Geometric and LP-based heuristics for the quadratic travelling salesman problem

Rostislav Staněk[*]   Peter Greistorfer[*]   Klaus Ladner[†]

Ulrich Pferschy[†]

Abstract

A generalization of the classical TSP is the so-called quadratic travelling salesman problem (QTSP), in which a cost coefficient is associated with the transition in every vertex, i.e. with every pair of edges traversed in succession. In this paper we consider two geometrically motivated special cases of the QTSP known from the literature, namely the angular-metric TSP, where transition costs correspond to turning angles in every vertex, and the angular-distance-metric TSP, where a linear combination of turning angles and Euclidean distances is considered.

At first we introduce a wide range of heuristic approaches, motivated by the typical geometric structure of optimal solutions. In particular, we exploit lens-shaped neighborhoods of edges and a decomposition of the graph into layers of convex hulls, which are then merged into a tour by a greedy-type procedure or by utilizing an ILP model. Secondly, we consider an ILP model for a standard linearization of QTSP and compute fractional solutions of a relaxation. By rounding we obtain a collection of subtours, paths and isolated points, which are combined into a tour by various strategies, all of them involving auxiliary ILP models. Finally, different improvement heuristics are proposed, most notably a matheuristic which locally reoptimizes the solution for rectangular sectors of the given point set by an ILP approach.

Extensive computational experiments for benchmark instances from the literature and extensions thereof illustrate the Pareto-efficient frontier of algorithms in a (running time, objective value)-space. It turns out that our new methods clearly dominate the previously published heuristics.

1 Introduction

The *travelling salesman problem* (TSP) is one of the best-known and most widely investigated combinatorial optimization problems with several famous books entirely devoted

---

[*]{peter.greistorfer, rostislav.stanek}@uni-graz.at. Department of Production and Operations Management, University of Graz, Universitätsstraße 15, 8010 Graz, Austria

[†]{klaus.ladner, pferschy}@uni-graz.at. Department of Statistics and Operations Research, University of Graz, Universitätsstraße 15, 8010 Graz, Austria

to its study [19, 25, 4, 16]. Plenty of variations of this problem have already been studied (see e. g. [16]).

In the *quadratic travelling salesman problem* (QTSP) we are still looking for a *tour*, i. e. a *Hamiltonian cycle*; however, different from the TSP, not only the costs of direct movements from one vertex to the next vertex are considered, but beyond that, costs arising for the successive use of two edges are taken into account. This means that we consider the *transition* in each vertex i which depends both on the predecessor and on the successor of i in the tour. Thus, transition costs, such as the effort of turning in path planning [1], changing the equipment in scheduling or the change of transportation means in logistic networks from one edge to another [3], can be modelled.

Mathematically, this aspect can be expressed by a quadratic objective function. The resulting optimization problem is known as the *(symmetric) quadratic travelling salesman problem* (QTSP) [12]. Note that the quadratic objective can also be extended to include costs for *every* pair of edges contained in the tour. Such a definition of quadratic costs was used for the *quadratic minimum spanning tree problem* (QMSTP) introduced in [5] and recently treated in [26] and [6] (see the references therein). Our version of costs only for successive edges is called *adjacent* QMSTP in these papers. The same terminology applies for the *quadratic shortest path problem* recently considered in [27] and [17].

In this paper we consider two relevant special cases of the QTSP, where the transition costs correspond to i) *turning angles* or to ii) a *linear combination of turning angles and Euclidean distances*. In both cases, vertices correspond to points in the Euclidean plane. The first problem was introduced by Aggarwal et al. [1] and is called the *angular-metric TSP* (AngleTSP). It arises from an application in robotics where changing the driving directions of a robot is more energy consuming for larger turning angles. Thus, one would prefer a tour which keeps the movement of the robot as closely as possible to a straight line. This is also relevant for steering-constrained robots and vehicles with high weight, were a straight movement is easier to control than narrow curves. Another application is the planning of trajectories of high-speed aircraft.

The second problem is the *angular-distance-metric TSP* (AngleDistanceTSP) which was first considered by Savla et al. [29] and later in [20]. It was introduced for an approximate solution of the *TSP for Dubins vehicles*, which also has applications in robotics.

In this paper we follow two lines of research: On one hand we focus on geometric properties of "good" solutions for both of these problems and derive a number of constructive heuristic solution algorithms. On the other hand we consider relaxations of a linearized ILP formulation for the QTSP studied in [2] and derive heuristic methods to generate feasible solutions from them. Some of our algorithms are *matheuristics* in the sense that they apply ILP models to (exactly) solve subproblems arising in the heuristic approach.

1.1 Related literature

The AngleTSP was introduced by Aggarwal et al. [1], who also showed that it is $\mathcal{NP}$-hard. Thus, the AngleDistanceTSP and the QTSP are $\mathcal{NP}$-hard as well, as they are generalizations of the AngleTSP. It was also shown that the AngleTSP allows a polynomial time approximation algorithm guaranteeing an approximation ratio within $O(\log n)$.

Several papers deal with the general QTSP in the context of applications in bioinformatics. The first contribution in this area is Jäger and Molitor [18] where the asymmetric version of QTSP is introduced. In that case, the cost of traversing vertices may differ between one direction and the other. The authors provide seven different heuristics and two exact solution approaches for the asymmetric version.

More recently, two papers extensively studied the symmetric QTSP and its bioinformatic application. The first one, due to Fischer et al. [13], gives three exact approaches and several heuristic algorithms. The exact approaches are a transformation to the TSP (which is then solved by an available software package), a Branch-and-Bound scheme and a Branch-and-Cut framework. However, these are able to solve only instances of moderate size within reasonable running time. Therefore, the authors also develop seven heuristics in [13]. The computational experiments do not give conclusive results, but suggest that the so-called *Cheapest-Insertion Heuristic* (generalizing the corresponding heuristic for the TSP, see Section 3.2) can be seen as the best currently known non-exact algorithm for the AngleTSP and the AngleDistanceTSP. Thus, it will be used as a reference method for our tests.

