When Hamming Meets Euclid: The Approximability of Geometric TSP and MST

[Extended Abstract]

Luca Trevisan*

Abstract

*We prove that the Traveling Salesperson Problem (*MIN TSP*) and the Minimum Steiner Tree Problem (*MIN ST*) are* Max SNP*-hard (and thus* NP*-hard to approximate within some constant $r > 1$) even if all cities (respectively, points) lie in the geometric space $\mathcal{R}^n$ (n is the number of cities/points) and distances are computed with respect to the l_1 (rectilinear) metric.*

The TSP hardness results also hold for any l_p metric, including the Euclidean metric, and in $\mathcal{R}^{\log n}$.

The running time of Arora's approximation scheme for geometric MIN TSP *in $\mathcal{R}^d$ is doubly exponential in d. Our results imply that this dependance is necessary unless* NP *has sub-exponential algorithms.*

We also prove, as an intermediate step, the hardness of approximating MIN TSP *and* MIN ST *in Hamming spaces. The reduction for* MIN TSP *uses error-correcting codes and random sampling; the reduction for* MIN ST *uses the integrality property of MIN-CUT. The only previous non-approximability results for* MIN TSP *and* MIN ST *involved metrics where all distances are 1 or 2.*

*Dept. of Computer Science, University of Geneva, Rue General-Dufour 24, CH 1211 Geneva, Switzerland. Email: `trevisan@cui.unige.ch`. Work partly done at the University of Rome "La Sapienza" and while visiting the Technical University of Catalunya (UPC).

1. Introduction

Given a metric space and a set U of points into it, the Traveling Salesperson Problem (MIN TSP) is to find a closed tour of shortest total length visiting each point exactly once, while the Minimum Steiner Tree Problem (MIN ST) is to find the minimum cost tree connecting all the points of U; the tree can possibly contain points not in U, that are called "Steiner points".

Both problems are among the most classical and most widely studied ones in Combinatorial Optimization, Operations Research and Computer Science during the past few decades, and before. Important special cases arise when the metric space is $\mathcal{R}^k$ and the distance is computed according to the ℓ_1 norm (the *rectilinear* case) or the ℓ_2 norm (the *Euclidean* case).

We establish the first non-approximability results for this class of problems. As an intermediate step, we prove that they are hard to approximate also in *Hamming spaces*. The approximability of the Hamming versions of MIN TSP and MIN ST seems to have never been considered before. Our main contributions are: (i) the identification of this class of metric spaces as the "right" one to prove hardness in more natural geometric spaces, and (ii) the derivation of combinatorial results that could have some independent interest.

Our techniques prove hardness of approximation for other problems mentioned in Arora's paper [Aro96] on approximation schemes for geometric problems.

We now state and discuss our results for MIN TSP and MIN ST.

1.1. The Traveling Salesperson Problem

Interest in the MIN TSP started during the 1930's. In 1966, the (already) long-standing failure of developing an efficient algorithm for the MIN TSP led Edmonds [Edm66] to conjecture that the problem is not in P: this is sometimes referred to as the first statement of the P $\neq$ NP conjecture. See the book of Lawler et al. [LLKS85] for a very complete survey on MIN TSP. Here we will only review the results that are relevant for the present paper. The MIN TSP is NP-hard even if the cities are restricted to lie in $\mathcal{R}^2$ and the distances are computed according to the ℓ_2 norm [GGJ76, Pap77]. Due to such a negative result, research concentrated on developing good heuristics. Recall that an r-approximate algorithm ($r > 1$) is a polynomial-time heuristic that is guaranteed to deliver a tour whose cost is at most r times the optimum cost. A 3/2-approximation algorithm that works for any metric space is due to Christofides [Chr76]. In twenty years of research no improvement of this bound had been found, even in the restricted case of geometric metrics.

In the late 1980's, the emergence of the theory of Max SNP-hardness [PY91] gave a means of possibly understanding this lack of results. Indeed, Papadimitriou and Yannakakis [PY93] proved that the MIN TSP is Max SNP-hard even when restricted to metric spaces (as we shall see later, the result also holds for a particularly restricted class of metric spaces), and thus a constant $\epsilon > 0$ exists such that metric MIN TSP cannot be approximated within a factor $(1+\epsilon)$ in polynomial time, unless P $=$ NP. The complexity of approximating MIN TSP in the case of geometric metrics remained a major open question. In his PhD thesis, Arora noted that proving the Max SNP-hardness of Euclidean MIN TSP in $\mathcal{R}^2$ should be very difficult, but that this could perhaps be done in $\mathcal{R}^{k(n)}$ for sufficiently large $k(n)$ [Aro94, Chapter 9]. In [GKP95], Grigni, Koutsopias and Papadimitriou proved that the restriction of the MIN TSP to shortest paths metrics of planar graphs can be approximated within $(1+\epsilon)$ in time $n^{O(1/\epsilon)}$. Such an approximation algorithm is called a *Polynomial Time Approximation Scheme (PTAS)*. This result led Grigni et al. [GKP95] to conjecture that Euclidean MIN TSP has a PTAS in $\mathcal{R}^2$. They again posed the question of determining the approximability of the problem for higher dimensions. In a very recent breakthrough, Arora [Aro96] developed a PTAS for the MIN TSP in $\mathcal{R}^2$ under any ℓ_p metric. Such an algorithm also works in higher dimensional spaces and, in particular, it runs in time $n^{\tilde{O}((\log^{d-2} n)/\epsilon^{d-1})}$ in $\mathcal{R}^d$. Note that the dependence of the running time on d is doubly exponential. In a preliminary version of [Aro96] Arora asked if it was possible to develop a PTAS for Euclidean MIN TSP in $\mathcal{R}^n$ or if, at least, it was possible to have the running time being singly exponential in d, e.g. $n^{O(d/\epsilon)}$.

