Hardness of Embedding Metric Spaces of Equal Size

Subhash Khot⋆ and Rishi Saket⋆⋆

Georgia Institute of Technology
{khot,saket}@cc.gatech.edu

Abstract. We study the problem embedding an n-point metric space into another n-point metric space while minimizing distortion. We show that there is no polynomial time algorithm to approximate the minimum distortion within a factor of $\Omega((\log n)^{1/4-\delta})$ for any constant $\delta > 0$, unless NP $\subseteq$ DTIME($n^{\text{poly}(\log n)}$). We give a simple reduction from the METRIC LABELING problem which was shown to be inapproximable by Chuzhoy and Naor [10].

1 Introduction

Given an *embedding* $f : X \mapsto Y$ from a finite metric space (X, d_X) into another metric space (Y, d_Y), define

$$\text{Expansion}(f) := \max_{i,j \in X, i \neq j} \frac{d_Y(f(i), f(j))}{d_X(i,j)},$$

$$\text{Contraction}(f) := \max_{i,j \in X, i \neq j} \frac{d_X(i,j)}{d_Y(f(i), f(j))}.$$

The *distortion* of f is the product of Expansion(f) and Contraction(f).

The problem of embedding one metric space into another has been well studied in Mathematics, especially in the context of bi-lipschitz embeddings. Metric embeddings have also played an increasing role in Computer Science; see Indyk's survey [12] for an overview of their algorithmic applications. Much of the research in this area has focussed on embedding finite metrics into a *useful* target space such as ℓ_2, ℓ_1 or distributions over tree metrics. For example, Bourgain's Theorem [8] shows that every n-point metric space embeds into ℓ_2 with distortion $O(\log n)$. Fakcharoenphol, Rao, and Talwar [11], improving on the work of Bartal [7], show that every n-point metric embeds into a distribution over tree metrics with distortion $O(\log n)$. Recently, the class of *negative type* metrics has received much attention, as they arise naturally as a solution of a SDP-relaxation for the SPARSEST CUT problem. Arora, Lee, and Naor [3], building on techniques from Arora, Rao, and Vazirani [4], show that every n-point negative type metric embeds into ℓ_2 with distortion $O(\sqrt{\log n} \log \log n)$.

⋆ Research supported in part by Sloan Foundation Fellowship, Microsoft New Faculty Fellowship and NSF CAREER award CCF-0643626.

⋆⋆ Research supported in part by Subhash Khot's Sloan Foundation Fellowship and Microsoft New Faculty Fellowship.

Given two spaces (X, d_X) and (Y, d_Y), let $c_Y(X)$ denote the minimum distortion needed to embed X into Y. A natural computational problem is to determine $c_Y(X)$, exactly or approximately (call it the MIN-DISTORTION problem). Of course, the complexity of the problem depends on the nature of the two spaces. It is known that $c_{\ell_2}(X)$ can be computed in polynomial time for any n-point metric X, using a straightfoward SDP-relaxation. It is however NP-hard to determine whether $c_{\ell_1}(X) = 1$. When the target space is the line metric, the problem is hard, even to approximate within a factor n^β for some constant $0 < \beta < 1$, as shown by Badoiu, Chuzhoy, Indyk, and Sidiropoulos [6].

One variant of the problem is when X and Y are of equal size and explicitly given (call it the MIN-DISTORTION$^=$ problem). Though a very natural problem, it has not been studied until recently, and not much is known about its complexity. The problem was formulated by Kenyon, Rabani, and Sinclair [14] motivated by applications to shape matching and object recognition. Note that when both are shortest path metrics on n-vertex graphs, determining whether $c_Y(X) = 1$ is same as determining whether the two graphs are isomorphic. On the positive side, Kenyon *et al.* give a dynamic programming based algorithm when the metrics are one-dimensional and the distortion is at most $2+\sqrt{5}$. In this case, the bijection f between the two point sets must obey certain properties regarding how the points are relatively ordered, which enables efficient computation of all such possible embeddings in a recursive manner. On the negative side, Papadimitriou and Safra [15] showed that $c_Y(X)$ (when $|X| = |Y|$) is hard to approximate within a factor of 3, and the result holds even when both point sets are in $\mathbb{R}^3$. They give a clever reduction from GRAPH 3-COLORING.