The second paper [14] by the same group of authors focuses on the bioinformatic details and tries to explore the boundaries of problem sizes that can still be solved to optimality by exact algorithms. To this purpose three different approaches are tested, namely Dynamic Programming, Branch-and-Bound and an improved Branch-and-Cut model, where the latter turns out to be vastly superior.

Polyhedral studies were done for the general symmetric QTSP by Fischer and Helmberg [12], who also derived several classes of strengthened subtour elimination constraints and proved that many of them are facet defining. In addition, they provided computational comparisons based on an ILP linearization and on the standard separation approach known from the TSP literature. Further polyhedral studies were provided in [11]. Recently, Fischer also analyzed the polytope of the asymmetric QTSP in [10] and presented involved lower bounding procedures based on an extended LP formulation which is solved by column generation in [28].

A simple ILP-based approach, based on integral separation of subtours, instead of the previously used standard fractional separation, was presented in [2] and performed remarkably well. Also a different linearization with a linear number of additional variables was developed in that paper. Moreover, the *maximization version* of the QTSP was introduced and some surprising structural properties for the case of the AngleTSP were shown.

Considering the QTSP from an approximation point of view, only few results are

known. It is easy to see that the standard TSP with distances d_{ij} can be represented as a special case of QTSP with identical solution value by defining costs $\frac{1}{2}(d_{ij} + d_{jk})$ for any transition from vertex i via j to k. Thus, the well-known fact that no constant-ratio approximation can exist for the TSP immediately carries over to the QTSP. For the special case of Halin graphs, it was shown in [31] that QTSP can be solved in linear time.

1.2 Our contribution

Looking at many optimal solutions of the AngleTSP plotted on the screen, we observed a certain recurrent pattern: tours tend to consist of large circular arcs trying to avoid sharp corners, resembling in some sense the structure of snail-shells, or the superposition of several circles. Moreover, solving an ILP model and relaxing the classical subtour elimination constraints, we often obtain a fairly moderate number of subtours, which is very different from the behaviour of the standard TSP.

In this paper we tried to exploit these empirical (although anecdotical) insights in a number of different ways:

1. Based on the observation that vertices near to an existing straight line (i. e. edge) of a preliminary path may be inserted into the path with little extra cost, the so-called *lens procedure* is introduced in Section 2, followed by some standard construction heuristics in Section 3.

2. The snail-shell shape of optimal tours mentioned above gives rise to several solution approaches based on the construction of convex hulls (Section 4). Peeling away convex hulls provides an onion-style set of subtours which can be merged together by different patching procedures, also with the help of ILP models.

3. Starting from existing ILP models, relaxing both integrality and subtour elimination constraints, but keeping the degree-two-constraint for every vertex, and rounding-up the obtained (fractional) solution, provides us with three types of components: subtours, paths and isolated points. We develop several procedures to transform these building blocks into one tour as described in Section 5. Some of the used (sub)procedures utilize again auxiliary ILP models of smaller size.

4. As for the standard TSP, it often makes sense to apply an improvement heuristic as a postprocessing step. We explore some ideas in this direction in Section 6 with particular emphasis on a matheuristic, which locally reoptimizes the solution with respect to all vertices lying in a rectangular sector of the given Euclidean plane; this procedure is then applied on different areas of the graph.

5. It is obvious that only extensive and carefully performed computational experiments can shed some light on the pros and cons of the individual algorithmic ideas and on the difficult question for their "best" combination. We pursued many more

possibilities than could be reasonable reported in the paper, but tried to narrow down the results to the most promising variants for each concept as reported in Section 7. All together, it turned out that the LP-based algorithms significantly dominate all other approaches for the AngleTSP. For the AngleDistanceTSP the situation is more diffuse: apart form the LP-based algorithms, a variant of the well known nearest-neighbour heuristic yields competitive solutions as well. Compared to the heuristics proposed in [13], our algorithms improve either the objective function values or the running times, and some of them dominate for both parameters.

Although our approaches were mainly motivated by the structure of AngleTSP, it turned out that many of them worked quite well also for the AngleDistanceTSP. Of course, the actual performances depends heavily on the weighting of the angle- and distance-component.

1.3 Formal problem definition

Let $G = (V, E)$ be an undirected complete graph with vertex set $V = \{1, 2, \ldots, n\}$ and edge set $E = \{\{i, j\} \colon i, j \in V, i \neq j\}$. A *path* P is an ordered sequence of vertices, i.e. $P = (p_1, p_2, \ldots, p_{|P|})$ with $p_i \in V$. Note that we will only consider simple paths, i.e. paths containing each vertex at most once. A *subtour* T is a path with the additional interpretation that all vertices should be visited in the given order and finally the edge from $p_{|T|}$ to p_1 is traversed. If T contains all vertices of V, we simply call T a *tour* in the graph, i.e. we have a *Hamiltonian cycle*. Alternatively, a tour T can be described by $T = (\sigma(1), \sigma(2), \ldots, \sigma(n))$ for a permutation σ of the vertices $1, 2, \ldots, n$. For a tour $T = \{t_1, t_2, \ldots, t_n\}$ we define the set of all *tour edges* as $E(T) \coloneqq \{\{t_1, t_2\}, \{t_2, t_3\}, \ldots, \{t_n, t_1\}\}$. Finally, for auxiliary graphs $\widehat{G} = (\widehat{V}, \widehat{E})$, we denote the *set of all neighbouring vertices* of a vertex $v \in \widehat{V}$ by $\delta(v)$.

