

The Approximation Ratio of the 2-Opt Heuristic for the Metric Traveling Salesman Problem

Stefan Hougardy Fabian Zaiser Xianghui Zhong

Research Institute for Discrete Mathematics, University of Bonn
Lennéstr. 2, 53113 Bonn, Germany

March 16, 2020

Abstract. The 2-Opt heuristic is one of the simplest algorithms for finding good solutions to the metric Traveling Salesman Problem. It is the key ingredient to the well-known Lin-Kernighan algorithm and often used in practice. So far, only upper and lower bounds on the approximation ratio of the 2-Opt heuristic for the metric TSP were known. We prove that for the metric TSP with n cities, the approximation ratio of the 2-Opt heuristic is $\sqrt{n/2}$ and that this bound is tight.

keywords: traveling salesman problem; metric TSP; 2-Opt; approximation algorithm

1 Introduction

In the Traveling Salesman Problem (TSP), we are given n cities with their pairwise distances. The task is to find a shortest tour that visits each city exactly once. The Traveling Salesman Problem is one of the most intensely studied problems in combinatorial optimization. It is well known to be NP-hard [4]. Without any additional assumptions, the Traveling Salesman Problem is also hard to approximate to any number that is polynomial in n [11]. The *metric TSP* is a special case of the TSP where the distance function satisfies the triangle inequality. The metric TSP is also NP-hard [6]. Therefore, a lot of time has been spent to find polynomial time algorithms with a small approximation ratio for the metric TSP. In 1976, Christofides [3] proposed an algorithm for the metric TSP with an approximation ratio of $3/2$. To date, no polynomial time algorithm with smaller approximation ratio is known.

For real-world instances appearing in practice, it turns out that many simple algorithms often find better solutions than Christofides' algorithm (see e.g. [5,1,10]). One of these algorithms is the *2-Opt heuristic*, which is

the key ingredient to the well-known Lin-Kernighan algorithm [8]. Starting with an arbitrary tour, the 2-Opt heuristic repeatedly replaces two edges of the tour by two other edges, as long as this yields a shorter tour. The 2-Opt heuristic stops when no further improvement can be made this way. A tour that the 2-Opt heuristic cannot improve is called *2-optimal*.

Experiments on real-world instances have shown that the 2-Opt heuristic applied to a greedy tour achieves much better results than Christofides' algorithm (see e.g. Bentley [1]). The exact approximation ratio of the 2-Opt heuristic for metric TSP was not known so far. In 1987, Plesník [9] proved a lower bound of $\sqrt{n/8}$. In 1999, Chandra, Karloff, and Tovey [2] presented a proof showing an upper bound of $4\sqrt{n}$. In 2013, Levin and Yovel [7] observed that this proof yields the value $2\sqrt{2n}$. This leaves a gap of factor 8 between the upper bound $2\sqrt{2n}$ and the lower bound $\sqrt{n/8}$. Our main result determines the exact approximation ratio of the 2-Opt heuristic:

Theorem 1. *The length of a 2-optimal tour in a metric TSP instance with n cities is at most $\sqrt{n/2}$ times the length of a shortest tour and this bound is tight.*

As the 2-Opt heuristic always returns a 2-optimal tour and the 2-Opt heuristic may start with any tour, we immediately get:

Corollary 1. *The 2-Opt heuristic for metric TSP instances with n cities has approximation ratio $\sqrt{n/2}$ and this result is tight.*

To prove Theorem 1, we show in Section 3 that the length of a 2-optimal tour in a metric TSP instance is bounded by $\sqrt{n/2}$ times the length of a shortest tour. In Section 4, we provide an infinite family of metric TSP instances and 2-optimal tours within these instances with length $\sqrt{n/2}$ times the length of a shortest tour. This proves the tightness stated in Theorem 1. Before proving the upper and the lower bound, we present in Section 2 some notation and background on the 2-Opt heuristic.