In this paper, we prove the following inapproximability result.

Theorem 1 *(Inapproximability of* MIN-DISTORTION$^=$*) For any constant $\delta > 0$, there is no polynomial time algorithm to approximate the distortion required to embed an n-point metric space into another n-point metric space within a factor of $(\log n)^{1/4-\delta}$, unless* NP $\subseteq$ DTIME$(n^{\mathrm{poly}(\log n)})$.

2 Preliminaries

In this section we formally state the problem MIN-DISTORTION$^=$, and some of the tools we require for the construction in the next section. We will be concerned only with finite metric spaces i.e. metrics on finite sets of points.

Definition 1 *The problem of* MIN-DISTORTION$^=$ *is, given two n point metric spaces (X, d_X) and (Y, d_Y), computing an embedding f of X into Y with the minimum distortion.*

A related problem is METRIC LABELING which was introduced by Kleinberg and Tardos [13].

Definition 2 *The problem of* METRIC LABELLING *is the following. Given a weighted graph $G(V, E, \{w_e\}_{e\in E})$ and a metric space (X, d), along with a cost function c :*

$V \times X \mapsto \mathbb{R}$, *the goal is to find a mapping* h *of* V *to* X *such that the following quantity is minimized,*

$$\sum_{(u,v) \in E(G)} w(u,v) \cdot d(h(u), h(v)) + \sum_{u \in V} c(u, f(u)).$$

The problem is essentially of finding a 'labeling' of the vertices of the graph G with the points in the metric space X, so as to minimize the sum of the connection costs (cost of labelling vertices in G with points in X) and weighted 'stretch' of the edges of G. A special case of METRIC LABELING is the $(0, \infty)$-EXTENSION problem in which, essentially, the connection costs are 0 or ∞. It is formally defined as follows.

Definition 3 *The problem of* $(0, \infty)$-EXTENSION *is the following. Given a weighted graph* $G(V, E, \{w_e\}_{e \in E})$ *and a metric space* (X, d)*, along with a subset of allowed labels* $s(u) \subseteq X$ *for every vertex* $u \in V(G)$*, the goal is to find a mapping* h *of* V *to* X*, satisfying* $h(u) \in s(u)$ *for all* $u \in V(G)$*, such that the following quantity is minimized,*

$$\sum_{(u,v) \in E(G)} w(u,v) \cdot d(h(u), h(v)).$$

It has been shown that METRIC LABELING is equivalent to its special case of $(0, \infty)$-EXTENSION [9]. For METRIC LABELING Kleinberg and Tardos [13] obtained an $O(\log n \log \log n)$ approximation, where n is the size of the metric X, which was improved to $O(\log n)$ [11]. Chuzhoy and Naor [10] give a hardness of approximation factor of $(\log n)^{1/2-\delta}$ for METRIC LABELING. The instance that they construct is the $(0, \infty)$-EXTENSION version of METRIC LABELING, and moreover G is unweighted.

Definition 4 *A* 3-SAT(5) *formula* ϕ *is a* 3-CNF *formula in which every variable appears in exactly* 5 *clauses.*

The reduction in [10] starts with an instance of the MAX-3-SAT(5) problem in which, given a 3-SAT(5) formula ϕ, the goal is to find an assignment to the variables that satisfies the maximum number of clauses. The following is the well known PCP theorem [2] [5].

Theorem 2 *(PCP theorem): There exists a positive constant* ε *such that, given an instance* ϕ *of* MAX-3-SAT(5) *there is no polynomial time algorithm to decide whether there is an assignment to the variables of* ϕ *that satisfies all the clauses (Yes instance) or that no assignment satisfies more than* $1 - \varepsilon$ *fraction of the clauses, unless* P $=$ NP.

Overview of reduction. The construction of [10] yields a $(0, \infty)$-EXTENSION version with graph G and target metric H, where both G and H are unweighted graphs, with the metric on H being the shortest path metric. They show that in the Yes case, there is an embedding of $V(G)$ into $V(H)$ such that the end points of an edge of G are mapped onto end points of some edge of H, so that the stretch of every edge in G is 1. However, in the No case in any such embedding, for at least a constant fraction of edges of G, their end points are mapped onto pairs of vertices in H which are $\Omega(k)$ distance apart, where k is a parameter used in the construction so that $k = \theta(\log |V(H)|)^{1/2-\delta}$. This yields a hardness of approximation factor of $(\log n)^{1/2-\delta}$ for METRIC LABELING for any $\delta > 0$.

