quadratic distance is $O(\sqrt{M})$.*

5 General Metrics

5.1 Integrality Gap on General Metrics

We now show that the integrality gap of the LP formulation is $O(\log k)$ on general metrics. This gives an alternative way of obtaining the result of Kleinberg-Tardos [26]. The latter algorithm uses the approach of first approximating the given metric probabilistically by a hierarchically well-separated tree (HST) metric [2] and then using an LP formulation to solve the problem on tree metrics. The Kleinberg-Tardos LP formulation has only an $O(1)$ integrality gap on HST metrics. Since any arbitrary k -point metric can be probabilistically approximated by an HST metric with an $O(\log k)$ distortion [16], their result follows. In contrast, our approach is based on directly using our LP formulation on the given general metric. As a first step, we use the LP solution to identify a deterministic HST metric approximation of the given metric such that the cost of our fractional solution on this HST metric is at most $O(\log k)$ times the LP cost on the original metric. This first step is done by using the following proposition from [16]. A weaker version was shown earlier in [9, 3] with a bound of $O(\log k \log \log k)$.

Proposition 1 *Let d be an arbitrary k -point metric and let α be a non-negative function defined over all pairs of points in the metric. Then d can be deterministically approximated by an HST metric d_T such that*

$$\sum_{i,j} \alpha(i, j) \cdot d_T(i, j) \leq O(\log k) \sum_{i,j} \alpha(i, j) \cdot d(i, j).$$

Given an optimal solution \bar{x} to our LP formulation, we apply Proposition 1 with the weight function $\alpha(i, j) = \sum_{(u,v) \in E} w(u, v) \cdot \bar{x}(u, i, v, j)$ for $1 \leq i, j \leq k$. Thus, $\alpha(i, j)$ is the fractional weight of edges between i and j . Let d_T denote the resulting HST metric. Since we are changing only the metric on the labels and not the assignments provided by the LP, the fractional solution is a feasible solution for this new metric and has cost at most $O(\log k) \cdot C^*$, where C^* is the optimal LP cost for the original metric. Thus, if we can now round our fractional solution on d_T by introducing only a constant factor increase in the solution cost, we will obtain an $O(\log k)$ -approximation algorithm. We prove this by showing that on any tree metric, our LP formulation is at least as strong as the Kleinberg-Tardos LP formulation (for tree metrics).

Given an edge weighted tree that defines the metric on the labels, we root the tree at some arbitrary vertex. Let T_a denote the subtree hanging off a vertex a in the rooted tree and let \mathcal{T} denote the set of all such trees. For any tree $T \in \mathcal{T}$ we denote by $\ell(T)$ the length of the edge leaving the root of T to its parent. Let $x_T(u)$ be compact notation for $\sum_{i \in T} x(u, i)$, the fractional assignment of the LP to labels in the subtree T . With this notation the LP formulation in [26] is as follows:

$$(KT) \quad \min \sum_{v \in V} \sum_{i=1}^k c(u, i) \cdot x(u, i) + \sum_{(u,v) \in E} w(u, v) \sum_{T \in \mathcal{T}} \ell(T) \cdot |x_T(u) - x_T(v)|$$

subject to:

$$\begin{aligned} \sum_{i=1}^k x(u, i) &= 1 & \forall u \in V \\ x(u, i) &\geq 0 & \forall u \in V \text{ and } i \in 1, \dots, k \end{aligned}$$

Let \bar{x} be a solution to our formulation (I). As we remarked in Section 3.1, for both (I) and (KT) above, the values $\bar{x}(u, i)$ completely determine the cost. We will show that cost of (KT) on \bar{x} is smaller than that of (I). Both linear programs (I) and (KT) coincide regarding the labeling cost. For each edge $(u, v) \in E$ we will show that LP distance for (u, v) is smaller in (I) than (KT). This is based on the following claim.