Our Results. In this paper we essentially answer negatively to both questions. We prove that MIN TSP in $\mathcal{R}^{\log n}$ is Max SNP-hard using any ℓ_p metric. It follows from our result that there cannot be a PTAS for these problems (unless P= NP) and that there cannot be $(1+\epsilon)$-approximate algorithms in $\mathcal{R}^d$ running in time $n^{O(d/\epsilon)}$ for any $\epsilon > 0$, unless NP $\subseteq$ DTIME$(n^{O(\log n)})$.

The Max SNP-hardness is proved by means of a reduction from the version of the metric MIN TSP that was shown to be Max SNP-hard in [PY93]. The reduction uses a mapping (see Lemma 7) of the metric spaces of [PY93] into Hamming spaces and the observation (see Proposition 3) that, for elements of $\{0,1\}^n$ a "gap" in the Hamming distance is preserved if distances are computed according to a ℓ_p metric. Our mapping of the metric spaces of [PY91] into Hamming spaces is *not* an *approximate isometry*, that is, it does *not* preserve distances up to negligible distorsion. We also suspect that such kind of mapping would be provably impossible. Instead, our mapping introduces a fairly high (yet constant) distorsion, but satisfies an additional condition: cities at distance one are mapped into cities at distance $\approx D_1$; cities at distance 2 are mapped into cities at distance $\approx D_2$, and D_2 is larger than D_1 by a multiplicative constant factor. This is sufficient to make the mapping be an *approximation preserving reduction*. Our mapping uses *error-correcting codes* (namely, Hadamard codes) to map cities into an $O(n)$-dimensional Hamming space, and then *random sampling* to reduce the number of dimensions to $O(\log n)$.

The Minimum k-Cities Traveling Salesman Problem (MIN k-TSP) and the Minimum Degree-Restricted Steiner Tree Problem (two problems mentioned in Arora's paper [Aro96] on approximation schemes for geometric problems) are generalizations of the MIN TSP. The hardness results that we prove for MIN TSP clearly extend to them.

1.2. The Minimum Steiner Tree Problem

The origins of the MIN ST problem seem to be even more remote than the MIN TSP's ones: the case when $|U| = 3$ and the metric space is $\mathcal{R}^2$ with the ℓ_2 norm has been studied by the Italian mathematician Torricelli (a student of Galilei's) in 17th century. Reportedly, Gauss had an interest to this problem as well. Recent results about this problem are similar to the ones for MIN TSP: exact optimization is NP-hard in $\mathcal{R}^2$ both in the Rectilinear (ℓ_1) case [GJ77] and in the Euclidean (ℓ_2) case [GGJ77]. Constant-factor approximation is achievable in any metric space (the best factor should be 1.644 due to Karpinski and Zelikovsky [KZ95]), in general metric spaces the problem is Max SNP-hard [BP89], Arora's algorithm achieves approximation $(1+\epsilon)$ in $\mathcal{R}^d$ in time $n^{\tilde{O}((\log^{d-2} n)/\epsilon^{d-1})}$. No non-approximability result was known for geometric versions of the problem.

Our Results. We prove the Max SNP-hardness of the problem in $\mathcal{R}^n$ under the ℓ_1 norm. As a preliminary step, we prove the hardness of the problem restricted to Hamming spaces. The latter hardness is proved via a reduction from the Minimum Vertex Cover problem (MIN VC) restricted to triangle-free graphs of maximum degree 3. The Max SNP-hardness of this very restricted version of MIN VC is proved in this paper and could be used as a starting point for other non-approximability results. The reduction from MIN VC to Hamming MIN ST uses a combinatorial result (Claim 14) stating that for an instance where all points have weight[1] 2 or 0, if a technical condition is satisfied, there exists an optimum solution where all Steiner points have weight 1. We remark that there exists an instance of Hamming Steiner Tree where all the points have weight 3 or 0 and such that an optimum solution must contain a Steiner point of weight at least 4. Thus, our combinatorial result cannot be generalized too much. Reducing from Hamming Steiner Tree to Rectilinear Steiner Tree requires another combinatorial result (Theorem 16): for an instance where all the points are in $\{0,1\}^n \subset \mathcal{R}^n$, there exists an optimum solution where all the Steiner points lie in $\{0,1\}^n$. We prove this fact using the *integrality property* of Min-CUT linear programming relaxations.

[1]For a vector $u \in \{0,1\}^n$, its weight is defined as the number of non-zero coefficients, e.g. the weight of $(0,1,1,0,1)$ is three.

1.3. Discussion

We give the first non-approximability results for geometric versions of network optimization problems. For Euclidean MIN TSP, it is still a major open question whether a PTAS exists in $\mathcal{R}^d$ for each *fixed* d. Note that the case of fixed d is not Max SNP-hard, unless $\mathsf{NP} \subseteq \mathsf{DTIME}(n^{\text{poly}\log n})$. A more general question is what is the best asymptotic relation between number of dimensions and running time. A running time $2^{2^d/\epsilon}\text{poly}(n)$ would be compatible with our results, but if we believe that NP does not admit sub-exponential algorithms (i.e. $\mathsf{NP} \not\subseteq \mathsf{DTIME}(2^{n^{o(1)}})$), then even a running time $2^{2^{o(d)}/\epsilon}\text{poly}(n)$ is unfeasible. Alternatively, our non-approximability result could be extended to $\mathcal{R}^{\log n/\log\log n}$. Much more consistent improvements are possible for MIN ST, however our results at least state very clearly that the number of dimensions *does matter* in the running time of an approximation scheme for these geometric problems.