The main idea behind our reduction is to start with the instance constructed in [10] and view the graph G also as a metric, and the mapping h as an embedding of G into H. This makes sense since their instance is a $(0, \infty)$-EXTENSION version and moreover G is unweighted, in which case the quantity to be optimized is just the average stretch of every edge in G by the embedding h.

The reduction proceeds in three steps where the first step constructs the instance of [10]. The second involves adding additional points in the two metrics in order to equalize the number of points in both the metrics and third step enforces the constraints of allowed labels, by adding some more points in the two metrics. The entire construction and analysis is presented in the following section.

3 Reduction from MAX-3-SAT(5)

Our reduction from MAX-3-SAT(5) is based on the construction in [10] for the hardness of METRIC LABELING. We modify the instance of METRIC LABELING to obtain a 'gap' instance of MIN-DISTORTION$^{=}$. In this section we describe the entire construction and analysis of the hardness factor.

3.1 Construction

We describe the k-prover protocol used in [10]. Let ϕ be a MAX-3-SAT(5) formula, and let n be the number of clauses in ϕ. Let $P_1, P_2, \ldots, P_k$ be k provers (the parameter k will be set to $\theta(poly(\log n))$). The protocol is as follows,

- For each (i, j), $1 \leq i < j \leq k$, the verifier chooses a clause C_{ij} from ϕ uniformly at random, and randomly selects x_{ij} a distinguished variable from C_{ij}. P_i is sent C_{ij} and returns an assignment to the variables of C_{ij}. P_j is sent x_{ij} and returns as assignment to the variable x_{ij}. Every other prover is sent both C_{ij} and x_{ij} and returns assignments to all the variables of the clause. Hence, the verifier sends $\binom{k}{2}$ coordinates to each prover.
- The verifier accepts if the answers of all the provers are consistent and satisfy all the clauses of the query.

We will denote the set of random strings used by the verifier as R. For $r \in R$, let $q_i(r)$ be the query sent to prover P_i when r is the random string chosen by the verifier. Let $Q_i = \cup_r q_i(r)$ be the set of all possible queries of to P_i. For $q \in Q_i$, let $\mathcal{A}_i(q)$ be the set of all answers to q that satisfy all the clauses in q. Let P_i and P_j ($1 \leq i < j \leq k$) be any two provers, and $q_i \in Q_i$ and $q_j \in Q_j$ be two queries such that for some $r \in R$, $q_i = q_i(r)$ and $q_j = q_j(r)$. Let $A_i \in \mathcal{A}_i(q_i)$ and $A_j \in \mathcal{A}_j(q_j)$ be respective answers to these queries. Then, A_i and A_j are called *weakly consistent* if the assignments to C_{ij} in A_i and x_{ij} in A_j are consistent and satisfy the clause C_{ij}. They are called *strongly* *consistent* if they are also consistent in all the other coordinates and also satisfy all the other clauses. The following theorem is due to Chuzhoy and Naor [10].

Theorem 3 *There is a constant* $0 < \varepsilon < 1$ *such that if* ϕ *is a Yes instance, then there is a strategy of the* k *provers such that the verifier accepts, and if* ϕ *is a No instance, then for any pair of provers* P_i *and*P_j*, the probability that their answer is weakly consistent is at most* $1 - \frac{\varepsilon}{3}$.

We now construct our instance of MIN-DISTORTION$^=$ starting from a MAX-3-SAT(5) formula ϕ. The reduction proceeds in three stages. At each stage we obtain two metric spaces, with the first stage yielding exactly the METRIC LABELING instance of [10] and at the end of the third stage we obtain two metric spaces of equal size which constitute the desired instance of MIN-DISTORTION$^=$.

STEP I. In this step we construct two graphs G_1 and H_1 which are same as the graphs constructed in the METRIC LABELING instance of [10]. The construction of the graph G_1 is as follows.