In the QTSP we associate costs with every pair of adjacent edges traversed in succession. So, using the (incident) edges e and f one after the other in a tour gives rise to a certain cost value $c_{ef} \in \mathbb{R}_0^+$, which is assigned to the edge pair (e, f). Note that in this paper we only consider the symmetric case where $c_{ef} = c_{fe}$ for all $e \neq f \in E$. Equivalently, we can state costs for every triple of vertices $(i, j, k) \in V \times V \times V$ by setting $c_{ijk} = c_{ef}$ for $e = \{i, j\}$ and $f = \{j, k\}$. Clearly, $c_{ijk} = c_{kji}$. The QTSP asks for a *tour* T minimizing the *objective function*

$$z(G, T) \coloneqq \left(\sum_{i=1}^{n-2} c_{\sigma(i)\sigma(i+1)\sigma(i+2)} \right) + c_{\sigma(n-1)\sigma(n)\sigma(1)} + c_{\sigma(n)\sigma(1)\sigma(2)}. \tag{1}$$

In analogy, the cost of a subtour is defined. For convenience, we define the contribution to the objective function of a path $P = (p_1, p_2, \ldots, p_{|P|})$:

$$\bar{z}(G, P) \coloneqq \begin{cases} 0 & \text{if } |P| \leq 2 \\ \sum_{i=1}^{|P|-2} c_{p_i, p_{i+1}, p_{i+2}} & \text{otherwise} \end{cases} \tag{2}$$

Although the costs $c_{ijk} \in \mathbb{R}_0^+$ can be arbitrary non-negative real numbers in general, in this paper we focus on a geometric scenario where vertices correspond to points in the Euclidean plane. However, some of our approaches also work in the general case.

For the AngleTSP we define costs representing the turning angle of a tour based on the coordinates $(x_i, y_i) \in \mathbb{R}^2$ of a vertex $i \in V$, i.e.

$$c_{ijk} = \alpha_{ijk} := \arccos_{[0,\pi]} \left(\frac{\begin{pmatrix} x_j \\ y_j \end{pmatrix} - \begin{pmatrix} x_i \\ y_i \end{pmatrix}}{\left\| \begin{pmatrix} x_j \\ y_j \end{pmatrix} - \begin{pmatrix} x_i \\ y_i \end{pmatrix} \right\|} \cdot \frac{\begin{pmatrix} x_k \\ y_k \end{pmatrix} - \begin{pmatrix} x_j \\ y_j \end{pmatrix}}{\left\| \begin{pmatrix} x_k \\ y_k \end{pmatrix} - \begin{pmatrix} x_j \\ y_j \end{pmatrix} \right\|} \right) \tag{3}$$

for all $i \neq j \neq k \in V$. Here $\cdot$ denotes the scalar product. For the AngleDistanceTSP a weighted sum of the above turning angles and the Euclidean distances between successive vertices is used in the objective function, i.e. for arbitrary non-negative real multipliers λ_1, λ_2 we define:

$$c_{ijk} = \lambda_1 \cdot \alpha_{ijk} + \frac{\lambda_2}{2} \left(\left\| \begin{pmatrix} x_j \\ y_j \end{pmatrix} - \begin{pmatrix} x_i \\ y_i \end{pmatrix} \right\| + \left\| \begin{pmatrix} x_k \\ y_k \end{pmatrix} - \begin{pmatrix} x_j \\ y_j \end{pmatrix} \right\| \right) \tag{4}$$

2 Lens procedure

Before dealing with the particular solution approaches, we introduce an extension procedure, which can be applied in different algorithms as a subroutine. Given a partial solution consisting of subtours and/or paths and isolated points, the *lens procedure* inserts some of the isolated vertices into the subtours and/or paths and thus moves the partial solution closer to feasibility. At the same time the resulting increase of the objective function value is upper-bounded.

Let $P = (p_1, p_2, \ldots p_{|P|})$ be a path in G and denote its non-empty complement as $F = V \setminus P \neq \emptyset$. Furthermore, let p_i and p_{i+1}, $1 \leq i < |P|$, be two vertices of P traversed in succession. In the lens heuristic we try to include one or more vertices $f \in F$ into the path P between the vertices p_i and p_{i+1}. In particular, we position a *lens* between the two vertices p_i and p_{i+1} in such a way that their two positive curvatures intersect at these vertices and consider the set of points $L \subseteq F$ contained in this lens. The motivation for this new idea comes from the fact that the inclusion of any point inside such a lens causes an additional turning angle that is upper-bounded by the angle γ of the curvatures' tangent in p_i (and p_{i+1} as well) and the line between p_i and p_{i+1}. Moreover, the distance added by inserting $l \in L$, namely $\Delta = d_{p_i,l} + d_{l,p_{i+1}} - d_{p_i,p_{i+1}}$, which is also part of the objective function value in the AngleDistanceTSP, is relatively small. In fact, one can bound $\Delta \leq d_{p_i,p_{i+1}} \left(\sqrt{\frac{2}{1-\cos(\pi - 2\gamma)}} - 1 \right)$. Using parameter γ, the thickness of a lens can be controlled and thus the number $|L|$ of potential insertion candidates: the higher γ, the larger $|L|$ can get.