We feel that one important contribution of this paper is the recognition of Hamming spaces as a class of metric spaces that seem to retain most of the hardness of general metrics while having a nice combinatorial structure. We believe that other non-approximability results could be established using Hamming spaces as intermediate steps. We also think that it should be worth trying to improve Christofides algorithm in Hamming spaces. While the well-behaved structure of Hamming spaces should not make this task impossible, it is likely that such an improved algorithm could give useful ideas for more general cases.

2. Preliminaries

We denote by $\mathcal{R}$ the set of real numbers. For an integer n we denote by $[n]$ the set $\{1,\ldots,n\}$. For a vector $\vec{a} \in \mathcal{R}^n$ and an index $i \in [n]$, we denote by $\vec{a}[i]$ the i-th coordinate of $\vec{a}$, Given an instance x of an optimization problem A, we will denote by $\mathsf{opt}_A(x)$ the cost of an optimum solution for x, we will also typically omit the subscript. For a feasible solution y (usually a tour or a tree) of an instance x of an optimization problem A, we denote its cost by $\mathsf{cost}_A(x,y)$

or, more often, as $\text{cost}(y)$. In this paper we will use the notions of L-reduction and Max SNP-hardness. Max SNP is a class of constant-factor approximable optimization problems that includes MAX 3SAT, we refer the reader to [PY91] for the formal definition.

Definition 1 (L-reduction) *An optimization problem A us said to be L-reducible to an optimization problem B if two constants α and β and two polynomial-time computable functions f and g exist such that*

1. *For an instance x of A, $x' = f(x)$ is an instance of B, and it holds $\text{opt}_B(x') \le \alpha \text{opt}_A(x)$.*
2. *For an instance x of A, and a solution y' feasible for $x' = f(x)$, $y = g(x, y')$ is a feasible solution for x and it holds $|\text{opt}_A(x) - \text{cost}_A(x,y)| \le \beta|\text{opt}_B(x') - \text{cost}_B(x', y')|$.*

We say that an optimization problem is Max SNP-hard if all Max SNP-problems are L-reducible to it. From [ALM+92] it follows that if a problem A is Max SNP-hard, then a constant $\epsilon > 0$ exists such that $(1+\epsilon)$-approximating A is NP-hard.

A function $d : U \times U \to \mathcal{R}$ is a *metric* if it is non-negative, if $d(u,v) = 0$ iff $u = v$, if it is symmetric (i.e. $d(u,v) = d(v,u)$ for any $u, v \in U$), and it satisfies the *triangle inequality* (i.e. $d(u,v) \le d(u,z) + d(z,v)$ for any $u, v, z \in U$).

Definition 2 ($(1,2) - B$ metrics) *A metric $d : U \times U \to \mathcal{R}$ is a $(1,2) - B$ metric if it satisfies the following properties:*

1. *For any $u, v \in U$, $u \ne v$, $d(u,v) \in \{1,2\}$.*
2. *For any u, at most B elements of U are at distance 1 from u.*

Papadimitriou and Yannakakis [PY93] have shown that a constant $B_0 > 0$ exists such that the MIN TSP is Max SNP-hard even when restricted to $(1,2) - B_0$ metrics.

For an integer $p \ge 1$, the ℓ_p norm in $\mathcal{R}^n$ is defined as $||(u_1, \ldots, u_n)||_p = \left(\sum_{i=1}^n |u_i|^p\right)^{(1/p)}$. The distance induced by the ℓ_p norm is defined as $d_p(\vec{u}, \vec{v}) = ||\vec{u} - \vec{v}||_p$. For a positive integer n, we denote by d_H^n the Hamming metric in $\{0,1\}^n$ (we will usually omit the superscripts). We will make some use of the following fact.

Proposition 3 *Let $\vec{u}, \vec{v} \in \{0,1\}^n \subseteq \mathcal{R}^n$. Then $d_p(\vec{u}, \vec{v}) = d_H(\vec{u}, \vec{v})^{1/p}$.*

Before starting with the presentation of our results, we make the following important caveat.

Remark 4 *In some of the proofs of this paper we implicitly make the (unrealistic) assumption that arbitrary real numbers can appear in an instance and that arithmetic operations (including squared roots) can be computed over them in constant time. However, our results still hold if we instead assume that numbers are rounded and stored in a floating point notation using $O(\log n)$ bits. This fact follows from a minor modification of the argument used in [Aro96] to reduce a general instance of Euclidean TSP or Steiner Tree into an instance where coordinates are positive integers whose value is $O(n^2)$.*

3. The MIN TSP

Recall that, for any $n = 2^h$ that is a power of 2, the *Hadamard code* $H_n \subset \{0,1\}^n$ is a set of n binary strings of length n whose pairwise Hamming distance is $n/2$. The elements of H_n can be seen as the set of liner functions $l : \{0,1\}^h \to \{0,1\}$.

We use Hadamard codes to prove a lemma relating $(1,2) - B$ metrics and Hamming metrics. The lemma gives a "distance preserving" embedding of $(1,2) - B$ metric spaces into Hamming spaces.

Lemma 5 *Let U be a finite set of cardinality n, where n is a power of two, and d be a $(1,2) - B$ metric over U. Then there exists an embedding $f : U \to \{0,1\}^{2Bn}$ such that for any $u, v \in U$:*

1. *If $d(u,v) = 2$, then $d_H(f(u), f(v)) = Bn$.*
2. *If $d(u,v) = 1$, then $d_H(f(u), f(v)) = (B - 1/2)n$.*

Such an embedding is computable in time polynomial in $|U|$.