– For every prover P_i and every $q \in Q_i$, there is a vertex $v(i, q)$.
– For every random string $r \in R$, there is a vertex $v(r)$.
– There is an edge of length 1 between $v(r)$ and $v(i, q)$ if $q = q_i(r)$.

Call the vertices of G_1 as 'query' vertices. The graph H_1 is constructed in the following manner.

– For every i ($1 \le i \le k$), every $q \in Q_i$, and answer $A_i \in \mathcal{A}_i(q)$, there is a vertex $v(i, q, A_i)$.
– For every $r \in R$ and every pairwise strongly consistent answers $A_1, A_2, \dots, A_k$ to the queries $q_1(r), q_2(r), \dots, q_k(r)$ respectively (where $A_i \in \mathcal{A}_i(q_i(r))$ for $1 \le i \le k$), there is a vertex $v(r, A_1, A_2, \dots, A_k)$.
– There is an edge of length 1 between $v(i, q, A_i)$ and $v(r, A_1, A_2, \dots, A_k)$ if $q = q_i(r)$ and the ith coordinate of the tuple $(A_1, A_2, \dots, A_k)$ is A_i.

The vertices of H_1 will be referred to as 'label' vertices. Figure 1 shows the local structure of G_1 and H_1. It can be seen that for every 'query' vertex u in G_1, there is a set of 'label' vertices $s(u)$ in H_1, such that $\{s(u)\}_{u \in V(G_1)}$ is a partition of $V(H_1)$. For the sake of convenience, we modify our notation to let $V(G_1) = \{u_1, u_2, \dots, u_N\}$, where $N = |V(G_1)|$ and $V(H_1) = \cup_{i=1}^{N} s(u_i)$, and let $\ell_1^i, \ell_2^i, \dots, \ell_{m_i}^i$ be the elements of $s(u_i)$, where $m_i = |s(u_i)|$. Also note that there is an edge between the sets $s(u_i)$ and $s(u_j)$ in H_1 only if there is an edge between u_i and u_j in G_1. The graphs G_1 and H_1 with the partition $\{s(u_i)\}_{u_i \in V(G_1)}$ constitute the instance of METRIC LABELING of [10]. They show that in the Yes case, there is a labeling of every $u_i \in V(G_1)$ with a label from $s(u_i)$ such that each edge in G_1 is mapped to an edge in H_1, while in the No case in every such labeling, a constant fraction of edges in G_1 are mapped to pairs of vertices $\Omega(k)$ distance apart in H_1. Figure 2 illustrates this structure of G_1 and H_1. Let d_{G_1} and d_{H_1} denote the shortest path metric on G_1 and H_1 respectively. Note that eventually we want two metric spaces of equal cardinality, in the next step of the

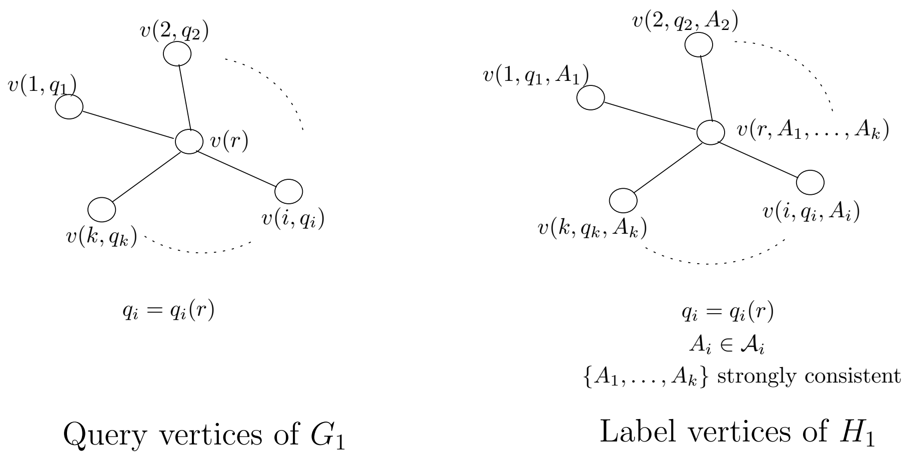

Fig. 1. Vertices of G_1 and H_1. All edges are of length 1.