Claim 5.1 *For any feasible solution \bar{x} of (I),*

$$\sum_{T \in \mathcal{T}} \ell(T) \cdot |\bar{x}_T(u) - \bar{x}_T(v)| \leq \sum_{i,j} d(i, j) \cdot \bar{x}(u, i, v, j).$$

The proof of the above claim is straight forward and simply relies on the fact that the distance between two vertices in a tree metric is defined by the unique path between them. We omit the details. We obtain the following theorem from the above discussion.

Theorem 5 *The integrality gap of LP (I) on a k -point metric is $O(\log k)$.*

5.2 Reduction to Zero-Infinity Assignment Costs

We now describe a transformation for the general problem that essentially allows us to eliminate the label assignment cost function. This transformation reduces an instance with arbitrary label assignment cost function c to one where each label assignment cost is either 0 or ∞ . We refer to an instance of this latter type as a *zero-infinity* instance. Our transformation exactly preserves the cost of any feasible solution but in the process increases the number of labels by a factor of n . This provides some evidence that the label cost assignment function does not play a strong role in determining the approximability threshold of the metric labeling problem. In particular,

existence of a constant factor approximation algorithm for zero-infinity instances would imply a constant factor approximation algorithm for general instances as well.

From an instance $I = \langle c, d, w, L, G(V, E) \rangle$ of the general problem, we create an instance of the zero-infinity variant $I' = \langle c', d', w, L', G(V, E) \rangle$ as follows. We define a new label set $L' = \{i_u \mid i \in L \text{ and } u \in V\}$, i.e. we make a copy of L for each vertex in G . The new label cost assignment function is given by $c'(u, i_v) = 0$ if $v = u$ and ∞ otherwise. Thus, each vertex has its own copy of the original label set, and any finite cost solution to I' would assign each vertex a label from its own private copy.

Let $W_u = \sum_{(u,v) \in E} w(u, v)$ for any vertex $u \in V$. The new distance metric on L' is defined in terms of the original distance metric as well as the original label cost assignment function. For $i \neq j$ or $u \neq v$,

$$d'(i_u, j_v) = d(i, j) + \frac{c(u, i)}{W_u} + \frac{c(v, j)}{W_v},$$

and $d'(i_u, i_u) = 0$. It can be verified that d' is indeed a metric, and that any solution to instance I can be mapped to a solution to instance I' , and vice versa, in a cost-preserving manner. The proof of the following theorem follows in a straightforward manner from the above construction.

Theorem 6 *If there exists a $f(n, k)$ -approximation algorithm for zero-infinity instances of the metric labeling problem, then there exists a $f(n, nk)$ -approximation algorithm for general instances.*

In fact, there is an even simpler reduction to zero-infinity instances, conveyed to us by Julia Chuzhoy [10], which does not change the input metric but changes the graph in a very simple way. For each vertex v and label j such that $c(v, j) > 0$, add a new vertex z_{vj} to the graph for which

$$c(z_{vj}, \ell) = \begin{cases} \infty & \text{if } \ell = j, \\ 0 & \text{if } \ell \neq j. \end{cases}$$

A new edge (v, z_{vj}) is added to the graph. Let $\ell \neq j$ be the label that minimizes $d(j, \ell)$, such that $d(j, \ell) > 0$. The weight of the edge (v, z_{vj}) is set to

$$w(v, z_{vj}) = \frac{c(v, j)}{d(j, \ell)}.$$

Set $c(v, j) = 0$ for $1 \leq j \leq k$. We now obtain a new instance of the metric labeling problem by the above transformation. Note that the metric has not been altered. It is not hard to verify that this reduction preserves the value of an optimal solution, and that an r -approximation to the new instance also yields an r -approximation to the original instance. Hence the following theorem is obtained.

