The x-and-y-axes travelling salesman problem

Eranda Çela* Vladimir Deineko†‡ Gerhard J. Woeginger§

Abstract

The x-and-y-axes travelling salesman problem forms a special case of the Euclidean TSP, where all cities are situated on the x-axis and on the y-axis of an orthogonal coordinate system of the Euclidean plane. By carefully analyzing the underlying combinatorial and geometric structures, we show that this problem can be solved in polynomial time. The running time of the resulting algorithm is quadratic in the number of cities.

Keywords. Combinatorial optimization; polynomial-time algorithm; computational complexity; Euclidean travelling salesman problem;

1 Introduction

The *travelling salesman problem* (TSP) is a fundamental and well-studied problem in combinatorial optimization; see for example the book by Lawler, Lenstra, Rinnooy Kan & Shmoys [10]. An instance of the TSP consists of n cities $1, \ldots, n$ together with the distances d_{ij} between all pairs i, j of cities. The goal is to find a round-trip through all cities of the smallest possible total length. The TSP is well known to be NP-hard, and hence one branch of algorithmic research has turned to the identification of well-solvable special cases; see the survey articles by Gilmore, Lawler & Shmoys [9] and Burkard & al [2] for a wealth of information on this area.

In the *Euclidean TSP* the cities are points in the two-dimensional plane, and their distances are measured according to the standard Euclidean metric. It can be seen that in the Euclidean situation the shortest TSP tour does not intersect itself (Flood [7]), and hence the underlying geometric structures make the problem somewhat easier. Nevertheless, the Euclidean TSP is still NP-hard (see Papadimitriou [11] or Chapter 3 in the TSP book [10]).

The literature contains a wide range [1, 3, 4, 5, 6, 8, 13, 14] of polynomially solvable special cases of the Euclidean TSP where the positions of the cities are restricted to certain families $\mathcal{C}$ of nice and well-behaved curves. One of the strongest and most general results in this area is due to the beautiful paper of Rubinstein, Thomas & Wormald [14], where the so-called

*cela@opt.math.tu-graz.ac.at. Institut für Optimierung und Diskrete Mathematik, TU Graz, Steyrergasse 30, A-8010 Graz, Austria

†V.Deineko@warwick.ac.uk. Warwick Business School, Coventry, CV4 7AL, United Kingdom

‡Corresponding author; tel. +44 02476524501

§gwoegi@win.tue.nl. Department of Mathematics and Computer Science, TU Eindhoven, P.O. Box 513, 5600 MB Eindhoven, Netherlands

Constrained TSP (CTSP) is introduced and fully investigated. In the CTSP the family $\mathcal{C}$ is a finite set of smooth (that is, continuously differentiable) and compact curves in the plane such that (i) each curve in $\mathcal{C}$ has finite length, (ii) the curves in $\mathcal{C}$ intersect and self-intersect in a finite number of points, and (iii) at any intersection point the different branches of the curves approach from different directions; we refer the reader to [14] for formal definitions. It is shown in [14] that the CTSP can be solved in polynomial time, where the degree of the polynomial heavily depends on the family $\mathcal{C}$ (and can be prohibitively large).

Other papers in this area investigate special cases of the Euclidean TSP where the curves in $\mathcal{C}$ have infinite length (and hence are not compact). The papers [3, 4, 5, 13] study cases where $\mathcal{C}$ consists of several parallel lines, or of almost parallel half-lines. In these cases the no-self-intersection property of optimal tours severely restricts the combinatorics of such tours, which makes the problem attackable by dynamic programming. The easiest case of the Euclidean TSP that is not covered by the results in [3, 4, 5, 13] has all the cities on two intersecting lines.

Contribution of this paper. In this paper we will solely deal with the so-called *x-and-y-axes* traveling salesman problem (XY-TSP), where all cities are situated on the x-axis or on the y-axis of an orthogonal coordinate system for the Euclidean plane. (Note that if the coordinates of all cities in the XY-TSP additionally are bounded by a constant, then the XY-TSP boils down to a special case of the CTSP of [14].) At first sight, the XY-TSP looks like a fairly primitive problem that should give way easily: There are just two curves in $\mathcal{C}$ (which furthermore are straight lines!), and there is only a single intersection, and this intersection is at an angle of 90^0. Nevertheless, the problem has remained unsolved for a long time. It was first formulated by Cutler [3] in 1980, and then kept circulating in the TSP community; see for instance [2, 5].

