Electronic Colloquium on Computational Complexity, Report No. 89 (2000)

Approximation Hardness of TSP with Bounded Metrics

Lars Engebretsen[1] and Marek Karpinski[2]

[1] MIT Laboratory for Computer Science
200 Technology Square, NE43-369
Cambridge, Massachusetts 02139-3594
E-mail: enge@mit.edu

[2] Department of Computer Science
University of Bonn
53117 Bonn
E-mail: marek@cs.uni-bonn.de

Abstract. The general asymmetric (and metric) TSP is known to be approximable only to within an $O(\log n)$ factor, and is also known to be approximable within a constant factor as soon as the metric is bounded. In this paper we study the asymmetric and symmetric TSP problems with bounded metrics and prove approximation lower bounds of 54/53 and 131/130, respectively, for these problems.

We prove also approximation lower bounds of 321/320 and 743/742 for the asymmetric and symmetric TSP with distances one and two, improving over the previous best lower bounds of 2805/2804 and 5381/5380 of Engebretsen for the case of distances one and two, by the order of magnitude. Furthermore, one of our constructions can be used to improve a recent lower bound of Papadimitriou and Vempala for the case of symmetric TSP with unbounded metric.

Key words. Approximation Ratios; Lower Bounds; Metric TSP; Bounded Metric.

1 Introduction

A common special case of the Traveling Salesman Problem (TSP) is the metric TSP, where the distances between the cities satisfy the triangle inequality. The decision version of this special case was shown to be **NP**-complete by Karp [8], which means that we have little hope of computing exact solutions in polynomial time. Christofides [4] has constructed an elegant algorithm approximating the metric TSP within 3/2, i.e., an algorithm that always produces a tour whose weight is at most a factor 3/2 from the weight of the optimal tour. For the case when the distance function may be asymmetric,

ISSN 1433-8092

the best known algorithm approximates the solution within $O(\log n)$, where n is the number of cities [6], although a constant factor approximation algorithm has recently been conjectured [3]. As for lower bounds, Papadimitriou and Yannakakis [10] have shown that there exists some constant, see also [1], such that it is **NP**-hard to approximate the TSP where the distances are constrained to be either one or two—note that such a distance function always satisfies the triangle inequality—within that constant. This lower bound was improved by Engebretsen [5] to $2805/2804 - \epsilon$ for the asymmetric and $5381/5380 - \epsilon$ for the symmetric, respectively, TSP with distances one and two. The instances produced in Engebretsen's construction also has the property that every city is close to at most four other cities, i.e., that there are at most four other cities at distance one from it; Trevisan [11] studies TSP for such metrics in greater detail.

It appears that the metric TSP lacks the good properties which have been needed (so far) for proving strong nonapproximability results. Therefore, any new insights into explicit lower bounds here seem to be of a considerable interest. Papadimitriou and Vempala [9] recently announced lower bounds of $42/41 - \epsilon$ and $129/128 - \epsilon$, respectively, for the asymmetric and symmetric versions, respectively, of the TSP with graph metric, but left the question of the approximability for the case with bounded metric open. Apart from being an interesting question on its own, it is conceivable that the special cases with bounded metric are easier to approximate than the cases when the distance between two points can grow with the number of cities in the instance. Indeed, the asymmetric TSP with distances bounded by B can be approximated within B by just picking any tour as the solution and the asymmetric TSP with distances one and two can be approximated within 17/12 [12]. The symmetric version of the latter problem can be approximated within 7/6 [10].

In this paper, we consider the case when the metric contains only integer distances between one and six and prove a lower bound of $54/53 - \epsilon$ for the asymmetric case and $131/130 - \epsilon$ for the symmetric case. This is an improvement of several orders of magnitude compared to the previous best known bounds of $2805/2804 - \epsilon$ and $5381/5380 - \epsilon$ for this case, respectively [5]. We also prove that it is **NP**-hard to approximate the asymmetric TSP with distances one and two within $321/320 - \epsilon$, for any constant $\epsilon > 0$. For the symmetric version of the latter problem we show a lower bound of $743/742 - \epsilon$.

The method of our proofs depends on explicit reductions from certain bounded dependency instances of linear equations satisfiability. The main idea is to construct certain uniform circles of equation gadgets and, in the second part, certain combined hybrid circle constructions.

Definition 1.1. *The* Asymmetric Traveling Salesman Problem (ATSP) *is the following minimization problem: Given a collection of cities and a matrix whose entries are interpreted as the distance from a city to another, find the shortest tour starting and ending in the same city and visiting every city exactly once.*

Definition 1.2. (1,B)-ATSP *is the special case of ATSP where the entries in the distance matrix obey the triangle inequality and the off-diagonal entries in the distance matrix are integers between 1 and B.* (1,B)-TSP *is the special case of (1,B)-ATSP where the distance matrix is symmetric.*

2 The hardness of (1,B)-ATSP

We reduce, similarly to Papadimitriou and Vempala [9], from Håstad's lower bound for E3-Lin mod 2 [7]. In fact, our gadgets for the (1,B)-ATSP case are syntactically identical to those of Papdimitriou and Vempala [9] but we use a slightly different accounting method. The construction consists of a circle of *equation gadgets* testing odd parity. This is no restriction since we can easily transform a test for even parity into a test for odd parity by flipping a literal. Three of the edges in the equation gadget correspond to the variables involved in the parity check. These edges are in fact gadgets, so called *edge gadgets*, themselves. Edge gadgets from different equation gadgets are connected to ensure consistency among the edges representing a literal. This requires the number of negative occurrences of a variable to be equal to the number of positive occurrences. This is no restriction since we can duplicate every equation a constant number of times and flip literals to reach this property.