Now we choose the locally best vertex $l' \in L$ for inclusion in the path by minimizing the increase of the objective function value and selecting

$$l' = \arg\min_{l \in L} \left\{ \bar{z}(G, P) \colon P = (p_1, \ldots, p_i, l, p_{i+1}, \ldots, p_{|P|}) \right\}, \tag{5}$$

and set $P = (p_1, \ldots, p_i, l', p_{i+1}, \ldots, p_{|P|})$. This procedure is repeated recursively on the edges $\{p_i, l'\}$ and $\{l', p_{i+1}\}$ until there exists no more candidate, which could be included. Note that non-recursively including all $l \in L$ at once could lead to an unwelcome increase of the objective function value since not only the total sum of all Euclidean distances, but also the sum of all turning angles has to be taken into account. In fact, not all vertices $l \in L$ are necessarily contained in the lenses, which are stretched out in the further recursion steps.

Example 1. *Let us consider the graph $G = (V, E)$ with $V = \{1, \ldots, 6\}$ depicted in Figure 1, let $P = (1, 2, 3, 4)$ be a path and let us assume that we are solving the AngleTSP. Starting from the edge $\{2, 3\}$ and choosing $\gamma = 30° (\approx 0.5236)$ we obtain a lens (dashed in Figure 1) containing both remaining vertices 5 and 6, i. e. $L = \{5, 6\}$. Since $\bar{z}(G, (1, 2, 6, 3, 4)) \approx 1.4201 < 1.8082 \approx \bar{z}(G, (1, 2, 5, 3, 4))$, vertex $l' = 6$ is chosen and the process starts recursively with the new path $P = (1, 2, 6, 3, 4)$ and the two new edges $\{2, 6\}$ and $\{6, 3\}$ (see Figure 2). The lens stretched out on the edge $\{2, 6\}$ contains just one vertex 5 and the other lens is empty. Thus the process stops with the new path $P = (1, 2, 5, 6, 3, 4)$.*


Figure 1: lens heuristic idea started from the edge $\{2, 3\}$: step 1

Figure 2: lens heuristic idea started from the edge $\{2, 3\}$: step 2

Although the approach described does not solve the problem itself, it can be used in combination with many other heuristics. Such combinations usually yield significantly better results than the original heuristics alone as documented in Section 7. We will denote the *lens procedure* by LENS.

3 Simple construction heuristics and related approaches

3.1 Nearest-neighbour heuristics

The *nearest-neighbour heuristic* is one of the most common constructive heuristics used for the TSP (for details see e. g. [25]). The original TSP-idea is to start from a vertex and move to the nearest neighbour (not yet visited) in every step until no vertex remains unvisited. Finally, the tour is closed. In our problem we cannot start just from a vertex since we need two (incident) edges e and f to be able to estimate the cost value c_{ef} of using them in succession. Thus we start from an arbitrary edge $e = \{i, j\} \in E$ and look for a

vertex $k \in V$, which minimizes the additional cost, i.e. $k = \operatorname{argmin}_{k \in V \setminus \{i,j\}} \{c_{kij}, c_{ijk}\}$. This vertex is added either at the beginning or at the end of the path (i, j) depending on the respective costs. Similarly to the TSP, this can be repeated until no unconnected vertex remains. In the last step, the tour is closed. This basic variant of the nearest-neighbour heuristic will be denoted by NN. Note that NN always looks for the best next vertex in both path directions; in the literature this variant is sometimes also called the *two-directional-nearest-neighbour heuristic*. Jäger and Molitor [18] and Fischer et al. [13], however, understand a slightly different algorithm under this name.

A common idea is to run NN from all possible starting points, in our case from every start edge $e = \{i, j\} \in E$, and to choose the best solution. This variant will be denoted as NN_S.

Based on NN_S, one can use an improvement heuristic on every generated NN-solution and then choose the best one. We used the well known 2-opt heuristic (for details see e.g. [25]), however, in order to reduce the computation time, we do not start from every possible edge $e = \{i, j\} \in E$, but from every edge $e = \{i, j\} \in E$ for a prespecified vertex i. We denote this NN_S-variant by NN_S^2.

An extended approach includes the lens procedure LENS, described in Section 2, in the nearest neighbour heuristic: Whenever a new vertex is selected, LENS is applied on the new edge to possibly include further vertices into the existing path. We denote the resulting extended approach NN^L (resp. NN_S^L) for the underlying basic heuristic NN (resp. NN_S). Because of the high running times of NN_S^2 together with LENS, we did not pursue this combination any further.

Since NN is highly myopic, a look-ahead strategy might be employed which takes the next few potential iterations into account. In particular, we consider not only the nearest neighbour, but the $w \in \mathbb{N}$ nearest neighbours; for each such candidate we tentatively select it and repeat the whole process for $h \in \mathbb{N}$ steps, i.e. we evaluate the h next steps of the algorithm and in every step we consider again the w best possible successors. In this way we construct a decision tree with *tree width w* and *tree depth h*. Finally, among all candidates generated in the first step, we choose the vertex, which led to the best solution in the last iteration, i.e. the best leaf of the decision tree. This idea has similarities with the *piloting* strategy described by Duin and Voß [8], where $w \in \mathbb{N}$ candidates are considered, but the whole algorithm is evaluated for every such candidate (i.e. no decision tree is constructed, but the heuristic is run to the end for every candidate generated in the first step). We tested this *look-ahead* strategy for all four variants NN, NN_S, NN^L and NN_S^L. However, it turned out that the objective function values were not better than the values produced by other "good" algorithms, while – as can be expected – the running times increased considerably. Thus, we did not follow this idea to the end.

3.2 Cheapest-insertion heuristic

The next construction approach, again well-known from the TSP-context, is the *cheapest-insertion heuristic* (for more details see e.g. [25]). A small tour is enlarged step by step

until it contains all vertices. In particular, among all not visited vertices and all possible insertion positions the cheapest possibility is chosen in every step. Theoretically, the whole process could be started, similarly to NN, with a random tour just having 3 vertices, Fischer et al. [13], however, offer a smarter rule: they start with the subtour (u', v', w'), where

$$(u', v') = \arg \min_{u \neq v \in V} \left(\min_{x \in V \setminus \{u,v\}} c_{xuv} + \min_{y \in V \setminus \{u,v\}} c_{uvy} \right) \text{ and} \tag{6}$$

$$w' = \arg \min_{w \in V \setminus \{u',v'\}} z(G, (u', v', w)). \tag{7}$$

We will denote this algorithm CIF.