2 Metric TSP and the 2-Opt Heuristic

Let $G = (V(G), E(G))$ be a complete undirected graph with $|V(G)| = n$. The set $E(G)$ contains all $\binom{n}{2}$ possible edges between the n vertices. The *distances* between the vertices are defined by a function $c : E(G) \rightarrow \mathbb{R}_{\geq 0}$. A *tour* in G is a cycle that contains all the vertices of G . The *length* of a

tour T in G is defined as $c(T) := \sum_{e \in E(T)} c(e)$. A *shortest tour* is a tour of minimum length among the tours in G . Given a graph $G = (V(G), E(G))$ and a function $c : E(G) \rightarrow \mathbb{R}_{\geq 0}$, the Traveling Salesman Problem is to find a shortest tour in G . To simplify the notation, we will denote the length of an edge $\{x, y\} \in E(G)$ simply by $c(x, y)$ instead of the more cumbersome notation $c(\{x, y\})$. In the metric TSP, the distance function c satisfies the triangle inequality, i.e. we have for any set of three vertices $x, y, z \in V(G)$:

$$c(x, y) + c(y, z) \geq c(x, z). \quad (1)$$

An algorithm for the traveling salesman problem has *approximation ratio* $\alpha(n) \geq 1$ if for every TSP instance with n vertices, it finds a tour that is at most $\alpha(n)$ times as long as a shortest tour.

The 2-Opt heuristic repeatedly replaces two edges from the tour by two other edges such that the resulting tour is shorter. Given a tour T and two edges $\{a, b\}$ and $\{x, y\}$ in T , there are two possibilities to replace these two edges by two other edges. Either we can choose the pair $\{a, x\}$ and $\{b, y\}$ or we can choose the pair $\{a, y\}$ and $\{b, x\}$. Exactly one of these two pairs will result in a tour again. Without knowing the other edges of T , we cannot decide which of the two possibilities we have to choose. Therefore, we will assume in the following that the tour T is an *oriented cycle*, i.e. the edges of T have an orientation such that each vertex has exactly one incoming and one outgoing edge. Using this convention, there is only one possibility to exchange a pair of edges such that the new edge set is a tour again: two directed edges (a, b) and (x, y) have to be replaced by the edges (a, x) and (b, y) . Note that to obtain an oriented cycle again, one has to reverse the direction of the segment between b and x , see Figure 1.




Fig. 1. A TSP tour (left) and the tour obtained after replacing the edges (a, b) and (x, y) with the edges (a, x) and (b, y) (right). The orientation of the tour segment between the vertices b and x has been reversed in the new tour.

If (a, b) and (x, y) are two edges in a tour T and we have

$$c(a, x) + c(b, y) < c(a, b) + c(x, y)$$

then we say that replacing the edges (a, b) and (x, y) in T by the edges (a, x) and (b, y) is an *improving 2-change*. Thus, the 2-Opt heuristic can be formulated as follows:

2-Opt heuristic ($G = (V(G), E(G)), c : E(G) \rightarrow \mathbb{R}_{\geq 0}$)

- 1 start with an arbitrary tour T
- 2 **while** \exists improving 2-change in T
- 3 perform an improving 2-change
- 4 **output** T

3 The Upper Bound on the Approximation Ratio

Chandra, Karloff, and Tovey [2] proved in 1999 that the 2-Opt heuristic has an approximation ratio of $4\sqrt{n}$ for metric TSP. In 2013, Levin and Yovel [7] observed that their proof yields the upper bound $2\sqrt{2n}$. Here we present a new proof which improves this bound by a factor of 4:

Theorem 2. *The approximation ratio of the 2-Opt heuristic on metric TSP is at most $\sqrt{\frac{n}{2}}$.*

Proof. Let $G = (V(G), E(G))$ with $c : E(G) \rightarrow \mathbb{R}_{\geq 0}$ and $|V(G)| = n$ be a metric TSP instance and let T be an optimal tour. We may assume that T has length 1. We fix an orientation of the tour T and choose two vertices $p, q \in V(G)$ arbitrarily. For each vertex $v \in V(G)$, let $i_p(v)$ be the length taken mod 1 of the unique shortest directed p - v path starting in p and using only edges of T . By our assumption, we have $i_p : V(G) \rightarrow [0, 1)$ and we define i_q similarly. For the following, it helps to think of $[0, 1)$ as the circle with circumference 1 and of i_p as an embedding of the optimal tour into this circle such that the arc distance of two consecutive vertices on the circle is the length of the edge between them.

Define the following metric d on the interval $[0, 1)$, interpreted as a circle: $d(x, y)$ is the length of the shorter of the two arcs between x and y on the circle, i.e., $d(x, y) := \min\{|x - y|, 1 - |x - y|\}$. For any points $x, y, z \in [0, 1)$ we have $d(x, y) + d(y, z) \geq d(x, z)$ since combining the two shortest arcs between x, y and y, z and deleting the overlap results in an arc between x, z .