PROOF: Let $U = \{u_1, \ldots, u_n\}$. Recall that a $(1,2) - B$ metric (U,d) can be represented as an undirected graph $G = (U, E)$, where $\{u,v\} \in E$ iff $d(u,v) = 1$ (see [PY93]). Note that G has maximum degree B. We say that two edges are *incident* if they share an endpoint.

Claim 6 *We can find in polynomial time a partition of E into $2B$ matchings $E_1, \ldots, E_{2B}$.*

PROOF: Repeatedly find a maximal matching and delete its edges. Let E_i be the maximal matching removed at the i-th phase. An edge $e \in E$ is not picked after i phases only if edges incident on e have been picked in each phase. Since G has maximum degree B, there can be at most $2(B-1)$ edges incident to given one. Thus, we will be always able to partition G into (at most) $2B-1$ matchings. The bound $2B$ is more convenient for notation. We may assume that one or more matchings are empty. □

Each node $u \in U$ is mapped into a string $f(u)$ that is the concatenation of $2B$ strings $\vec{a}_u^1, \ldots, \vec{a}_u^{2B} \in H_n$:

$$f(u) = \vec{a}_u^1 \circ \ldots \circ \vec{a}_u^{2B} \ .$$

For a fixed $i \in \{1, \ldots, 2B\}$, the strings $\{\vec{a}_u^i\}_{u \in U}$ are chosen arbitrarily in H_n such that $\vec{a}_u^i = \vec{a}_v^i$ if and only if $\{u, v\} \in E_i$. Since H_n can be generated in polynomial time in n, the overall construction can be carried out in poly(n) time.

Let us now compute the distance between two strings $f(u)$ and $f(v)$. There are two cases to be considered.

1. If $\{u, v\} \notin E$, then $\vec{a}_u^i \neq \vec{a}_v^i$ for all $i = 1, \ldots, 2B$, and so $d_H(f(u), f(v)) = 2B \cdot n/2$.
2. If $\{u, v\} \in E$, then $\{u, v\} \in E_j$ for some j, and we have $\vec{a}_u^j = \vec{a}_v^j$ and $\vec{a}_u^i \neq \vec{a}_v^i$ for $i \neq j$. It follows that $d_H(f(u), f(v)) = (2B-1) \cdot n/2$.

□

The embedding described in the previous corollary uses somehow *too many dimensions* in the target Hamming space. We can reduce them using *random sampling*. The idea is as follows: let $b_1, \ldots, b_n \in \{0, 1\}$ be unknown values. If we pick a random subset $b_{i_1}, \ldots, b_{i_m}$ of m elements, where $m = O((\log 1/\delta)/\epsilon^2)$, then with probability $1-\delta$ it holds

$$\left| \sum_{i=1}^{n} b_i - (n/m) \sum_{j=1}^{m} b_{i_j} \right| \leq \epsilon n \ .$$

Now, if we pick $O((\log n)/\epsilon^2)$ coordinates from the target Hamming space of the previous reduction, the distance between two fixed cities will suffer a distorsion at most $O(\epsilon Bn)$ with probability $(1 - 1/\text{poly}(n))$. In particular, there is a constant probability that all the pairwise distances are simultaneously distorced by at most $O(\epsilon Bn)$. Using the oblivous sampler of Bellare and Rompel [BR94] (or alternatively, the Chernoff bound for random walks on expander graphs [Gil93]) we can find such a set of $O((\log n)/\epsilon^2)$ coordinates *deterministically* in polynomial time.

Lemma 7 *Let U be a finite set of cardinality n, where n is a power of two, d be a $(1,2)-B$ metric over U, $\epsilon > 0$ be a positive constant. Then there exists an embedding $f : U \to \{0,1\}^m$ (where $m = O((\log Bn)/\epsilon^2)$) such that for any $u, v \in U$:*

1. *If $d(u, v) = 2$, then $(1-\epsilon)m/2 \leq d_H(f(u), f(v)) \leq (1+\epsilon)m/2$.*
2. *If $d(u, v) = 1$, then $(1-\epsilon)(1-1/2B)m/2 \leq d_H(f(u), f(v)) = (1+\epsilon)(1-1/2B)m/2$.*

Such an embedding is computable in time polynomial in $|U|$.

The following simple corollary of Proposition 3 and of Lemma 7 is required in the proof of our hardness result.

Corollary 8 *Let $p \geq 1$ be fixed. Let U be a finite set and d be a $(1,2)-B$ metric over U. Then there exist a constant δ (depending on B) and an embedding $f : U \to \mathcal{R}^{O(\log n/\epsilon^2)}$ such that for any $u, v \in U$:*

1. *If $d(u, v) = 1$, then $1 - \epsilon \leq d_p(f(u), f(v)) \leq 1$.*
2. *If $d(u, v) = 2$, then $1 + \delta - \epsilon \leq d_p(f(u), f(v)) \leq 1 + \delta$.*

Such an embedding is computable in time polynomial in $|U|$.

The main result of this section is now only a matter of standard calculations.

Theorem 9 *For any fixed $p \geq 1$, a constant $\epsilon^{(p)} > 0$ exists such that the* MIN TSP *is* NP*-hard to approximate within $1 + \epsilon^{(p)}$, even when restricted to the ℓ_p metric in $\mathcal{R}^{\log n}$ (n is the number of cities).*

PROOF: From [PY93] and [ALM$^+$92] we have the following result: constants $B_0 > 0$ and $r_0 > 1$ exist such that, given an instance x of MIN TSP with a $(1,2) - B_0$ metric and n cities, and given the promise that either $\mathsf{opt}(x) = n$ or $\mathsf{opt}(x) \geq r_0 n$, it is NP-hard to distinguish which of the two cases holds.