Definition 2.1. E3-Lin mod 2 *is the following problem: Given an instance of n variables and m equations over $\boldsymbol{Z}_2$ with exactly three unknowns in each equation, find an assignment to the variables that satisfies as many equations as possible.*

Theorem 2.2 [**7**]**.** *There exists instances of E3-Lin mod 2 with $2m$ equations such that, for any constant $\epsilon > 0$, it is **NP**-hard to decide if at most ϵm or at least $(1-\epsilon)m$ equations are left unsatisfied by the optimal assignment. Each variable in the instance occurs a constant number of times, half of them negated and half of them unnegated.*

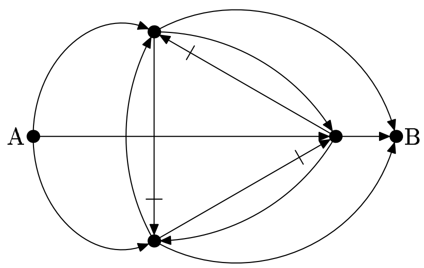

Figure 1. The gadget for equations of the form $x + y + z = 0$. There is a Hamiltonian path from A to B only if zero or two of the ticked edges are traversed.

To describe an instance of (1,B)-ATSP, it is enough to specify the edges of weight one. We do this by constructing an unweighted directed graph G and then let the (1,B)-ATSP instance have the nodes of G as cities. For two nodes u and v in G, let $\ell(u, v)$ be the length of the shortest path from u to v in G. The distance between two cities u and v is the (1,B)-ATSP instance is then defined to be $\min\{B, \ell(u, v)\}$.

2.1 The gadgets

The equation gadget for equations of the form $x + y + z = 0$ is shown in Fig. 1. The key property of this gadget is that there is a Hamiltonian path through the gadget only if zero or two of the ticked edges are traversed. To form the circle of equation gadgets, vertex A in one gadget coincides with vertex B in another gadget.

The ticked edges in Fig. 1 are gadgets themselves. This gadget is shown in Fig. 2. Each of the bridges, i.e., the pairs of undirected edges in the gadget, is shared between two different edge gadgets, one corresponding to a positive occurrence of the literal and one corresponding to a negative occurrence. The precise coupling is provided by a perfect matching in a 5-regular bipartite multigraph with the following property: For any partition of the left k vertices into subsets S_1, U_1 and T_1 and any partition of the right k vertices into subsets S_2, U_2 and T_2 such that there are no edges from U_1 to U_2, $|S_1| + |S_2| \leq k$ and $|U_1| + |T_2| \leq |T_1| + |U_2|$, the total number of edges from vertices in T_1 to vertices in T_2 is greater than

$$\frac{1}{2} \min\{|U_1| + |T_2|, k - |S_1| - |S_2|\}.$$

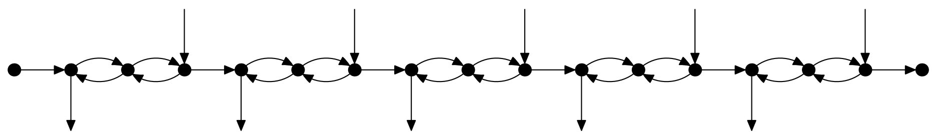

Figure 2. The edge gadget consists of five bridges—each of the bridges are shared between two different edge gadgets.

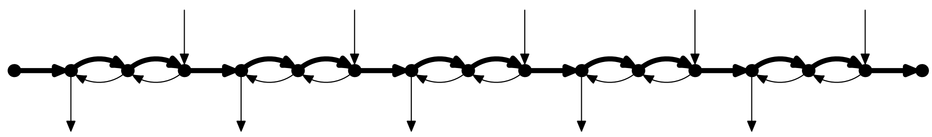

Figure 3. A traversed edge gadget represents the value 1.

(We sketch the proof that such graphs exist in Sec. 2.3.) The purpose of this construction is to ensure that it is always optimal for the tour to traverse the graph in such a way that all variables are given consistent values. The edge gadget gives an assignment to an occurrence of a variable by the way it is traversed.

Definition 2.3. *We call an edge gadget where all bridges are traversed from left to right in Fig. 2* traversed *and an edge gadget where all bridges are traversed from right to left* untraversed*. All other edge gadgets are called* semitraversed.

2.2 Proof of correctness

If we assume that the tour behaves nicely, i.e., that the edge gadgets are either traversed or untraversed, it is straightforward to establish a correspondence between the length of the tour and the number of unsatisfied equations.

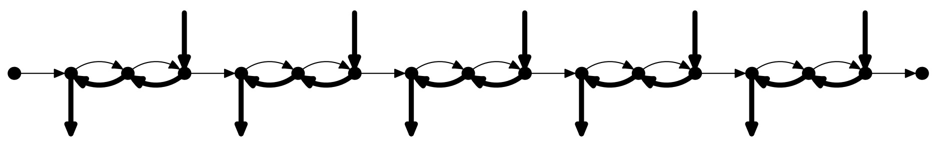

Figure 4. An untraversed edge gadget represents the value 0.

Lemma 2.4. *The only way to traverse the equation gadget in Fig. 1 with a tour of length 4—if the edge gadgets count as length one for the moment—is to traverse an odd number of edge gadgets. All other locally optimal traversals have length 5.*

Proof. It is easy to see that any tour traversing two ticked edges and leaving the third one untraversed has length 4. Any tour traversing one ticked edge and leaving the other two ticked edges untraversed has length at least 5. Strictly speaking, it is impossible to have three traversals since this does not result in a tour. However, we can regard the case when the tour leaves the edge gadget by jumping directly to the exit node of the equation gadget as a tour with three traversals; such a tour gives a cost of 5. □

Lemma 2.5. *In addition to the length 1 attributed to the edge gadget above, the length of a tour traversing an edge gadget in the intended way is 15.*

Proof. Each bridge has length 2 and every bridge must have one of the outgoing edge traversed. Thus, the total cost is $5 \cdot (2 + 1) = 15$. □

Lemma 2.6. *Suppose that there are $2m$ equations in the E3-Lin mod 2 instance. If the tour is shaped in the intended way, i.e., every edge gadget is either traversed or untraversed, the length of the tour is $53m + u$, where u is the number of unsatisfied equations resulting from the assignment represented by the tour.*

Proof. The length of the tour on an edge gadgets is 15. There are three edge gadgets corresponding to every equation and every bridge in the edge gadget is shared between two equation gadgets. Thus, the length of the tour on the edge gadgets is $2m \cdot 3 \cdot 15/2 = 45m$. The length of the tour on an equation gadgets is 4 if the equation is satisfied and 5 otherwise. Thus, the total length is $53m + u$. □

The main challenge now is to prove that the above correspondence between the length of the optimum tour and the number of unsatisfied equation holds also when we drop the assumption that the tour is shaped in the intended way.