Let T' be a 2-optimal tour. As usual, we assume that it is directed. Now consider for each edge (u, v) of T' the set

$$S_{p,q}(u, v) = \{(x, y) \in [0, 1) \times [0, 1) : d(x, i_p(u)) + d(y, i_q(v)) < c(u, v)\},$$

as shown in Figure 2. We claim that all these sets are pairwise disjoint for distinct edges $(u_1, v_1), (u_2, v_2) \in E(T')$. Suppose that $S_{p,q}(u_1, v_1)$ and $S_{p,q}(u_2, v_2)$ intersect in (x, y) . Then, by the triangle inequality for c and d , we have

$$\begin{aligned} c(u_1, u_2) + c(v_1, v_2) &\leq d(i_p(u_1), i_p(u_2)) + d(i_q(v_1), i_q(v_2)) \\ &\leq d(i_p(u_1), x) + d(x, i_p(u_2)) + d(i_q(v_1), y) + d(y, i_q(v_2)) \\ &< c(u_1, v_1) + c(u_2, v_2). \end{aligned}$$

This contradicts the 2-optimality of T' . Hence, all these sets $S_{p,q}(u, v)$ are disjoint.




Fig. 2. The sets $S_{p,q}(a, b)$ (red) and $S_{p,q}(u, v)$ (green) assigned to the edges (a, b) and (u, v) of a 2-optimal tour. The sets are taken modulo the unit square and thus may consist of up to four parts.

Next, we want to show that the area of each set is independent of the choice of p and q . Let p' and q' be a different choice. Note that for all

vertices u , we have $i_{p'}(u) = i_{p'}(p) + i_p(u) \bmod 1$. In particular, we find $d(x, i_p(u)) = d(x + i_{p'}(p) \bmod 1, i_{p'}(u))$ because both points are shifted by $i_{p'}(p)$ on the circle $[0, 1)$. By the definition of $S_{p,q}(u, v)$, this means that the map

$$\begin{aligned} t : [0, 1) \times [0, 1) &\rightarrow [0, 1) \times [0, 1) \\ (x, y) &\mapsto (x + i_{p'}(p) \bmod 1, y) \end{aligned}$$

bijectionally sends $S_{p,q}(u, v)$ to $S_{p',q}(u, v)$. In other words, we obtain $S_{p',q}(u, v)$ from $S_{p,q}(u, v)$ by cutting the unit square vertically at $1 - i_{p'}(p) = i_p(p')$ into two rectangles and reassembling them, as described by the following two translations:

$$\begin{aligned} t_1 : [0, i_p(p')) \times [0, 1) &\rightarrow [i_{p'}(p), 1) \times [0, 1) \\ (x, y) &\mapsto (x + i_{p'}(p), y) \\ t_2 : [i_p(p'), 1) \times [0, 1) &\rightarrow [0, i_p(p')) \times [0, 1) \\ (x, y) &\mapsto (x - i_p(p'), y) \end{aligned}$$

Since they have disjoint domains and disjoint images, their union $t = t_1 \cup t_2$ is a bijection $[0, 1) \times [0, 1) \rightarrow [0, 1) \times [0, 1)$; sends $S_{p,q}(u, v)$ bijectionally to $S_{p',q}(u, v)$; and preserves the area of this set because it consists of translations. Analogously, we can cut the square horizontally at $i_q(q')$ to obtain $S_{p',q'}(u, v)$ from $S_{p',q}(u, v)$, again preserving its area. We conclude that the area of $S_{p,q}(u, v)$ is independent of the choice of p and q .

Now we want to show that the area of $S_{p,q}(u, v)$ is $2c(u, v)^2$ for any edge $(u, v) \in E(T')$. By the previous paragraph, we can choose $p = u$ and $q = v$. Then $S_{u,v}(u, v) = \{(x, y) \in [0, 1) \times [0, 1) : d(x, 0) + d(y, 0) < c(u, v)\}$. This set consists of four congruent isosceles right-angled triangles whose legs have length $c(u, v)$. Note that they do not overlap because the metric property ensures $c(u, v) \leq \frac{1}{2}$. Hence we have: $\text{area}(S_{p,q}(u, v)) = 4 \cdot \frac{c(u, v)^2}{2} = 2c(u, v)^2$.