Given an instance x of $(1,2) - B_0$ MIN TSP, we apply the mapping of Corollary 8 with $B = B_0$ and $\epsilon = \delta(r_0 - 1)/2$ (where δ is the constant of Corollary 8 relative to B_0). We also let $\epsilon^{(p)} = \delta(r_0 - 1)/3$. In this way, we obtain an instance x' of geometric MIN TSP in $\mathcal{R}^{O(\log n)}$. It is easily seen that if $\mathsf{opt}(x) = n$, then $\mathsf{opt}(x') \leq n$. On the other hand, if $\mathsf{opt}(x) \geq r_0 n$, then $\mathsf{opt}(x') \geq (1-\epsilon)n + \delta(r_0 - 1)n$. An approximation better than $(1 - \epsilon + \delta(r_0 - 1))$ (e.g. an approximation $1 + \epsilon^{(p)}$) is sufficient to distinguish between the two cases, and so is NP-hard to achieve.

The claim of the Theorem asks for the cities to be in $\mathcal{R}^{\log n}$, rather than in $\mathcal{R}^{c \log n}$ as in the previous construction. However, we can add $(n^c - n)$ new cities, all at distance $1/n^{c+1}$ from a given one. This perturbs the optimum in a negligible way, and gives an instance with $N = n^c$ cities in $\mathcal{R}^{\log N}$. □

Using techniques of Khanna et al. [KMSV94], the non-approximability result of Theorem 9 implies that geometric MIN TSP in $\mathcal{R}^{\log n}$ under any ℓ_p norm is APX PB-hard (in particular, Max SNP-hard) under E-reductions and APX-complete under AP-reductions [CKST95].

4. The MIN ST Problem

The hardness of approximating MIN ST will be established with a longish chain of reductions. The starting point is the following hardness result, that may have a little independent interest. Recall that in the Minimum Vertex Cover (MIN VC) problem one is given a graph $G = (V, E)$ and looks for the smallest set $C \subseteq V$ such that C contains at least one endpoint of any edge in E.

Theorem 10 *The* MIN VC *problem is* Max SNP*-hard even when restricted to triangle-free graphs with maximum degree 3 (we call this restriction* MIN TF VC-3*).*

PROOF: The MAX 2SAT problem is Max SNP-hard even when restricted to instances where each variable occurs in at most 3 clauses (apply to MAX 2SAT the reduction from MAX 3SAT to MAX 3SAT-3 described in [Pap94]). One can assume without loss of generality that the 3 occurrences of each variable are either one positive occurrence and two negative occurrences, or vice versa. We reduce MAX 2SAT-3 to MIN VC using the reduction of [PY91]: we create a graph with a node for any occurrence of any literal, putting an edge between two nodes if they represent literals that occur in the same clause or if they are one the complement of the other. See [PY91] for the proof that this is an L-reduction. The obtained graph has maximum degree 3: each literal is adjacent to the fellow literal occurring in the same clause and to the (at most) two occurrences of its complement. Also, the graph is triangle-free: let l_1, l_2 and l_3 be any three occurrences of literals. Since clauses contain only two literals, from pigeonhole principle it follows that one of the three occurrences (say, l_1) does not occur in the same clause with any of other two. Then, if l_1, l_2 and l_3 form a triangle it follows that l_2 and l_3 are both the complement of l_1. Being adjacent, they also have to occur in the same clause, but this is a contraddiction since the literals occurring in a clause have to be different. □

We note in passing that, as a corollary, we obtain that the MAX INDEPENDENT SET problem is Max SNP-hard in the same, very restricted class of graphs. We now move to the restriction of MIN ST to Hamming spaces.

NOTATION: For a pair of indices $i, j \in [n]$ we define $\vec{a}^n_{i,j} \in \{0,1\}^n$ as the n-dimensional boolean vector all whose coordinates are zero but the i-th and the the j-th, e.g. $\vec{a}^5_{1,4} = (1,0,0,1,0)$. Similarly, we let $\vec{a}^n_i$ be the the vector in $\{0,1\}^n$ whose only non-zero coordinate is the i-th, e.g. $\vec{a}^4_3 = (0,0,1,0)$. For a vector $\vec{a} \in \{0,1\}^n$ and indices $i, j \in [n]$, we let $\mathsf{red}_{i,j}(\vec{a}) \in \{0,1\}^n$ be the vector defined as follows

$$\mathsf{red}_{i,j}(\vec{a})[h] = \begin{cases} 0 & \text{if } (h = i \vee h = j) \wedge \\ & \quad \vec{a}[i] = \vec{a}[j] = 1 \\ a[h] & \text{otherwise.} \end{cases}$$

In other words, $\mathsf{red}_{i,j}(\vec{a})$ is equal to $\vec{a}$ unless a has a one in the i-th and the j-th coordinate. In this latter case, the i-th and the j-th coordinate of $\mathsf{red}_{i,j}(\vec{a})$ are set to zero. For example $\mathsf{red}_{1,3}(0,1,1,1) = (0,1,1,1)$, while $\mathsf{red}_{2,3}(0,1,1,1) = (0,0,0,1)$. We will make use of the following simple combinatorial lemma.

Lemma 11 *For any $\vec{a}, \vec{b} \in \{0,1\}^n$, for any $i, j \in [n]$, $d_H(\mathsf{red}_{i,j}(\vec{a}), \mathsf{red}_{i,j}(\vec{b})) \leq d_H(\vec{a}, \vec{b})$.*

PROOF: Case analysis. □

Theorem 12 *The* MIN ST *problem is* Max SNP*-hard when restricted to Hamming spaces.*

PROOF: We give an L-reduction from MIN TF VC-3. Let $G = (V, E)$ be a triangle-free graph of maximum degree 3, assume $V = [n]$ and let $m = |E|$. We define an instance of Hamming MIN ST as follows: the number of dimensions is n and the set of points is

$$U = \{\vec{0}\} \cup \{\vec{a}^n_{ij} : \{i,j\} \in E\}$$

where $\vec{0}$ is the vector with all zero entries.