This algorithm exactly corresponds to the *Cheapest-Insert Heuristic* (CI) in [13], where it can be seen as the winner under the presented non-exact algorithms if the instance classes are restricted to the AngleTSP and the AngleDistanceTSP; thus it will be used as a "reference method" in this paper.

4   Approaches based on convex hulls

Although the majority of the test instances used in [12, 9, 13] are AngleTSP- or AngleDistanceTSP-instances (for exact definitions of these instance-types see Section 7.1.1), all heuristics presented in these papers are designed for the general QTSP setting and hardly make use of the geometric properties of AngleTSP and AngleDistanceTSP. Focusing on the AngleTSP we observed that optimal solutions consist of collections of large "circles" or "spiral" shapes (compare Aichholzer et al. [2]). We try to use this observation to design the following approach for the AngleTSP (note that we do not target the AngleDistanceTSP in this section). Recall that vertices are points in the Euclidean plane.

If all vertices $v \in V$ of G lay on a convex hull, an optimal AngleTSP-tour is obvious: it simply corresponds to the convex hull. Thus the main idea of our approach can be described as follows. First, the convex hull $H^1 \subseteq V$ of all vertices in V is computed and defines the first (outermost) subtour T^1. In the next step, all vertices in H^1 are removed from the graph and a new convex hull $H^2 \subseteq V \setminus H^1$ is computed on the remaining vertex set, defining a new subtour T^2. This process is repeated until at most two unused vertices remain; these vertices are included into the innermost (last) subtour by the CIF-idea. Finally, all τ obtained subtours $T^1, T^2, \ldots, T^\tau$ are merged into one single tour. We tested two different merging approaches.

4.1   Greedy merging approach

Consider two different arbitrarily chosen subtours T^i and T^ℓ with $1 \leq i \neq \ell \leq \tau$ and an arbitrary edge from each of them denoted by $e = \{u_e, v_e\} \in E\left(T^i\right)$ and $f = \{u_f, v_f\} \in E\left(T^\ell\right)$. There are two possibilities of merging the subtours T^i and T^ℓ: either we connect

them by the new edges $\{u_e, u_f\}$ and $\{v_e, v_f\}$ or by the new edges $\{u_e, v_f\}$ and $\{u_f, v_e\}$. In our algorithm we try both possibilities for all edge pairs $e \in E\left(T^i\right)$ and $f \in E\left(T^\ell\right)$ and for all pairs of subtours T^i and T^ℓ. Among all possibilities, we choose the one with the smallest sum of objective function values over all subtours contained in the respective solution. The number of subtours decreases by one in this step. This procedure is repeated until a single tour is obtained. The algorithm in this form (i.e. convex-hull approach combined with the described Greedy procedure) is abbreviated CH.

This simple Greedy idea can be enhanced by using *piloting*, i.e. a look-ahead probing strategy, applied in an analogous way as it was done to NN-based approaches in Section 3.1. However, in most cases the resulting algorithm yields the same tours as the basic *Greedy approach* while considerably worsening the computational time, so we did not follow this idea.

4.2 ILP merging approach

Another possibility is to merge all subtours $T^1, T^2, \ldots, T^\tau$, $k \geq 2$, at once by using a single ILP model. The idea is to remove exactly one edge in every subtour and then to connect the resulting paths into a single tour in an optimal way.

In order to do so, we define an auxiliary graph $\widehat{G} = (\widehat{V}, F)$ where each vertex $i \in V$ is represented by a pair of vertices $a, b \in \widehat{V}$. The edge set $F := F_L \cup F_S \cup F_C$ contains three types of edges: F_L, called *long edges*, replaces the original edges in a subtour, F_S, called *short edges*, joins two vertices a, b arising from one vertex $i \in V$ and F_C, called *connecting edges*, joins vertices of different subtours. Formally, each subtour T^i consists of $\left(v_1^i, v_2^i, \ldots, v_{|T^i|}^i\right)$ for $1 \leq i \leq \tau$ with $v_j^i \in V$. We define for $1 \leq i \leq \tau$:

$$\widehat{V}^i = \left\{a_j^i, b_j^i \colon 1 \leq j \leq \left|T^i\right|\right\}$$
$$F_S^i = \left\{\{a_j^i, b_j^i\} \colon 1 \leq j \leq \left|T^i\right|\right\}$$
$$F_L^i = \left\{\{b_j^i, a_{j+1}^i\} \colon 1 \leq j \leq \left|T^i\right| - 1\right\} \cup \left\{\{b_{|T^i|}^i, a_1^i\}\right\}$$

For every pair of subtours T^i and T^ℓ with $1 \leq i \neq \ell \leq \tau$ we define all possible connecting edges:

$$F_C^{i\ell} = \left\{\{u, v\} \colon u \in \widehat{V}^i, v \in \widehat{V}^\ell\right\}$$

Finally, we have

$$\widehat{V} := \cup_{i=1}^\tau \widehat{V}^i, \tag{8}$$
$$F_S := \cup_{i=1}^\tau F_S^i, \ \ F_L := \cup_{i=1}^\tau F_L^i, \ \ F_C := \cup_{i=1}^\tau \cup_{\ell=1, \ell \neq i}^\tau F_C^{i\ell} \tag{9}$$
$$F := F_S \cup F_L \cup F_C \tag{10}$$