To count the excessive cost due to traversed edges of weight more than one, we note that every traversed edge of weight $w > 1$ corresponds to a path of length $\min\{w, B\}$ on edges of weight one. To ease the analysis of the impact of such tours, we reroute every such tour its corresponding path if $w \leq B$; if $w > B$ we make the tour follow the first $B/2$ and last $B/2$ edges of the path and then pretend that the tour does a jump of zero

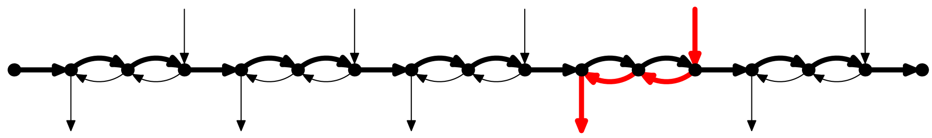

Figure 5. A traversed edge gadget that shares a bridge with another traversed edge gadget..

cost between these two vertices. For clarity we color these new traversals red. This produces something which is not a tour—we call it a *pseudo-tour*—since some edges are traversed more than once and some vertices are connected to more than two traversed edges. From now on, most of the reasoning concerns this pseudo-tour. The following sequence of lemmas give a lower bound on the extra cost, not counting the "normal" cost of 15 per edge gadget and 4 per equation gadget, that results from a non-standard behavior of the tour.

We have already seen that an unsatisfied equation adds an extra cost of 1. Edge gadgets that are either traversed or untraversed do not add any extra cost, except for the case when two traversed equation gadgets share a bridge; this results in a bridge being traversed in both directions by the pseudo-tour. A pseudo-tour resulting from a proper TSP tour can never result in two untraversed edge gadgets sharing a bridge; this would imply a cycle containing three vertices in the original TSP tour.

Lemma 2.7. *Two traversed edge gadgets that share a bridge give an extra cost of 2 to the length of the tour.*

Proof. If two traversed edge gadgets are connected, there must be a bridge that is traversed in both directions. Such a bridge gives an extra cost of 2. □

So far we have dealt with the traversed and the untraversed edge gadgets. What remains is the difficult case—the semitraversed edge gadgets.

Lemma 2.8. *Suppose that $B \geq 6$. Then every semitraversed edge gadget adds an extra cost of at least one to the length of the tour.*

Proof sketch. We call a bridge balanced with respect to a pseudo-tour if there is at least one edge of the pseudo-tour adjacent to each endpoint of the bridge. Note that an unbalanced bridge always gives an extra cost of two, since the bridge must be traversed in both directions by the pseudo-tour. Thus, we always obtain an extra cost of two if any of the bridges

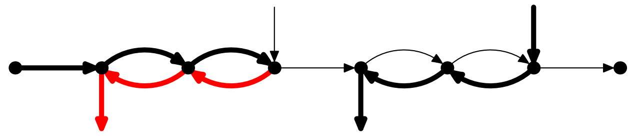

Figure 6. An unbalanced bridge always gives an extra cost of 2.

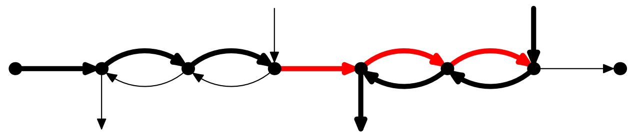

Figure 7. A balanced bridge always gives an extra cost of 2.

are unbalanced. This cost can be divided between two semitraversed edge gadgets, resulting in a cost of at least one per semitraversed edge gadget. We show one such case in Fig. 6, the other cases are handled similarly.

Now assume that all bridges are balanced. Since the edge gadget is semitraversed, all bridges cannot be traversed in the same direction. Thus, there are two adjacent bridges that are traversed in different directions. When $B \geq 6$ this gives an extra cost of two that may be shared by two different semitraversed edge gadgets. We show one such case in Fig. 7, the other cases are handled similarly. □

Lemma 2.9. *Suppose that there are $2k$ occurrences of the variable x. Then at most k of the equation gadgets corresponding to x can be semitraversed.*

Proof. Assume that there are e semitraversed edge gadgets corresponding to x. Then it is possible to change the length of the tour by making all edge gadgets corresponding to x either traversed or untraversed—whatever satisfies the largest number of equations. This transformation itself decreases the length of the tour by at least e, but since we may introduce up to k unsatisfied equations in the process, we may also get an increase of at most k. Summing up, the length of the tour decreases by at least $e-k$. This number is positive when $e > k$. □

Lemma 2.10. *There exists a coupling of the equation gadgets with the property that there can never be advantageous to have inconsistently traversed equation gadgets.*

Proof. For any variable x, the number of semitraversed occurrences is at most k. Consider now the bipartite graph with occurrences of x at one side and occurrences of $\bar{x}$ on the other side. Each vertex in this graph can be labeled as T, U or S, depending on whether it is traversed, untraversed or semitraversed. Let T_1 be the set of traversed positive occurrences and T_2 be the set of traversed negative occurrences. Define U_1, U_2, S_1, and S_2 similarly. We can assume that $|U_1| + |T_2| \leq |U_2| + |T_1|$—otherwise we just change the indexing convention.

We now consider a modified tour where the positive occurrences are traversed and the negative occurrences are untraversed. This decreases the cost of tour by at least $|S_1| + |S_2| + 2|(T_1, T_2)|$, where $|(T_1, T_2)|$ denotes the number of edges between T_1 and T_2, and increases it by $\min\{k, |S_1| + |S_2| + |U_1| + |T_2|\}$. But the bipartite graph has the property that

$$2|(T_1, T_2)| \geq \min\{|U_1| + |T_2|, k - |S_1| - |S_2|\},$$

which implies that the cost of tour decreases by this transformation. Thus, we can assume that x is given a consistent assignment by the tour. □

Theorem 2.11. *For any constant $\epsilon > 0$, it is **NP**-hard to approximate (1,6)-ATSP within $54/53 - \epsilon$.*

Proof. Given an instance of E3-Lin mod 2 with $2m$ equations where every variable occurs a constant number of times, we construct the corresponding instance of (1,6)-ATSP. This can be done in polynomial time. By the above lemma, we can assume that all edge gadgets are traversed consistently in this instance. The assignment obtained from this traversal satisfies $2m - u$ equations if the length of the tour is $53m + u$. If we could decide if the length of the optimum tour is at most $(53 + \epsilon_1)m$ or at least $(54 - \epsilon_2)m$, we could decide if at most $\epsilon_1 m$ or at most $(1 - \epsilon_2)m$ of the equations are let unsatisfied by the corresponding assignment. But to decide this is **NP**-hard by Theorem 2.2. □

2.3 The bipartite graph

In this section we sketch the proof that there exist bipartite graphs with good enough expansion properties for the particular set of parameters we have in our case. The exact statement on all parameters involved is given in Theorem 2.12 below.