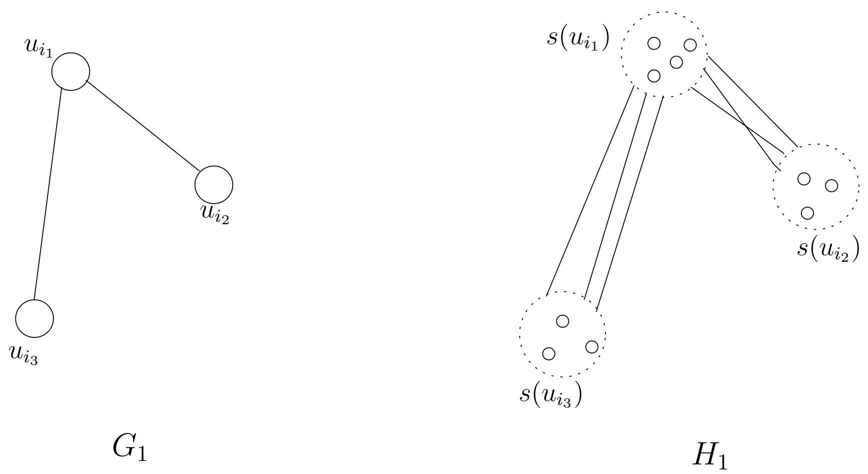

Fig. 2. Structure of G_1 and H_1.

construction we achieve that goal.

STEP II. We first modify G_1 as follows. For every vertex u_i ($1 \le i \le N$) in G_1, we add $m_i - 1$ vertices $t_1^i, t_2^i, \ldots t_{m_i-1}^i$ and add an edge from each t_j^i to u_i of length $\sqrt{k}$. Let G_2 be the new graph created in this process and d_{G_2} be the shortest path metric on G_2. The transformation is shown in Figure 3.

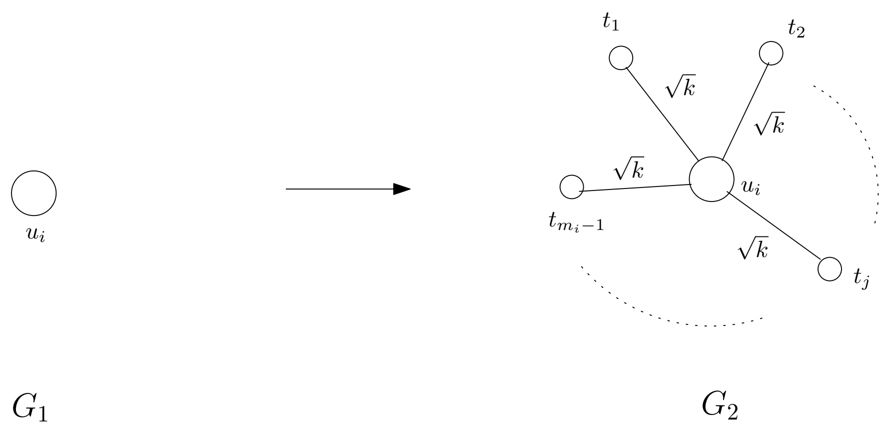

Fig. 3. Transformation from G_1 to G_2.

We truncate d_{H_1} from below to $\sqrt{k}$ and from above to $10k$, i.e. if the distance between a particular pair of points is less than $\sqrt{k}$ then it is set to $\sqrt{k}$ and if it is greater than $10k$ then it is set to $10k$, otherwise it remains the same. We also truncate d_{G_1} from above to $10k$, but not from below. Let H_2 be the new resultant graph and d_{H_2} be this new metric on $V(H_2)$. Observe that G_2 and H_2 have the same number of vertices.