Example 2. *Consider the graph $G = (V, E)$ depicted in Figure 3 and let $T^1 = \left(v_1^1, v_2^1, v_3^1, v_4^1\right)$ and $T^2 = \left(v_1^2, v_2^2, v_3^2, v_4^2\right)$ be two subtours. The corresponding graph*

$\widehat{G} = (\widehat{V}, F)$ *is depicted in Figure 4. The edges* $\{a_j^i, b_j^i\} \in F_S^i$, $i = 1, 2$, *are depicted densely dotted, the edges in* F_L^i, $i = 1, 2$, *solid and some of the edges in* F_C *incident to* a_2^1 *dashed.*


Figure 3: graph $G = (V, E)$ with subtours T^1 and T^2

Figure 4: graph $\widehat{G} = (\widehat{V}, F)$ (not all edges F_C are depicted)

The goal of this auxiliary graph $\widehat{G} = (\widehat{V}, F)$ is to transform the *quadratic costs* (depending on two edges) into *linear weights* (depending on just one particular edge). Obviously, by opening for each tour T^i one long edge $l^i \in F_L^i$ and its two incident short edges $s_1^i, s_2^i \in F_S^i$, and choosing exactly one connecting edge $r^i \in \cup_{\ell=1, \ell \neq i}^k F_C^{i\ell}$ for each tour T^i, $1 \leq i \leq \tau$, we can build one single subtour in the auxiliary graph $\widehat{G} = (\widehat{V}, F)$, which corresponds to a tour in the original graph $G = (V, E)$.

Example 3. *Let us again consider the graphs depicted in Figures 3 and 4. By taking all long and short edges except for the long edges* $\{b_2^1, a_3^1\}$ *and* $\{b_4^2, a_1^2\}$ *(one for every* $i = 1, 2$*) and the incident short edges* $\{a_2^1, b_2^1\}$, $\{a_3^1, b_3^1\}$, $\{a_1^2, b_1^2\}$ *and* $\{a_4^2, b_4^2\}$ *(two for every* $i = 1, 2$*) and by taking the two connecting edges* $\{a_2^1, b_1^2\}$ *and* $\{a_4^2, b_1^1\}$ *(one for every* $i = 1, 2$*) we clearly obtain the single subtour* $T = (a_1^1, b_1^1, a_2^1, b_1^2, a_2^2, b_2^2, a_3^2, b_3^2, a_4^2, b_3^1, a_4^1, b_4^1)$ *in the auxiliary graph* $\widehat{G} = (\widehat{V}, F)$ *corresponding to the tour* $T = (v_1^1, v_2^1, v_1^2, v_2^2, v_3^2, v_4^2, v_3^1, v_4^1)$ *in the original graph* $G = (V, E)$. *Note also that the vertices* b_2^1, a_3^1, b_4^2 *and* a_1^2 *(i.e. two vertices for every* $i = 1, 2$*) remain unconnected in the auxiliary graph* $\widehat{G} = (\widehat{V}, F)$ *after our merging procedure.*

Now, we assign linear weights $w_e \in \mathbb{R}_0^+$ to all edges $e \in F$. We also write w_{uv} for w_e if $e = \{u, v\}$. For the short edges $s \in F_S$, we set

$$w_{a_j^i b_j^i} = c_{v_{j-1}^i v_j^i v_{j+1}^i}$$

for all $1 \leq i \leq \tau$ with the obvious cyclic completion around $j = |T^i|$. This means that the weight of a short edge represents the quadratic cost arising from the turn in the corresponding vertex v_j^i in the particular subtour T^i for $1 \leq i \leq \tau$. For the long edges

$l \in F_L$ we set $w_l = 0$ and for the connecting edges F_C we set

$$w_{a_j^i a_k^\ell} = c_{v_{j-1}^i v_j^i v_k^\ell} + c_{v_j^i v_k^\ell v_{k-1}^\ell}$$

$$w_{a_j^i b_k^\ell} = c_{v_{j-1}^i v_j^i v_k^\ell} + c_{v_j^i v_k^\ell v_{k+1}^\ell}$$

$$w_{b_j^i a_k^\ell} = c_{v_{j+1}^i v_j^i v_k^\ell} + c_{v_j^i v_k^\ell v_{k-1}^\ell}$$

$$w_{b_j^i b_k^\ell} = c_{v_{j+1}^i v_j^i v_k^\ell} + c_{v_j^i v_k^\ell v_{k+1}^\ell},$$

i. e. the weights reflect the additional costs arising by the new connections. Again, the wrapping around of indices at the beginning and end of each subtour follows in a natural way.

Example 4. *Let us again consider the graphs depicted in Figures 3 and 4. Then*

$w_{a_2^1 b_2^1} = c_{v_1^1 v_2^1 v_3^1}$, $w_{a_3^1 b_3^1} = c_{v_2^1 v_3^1 v_4^1}$, $w_{a_4^2 b_4^2} = c_{v_3^2 v_4^2 v_1^2}$, $w_{a_1^2 b_1^2} = c_{v_4^2 v_1^2 v_2^2}$, $w_{a_2^2 b_2^2} = c_{v_1^2 v_2^2 v_3^2}$,
$w_{a_1^1 b_1^2} = c_{v_1^1 v_2^1 v_1^2} + c_{v_2^1 v_1^2 v_2^2}$, $w_{a_2^1 a_1^2} = c_{v_1^1 v_2^1 v_1^2} + c_{v_2^1 v_1^2 v_4^2}$, *etc.*