Theorem 2.12. *For $d \geq 5$ and a large enough constant k, there exists a bipartite d-regular multigraph on $2k$ vertices with the following property:*

For any partition of the left k vertices into subsets S_1, U_1 and T_1 and any partition of the right k vertices into subsets S_2, U_2 and T_2 such that there are no edges from U_1 to U_2, $|S_1|+|S_2| \leq k$ and $|U_1|+|T_2| \leq k-(|S_1|+|S_2|)/2$, the total number of edges from vertices in T_1 to vertices in T_2 is greater than

$$\frac{1}{2}\min\{|U_1|+|T_2|, k-|S_1|-|S_2|\}.$$

We view a d-regular bipartite graph as a perfect matching on a $dk \times dk$ bipartite graph—every node in the d-regular graph corresponds to d nodes in the larger graph. We select a matching uniformly at random and want to estimate the probability of failure. To do this, we upper bound the number of "bad" matchings and divide by $(dk)!$, the total number of matchings.

First note that for any choice of $s_1, t_1, u_1, s_2, t_2, u_2$ such that $s_1+t_1+u_1=k$ and $s_2+t_2+u_2=k$, there are less than 9^k ways to partition the nodes of the d-regular bipartite graph into sets $S_1, T_1, U_1, S_2, T_2, U_2$ with sizes $s_1, t_1, u_1, s_2, t_2, u_2$, respectively. Given such a partition, there are

$$\begin{aligned}
&P(s_1, t_1, u_1, s_2, t_2, u_2, e_{s_1 s_2}, e_{t_1 t_2}) \\
&\quad = \binom{ds_1}{e_{s_1 s_2}}\binom{ds_2}{e_{s_1 s_2}}(e_{s_1 s_2})!\binom{dt_1}{e_{t_1 t_2}}\binom{dt_2}{e_{t_1 t_2}}(e_{t_1 t_2})! \\
&\quad \times (d(k-u_1-u_2)-e_{t_1 t_2}-e_{s_1 s_2}))!(du_1)!(du_2)!
\end{aligned}$$

different matchings such that there are $e_{s_1 s_2}$ edges from S_1 to S_2 and $e_{t_1 t_2}$ edges from T_1 to T_2. Thus, we can bound the probability of failure by

$$\frac{9^k}{(dk)!}\sum P(s_1, t_1, u_1, s_2, t_2, u_2, e_{s_1 s_2}, e_{t_1 t_2}).$$

Since the number of terms in the above sum is polynomial in k, we can approximate the bound by

$$\text{poly}(k)\frac{9^k}{(dk)!}P(s_1, t_1, u_1, s_2, t_2, u_2, e_{s_1 s_2}, e_{t_1 t_2}),$$

where the parameters are chosen to maximize P. Since

$$\binom{ds_1}{e_{s_1 s_2}}\binom{ds_2}{e_{s_1 s_2}}(e_{s_1 s_2})!\binom{dt_1}{e_{t_1 t_2}}\binom{dt_2}{e_{t_1 t_2}}(e_{t_1 t_2})!$$

is at most

$$\binom{d(s_1+t_1)}{e_{s_1 s_2}+e_{t_1 t_2}}\binom{d(s_2+t_2)}{e_{s_1 s_2}+e_{t_1 t_2}}(e_{s_1 s_2}+e_{t_1 t_2})!,$$

it suffices to consider $P(0, t_1, u_1, 0, t_2, u_2, 0, e_{t_1 t_2})$. Furthermore, $s_1 = s_2 = 0$ implies $t_1 = t_2 \wedge u_1 = u_2$, which, in turn, implies that $e_{t_1 t_2} = d(t_1 - u_2) =$

$d(t_2 - u_1)$. Thus, it suffices to consider $P(0, k-u, u, 0, k-u, u, 0, d(k-2u))$ for all u. Write $u = \alpha k$, then the probability of failure can be bound from above by

$$\operatorname{poly}(k) 9^k \frac{\binom{d(1-\alpha)k}{d(1-2\alpha)k}^2 (d(1-2\alpha)k)!((d\alpha)!)^2}{(dk)!}$$

which is equal to

$$\operatorname{poly}(k) 9^k \frac{(d(1-\alpha)k)!(d(1-\alpha)k)!}{(d(1-2\alpha)k)!(dk)!}.$$

By Stirling's formula, this can be written as

$$\operatorname{poly}(k) \left(\sqrt[d]{9} \frac{(1-\alpha)^{2(1-\alpha)}}{(1-2\alpha)^{1-2\alpha}} \right)^{dk}.$$

The latter quantity is strictly less than one for large enough k if

$$\sqrt[d]{9} \frac{(1-\alpha)^{2(1-\alpha)}}{(1-2\alpha)^{1-2\alpha}} < 1.$$

When $d \geq 5$ this is true for $2\alpha \geq 1 - 1/(2d)$, which translates into $e_{t_1 t_2} \leq k/2$.

3 The hardness of (1,B)-TSP

To adapt the construction from the previous section for the symmetric case we need to change some of the gadgets. Most changes in the equation gadgets are minor—the main change being that we test odd instead of even parity for equations with three variables (Fig. 8). There is a more substantial change in the edge gadget; it is changed according to Fig. 9.

If we assume that the tour behaves nicely, it is straightforward to prove a correspondence between the length of a tour and the number of equations left unsatisfied by the corresponding assignment.