Claim 13 *Given a vertex cover $C \subseteq V$ in G it is possible to find a Steiner tree for U of cost $m + C$.*

PROOF:[Of Claim 13] Let $S = \{\vec{a}^n_i : i \in C\}$. Consider the graph whose vertex set is $S \cup U$ and such that two vertices are adjacent iff their Hamming distance is one. We claim that this graph is connected: indeed all the nodes of S are clearly adjacent to $\vec{0}$; furthermore any node in U is adjacent to some node in S (since C is a vertex cover), thus all the nodes are connected to $\vec{0}$. Since the graph is connected it admits a spanning tree, that is also a Steiner tree for U. All the edges of such Steiner tree have cost 1, and there are $|C| + m$ of them (because the tree has $|S| + |U| = |C| + m + 1$ nodes), so the claim follows. □

From the above claim it follows that $\mathsf{opt}(U) \leq m + \mathsf{opt}(G) \leq 4\mathsf{opt}(G)$, and we have established the first condition of the L-reducibility. As usual, the other condition is more difficult to prove.

Claim 14 *Given a Steiner tree T for U it is possible to find in polynomial time another Steiner tree T' such that: (i) $\mathsf{cost}(T') \leq \mathsf{cost}(T)$ and (ii) all the edes of T' have cost one and all the Steiner nodes of T' are weight-one vectors.*

PROOF:[Of Claim 14] We first make sure that all edges have cost 1: any edge of cost $d > 1$ is broken into a length-d path using $d - 1$ additional Steiner nodes. Let S be the new set of Steiner vertices. We now reduce the number of non-zero coordinates of Steiner vertices. For any $\{i,j\} \notin E$ we map each point $\vec{a} \in S \cup U$ into $\mathsf{red}_{i,j}(\vec{a})$; this mapping only changes Steiner points (by definition of $\mathsf{red}_{i,j}$, definition of U, and the fact that $\{i,j\} \notin E$). From Lemma 11 we also have that any phase does not increase the cost of the tree. At the end of this set of transformations, we run a "clean-up" phase that does the following: if some transformation has collapsed one node onto another, we take only one node (if a Steiner node is collapsed onto a node in U we clearly take the node in U). If the transformation creates cycles, we break them (e.g. finding a spanning tree of the final graph), and, again, this does not increase the cost. It remains to see that, after this process, no Steiner node can have more than one non-zero coordinate. If a Steiner node has some set of k non-zero coordinates, then they must correspond to a clique in G (otherwise, at some phase, some of them would have been changed by the application of the red operator): since G is triangle-free, $k \leq 2$, but if $k = 2$ then the Steiner node would be equal to a node of U, and thus would have been removed in the clean-up phase. It follows that $k = 1$. □

From the above claim, the next one, whose proof we omit, follows quite easily.

Claim 15 *Given a Steiner tree T for U it is possible to find in polynomial time a vertex cover C for G such that $|C| \leq \mathsf{cost}(T) - m$*

If T is any Steiner tree of U, the vertex cover C for G computed according the previous claim satisfies

$$\begin{aligned} & \mathsf{cost}(C) - \mathsf{opt}(G) \\ \leq \; & (\mathsf{cost}(T) - m) - (\mathsf{opt}(U) - m) \\ = \; & \mathsf{cost}(T) - \mathsf{opt}(U) \end{aligned}$$

and so also the second condition of the L-reduction is satisfied. □

If the following conjecture holds, then we can reduce MIN VC-B to Hamming MIN ST (without imposing the triangle-free restriction).

Conjecture 1 *Let $U \subset \{0,1\}^n$ be an instance of Hamming* MIN ST *such that $\vec{0} \in U$ and all vectors of U have weight at most 2. Then there exists an optimum solution where all the Steiner nodes have weight at most 2.*

Janos Körner proposed a further generalization: if U is contained in the Hamming sphere centered in some $\vec{u} \in U$ and of radius k, then there exists an optimum solution all whose Steiner nodes lie in the same sphere. This seemed to be a reasonable combinatorial analog of the fact that if the points are in $\mathcal{R}^k$ and distances are computed according to the Euclidean metric, the Steiner points of an optimum solution will be in the *convex hull* of the points of the instance. Subsequently, Janos refuted the generalized conjecture. The instance $U = \{(0,0,0,0),(0,1,1,1),(1,0,1,1),$ $(1,1,0,1),(1,1,1,0)\}$ refutes the generalized conjecture even for $k = 3$. An optimum solution of cost 7 uses the Steiner node $(1,1,1,1)$. Computational experiments show that any solution without $(1,1,1,1)$ has cost at least 8.

To approach the MIN ST in ℓ_1 normed spaces we use a reduction from the Hamming case. Note that for points in $\{0,1\}^n$ the ℓ_1 distance equals the Hamming distance. However, the reduction is non-trivial since $\mathcal{R}^n$ contains so many points that are not in $\{0,1\}^n$ and we have to argue that having much more choice for the Steiner nodes does not make the problem easier. The Rectilinear MIN ST problem looks very much like a *relaxation* of the Hamming MIN ST problem; our reduction makes use of a *rounding scheme* proving that the relaxation does not change the optimum.

Theorem 16 *Let $U \subseteq \{0,1\}^n \subset \mathcal{R}^n$ be an instance of Rectilinear* MIN ST *all whose points are in the Boolean cube. Let T be a feasible solution for U. Then it is possible to find in polynomial time (in the size of T) another solution T' such that $\mathsf{cost}(T') \leq \mathsf{cost}(T)$ and all the Steiner nodes of T' are in $\{0,1\}^n$.*

Before proving the theorem, we note the following relevant consequence.