STEP III. We now have two graphs G_2 and H_2 with equal number of vertices. However, we desire that any 'good' embedding of vertices of G_2 into H_2 map the set $\{u_i, t_1^i, \ldots, t_{m_i-1}^i\}$ onto the set $s(u_i)$. For this we add new vertices to both the graphs in the following manner. Let $\eta_i = 2^{-iN}$. For every i ($1 \le i \le N$), we add m_i vertices $\alpha_1^i, \alpha_2^i, \ldots, \alpha_{m_i}^i$ to G_2, where $\alpha_{m_i}^i$ is at a distance η_i from u_i and similarly α_j^i is at a distance of η_i from t_j^i for each $j = 1, 2, \ldots, m_i - 1$. Similarly, in the graph H, we add vertices vertices $\beta_1^i, \beta_2^i, \ldots, \beta_{m_i}^i$, such that β_j^i is at a distance of η_i from ℓ_j^i for $j = 1, 2, \ldots, m_i$. Let G_3 and H_3 be the graphs formed in this manner (refer to Figure 4), with the metrics d_{G_3} and d_{H_3} extended to the new vertices according to the distances defined above. Note that we have ensured that G_3 and H_3 have equal number of vertices. We output G_3 and H_3 along with the respective metrics d_{G_3} and d_{H_3} on $V(G_3)$ and $V(H_3)$ respectively, as the instance of MIN-DISTORTION$^=$.

Remark. Note that all the distances in G_2 and H_2 were at least 1. In G_3 and H_3 each, we added m_i edges of length 2^{-iN}, for $i = 1, \ldots, N$. So, any bijection between $V(G_3)$

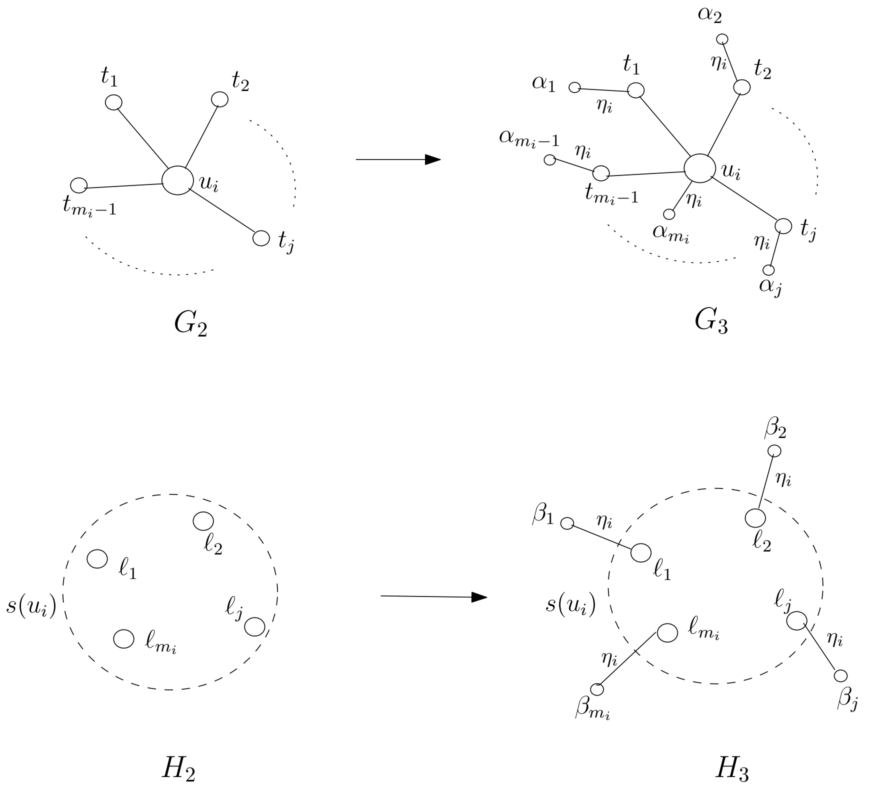

Fig. 4. Transformation from G_2 to G_3 and H_2 to H_3.

and $V(H_3)$ must map edges of length 2^{-iN} in G_3 to edges of length 2^{-iN} in H_3 for $i = 1, \ldots, N$, otherwise at least one edge will be stretched or contracted by a factor of at least 2^N. Therefore, this ensures that in any embedding of G_3 into H_3, for any given i ($1 \leq i \leq N$) the set $\{\alpha_1^i, \ldots, \alpha_{m_i}^i\}$ is mapped onto the set $\{\beta_1^i, \ldots, \beta_{m_i}^i\}$ and the set $\{u_i, t_1^i, \ldots, t_{m_i-1}^i\}$ is mapped onto the set $\{\ell_1^i, \ldots, \ell_{m_i}^i\}$, otherwise the distortion is at least 2^N.