Lemma 3.1. *The only way to traverse the equation gadget in Fig. 8 with a tour of length 5—if the edge gadgets count as length one for the moment—is to traverse an odd number of edge gadgets. All other locally optimal traversals have length 6.*

Proof. It is easy to see that any tour traversing either one or three of the ticked edges and leaving the third one untraversed has length 5. Any tour traversing zero or two ticked edges end up on the wrong side of the gadget and needs an extra cost of at least one to get back to the other side. □

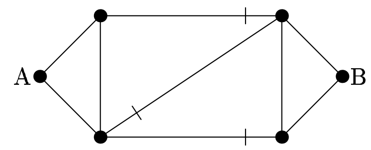

Figure 8. The gadget for equations of the form $x + y + z = 1$. There is a Hamiltonian path from A to B only if one or three of the ticked edges are traversed.

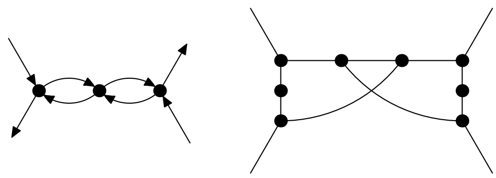

Figure 9. To transform the edge gadget from Fig. 2 into a gadget that can be used in the symmetric case, all occurrences of the structure to the left above are replaced with the structure to the right above.

Lemma 3.2. *In addition to the length 1 attributed to the edge gadget above, the length of a tour traversing an edge gadget in the intended way is 48.*

Proof. The total cost is $5 \cdot (7 + 1) = 40$. □

Lemma 3.3. *Suppose that there are $2m$ equations in the E3-Lin instance. If the tour is shaped in the intended way, i.e., every edge gadget is either traversed or untraversed, the length of the tour is $130m + u$, where u is the number of unsatisfied equations resulting from the assignment represented by the tour.*

Proof. The length of the tour on the edge gadgets is 18. There are three edge gadgets corresponding to every equation and every bridge in the edge gadget is shared between two equation gadgets. Thus, the length of the tour on the edge gadgets is $2m \cdot 3 \cdot 40/2 = 120m$. The length of the tour in the equation gadgets is 5 if the equation is satisfied and 6 otherwise. Thus, the total length is $130m + u$. □

In the same way as in the asymmetric case, it can be shown that the tour can be assumed to behave in the intended way. This gives the following lemma (we omit the proof):

Lemma 3.4. *Two traversed edge gadgets that share a bridge give an extra cost of at least 2 to the length of the tour.*

Suppose that $B \geq 6$. Then every semitraversed edge gadget adds an extra cost of at least one to the length of the tour.

Suppose that there are $2k$ occurrences of the variable x. Then at most k of the equation gadgets corresponding to x can be semitraversed.

There exists a coupling of the equation gadgets with the property that there can never be advantageous to have inconsistently traversed equation gadgets.

Given the above lemma, the main theorem follows in exactly the same way as in the asymmetric case.

Theorem 3.5. *For any constant $\epsilon > 0$, it is **NP**-hard to approximate (1,6)-TSP within $131/130 - \epsilon$.*

Proof. Given an instance of E3-Lin mod 2 with $2m$ equations where every variable occurs a constant number of times, we construct the corresponding instance of (1,6)-TSP. This can be done in polynomial time. By the above lemma, we can assume that all edge gadgets are traversed consistently in this instance. The assignment obtained from this traversal satisfies $2m - u$ equations if the length of the tour is $130m + u$. If we could decide if the length of the optimum tour is at most $(130 + \epsilon_1)m$ or at least $(131 - \epsilon_2)m$, we could decide if at most $\epsilon_1 m$ or at most $(1 - \epsilon_2)m$ of the equations are let unsatisfied by the corresponding assignment. But to decide this is **NP**-hard by Theorem 2.2. □

4 The hardness of (1,2)-ATSP

To prove a lower bound for (1,2)-ATSP we apply the construction used by Berman and Karpinski [2], a reduction from systems of linear equations mod 2 with exactly three unknowns in each equation to a problem called *Hybrid*, to prove hardness results for instances of several combinatorial optimization problems where the number of occurrences of every variable is bounded by some constant.

Definition 4.1. Hybrid *is the following maximization problem: Given a system of linear equations mod 2 containing n variables, m_2 equations with exactly two unknowns, and m_3 equations exactly with three unknowns, find an assignment to the variables that satisfies as many equations as possible.*

Theorem 4.2 [2]. *There exists instances of Hybrid with 42ν variables, 60ν equations with two variables, and 2ν equations with three variables such that:*

1. *Each variable occurs exactly three times.*
2. *For any constant $\epsilon > 0$, it is **NP**-hard to decide if at most $\epsilon\nu$ or at least $(1-\epsilon)\nu$ equations are left unsatisfied.*

Since we adopt the construction of Berman and Karpinski [2], we can partly rely on their main technical lemmas, which simplifies our proof of correctness.

On a high level, the (1,2)-ATSP instance in our reduction consists of a circle formed by *equation gadgets* representing equations of the form $x+y+z=0$ and $x+y=1$. These gadgets are coupled in a way ensuring that the three occurrences of a variable are given consistent values. In fact, the instances of Hybrid produced by the Berman-Karpinski construction have a very special structure. Every variable occurs in at least two equations with two unknowns, and those equations are all equivalences, i.e., equations of the form $x+y=0$. Since our gadget for equations with two unknowns tests odd parity, we have to rewrite those equations as $x+\bar{y}=1$ instead. Similarly, the equations of the form $x+y+z=1$ must be rewritten with one variable negated since our gadgets for equations with three unknowns only test even parity. Turning to the coupling needed to ensure consistency, we have three occurrences of every variable. Since we do not have any gadgets testing odd parity for three variables or even parity for two variables, we may have to negate some of the occurrences. We now argue that there are either one or two negated occurrences of every variable. The Hybrid instance produced by the Berman-Karpinski construction can be viewed as a collection of wheels where the nodes correspond to variables and edges to equations. The edges within a wheel all represent equations with two unknowns, while the equations with three unknowns are represented by hyperedges connecting three different wheels. Figure 10 gives an example of one such wheel. The equations corresponding to the edges forming the perimeter of the wheel can be written as $x_1+\bar{x}_2=1$, $x_2+\bar{x}_3=1$, $\ldots$, $x_{k-1}+\bar{x}_k=1$, and $x_k+\bar{x}_1=1$, which implies that there is at least one negated and at least one unnegated occurrence of each variable.