If we would merge the original two subtours $T^1 = \left(v_1^1, v_2^1, v_3^1, v_4^1\right)$ *and* $T^2 = \left(v_1^2, v_2^2, v_3^2, v_4^2\right)$ *to one tour* $T = \left(v_1^1, v_2^1, v_1^2, v_2^2, v_3^2, v_4^2, v_3^1, v_4^1\right)$ *(in the same way as in Example 3), then we would decrease the total costs by* $c_{v_1^1 v_2^1 v_3^1}$, $c_{v_2^1 v_3^1 v_4^1}$, $c_{v_4^2 v_1^2 v_2^2}$ *and* $c_{v_3^2 v_4^2 v_1^2}$ *and increase them by* $c_{v_1^1 v_2^1 v_1^2}$, $c_{v_2^1 v_1^2 v_2^2}$, $c_{v_3^2 v_4^2 v_3^1}$ *and* $c_{v_4^2 v_3^1 v_4^1}$. *In the auxiliary graph* $\widehat{G} = (\widehat{V}, F)$, *this exactly corresponds to omitting the long and short edges* $\{b_2^1, a_3^1\}$, $\{b_4^2, a_1^2\}$, $\{a_2^1, b_2^1\}$, $\{a_3^1, b_3^1\}$, $\{a_1^2, b_1^2\}$ *and* $\{a_4^2, b_4^2\}$ *and inserting the connecting edges* $\{a_2^1, b_1^2\}$ *and* $\{a_4^2, b_3^1\}$. *Note that we omit exactly 6 edges (one long and two incident short edges for every* $i = 1, 2$*) and insert 2 connecting edges (one for every* $i = 1, 2$*).*

After having defined the weights, for every $1 \le i \le \tau$ in the graph $\widehat{G} = (\widehat{V}, F)$, we want to remove one long edge $l^i \in F_L^i$ and its two incident short edges $s_1^i, s_2^i \in F_S^i$, and choose exactly one new connecting edge $r^i \in \cup_{\ell=1, \ell \ne i}^k F_C^{i\ell}$ in order to obtain one single subtour $\widehat{T}$ (which can be subsequently transformed into one single tour T in the original graph $G = (V, E)$). The edges of this process are chosen such that the total costs of the merging are minimized. This is realized by the following ILP. Let $x_f \in \{0, 1\}$ be a binary variable indicating whether the edge $f \in F$ is included in the tour $\widehat{T}$ or not. Then we can model the described problem as follows.

$$\min \sum_{f \in F} w_f x_f \tag{11}$$

$$\text{s.t.} \sum_{u \in \delta(v)} x_{uv} \le 2 \qquad \forall\, v \in \widehat{V}, \tag{12}$$

$$\sum_{u \ne v \in S} x_{uv} \le |S| - 1 \qquad \forall\, S \subset \widehat{V},\ S \ne \emptyset, \tag{13}$$

$$\sum_{l \in F_L^i} x_l = \left|T^i\right| - 1 \qquad \forall\ 1 \le i \le \tau, \qquad (14)$$

$$\sum_{s \in F_S^i} x_s = \left|T^i\right| - 2 \qquad \forall\ 1 \le i \le \tau, \qquad (15)$$

$$2x_l \ge \sum_{\substack{\{t,u\} \in F_S \cup F_C \\ t \in \delta(u)}} x_{tu} + \sum_{\substack{\{v,t\} \in F_S \cup F_C \\ t \in \delta(v)}} x_{vt} \quad \forall\ l = \{u,v\} \in F_L, \qquad (16)$$

$$x_l \le x_{s_1} + x_{s_2} \qquad \forall\ l \in F_L^i,\ \forall\ 1 \le i \le \tau,$$
$$s_1 \ne s_2 \in F_S^i,\ |s_1 \cap l| = |s_2 \cap l| = 1, \quad (17)$$

$$\sum_{u \in \widehat{V}^i,\ v \in \widehat{V} \backslash \widehat{V}^i} x_{u,v} = 2 \qquad 1 \le i \le \tau, \qquad (18)$$

$$x_f \in \{0,1\} \qquad \forall\ f \in F. \qquad (19)$$

The constraints (12) ensure that the vertex degree is at most 2. Note that we cannot enforce equality since we leave some vertices unconnected in graph $\widehat{G}$ (see also Example 3). (13) define the standard subtour elimination constraints while (14) and (15) guarantee that exactly one long edge $l \in F_L^i$ and two short edges $s_1^i \ne s_2^i \in F_S^i$ are excluded for every $1 \le i \le \tau$. (16) and (17) imply that each of the two excluded short edges s_1^i, s_2^i shares a vertex with the removed long edge l. Moreover, the constraints (16) ensure that all vertices, which are not included in the final tour $\widehat{T}$, are isolated and not connected among each other. The cut constraints (18) make sure that every vertex set $\widehat{V}^i$, $1 \le i \le \tau$, is connected to its complement $\widehat{V} \backslash \widehat{V}^i$ by exactly two connecting edges $e \in F_C$. Finally, the integrality constraints are expressed in (19). Note also that (18) does not follow from (13) since we do not ensure equality in (12).

This model has similarities to the ILP known for the TSP (the objective function (11) and the constraints (13) and (19) are the same and the constraints (12) are similar). To solve the above ILP, which has an exponential number of subtour elimination constraints (13), we adapt the solution approach proposed by Pferschy and Staněk in [21]. In particular, we relax the constraints (13), solve the remaining ILP to integrality, add one subtour elimination constraint for every subtour and repeat the whole process until just one tour $\widehat{T}$ exists.

Unfortunately, it turns out that the *ILP approach* for merging the subtours $T^1, T^2, \dots, T^\tau$ representing the collection of convex hulls does not outperform the *Greedy approach* (see Section 4.1) since the objective function values are fairly similar and the ILP approach always needs more time. Nevertheless, the ILP model introduced is used later in Section 5.1 in a different context, where it yields very good results in competitive computational times. Note that the above model contains much more restrictions than the related ILP for the standard TSP and therefore can be solved much faster than a pure TSP of the same size.