3.2 Analysis

Yes instance. Suppose that the MAX-3-SAT(5) formula had a satisfying assignment, say σ to the variables. We will construct an embedding f of $V(G_3)$ into $V(H_3)$ such that no edge in G_3 is contracted, whereas any edge in G_3 is stretched by at most a factor of $O(\sqrt{k})$, which leads to an embedding with distortion at most $O(\sqrt{k})$. For any i ($1 \leq i \leq N$), consider the vertex u_i. We recall that u_i was a 'query' vertex, and $s(u_i)$ is the set of 'label' vertices corresponding to u_i. We set $f(u_i)$ to be a label vertex in $s(u_i)$ given by the satisfying assignment σ. The map f takes $t_1^i, t_2^i, \ldots, t_{m_i-1}^i$ arbitrarily to the remaining $m_i - 1$ vertices in $s(u_i)$. The vertex α_j^i ($1 \leq j \leq m_i - 1$) is mapped to $\beta_{j'}^i$ such that $\beta_{j'}^i$ at distance η_i to $f(t_j^i)$. Similarly, the vertex $\alpha_{m_i}^i$ is mapped to $\beta_{j''}^i$ which is at a distance of η_i from $f(u_i)$.

Consider any two vertices u_i and $u_{i'}$ in in G_1 that are adjacent. Since σ is a satisfying assignment, we have that the vertices $f(u_i)$ and $f(u_{i'})$ are also adjacent in H_1. And since adjacent vertices remain adjacent in G_3 and adjacent vertices in H_1 have distance $\sqrt{k}$ in H_3, therefore the stretch of the edge (u_i, u_j) in G_3 is $\sqrt{k}$. The edge (u_i, t_j^i) ($1 \leq j \leq m_i - 1$) has length $\sqrt{k}$ and since the diameter of the metric on H_3 (and G_3) is $O(k)$, the stretch of the edge is at most $O(\sqrt{k})$. Also, clearly the distances of the edges $(t_1^i, \alpha_1^i), \ldots, (t_{m_i-1}^i, \alpha_{m_i-1}^i), (u_i, \alpha_{m_i})$ are not stretched or contracted. One key fact that we have utilized is that there is an edge (in H_1) between $s(u_i)$ and $s(u_{i'})$ only if u_i and $u_{i'}$ are adjacent in G_1. This ensures that there is no edge (in G_1) contracts, which guarantees that no distance in G_3 contracts. Therefore, there is an embedding of the vertices of G_3 into H_3 with distortion at most $O(\sqrt{k})$.

No instance. Suppose that the MAX-3-SAT(5) formula has no satisfying assignment that satisfies more that $1 - \varepsilon$ fraction of the clauses. In this case, as a consequence of Proposition 4.4 and Lemma 4.5 of [10], we have that in every mapping g of $V(G_1)$ into $V(H_1)$, such that $g(u) \in s(u)$, for all $u \in V(G)$, there is a constant fraction of edges of G_1 whose end points are mapped to pairs of vertices of H_1 that are $\theta(k)$ distance apart in H_1. Since we truncate the metric on d_{H_1} from above to $\Omega(k)$, it is also true for the truncated metric on H_1.

As noted in the remark after step III of the construction, any embedding f of $V(G_3)$ into $V(H_3)$ which does not satisfy the condition that $f(u_i) \in s(u_i)$ for all $i = 1, 2, \ldots, N$ incurs a distortion of at least $2^N > k$. Therefore we may assume that the above condition holds for f. Now, using what we stated above, we get that there is a unit distance edge $(u_i, u_{i'})$ is stretched by a factor of $\Omega(k)$, i.e. $d_{H_3}(f(u_i), f(u_{i'}) \geq$

$\Omega(k) \cdot d_{G_1}(u_i, u_{i'})$. Note that unit distances in G_1 are preserved in G_3 and the truncation of distances of H_1 to $\sqrt{k}$ from below only helps us. Therefore, in this case the distortion is at least $\Omega(k)$.

