

The Approximation Ratio of the Greedy Algorithm for the Metric Traveling Salesman Problem

Judith Brecklinghaus and Stefan Hougardy

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

December 24, 2014

Abstract. We prove that the approximation ratio of the greedy algorithm for the metric Traveling Salesman Problem is $\Theta(\log n)$. Moreover, we prove that the same result also holds for graphic, euclidean, and rectilinear instances of the Traveling Salesman Problem. Finally we show that the approximation ratio of the Clarke-Wright savings heuristic for the metric Traveling Salesman Problem is $\Theta(\log n)$.

keywords: traveling salesman problem; greedy algorithm, Clarke-Wright savings heuristic, approximation algorithm

1 Introduction

Let $G = (V, E)$ be an undirected complete graph and $c : E(G) \rightarrow \mathbb{R}_{\geq 0}$ be a length function on the edges. A *tour* is a cycle in G that visits each vertex exactly once. The *length* of a tour T is the sum of the lengths of the edges in T . The Traveling Salesman Problem (TSP) is to find a tour of minimum length. In this paper we study the *metric* Traveling Salesman Problem which is the special version of the Traveling Salesman Problem where the function c is metric, i.e., we have $c(\{x, y\}) + c(\{y, z\}) \geq c(\{x, z\})$ for any three vertices $x, y, z \in V(G)$. The metric Traveling Salesman Problem is known to be NP-hard [5] and therefore much effort has been spent to design polynomial time algorithms that find good tours.

An algorithm A for the traveling salesman problem has *approximation ratio* c if for every TSP instance it finds a tour that is at most c times longer than a shortest tour. The *greedy algorithm* is one of the simplest algorithms to find a TSP tour. It starts with an empty edge set and adds in each iteration the cheapest edge (with respect to the length function c), such that the resulting graph is a subgraph of a tour. Because of its simplicity and because the greedy algorithm achieves quite good results on real world instances (see e.g. [7, page 98]) it is often used in practice.

In Section 2 we will prove that the approximation ratio of the greedy algorithm for metric TSP is $\Theta(\log n)$ for instances with n vertices. Our result closes the long-standing gap between the so far best known lower bound of $\Omega(\log n / \log \log n)$ and the upper bound of $O(\log n)$. Both these bounds were proved in 1979 by Frieze [2].

Our result also holds for the *euclidean*, the *rectilinear*, and the *graphic* TSP. These are well studied special cases of the metric TSP. In the euclidean and the rectilinear TSP the cities are points in the plane and the distance between two cities is defined as the euclidean respectively rectilinear distance. A graphic TSP is obtained from an unweighted, undirected, and connected graph G which has as vertices all the cities. The distance between two cities is then defined as the length of a shortest path in G that connects the two cities.

Another well established approximation algorithm for the traveling salesman problem that achieves good results in practice (see e.g. [7, page 98]) is the Clarke-Wright savings heuristic [1]. This heuristic selects one city x and connects it by two parallel edges to each other city. This way one obtains a Eulerian tour that is transformed into a TSP tour as follows. For each pair of cities a, b let the *savings* be the amount by which the edge $\{a, b\}$ is shorter than the sum of the lengths of the two edges $\{a, x\}$ and $\{b, x\}$. The replacement of the edges $\{a, x\}$ and $\{b, x\}$ by the edge $\{a, b\}$ is called a *short-cut*. Now go through all city pairs in nonincreasing order of savings and short-cut the pair if this does not close a cycle on the cities different from x and if both cities in the pair do not get adjacent to more than two other cities. After going through all city pairs the result will be a TSP tour.

In Section 3 we will prove that the approximation ratio of the Clarke-Wright savings heuristic for metric TSP is $\Theta(\log n)$. Again, we close a long-standing gap between the so far best known lower bound of $\Omega(\log n / \log \log n)$ and the upper bound of $O(\log n)$. Both these bounds were proved in 1984 by Ong and Moore [6].

2 The Approximation Ratio of the Greedy Algorithm

In this section we describe the construction of a family of metric TSP instances G_k on which the greedy algorithm can find a TSP tour that is by a factor of $\Omega(k)$ longer than an optimum tour. Our construction is similar to the approach in [3].

(d) the edges in P have length 3^i with $0 \leq i \leq k$.

Proof. We use induction on k to prove the lemma. For $k = 0$ a partial greedy tour of length $8 \cdot 3^0 - 1 = 7$ that satisfies (a)–(d) is shown in Figure 1. Now assume we already have defined a partial greedy tour for G_k . Then we define a partial greedy tour for G_{k+1} recursively as follows. By (1) we have $|V_{k+1}| = 3^{k+3} - 1 = 3 \cdot (3^{k+2} - 1) + 2$. Therefore, we can think of G_{k+1} to be the disjoint union of three copies G'_k , G''_k , and G'''_k of G_k such that G'_k and G''_k are separated by a 2×1 grid. Moreover, we assume that G''_k is embedded after mirroring at the y-axis. This is shown in Figure 2.