Corollary 17 *For any instance $U \subseteq \{0,1\}^n$ of Rectilinear* MIN ST, *an optimum solution exists all whose Steiner points are in $\{0,1\}^n$.*

We now prove Theorem 16.

PROOF:[Of Theorem 16] Let $S = \{\vec{s}_1,\ldots,\vec{s}_m\}$ be the set of Steiner points of T, and let E be the set of edges of T. For any $\vec{s}_j \in S$ we will find a new point $\vec{s}'_j \in \{0,1\}^n$, so that if we let T' be the tree obtained from T by substituting the $\vec{s}$ points with the corresponding $\vec{s}'$ points, the cost of T' is not greater than the cost of T. The latter statement is equivalent to

$$\begin{aligned}&\sum_{(\vec{s}_j,\vec{u})\in E, \vec{u}\in U} ||\vec{s}_j - \vec{u}||_1 + \sum_{(\vec{s}_j,\vec{s}_h)\in E} ||\vec{s}_j - \vec{s}_h||_1 \\ \geq &\sum_{(\vec{s}'_j,\vec{u})\in E, \vec{u}\in U} ||\vec{s}'_j - \vec{u}||_1 + \sum_{(\vec{s}'_j,\vec{s}'_h)\in E} ||\vec{s}'_j - \vec{s}'_h||_1\end{aligned}$$

We will indeed prove something stronger, namely, that for any $i \in [n]$ it holds

$$\begin{aligned}&\sum_{(\vec{s}_j,\vec{u})\in E, \vec{u}\in U} |\vec{s}_j[i] - \vec{u}[i]| + \sum_{(\vec{s}_j,\vec{s}_h)\in E} |\vec{s}_j[i] - \vec{s}_h[i]| \\ \geq &\sum_{(\vec{s}'_j,\vec{u})\in E, \vec{u}\in U} |\vec{s}'_j[i] - \vec{u}[i]| + \sum_{(\vec{s}'_j,\vec{s}'_h)\in E} |\vec{s}'_j[i] - \vec{s}'_h[i]|\end{aligned} \tag{1}$$

Let $i \in [n]$ be fixed, we now see how to find values of $\vec{s}'_1[i],\ldots,\vec{s}'_m[i] \in \{0,1\}$ such that (1) holds. We express as a linear program the problem of finding values of $\vec{s}'_1[i],\ldots,\vec{s}'_m[i]$ that minimize the right-hand side of (1). For any $j \in [m]$ we have a variable x_j (representing the value to be given to $\vec{s}'_j[i]$) and for any edge $e = (\vec{a},\vec{b})$ such that at least one endpoint is in S we have a variable y_e, representing the lenght $|\vec{a}[i] - \vec{b}[i]|$. The linear program is as follows

$$\begin{array}{lll}\min & \sum_e y_e & \\ \text{s.t.} & & \\ & y_e \geq x_j - x_h & \forall e = (\vec{s}_j,\vec{s}_h) \in E \\ & y_e \geq x_h - x_j & \forall e = (\vec{s}_j,\vec{s}_h) \in E \\ & y_e \geq x_j & \forall e = (\vec{s}_j,\vec{u}_h) \in E . \vec{u}_h[i] = 0 \\ & y_e \geq 1 - x_j & \forall e = (\vec{s}_j,\vec{u}_h) \in E . \vec{u}_h[i] = 1 \\ & x_j \geq 0 & \\ & y_e \geq 0 & \end{array} \tag{LP}$$

Setting $x_j = s_j[i]$ and setting $y_{(\vec{a},\vec{b})} = |\vec{a}[i] - \vec{b}[i]|$ yields a feasible solution, and its cost is the left-hand side of (1). Let $(\vec{x}^*, \vec{y}^*)$ be an optimum solution for (LP). From the previous observation we have that setting $\vec{s}'_j[i] = x^*_j$ we satisfy (1). It remains to be seen that (LP) has an optimum solution where all variables take value from $\{0, 1\}$. This follows from the fact that (LP) is the linear programming relaxation of an undirected Min-CUT problem, where all the $\vec{u}$ such that $\vec{u}[i] = 0$ (respectively, $\vec{u}[i] = 1$) are identifed with the source (respectively, the sink), each $\vec{s}_j$ is a node, and the edges are like in T. It is well known (see e.g. [PS82]) that a Min-CUT linear programming relaxation has optimum 0/1 solutions, and that such a solution can be found in polynomial time. □

Remark 18 *There seems to be no natural analog of Theorem 16 in other norms. Even in $\mathcal{R}^2$, using the Eulcidean metric, we have that the optimum solution of the instance $\{(0, 0), (1, 0), (0, 1)\}$ must use a Steiner point not in $\{0, 1\}^2$.*

Theorem 19 *Rectilinear* MIN ST *is* Max SNP*-hard.*

PROOF: We reduce from Hamming MIN ST. The reduction leaves the instance unchanged. For an instance $U \subseteq \{0, 1\}^n$, we let $\mathsf{opt}_H(U)$ (respectively, $\mathsf{opt}_R(U)$) be the cost of an optimum solution for U, when seen as an instance of Hamming MIN ST (respectively, of Rectilinear MIN ST). Clearly, we have that $\mathsf{opt}_R(U) \leq \mathsf{opt}_H(U)$. Given a solution T for U, we find a solution T' as in Theorem 16. Since in $\{0, 1\}^n$ the distance induced by the ℓ_1 norm equals the Hamming distance, we have that $\mathsf{cost}_H(T') = \mathsf{cost}_R(T') \leq \mathsf{cost}_R(T)$. We have an L-reduction with $\alpha = \beta = 1$. □

5. Conclusions and Open questions

We do not know how to extend our non-approximability result for MIN ST to the Euclidean case. Arora [Aro96] notes that, by inspecting the way his algorithm works, it is possible to claim that, for any instance of Euclidean MIN ST, there exists a near-optimal solution where the Steiner points lie in some well-specified positions (either at "portals" or in positions chosen at the bottom of the recursion). This observation could perhaps be a starting point.