Corollary 4.3. *There are instances of Hybrid with 42ν variables, 60ν equations of the form $x+\bar{y}=1 \bmod 2$, and 2ν equations of the form $x+y+z=0 \bmod 2$ or $x+y+\bar{z}=0 \bmod 2$ such that:*

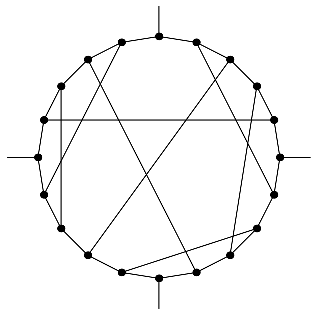

Figure 10. The Hybrid instance produced by the Berman-Karpinski construction can be viewed as a collection of wheels where the nodes correspond to variables and edges to equations.

1. *Each variable occurs exactly three times.*
2. *There is at least one positive and at least one negative occurrence of each variable.*
3. *For any constant $\epsilon > 0$, it is* ***NP****-hard to decide if at most $\epsilon\nu$ or at least $(1-\epsilon)\nu$ equations are left unsatisfied.*

To prove our hardness result for (1,2)-ATSP, we reduce instances of Hybrid of the form described in Corollary 4.3 to instances of (1,2)-ATSP and prove that, given a tour in the (1,2)-ATSP instance, it is possible to construct an assignment to the variables in the original Hybrid instance with the property that the number of unsatisfied equations in the Hybrid instance is related to the length of the tour in the (1,2)-ATSP instance.

To describe a (1,2)-TSP instance, it is enough to specify the edges of weight one. We do this by constructing a graph G and then let the (1,2)-TSP instance have the nodes of G as cities. The distance between two cities u and v is defined to be one if (u, v) is an edge in G and two otherwise. To compute the weight of a tour, it is enough to study the parts of the tour traversing edges of G. In the asymmetric case G is a directed graph.

Definition 4.4. *We call a node where the tour leaves or enters G an* end-point. *A node with the property that the tour both enters and leaves G in that particular node is called a* double endpoint *and counts as two endpoints.*

If c is the number of cities and $2e$ is the total number of endpoints, the weight of the tour is $c+e$ since every edge of weight two corresponds to two endpoints.

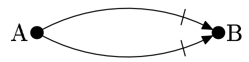

Figure 11. The gadget for equations of the form $x + y = 1$. There is a Hamiltonian path from A to B only if one of the ticked edges is traversed.

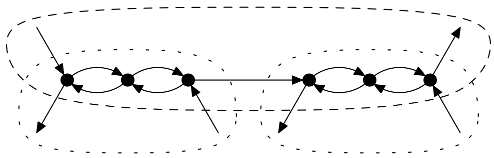

Figure 12. The gadget ensuring consistency for a variable. If there are two positive occurrences of the variable, the ticked edges corresponding to those occurrences are represented by the parts enclosed in the dotted curves and the ticked edge corresponding to the negative occurrence is represented by the part enclosed in the dashed curve. If there are two negative occurrences, the rôles are reversed.

4.1 The gadgets

The equation gadget for equations of the form $x + y + z = 0$ is shown in Fig. 1—the same gadget as in the $(1,B)$ case. However, the ticked edges now represent a different structure.

The equation gadget for equations of the form $x + y = 1$ is shown in Fig. 11. The key property of this gadget is that there is a Hamiltonian path through the gadget only if one of the ticked edges is traversed.

The ticked edges in the equation gadgets are syntactic sugar for a construction ensuring consistency among the three occurrences of each variable. As we noted above, either one or two of the occurrences of a variable are negated. The construction in Fig. 12 ensures that the occurrences are given consistent values, i.e., that either $x = 0$ and $\bar{x} = 1$, or $x = 1$ and $\bar{x} = 0$. If there is one negated occurrence of a variable, the upper part of the gadget connects with that occurrence and the lower part connects with the two unnegated occurrences. If there are two negated occurrences, the situation is reversed.

4.2 Proof of correctness

We want to prove that every unsatisfied equation has an extra cost of one associated with it. At first, it would seem that this is very easy—the gadget in Fig. 1 is traversed by a path of length four if the equation is satisfied and a path of length at least five otherwise; the gadget in Fig. 11 is traversed by a path of length one if the equation is satisfied and a path of length at least two otherwise; and the gadget in Fig. 12 ensures consistency and is traversed by a tour of length six, not counting the edges that were accounted for above. Unfortunately, things are more complicated than this. Due to the consistency gadgets, the tour can leave a ticked edge when it is half-way through it, which forces us to be more careful in our analysis.

We count the number of *endpoints* that occur within the gadgets; each endpoint gives an extra cost of one half. We say that an occurrence of a literal is *traversed* if both of its connected edges are traversed, *untraversed* if none of its connecting edges are traversed, and *semitraversed* otherwise. To construct an assignment to the literals, we use the convention that a literal is true if it is either traversed or semitraversed. We need to show that there are two endpoints in gadgets that are traversed in such a way that the corresponding assignment to the literals makes the equation unsatisfied. The following lemmas are easy, but tedious, to verify by case analysis:

Lemma 4.5. *It is locally optimal to traverse both* bridges*, i.e., both pairs of undirected edges, in the consistency gadget.*

Proof. By case analysis. □

Lemma 4.6. *Every semitraversed occurrence introduces at least one endpoint.*

Proof. By case analysis on traversed connection edges. □

Lemma 4.7. *It is always possible to change a semitraversed occurrence into a traversed one without introducing any endpoints in the consistency gadget.*

Proof. By case analysis on traversed connection edges. □

Given the above lemmas, the following two lemmas prove the properties we need regarding the equation gadgets.

Lemma 4.8. *A "satisfying traversal" of the gadget in Fig. 11 has length 1, all other locally optimal traversals have length at least 2, i.e., contain at least two endpoints within the gadget.*

Proof. If one of the ticked edges is traversed and the other is untraversed, the gadget is traversed by a tour of length 1. It is suboptimal to have one semitraversed and one untraversed edge, in this case it is possible to shorten the tour by transforming the semitraversed edge into a traversed one.