Theorem 7 (Chuzhoy [10]) *If there is a $f(n, k)$ -approximation algorithm for zero-infinity instances of metric labeling, then there is a $f(n + nk, k)$ -approximation algorithm for general instances.*

6 Conclusions

As we mentioned in Section 1, our LP is integral when G is a tree. We give a brief sketch of the idea. Consider an edge (u, v) in G where u is a leaf connected to v . If v is assigned a label i in some solution, then it is easy to see that an optimal assignment to u is to assign it a label j where $c(u, j) + w(u, v)d(i, j) = \min_k (c(u, k) + w(u, v)d(i, k))$. Hence the assignment to u is completely fixed by the assignment to v . We can eliminate u from G and incorporate the cost of u in the assignment cost of v as follows: set $c'(v, i) = c(v, i) + \min_k (c(u, k) + w(u, v)d(i, k))$. This transformation can be repeated until there is only one node left and the optimal solution

then is trivial. The same argument above can be used to show optimality of our LP formulation for trees. We leave the details to the reader.

Chuzhoy and Naor [11] recently obtained an $\Omega(\sqrt{\log k})$ -factor hardness of approximation for the metric labeling problem. The 0-extension problem generalizes the multi-way cut problem and is a special case of the metric labeling problem. As mentioned earlier, Călinescu, Karloff, and Rabani [8] established an $\Omega(\sqrt{\log k})$ lower bound on the integrality gap of the metric relaxation for the 0-extension problem. Fakcheroenphol et al. [15] improved the upper bound on the integrality gap of the metric relaxation to $O(\log k / \log \log k)$. It is worthwhile to study the integrality gap of our formulation for this restricted problem. See [1] for results in this direction.

The truncated quadratic distance function is of particular interest to applications in vision. Though this distance function does not form a metric, it is quite possible that a constant factor approximation is achievable. Here too our formulation might be of use in developing improved algorithms.

Acknowledgments We thank Olga Veksler and Mihalis Yannakakis for useful discussions. We thank Julia Chuzhoy for allowing us to include her reduction in Section 5.2.

References

- [1] A. Archer, J. Fakcheroenphol, C. Harrelson, R. Krauthgamer, K. Talwar, and É. Tardos. Approximate classification via earthmover metrics. In *Proceedings of the 15th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 1079–1087, 2004.
- [2] Y. Bartal. Probabilistic approximation of metric spaces and its algorithmic applications. In *Proceedings of the 37th Annual IEEE Symposium on Foundations of Computer Science*, pages 184–193, October 1996.
- [3] Y. Bartal. On approximating arbitrary metrics by tree metrics. In *Proceedings of the 30th Annual ACM Symposium on Theory of Computing*, pages 161–168, May 1998.
- [4] Y. Boykov, O. Veksler, and R. Zabih. Markov random fields with efficient approximations. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 648–655, 1998.
- [5] 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*, pages 377–384, 1999.
- [6] Y. Boykov, O. Veksler, and R. Zabih. A new algorithm for energy minimization with discontinuities. In *International Workshop on Energy Minimization Methods in Computer Vision and Pattern Recognition*, 1999.
- [7] G. Călinescu, H. Karloff, and Y. Rabani. An improved approximation algorithm for multiway cut. *J. Comput. System Sci.*, Vol. 60, pp. 564–574, 2000.
- [8] G. Călinescu, H. Karloff, and Y. Rabani. Approximation algorithms for the 0-extension problem. In *Proceedings of the Twelfth Annual Symposium on Discrete Algorithms*, pages 8–16, 2001.
- [9] M. Charikar, C. Chekuri, A. Goel, and S. Guha. Rounding via trees: deterministic approximation algorithms for group steiner trees and k -median. In *Proceedings of the 30th ACM Symposium on Theory of Computing*, pages 114–123, 1998.
- [10] J. Chuzhoy. Private communication, January 2001.