Since the sets $S_{p,q}(u, v)$ for $(u, v) \in E(T')$ are pairwise disjoint, their combined area cannot exceed that of the unit square:

$$2 \sum_{e \in E(T')} c(e)^2 = \sum_{(u, v) \in E(T')} \text{area}(S_{p,q}(u, v)) \leq \text{area}([0, 1) \times [0, 1)) = 1.$$

Then the inequality of arithmetic and quadratic means implies

$$\frac{\sum_{e \in E(T')} c(e)}{n} \leq \sqrt{\frac{\sum_{e \in E(T')} c(e)^2}{n}} \leq \frac{1}{\sqrt{2n}}.$$

Hence, the length of the 2-optimal tour T' satisfies $\sum_{e \in E(T')} c(e) \leq \sqrt{\frac{n}{2}}$. \square

4 The Lower Bound on the Approximation Ratio

To prove a lower bound α on the approximation ratio of the 2-Opt heuristic for the metric TSP, one has to show that for infinitely many n , there exists a metric TSP instance with n cities that contains a 2-optimal tour which is α times longer than a shortest tour.

In 1999, Chandra, Karloff, and Tovey [2] provided such a construction for all n of the form $4 \cdot k^2$ for positive integers k , which shows a lower bound of $\frac{1}{4}\sqrt{n}$. Several years earlier, Plesník [9] had given another construction without explicitly stating a lower bound. It turns out that his construction yields a lower bound of $\frac{1}{\sqrt{8}}\sqrt{n}$ and works for all n of the form $8 \cdot k^2 - 8 \cdot k + 3$ for positive integers k .

The following result improves Plesník's lower bound by a factor of 2, and yields the tight result stated in Theorem 1.

Theorem 3. *The approximation ratio of the 2-Opt heuristic on the metric TSP is at least $\sqrt{\frac{n}{2}}$.*

Proof. Let G be a complete graph on $n := 2 \cdot k^2$ nodes with vertex set $V(G) := \{v_{i,j} : 1 \leq i, j \leq k\} \cup \{w_{i,j} : 1 \leq i, j \leq k\}$. For each i with $1 \leq i \leq k$, we call $V_i := \{v_{i,j} : 1 \leq j \leq k\}$ and $W_i := \{w_{i,j} : 1 \leq j \leq k\}$ a *section* of $V(G)$ and the v -vertices and w -vertices the two *halves* of $V(G)$.

We define a distance function $c : E(G) \rightarrow \mathbb{R}_{\geq 0}$ as follows:

$$\begin{aligned} c(v_{i,j}, w_{i',j'}) &= 1 && \text{for all } 1 \leq i, i', j, j' \leq k \\ c(v_{i,j}, v_{i',j'}) &= \begin{cases} 0 & i = i' \\ 2 & i \neq i' \end{cases} && \text{for all } 1 \leq j, j' \leq k \\ c(w_{i,j}, w_{i',j'}) &= \begin{cases} 0 & i = i' \\ 2 & i \neq i' \end{cases} && \text{for all } 1 \leq j, j' \leq k \end{aligned}$$

It is not hard to see that the function c satisfies the triangle inequality: Let u, v, w be any three vertices in $V(G)$. We want to show that $c(u, w) \leq c(u, v) + c(v, w)$. As c takes only the values 0, 1, 2, this is obvious if $c(u, v) \geq 1$ and $c(v, w) \geq 1$. Otherwise, without loss of generality, we may assume $c(u, v) = 0$. i.e., u and v are in the same section of $V(G)$. But then the definition of c implies $c(u, w) = c(v, w)$ and the triangle inequality is satisfied. Therefore, the graph G with cost function c is a metric TSP instance.




Fig. 3. The optimal tour T (left) and the 2-optimal tour T' (right) for $k = 4$. Note that the w -vertices on the right are mirrored at the diagonal compared to the w -vertices on the left. Thus, on the left, vertices within the sections V_i and W_i are in a row. On the right, the vertices in the sections V_i are in a row while the vertices in a section W_i are within a column. The colored bars contain the vertices belonging to the same section.

In the following, we will construct two special tours in G , which are depicted in Figure 3. Let T be the tour consisting of the edges

$$\begin{aligned}
 E(T) = & \{(v_{i,j}, v_{i,j+1}) : 1 \leq i \leq k, 1 \leq j < k\} \cup \\
 & \{(w_{i,j}, w_{i,j+1}) : 1 \leq i \leq k, 1 \leq j < k\} \cup \\
 & \{(v_{i,k}, w_{i,1}) : 1 \leq i \leq k\} \cup \\
 & \{(w_{i,k}, v_{i+1,1}) : 1 \leq i < k\} \cup \\
 & \{(w_{k,k}, v_{1,1})\}.
 \end{aligned}$$

The edges in the first two sets have length 0; the $2k$ edges in the other three sets have length 1. Therefore, we have $c(T) = 2k$. This tour is optimal because any tour has to visit all $2k$ sections of $V(G)$ and the distance of two vertices from different sections is at least 1.