Fig. 2. The recursive construction of a partial greedy tour for the instance G_{k+1} . The dashed lines indicate partial greedy tours in G'_k , G''_k , and G'''_k .

By definition s_{k+1} is the vertex at position $\frac{1}{2}(3^{k+2} + 1)$ in the top row of V_{k+1} . As we have $\frac{1}{2}(3^{k+2} + 1) = \frac{1}{2}(3^{k+2} - 1) + 1 = \frac{1}{2}|V_k| + 1$ this is the first vertex in the top row behind G'_k .

We now claim that a partial greedy tour for G_{k+1} may look as follows (see the bold edges in Figure 2): it contains the three recursively defined partial greedy tours in G'_k , G''_k , and G'''_k plus four additional edges. These four additional edges are the two edges of length one that leave r'_k and r''_k to the right respectively to the left plus the two edges $\{s'_k, s_{k+1}\}$ and $\{s''_k, s'''_k\}$. As s'_k is at position $\frac{1}{2}(3^{k+1} + 1)$ in the top row and s_{k+1} is at position $\frac{1}{2}(3^{k+2} + 1)$ in the top row, we have that the edge $\{s'_k, s_{k+1}\}$ has length

$$\frac{1}{2}(3^{k+2} + 1) - \frac{1}{2}(3^{k+1} + 1) = 3^{k+1}.$$

The edge $\{s''_k, s'''_k\}$ has length

$$\frac{1}{2}(3^{k+1} + 1) + \frac{1}{2}(3^{k+1} + 1) - 1 = 3^{k+1}.$$

Thus, we have constructed a partial tour that satisfies (a), (b), and (d) and its total length is

$$3 \cdot ((2k + 8) \cdot 3^k - 1) + 2 + 2 \cdot 3^{k+1} = (2(k + 1) + 8) \cdot 3^{k+1} - 1$$

and therefore also (c) holds. We still have to prove that this partial tour is contained in some greedy tour. The two edges of length one that leave r'_k and r''_k to the right respectively to the left may be chosen by the greedy algorithm at the very beginning. Afterwards, by induction the partial greedy tours within G'_k , G''_k , and G'''_k may be chosen. By induction, the longest edge that has to be considered by the greedy algorithm up to this step has length 3^k . We now claim that all edges that are incident to the vertices s'_k, s''_k, s'''_k , and s_{k+1} and that may be added by the greedy algorithm have length at least 3^{k+1} . For the vertices s''_k, s'''_k , and s_{k+1} this is immediately clear, as any vertex not belonging to G_{k+1} has distance at least 3^{k+1} to these vertices. We still have to rule out that s'_k has a neighbor to the left of G_{k+1} within a distance smaller than 3^{k+1} . But if there exists a neighbor to the left of G_{k+1} then by the recursive construction G_{k+1} is contained in an embedding of G_{k+2} and this implies that to the left of G'_k there is a mirrored copy of G_k . Thus the next neighbor to the left of s'_k has distance at least 3^{k+1} .

Theorem 1. *On graphic, euclidean, and rectilinear TSP instances with n cities the approximation ratio of the greedy algorithm is $\Theta(\log n)$.*

Proof. The upper bound is proven in [2]. For the lower bound we apply Lemma 1. The instance G_k defined above has $n := 3^{k+2} - 1$ cities and an optimum TSP tour in G_k has length n . As shown in Lemma 1 there exists a partial greedy tour in G_k of length at least $(2k + 8) \cdot 3^k - 1$. Thus, for $k \geq 1$ the approximation ratio of the greedy algorithm is greater than

$$\frac{(2k + 8) \cdot 3^k - 1}{3^{k+2} - 1} \geq \frac{(2k + 8) \cdot 3^k}{3^{k+2}} = \frac{2k + 8}{9} \geq \frac{2}{9} \cdot \log_3(n + 1).$$

Conditions (i) and (ii) are satisfied whenever the distances in G_k are defined by an L^p -norm. Thus Theorem 1 not only holds for the L^2 - and the L^1 -norm but for all L^p -norms.

2.1 The 1-2-TSP

For completeness we also provide the approximation ratio of the greedy algorithm for the 1-2-TSP.