- [11] J. Chuzhoy and J. Naor. The Hardness of Metric Labeling. To appear, In *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science*, 2004.
- [12] F. S. Cohen. *Modeling and Applications of Stochastic Processes*, chapter Markov random fields for image modeling and analysis. Kluwer, 1986.
- [13] E. Dalhaus, D. Johnson, C. Papadimitriou, P. Seymour, and M. Yannakakis. The complexity of multi-terminal cuts. *SIAM J. Comput.*, 23(4):864–94, 1994.
- [14] R. Dubes and A. Jain. Random field models in image analysis. *J. of Applied Statistics*, 16, 1989.
- [15] J. Fakcheroenphol, C. Harrelson, S. Rao, and K. Talwar. An improved approximation algorithm for the 0-extension problem. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 257–265, 2003.
- [16] J. Fakcheroenphol, S. Rao, and K. Talwar. A tight bound on approximating arbitrary metrics by tree metrics. In *Proceedings of 35th Annual ACM Symposium on Theory of Computing (STOC)*, pp. 448–455, 2003.
- [17] D. Greig, B. Porteous, and A. Seheult. Exact maximum *a posteriori* estimation for binary images. *J. of Royal Statistical Society B*, 48, 1989.
- [18] A. Gupta and E. Tardos. Constant factor approximation algorithms for a class of classification problems. In *Proceedings of the 32nd Annual ACM Symposium on the Theory of Computing*, pages 652–658, 2000.
- [19] D. Hochbaum. An efficient algorithm for image segmentation, markov random fields and related problems. *Journal of the ACM*, Vol. 48, pp. 686-701, 2001.
- [20] A. J. Hoffman. On simple linear programming problems. *Convexity: Proc. of the 7th Symposium in Pure Mathematics*, volume 7, American Mathematical Society, 317–327, Providence, RI, 1961.
- [21] H. Ishikawa and D. Geiger. Segmentation by grouping junctions. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 125–131, 1998.
- [22] H. Ishikawa and D. Geiger. Mapping Image Restoration to a Graph Problem. In *Proceedings of the 1999 IEEE-EURASIP Workshop on Nonlinear Signal and Image Processing*, 1999.
- [23] D. Karger, P. Klein, C. Stein, M. Thorup, and N. Young. Rounding algorithms for a geometric embedding of multiway cut. In *Proceedings of the Thirtieth Annual ACM Symposium on Theory of Computing*, pages 668–678, 1999.
- [24] A. Karzanov. Minimum 0-extension of graph metrics. *Europ. J. Combinat.*, 19:71–101, 1998.
- [25] A. Karzanov. A combinatorial algorithm for the minimum $(2, r)$ -metric problem and some generalizations. *Combinatorica*, 18(4):549–569, 1999.
- [26] J. Kleinberg and E. Tardos. Approximation algorithms for classification problems with pairwise relationships: metric labeling and markov random fields. *Journal of the ACM*, Vol. 49, pp. 616-630, 2002.
- [27] S. Z. Li. *Markov Random Field Modeling in Computer Vision*. Springer-Verlag, 2nd Edition, 2001.
- [28] I. Milis. Task assignment in distributed systems using network flow methods. In *Proc. Conf. on Combinatorics and Computer Science (CCS)*. Lecture Notes in Computer Science, LNCS-1120, pp. 396-405, 1996.

- [29] G. Monge. Déblai et Ramblai. *Mem. de l'Academie des Sciences*, 1781.
- [30] S. Peleg, M. Werman, and H. Rom. A unified approach to the change of resolution: space and gray-level. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11:739-742, 1989.
- [31] S. Rachev. The Monge-Kantorovich mass transference problem and its stochastic applications. *Theory of Probability and its Applications*, 29(4):647–676, 1984.
- [32] S. Roy and I. Cox. A maximum-flow formulation of the N-camera stereo correspondence problem. In *Proceedings of the International Conference on Computer Vision*, pages 492–499, 1998.
- [33] Y. Rubner, C. Tomasi, and L. Guibas. The earth mover's distance as a metric for image retrieval. *International Journal of Computer Vision*, 40(2): 99–121, 2000.