We don't have explicit estimations of the constants to within which it is hard to approximate geometric MIN TSP and rectilinear MIN ST. The constant for MIN TSP should be only slightly smaller than the corresponding constant for the $(1, 2) - B$ case (estimated around $1 + 10^{-5}$). The constant for MIN ST is more likely to be around $1 + 10^{-4}$. Finding much stronger estimations (comparable to the $3/2$ bound of Christofides and the 1.644 bound of Karpinski and Zelikovsky) is an open and challenging question. It appears that MIN TSP and MIN ST lack the nice logical definability that allows to prove very strong non-approximability results for MAX CUT and MAX 3SAT using so-called "gadget reductions" [BGS96, TSSW96, Hås97].

Acknowledgments. I'm grateful to Sanjeev Arora, Madhu Sudan, Laszlo Babai and Janos Körner for very interesting and helpful discussions. In particular, Lemma 11 is due to Madhu. I thank Ray Greenlaw for having read a preliminary version of this paper and given several useful remarks.

References

[ALM+92] S. Arora, C. Lund, R. Motwani, M. Sudan, and M. Szegedy. Proof verification and hardness of approximation problems. In *Proc. of the 33rd IEEE FOCS*, pages 14–23, 1992.

[Aro94] S. Arora. *Probabilistic Checking of Proofs and Hardness of Approximation Problems*. PhD thesis, University of California at Berkeley, 1994.

[Aro96] S. Arora. Polynomial time approximation schemes for Euclidean TSP and other geometric problems. In *Proc. of the 37th IEEE FOCS*, 1996.

[BGS96] M. Bellare, O. Goldreich, and M. Sudan. Free bits, PCP's and non-approximability – towards tight results (3rd version). Technical Report ECCC TR95-24, 1996. Preliminary version in *Proc. of FOCS'95*.

[BP89] M. Bern and P. Plassmann. The Steiner tree problem with edge lengths 1 and 2. *Information Processing Letters*, 32:171–176, 1989.

[BR94] M. Bellare and J. Rompel. Randomness-efficient oblivious sampling. In *Proc. of the 35th IEEE FOCS*, pages 276–287, 1994.

[Chr76] N. Christofides. Worst-case analysis of a new heuristic for the travelling salesman problem. Technical report, Carnegie-Mellon University, 1976.

[CKST95] P. Crescenzi, V. Kann, R. Silvestri, and L. Trevisan. Structure in approximation classes. Technical Report ECCC TR96-66, 1996. Preliminary version in *Proc. of COCOON'95*.

[Edm66] J. Edmonds. Optimum branchings. *Journal of Research National Bureau of Standards, Part B*, 17B(4):233–240, 1966.

[GGJ76] M.R. Garey, R.L. Graham, and D.S. Johnson. Some NP-complete geometric problems. In *Proc. of the 8th ACM STOC*, pages 10–22, 1976.

[GGJ77] M.R. Garey, R.L. Graham, and D.S. Johnson. The complexity of computing Steiner minimal trees. *SIAM J. of Applied Mathematics*, 34:477–495, 1977.

[Gil93] D. Gillman. A Chernoff bound for random walks on expander graphs. In *Proc. of the 34th IEEE FOCS*, pages 680–691, 1993.

[GJ77] M.R. Garey and D.S. Johnson. The rectilinear Steiner tree problem is NP-complete. *SIAM J. on Applied Mathematics*, 32(4):826–834, 1977.

[GKP95] M. Grigni, E. Koutsoupias, and C.H. Papadimitriou. An approximation scheme for planar graph TSP. In *Proc. of the 36th IEEE FOCS*, pages 640–645, 1995.

[Hås97] J. Håstad. Some optimal inapproximability results. These proceedings, 1997.

[KMSV94] S. Khanna, R. Motwani, M. Sudan, and U. Vazirani. On syntactic versus computational views of approximability. In *Proc. of the 35th IEEE FOCS*, pages 819–830, 1994.

[KZ95] M. Karpinski and A. Zelikovsky. New approximation algorithms for the Steiner tree problems. Technical Report ECCC TR95-030, 1995.

[LLKS85] E.L. Lawler, J.K. Lenstra, A.H.G. Rinnooy Kan, and D.B. Shmoys. *The traveling salesman problem*. John Wiley, 1985.

[Pap77] C.H. Papadimitriou. Euclidean TSP is NP-complete. *Theoretical Computer Science*, 4:237–244, 1977.

[Pap94] C.H. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.

[PS82] C.H. Papadimitriou and K. Steiglitz. *Combinatorial Optimization: Algorithms and Complexity*. Prentice-Hall, 1982.

[PY91] C. H. Papadimitriou and M. Yannakakis. Optimization, approximation, and complexity classes. *Journal of Computer and System Sciences*, 43:425–440, 1991. Preliminary version in *Proc. of STOC'88*.

[PY93] C.H. Papadimitriou and M. Yannakakis. The travelling salesman problem with distances one and two. *Mathematics of Operations Research*, 18:1–11, 1993.

[TSSW96] L. Trevisan, G.B. Sorkin, M. Sudan, and D.P. Williamson. Gadgets, approximation, and linear programming. In *Proc. of the 37th IEEE FOCS*, 1996.