Two untraversed edges give a total cost of at least 2. It is impossible to have either two traversed edges or one traversed and one semitraversed ticked edge, since that gives a traversal which is not a tour. Two semitraversed edges give an extra cost of $\frac{1}{2}$ each, giving a total cost of at least 2. □

Lemma 4.9. *A "satisfying traversal" of the gadget in Fig. 1 has length 4, all other locally optimal traversals have length at least 5, i.e., contain at least two endpoints within the gadget.*

Proof. It is easy to see that any tour traversing two ticked edges and leaving the third one untraversed has length 4. The case with two semitraversed occurrences and one untraversed is suboptimal since a shorter tour can be produced in this way: Make the semitraversed occurrences traversed and then adjust the tour on the non-ticked edges to get a tour of length 4. Similarly, the case with one traversed and one semitraversed occurrence can be transformed into two semitraversed occurrences.

Any tour traversing one ticked edge and leaving the other two ticked edges untraversed has length at least 5. A tour semitraversing one ticked edge and leaving the other ticked edges untraversed can be transformed into a tour with one traversal and two non-traversals. It is impossible to have three traversals since this does not result in a tour. The case with two traversals and one semitraversal gives a cost of 5, and so does case with one traversal and two semitraverals, since each semitraversal has an extra cost of $\frac{1}{2}$ associated with it. □

When the above lemmas have been proven, we only need to prove that the gadget we use for consistency actually implements consistency.

Lemma 4.10. *The gadget in Fig. 12 ensures consistency and is traversed by a tour of length 6, not counting the edges or endpoints that were accounted for in the above lemmas.*

Proof. If there are no semitraversed occurrences, the gadget implements consistency correctly.

Suppose that the upper occurrence in Fig. 12 is semitraversed in such a way that the leftmost connecting edge is traversed but the rightmost is not. Then it is possible to have the lower left occurrence untraversed and the lower right occurrence traversed. Since a semitraversed occurrence is always part of an unsatisfied equation gadget, the following procedure produces a tour with equal cost: Make the upper occurrence untraversed and the lower left occurrence traversed. This makes the equation gadget that the upper occurrence is connected to satisfied and may make the equation gadget that the lower left occurrence is connected to unsatisfied.

Suppose that the lower left occurrence in Fig. 12 is semitraversed in such a way that the leftmost connecting edge is traversed but the rightmost is not. Then it is possible to have the lower right occurrence untraversed and the upper occurrence semitraversed. Since a semitraversed occurrence is always part of an unsatisfied equation gadget, the following procedure produces a tour with equal cost: Make the upper occurrence untraversed and the lower right occurrence traversed. This makes the equation gadget that the upper occurrence is connected to satisfied and may make the equation gadget that the lower right occurrence is connected to unsatisfied.

With similar arguments it can be shown that the lemma holds for all other possible cases. □

By combining the above lemmas, we have shown the following connection between the length of an optimum tour and the number of unsatisfied equations in the corresponding instance of Hybrid.

Theorem 4.11. *Suppose that we are given an arbitrary instance of Hybrid with n variables, m_2 equations of the form $x+\bar{y} = 1 \bmod 2$, and m_3 equations of the form $x + y + z = 0 \bmod 2$ or $x + y + \bar{z} = 0 \bmod 2$ such that:*

1. *Each variable occurs exactly three times.*
2. *There is at least one positive and at least one negative occurrence of each variable.*

Then we can construct an instance of (1,2)-ATSP with the property that a tour of length $6n + m_2 + 4m_3 + u$ corresponds to an assignment satisfying all but u of the equations in the Hybrid instance.

Corollary 4.12. *For any constant $\epsilon > 0$, it is* **NP**-*hard to approximate (1,2)-ATSP within $321/320 - \epsilon$.*

Proof. We connect Theorem 4.11 with Corollary 4.3 and obtain an instance of (1,2)-ATSP with the property that a tour of length

$$6n + m_2 + 4m_3 + u = 6 \cdot 42\nu + 60\nu + 4 \cdot 2\nu + u = 320\nu + u$$

corresponds to an assignment satisfying all but u of the equations in the Hybrid instance. Since, for any constant $\epsilon' > 0$, it is **NP**-hard to distinguish the cases $u \leq \epsilon'$ and $u \geq 1 - \epsilon'$, it is **NP**-hard to approximate (1,2)-ATSP within $321/320 - \epsilon$ for any constant $\epsilon > 0$. □

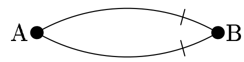

Figure 13. The gadget for equations of the form $x + y = 1$. There is a Hamiltonian path from A to B only if one of the ticked edges is traversed.

5 The hardness of (1,2)-TSP

It is possible to adapt the above construction for (1,2)-ATSP to prove a lower bound also for (1,2)-TSP. The equation gadget for equations containing three variables is changed in the same way as in the (1,B) case, the consistency gadget is change in a similar way.

5.1 The gadgets

The equation gadget for equations of the form $x+y = 1$ is shown in Fig. 13. The key property of this gadget is that there is a Hamiltonian path through the gadget only if one of the ticked edges is traversed.

The equation gadget for equations of the form $x + y + z = 1$ is shown in Fig. 8—the same gadget as in the (1,B) case.

The ticked edges in the equation gadgets are syntactic sugar for a construction ensuring consistency among the three occurrences of each variable. As we noted above, either one or two of the occurrences of a variable are negated. The construction in Fig. 14 ensures that the occurrences are given consistent values, i.e., that either $x = 0$ and $\bar{x} = 1$, or $x = 1$ and $\bar{x} = 0$. If there is one negated occurrence of a variable, the upper part of the gadget connects with that occurrence and the lower part connects with the two unnegated occurrences. If there are two negated occurrences, the situation is reversed.