Theorem 2. *The approximation ratio of the greedy algorithm for the 1-2-TSP is $\frac{3}{2} - \frac{1}{2n}$ on instances with n cities.*

Proof. We first present an example showing that the approximation ratio is at least $\frac{3}{2} - \frac{1}{2n}$. For an odd number n with $n > 4$ consider the cycle

on vertices $\{1, \dots, n\}$. Add to this cycle all edges $\{i, j\}$ with $|i - j| = 2$ and i and j odd. All these edges get length one while all other edges get length 2. Now the greedy algorithm may choose $(n + 1)/2$ edges of length one by taking the edges $\{2, 3\}$ and $\{1, n\}$ in addition to the edges $\{3, 5\}, \{5, 7\}, \dots, \{n - 2, n\}$. No other edge of length one can be chosen by the greedy algorithm. Thus the total length of the tour returned by the greedy algorithm is $(n + 1)/2 + 2(n - 1)/2 = (3n - 1)/2$. As an optimum tour obviously has length n , the approximation ratio of the greedy algorithm is at least $\frac{3}{2} - \frac{1}{2n}$.

Now we prove that the approximation ratio of the greedy algorithm is never worse than $\frac{3}{2} - \frac{1}{2n}$. Let T be a tour found by the greedy algorithm and denote by G_1 the subgraph of T that contains all edges of length one. Let m be the number of edges in G_1 . We may assume that T is not an optimum tour and therefore G_1 will consist of several connected components, each of which must be a path. Let e be an edge of length one contained in an optimum TSP tour. As e was not added by the greedy algorithm, it must either be incident to a vertex of degree two in G_1 or it connects two vertices of degree one in the same connected component of G_1 . For each connected component C of G_1 of size at least two we can have at most $2(|C| - 2) + 1 = 2(|C| - 1) - 1$ edges of the optimum tour that have length one and are either incident to a vertex of degree two in C or connect two vertices of degree one in C . By summing over all connected components of size at least two in G_1 we obtain $k \leq 2m - 1$ where k is the number of length one edges in an optimum TSP tour. As the greedy tour has length $2n - m \leq 2n - (k + 1)/2$ and the optimum tour has length $2n - k$ we conclude that the approximation ratio of the greedy algorithm is at most

$$\frac{2n - (k + 1)/2}{2n - k} \leq \frac{2n - (n + 1)/2}{2n - n} = \frac{3}{2} - \frac{1}{2n}.$$

3 The Approximation Ratio of the Clarke-Wright Savings Heuristic

Theorem 3. *For metric TSP instances with n cities the approximation ratio of the Clarke-Wright savings heuristic is $\Theta(\log n)$.*

Proof. The upper bound is proven in [6]. For the lower bound extend the construction used to prove Lemma 1. Let G_k be the graph as defined above and let x be a new vertex that is connected to all vertices in G_k . Take the graphic instance induced by G_k and define the length of all edges incident

to x as $\frac{1}{2} \cdot 3^{k+2}$. Note that this value is larger than any distance within the graph G_k .

Now start the Clarke-Wright savings heuristic on this instance with city x . As all edges incident to x have the same length, the short-cutting of a pair results in a reduction of the edge length that is only dependent on the length of the edge between the two cities in the pair. Therefore, the Clarke-Wright savings heuristic may choose the pairs to short-cut exactly in the same order as the greedy algorithm applied to G_k . Thus, we just have taken into account that the greedy tour and an optimum tour contain two edges of total length 3^{k+2} incident with city x . Similarly as in the proof of Theorem 1 we get that for $k \geq 1$ the approximation ratio of the Clarke-Wright savings heuristic is greater than

$$\frac{(2k+8) \cdot 3^k - 1 + 3^{k+2}}{3^{k+2} - 1 + 3^{k+2}} \geq \frac{(2k+17) \cdot 3^k}{2 \cdot 3^{k+2}} = \frac{2k+17}{18} \geq \frac{1}{9} \cdot \log_3 n.$$

References

1. G. Clarke and J. W. Wright. Scheduling of vehicles from a central depot to a number of delivery points. *Operations Research*, 12(4):568–581, 1964.
2. A. M. Frieze. Worst-case analysis of algorithms for travelling salesman problems. In *Third Symposium on Operations Research (Univ. Mannheim, Mannheim, 1978), Section 2*, volume 32 of *Operations Res. Verfahren*, pages 93–112. Hain, Königstein/Ts., 1979.
3. Stefan Hougardy and Mirko Wilde. On the nearest neighbor rule for the metric traveling salesman problem. *Discrete Applied Mathematics*, to appear.
4. Cor A.J. Hurkens and Gerhard J. Woeginger. On the nearest neighbor rule for the traveling salesman problem. *Operations Research Letters*, 32(1):1–4, 2004.
5. 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, 1972.
6. Hoon Liong Ong and J. B. Moore. Worst-case analysis of two travelling salesman heuristics. *Operations Research Letters*, 2(6):273–277, 1984.
7. Gerhard Reinelt. *The Traveling Salesman. Computational Solutions for TSP Applications*. Springer-Verlag, 1994.