Our contribution is a fast $O(n^2)$ algorithm for the XY-TSP, which ultimately is based on the main insights of [14]. We first carefully analyze the underlying combinatorial and geometric structures of the problem in Sections 3 through 6; each of these four sections is centered around one structural insight. In Section 7 we then combine all our insights into a fast polynomial time algorithm that is based on dynamic programming.

2 Preliminaries

Whenever we speak of the *origin*, we are referring to the origin $(0, 0)$ of the underlying orthogonal coordinate system in the Euclidean plane. Throughout the paper we will make the technical assumption that every instance of the XY-TSP satisfies the following:

> **Property P.**
> Either one of the cities is situated in the origin, or no edge in any optimal tour does contain the origin.

Let us justify this assumption. The assumption is trivially true for instances that have a city in the origin. Hence consider an instance I without any city in the origin, and define a new instance I' by adding the origin as a city. If I has an optimal tour with an edge containing the origin, then this induces an optimal tour of the same length for instance I'. If I has no optimal tour with an edge containing the origin, then the optimal objective value for I' is strictly larger than the optimal objective value for I.

Therefore, we may simply first solve instance I under property P and then solve instance I' under property P, and output the shorter of the two solution tours. This only doubles the running time.

Throughout the paper the notation $[a, b]$ will denote an edge connecting vertices a and b in a TSP tour. In the case where a and b lie on the same axis, $[a, b]$ will also be used to denote a path connecting a and b, where all the intermediate vertices along the path lie on the same axis as a and b. In the latter case we refer to such a path $[a, b]$ as an edge in the tour.

The following two simple lemmas will be useful in our proofs.

Lemma 2.1 *Let b and B be real numbers with $0 \le b < B$. Then $f(x) = \sqrt{x^2 + B} - \sqrt{x^2 + b}$ is a strictly decreasing function on the interval $[0, +\infty)$.*

Proof. One easily checks that f is continuous on $[0, +\infty)$, and that its derivative is negative on $(0, \infty)$. □

Lemma 2.2 $\sqrt{x_1^2 + y_2^2} + \sqrt{x_2^2 + y_1^2} > \sqrt{x_1^2 + y_1^2} + \sqrt{x_2^2 + y_2^2}$ *for* $0 \le y_1 < y_2$ *and* $0 \le x_1 < x_2$.

Proof. The statement follows from Lemma 2.1 with $B = y_2$ and $b = y_1$. □

3 The structure of optimal tours: implied edges

In our investigation, we will exploit one of the central ideas from [14] which in the context of the XY-TSP can be formulated as follows: Let $\mathcal{R}$ be a fixed circle centered at the origin. Then in any optimal tour, the number of edges that leave circle $\mathcal{R}$ is bounded by a small constant. Whereas for the general case of the CTSP considered in [14] it was not possible to explicitly calculate this constant, it can be easily done for the XY-TSP.

Theorem 3.1 *Consider an instance of the XY-TSP, and let $\mathcal{R}$ be a fixed circle centered at the origin. Let r_1 and r_2 be two cities on or inside $\mathcal{R}$. Let x_1 and x_2 be two cities that lie on or outside of $\mathcal{R}$, and on one of the half-axes, say on a half-axis of X.*

If an optimal tour contains both edges $[x_1, r_1]$ and $[x_2, r_2]$ then it also contains edge $[x_1, x_2]$.

This entire section will be dedicated to the proof of Theorem 3.1. Recall that an optimal tour of the Euclidean TSP can not intersect itself (Flood [7]). Thus in an optimal tour of the XY-TSP as described in Theorem 3.1 the endpoints r_1 and r_2 can not both lie on the same axis as the endpoints x_1 and x_2, as otherwise the tour would intersect itself. Thus, in order to prove the theorem it is enough to consider the four cases depicted in Figure 1 related to the positions of the endpoints r_1, r_2 with respect to the (given) half-axis containing x_1 and x_2. However, if we take property P into account (see Section 2), then case D becomes redundant: The only possible configuration in this case would be that r_1 coincides with the origin of the coordinates, and this configuration is already covered by case C.