Construction size. Since each query consists of at most k^2 clauses, the size of R is at most $3 \cdot n^{k^2}$ and there are at most 7^{k^2} answers to each query. Also, since in steps II and III we blow up the size only polynomially, we have that the total size of the construction is at most $n^{O(k^2)}$.

Hardness factor. In the Yes case we have an embedding of G_3 into H_3 with distortion at most $O(\sqrt{k})$, while in the No case any embedding has distortion at least $\Omega(k)$. Therefore, we have a hardness factor $\Omega(\sqrt{k})$ and choosing $k = \text{poly}(\log n)$, we have $k = \Omega((\log |V(G_3)|)^{1/2-\delta})$. Therefore, we get that there is no polynomial time algorithm to approximate the distortion of embedding two n-point metrics within a factor of $(\log(n))^{1/4-\delta}$ of the optimal for any positive constant δ, unless NP $\subseteq$ DTIME($n^{\text{poly}(\log n)}$). This proves Theorem 1.

4 Conclusion

In this paper we have shown a hardness factor of $(\log n)^{1/4-\delta}$ for approximating the distortion required to embed two general n point metrics. An interesting question is whether such a superconstant lower bound also holds for constant dimensional metrics. Papadimitriou and Safra [15] prove a 3 factor hardness of approximating the distortion for 3-dimensional metrics. Extending such a constant factor hardness to 2-dimensional, or 1-dimensional metrics is an important question. On the algorithmic side the only positive result is for the case when the metrics are 1-dimensional and the optimum distortion is at most $2 + \sqrt{5}$ [14], and extending it to higher dimensions remains an open question. For general n point metrics, no non-trivial upper bound is known, and it would be interesting if a reasonable upper bound can be derived for this problem.

References

1. S. Arora. Polynomial time approximation schemes for Euclidean Traveling Salesman and other Geometric problems. *Journal of the ACM*, 45(5):753–782, 1998.
2. S. Arora, C. Lund, R. Motawani, M. Sudan, and M. Szegedy. Proof verification and the hardness of approximation problems. *Journal of the ACM*, 45(3):501–555, 1998.
3. S. Arora, J. R. Lee, and A. Naor. Euclidean distortion and the sparsest cut. In *Proc. STOC*, 553–562, 2005.
4. S. Arora, S. Rao, and U. V. Vazirani. Expander flows, geometric embeddings and graph partitioning. In *Proc. STOC*, 222–231, 2004.
5. S. Arora and S. Safra. Probabilistic checking of proofs : A new characterization of NP. *Journal of the ACM*, 45(1):70–122, 1998.
6. M. Badoiu, J. Chuzhoy, P. Indyk, and A. Sidiropoulos. Low-distortion embeddings of general metrics into the line. In *Proc. STOC*, 225–233, 2005.
7. Y. Bartal. On approximating arbitrary metrices by tree metrics. In *Proc. STOC*, 161–168, 1998.

8. J. Bourgain. On Lipschitz embeddings of finite metrics in Hilbert space. *Israel Journal of Mathematics*, 52:46–52, 1985.
9. C. Chekuri, S. Khanna, J. Naor, and L. Zosin. Approximation algorithms for the metric labeling problem via a new linear programming formulation. In *Proc. SODA*, 109–118, 2001.
10. J. Chuzhoy and J. Naor. The hardness of Metric Labeling. In *Proc. FOCS*, 108–114, 2004.
11. J. Fakcharoenphol, S. Rao, and K. Talwar. A tight bound on approximating arbitrary metrics by tree metrics. *J. Comput. Syst. Sci*, 69(3):485–497, 2004.
12. P. Indyk. Algorithmic applications of low-distortion embeddings. In *Proc. FOCS*, 10–33, 2001.
13. J. Kleinberg and E. Tardos. Approximation algorithms for classification problems with pairwise relationships: metric labeling and Markov random fields. *Journal of the ACM*, 49:616–630, 2002.
14. C. Kenyon, Y. Rabani, and A. Sinclair. Low distortion maps between point sets. In *Proc. STOC*, 272–280, 2004.
15. C. H. Papadimitriou and S. Safra. The complexity of low-distortion embeddings between point sets. In *Proc. SODA*, 112–118, 2005.