Next we consider the tour T' with

$$\begin{aligned}
 E(T') = & \{(v_{i,j}, w_{j,i}) : 1 \leq i, j \leq k\} \cup \\
 & \{(w_{j,i}, v_{i,j+1}) : 1 \leq i \leq k, 1 \leq j < k\} \cup \\
 & \{(w_{k,i}, v_{i+1,1}) : 1 \leq i < k\} \cup \\
 & \{(w_{k,k}, v_{1,1})\}.
 \end{aligned}$$

Each edge of T' has length 1. Thus we have $c(T') = 2k^2$. We claim that the tour T' is 2-optimal. Assume by contradiction that T' is not 2-optimal. Consider a pair of edges $(a, b), (x, y)$ that allows an improving 2-change to $(a, x), (b, y)$. Hence $c(a, x) + c(b, y) < c(a, b) + c(x, y) = 2$ and one of $c(a, x)$ or $c(b, y)$ must be zero. This means a and x or b and y must be in the same section. But since a and b are in opposite halves of $V(G)$ (just like x and y), this means that a and x are in one half of $V(G)$ and b and y in the other. Hence $c(a, x), c(b, y) \in \{0, 2\}$. For an improving 2-change, we must have $c(a, x) = c(b, y) = 0$. This implies that a and x lie in the same section of $V(G)$ and b and y lie in the same section of $V(G)$. Thus there must exist indices i and j with $1 \leq i, j \leq k$ such that $a, x \in V_i$ and $b, y \in W_j$ or such that $a, x \in W_i$ and $b, y \in V_j$. This implies that there must exist two different edges from V_i to W_j or from W_i to V_j . However, this is a contradiction as by definition of T' , for any pair i, j with $1 \leq i, j \leq k$, there exists exactly one edge directed from V_i to W_j (namely the edge $(v_{i,j}, w_{j,i})$) and exactly one edge directed from W_j to V_i . This proves the 2-optimality of T' .

Combining the above findings we get

$$\frac{c(T')}{c(T)} = \frac{2k^2}{2k} = k = \sqrt{\frac{2k^2}{2}} = \sqrt{\frac{n}{2}}.$$

□

References

1. Jon Louis Bentley. Fast algorithms for geometric Traveling Salesman Problems. *ORSA Journal on Computing*, 4:387–411, 1992.
2. Barun Chandra, Howard Karloff, and Craig Tovey. New results on the old k -opt algorithm for the Traveling Salesman Problem. *SIAM J. Comput.*, 28:1998–2029, 1999.
3. Nicos Christofides. Worst-case analysis of a new heuristic for the Travelling Salesman Problem. Technical Report 388, Carnegie-Mellon University, 1976.
4. Michael R. Garey and David S. Johnson. *Computers and Intractability. A Guide to the Theory of NP-Completeness*. W. H. Freeman and Company, 1979.
5. David S. Johnson. Local optimization and the Traveling Salesman Problem. In *International Colloquium on Automata, Languages, and Programming. ICALP 1990*, pages 446–461, 1990.
6. Richard M. Karp. Reducibility among combinatorial problems. In Raymond E. Miller and James W. Thatcher, editors, *Complexity of Computer Computations*, pages 85–103, New York – London, 1972. Plenum Press.
7. Asaf Levin and Uri Yovel. Nonoblivious 2-opt heuristics for the Traveling Salesman Problem. *Networks*, 62:201–219, 2013.
8. S. Lin and B. W. Kernighan. An effective heuristic algorithm for the Traveling-Salesman Problem. *Operations Research*, 21:498–516, 1973.
9. Ján Plesník. Bad examples of the metric Traveling Salesman Problem for the 2-change heuristic. *Acta Mathematica Universitatis Comenianae*, 48–49:203–207, 1987.

10. Gerhard Reinelt. *The Traveling Salesman. Computational Solutions for TSP Applications*. Springer, 1994.
11. Sartaj Sahni and Teofilo Gonzalez. P-complete approximation problems. *Journal of the Association for Computing Machinery*, 23(3):555–565, July 1976.