5.2 Proof of correctness

In the same way as in the asymmetric case, it can be shown that the tour can be assumed to behave in the intended way. When this result is combined with the lower bound on the approximability of Hybrid, we obtain the following theorem:

Theorem 5.1. *Suppose that we are given an instance of Hybrid with n variables, m_2 equations of the form $x + \bar{y} = 1 \bmod 2$, and m_3 equations of the form $x + y + z = 0 \bmod 2$ or $x + y + \bar{z} = 0 \bmod 2$ such that:*

1. *Each variable occurs exactly three times.*

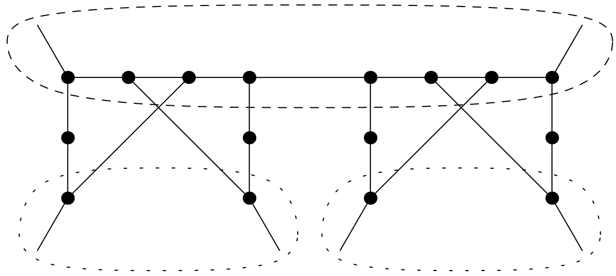

Figure 14. The gadget ensuring consistency for a variable. If there are two positive occurrences of the variable, the ticked edges corresponding to those occurrences are represented by the parts enclosed in the dotted curves and the ticked edge corresponding to the negative occurrence is represented by the part enclosed in the dashed curve. If there are two negative occurrences, the rôles are reversed.

2. There is at least one positive and at least one negative occurrence of each variable.

Then we can construct an instance of (1,2)-TSP with the property that a tour of length $16n + m_2 + 5m_3 + u$ *corresponds to an assignment satisfying all but u of the equations in the Hybrid instance.*

Theorem 5.2. *For any constant* $\epsilon > 0$*, it is* ***NP****-hard to approximate (1,2)-TSP within* $743/742 - \epsilon$.

Proof. We connect Theorem 5.1 with Corollary 4.3 and obtain an instance of (1,2)-TSP with the property that a tour of length

$$16n + m_2 + 5m_3 + u = 16 \cdot 42nu + 60\nu + 5 \cdot 2\nu + u = 742\nu + u$$

corresponds to an assignment satisfying all but u of the equations in the Hybrid instance. Since, for any constant $\epsilon' > 0$, it is **NP**-hard to distinguish the cases $u \leq \epsilon'$ and $u \geq 1 - \epsilon'$, it is **NP**-hard to approximate (1,2)-TSP within $743/742 - \epsilon$ for any constant $\epsilon > 0$. □

6 The unbounded metric case

Finally, we note that the equation gadget in Fig. 8 can be used to improve the bound given by Papadimitriou and Vempala [9] for the symmetric TSP with graph metric. Their construction has the property that the cost of a standard tour that corresponds to an assignment satisfying all but K equations is

$$4(8+1) \cdot \frac{3n}{2} + 10n + K,$$

where the first term corresponds to the cost of traversing the edge gadgets and the other two terms correspond to the cost of traversing the equation gadgets [9]. By removing the center node in every bridge in the edge gadgets and replacing the equation gadget with our gadget from Fig. 8, the cost becomes

$$4(7+1) \cdot \frac{3n}{2} + 5n + K = 53n + K,$$

which gives a lower bound of $107/106 + \epsilon$.

7 Conclusions

It should be possible to improve the reduction by eliminating the vertices that connect the equation gadgets for $x + y + z = \{0, 1\}$ with each other. This reduces the cost of those equation gadgets by one, which improves our bounds—but only by a miniscule amount. The big bottleneck, especially in the (1,2) case, is the consistency gadgets. If, for the asymmetric case, we were able to decrease the cost of them to four instead of six, we would improve the bound to $237/236 - \epsilon$; if we could decrease the cost to three, the bound would become $195/194 - \epsilon$. We conjecture that some improvement for the (1,2) case is still possible along these lines.

Acknowledgments

We thank Santosh Vempala for many clarifying discussions on the subject of this paper.

References

1. Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Márió Szegedy. Proof verification and the hardness of approximation problems. *Journal of the ACM*, 45(3):501–555, 1998.
2. Piotr Berman and Marek Karpinski. On some tighter inapproximability results. In *Proceedings of 26th International Colloquium on Automata, Languages and Programming*, volume 1644 of *Lecture Notes in Computer Science*, pages 200–209, 1999.

3. Robert D. Carr and Santosh Vempala. Towards a 4/3 approximation for the asymmetric traveling salesman problem. In *Proceedings of the Eleventh Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 116–125, 2000.
4. Nicos Christofides. Worst-case analysis of a new heuristic for the traveling salesman problem. Technical Report CS-93-13, G.S.I.A., Carnegie Mellon University, Pittsburgh, 1976.
5. Lars Engebretsen. An explicit lower bound for TSP with distances one and two. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science*, volume 1563 of *Lecture Notes in Computer Science*, pages 373–382, Trier, 1999.
6. Alan Frieze, Giulia Galbiati, and Francesco Maffioli. On the worst-case performance of some algorithms for the asymmetric traveling salesman problem. *Networks*, 12(1):23–39, 1982.
7. Johan Håstad. Some optimal inapproximability results. In *Proceedings of the Twenty-Ninth Annual ACM Symposium on Theory of Computing*, pages 1–10, 1997. Accepted for publication in *Journal of the ACM*.
8. Richard M. Karp. Reducibility among combinatorial problems. In Raymond E. Miller and James W. Thatcher, editors, *Complexity of Computer Computations*, pages 85–103. Plenum Press, New York, 1972.
9. Christos H. Papadimitriou and Santosh Vempala. On the approximability of the traveling salesman problem. In *Proceedings of the Thirty-second Annual ACM Symposium on Theory of Computing*, pages 126–133, 2000.
10. Christos H. Papadimitriou and Mihalis Yannakakis. The traveling salesman problem with distances one and two. *Mathematics of Operations Research*, 18(1):1–11, 1993.
11. Luca Trevisan. When Hamming meets Euclid: The approximability of geometric TSP and MST. In *Proceedings of the Twenty-Ninth Annual ACM Symposium on Theory of Computing*, pages 21–29, 1997.
12. Sundar Vishwanathan. An approximation algorithm for the asymmetric travelling salesman problem with distances one and two. *Information Processing Letters*, 44(6):297–302, 1992.

ECCC ISSN 1433-8092
http://www.eccc.uni-trier.de/eccc
ftp://ftp.eccc.uni-trier.de/pub/eccc
ftpmail@ftp.eccc.uni-trier.de, subject 'help eccc'