4.3 Enhanced convex hull strategies

The solutions generated by the merging of convex hulls as described above yield fairly well-structured solutions in general. However, this positive aspect deteriorates towards the inner vertices, i.e. for the convex hulls computed in the last few convex hull building iterations. It can be observed that the outer convex hulls (from the first iterations) usually contain a larger number of vertices, while the inner convex hulls (from the last iterations) often consist of just a few vertices. Since the sum of all turning angles in a convex polygon is always 2π, the average contribution to the objective function value of each vertex is usually much larger for these subtours with a small number of vertices.

Therefore we introduce the following matheuristic: We stop the convex hull building process as soon as the number of remaining vertices is at most a prespecified parameter $C \in \mathbb{N}$. Thus, we obtain at the end of the process the convex hulls H^1, H^2, ..., $H^{\tau'}$, $\tau' \in \mathbb{N}$ with the corresponding subtours, and a set of remaining vertices $R = V \setminus \{H^1 \cup H^2 \cup \ldots \cup H^{\tau'}\}$. Then we compute an optimal subtour $T^{\tau'+1}$ on the induced subgraph $G[R]$. This is done by the *integral ILP based approach* described in [2]. By choosing the value of parameter C the trade-off between running time and solution quality can be controlled. This approach is denoted by $\mathrm{CH_C}$. Clearly, for $C = 2$ we obtain CH.

A second enhanced strategy incorporates LENS procedure: After computing the vertices H^i of the current convex hull and the corresponding subtour T^i in iteration i, $1 \leq i \leq \tau$, we run LENS procedure on every edge in T^i and thus possibly enlarge the subtour with an only moderate increase of the objective function value. We will denote this variant by $\mathrm{CH^L}$. The combined variant including the features of both $\mathrm{CH_C}$ and $\mathrm{CH^L}$ will be denoted by $\mathrm{CH_C^L}$.

5 LP-based approaches

In this section we present a different solution strategy based on the utilization of LP relaxations. The general approach can be described as follows. First, QTSP is modelled via an ILP. In the next step, the LP relaxation of this model is solved and the fractional values of variables are rounded to integrality. In this way, one obtains either a partial solution or an infeasible solution with only limited violations of feasibility. Finally, the result of the rounding process is modified to either complete a partial solution or to reach feasibility. This is a standard approach for designing approximation algorithms (see e.g. [30, 15]) and heuristics.

The QTSP can be written as the following quadratic integer program with binary

edge variables x_{uv} for $\{u, v\} \in E$.

$$\min \sum_{u \neq v \neq t \in V} c_{uvt} x_{uv} x_{vt} \tag{20}$$

$$\text{s.t.} \quad \sum_{u \in \delta(v)} x_{uv} = 2 \qquad \forall\, v \in V, \tag{21}$$

$$\sum_{u \neq v \in S} x_{uv} \leq |S| - 1 \quad \forall\, S \subset V,\ S \neq \emptyset, \tag{22}$$

$$x_{uv} \in \{0, 1\} \quad \forall\, \{u, v\} \in E. \tag{23}$$

In the objective function (20) costs c_{uvt} are taken into account for the pair of consecutive edges $\{u, v\}$, $\{v, t\} \in E$, $u \neq v \neq t$, if both edges $\{u, v\}$ and $\{v, t\}$ are contained in the tour. Equations (21) are the *degree constraints* ensuring that each vertex is visited exactly once, (22) are the well-known *subtour elimination constraints* and, finally, (23) are the integrality constraints on the edge variables. In comparison to the standard model for the TSP by Dantzig et al. [7] only the objective function is changed.

This quadratic integer program can be linearized as suggested in [12]. We introduce a cubic number of additional integer variables y_{uvt} for all pairs of incident edges $\{u, v\}$, $\{v, t\} \in E$, $u \neq v \neq t$, where $y_{uvt} = 1$ if and only if the vertices u, v and t are visited in the tour in consecutive order.

$$\min \sum_{u \neq v \neq t \in V} c_{uvt} y_{uvt} \tag{24}$$

$$\text{s.t.} \qquad (21), (22), (23),$$

$$x_{uv} = \sum_{t \in V \setminus \{u,v\}} y_{uvt} = \sum_{t \in V \setminus \{u,v\}} y_{tuv} \quad \forall\, \{u, v\} \in E, \tag{25}$$

$$y_{uvt} \in \{0, 1\} \qquad\qquad\qquad \forall\, u \neq v \neq t \in V. \tag{26}$$

The x-variables have to correspond to a tour due to the constraints (21)–(23). Apart from that, this model has a linear objective function (24). Constraints (25) couple the x- and the y-variables. Finally, conditions (26) ensure the integrality of the y-variables.

Our approach works as follows. Based on this ILP, we relax the integrality constraints (23) and (26) and the subtour elimination constraints (22) and solve the obtained LP (obviously having polynomial size) to optimality. Let us denote by x'_e the obtained solution value of the x-variable for each edge $e \in E$. To construct a partial solution we first set $x_e = 0$ for all $e \in E$ and choose a rounding parameter ρ. Then we consider all x'-variables with $x'_e > \rho$ in decreasing order and set $x_e = 1$ if and only if the following two conditions are fulfilled:

C1. The degree of each vertex remains at most 2, i.e. $\sum_{u \in \delta(v)} x_{uv} \leq 2$ for all $v \in V$.

C2. The rounding does not imply any cycle in the resulting partial solution.

By this rounding procedure we obviously obtain a set of paths and isolated points. From these building blocks we construct a tour, i. e. a feasible QTSP-solution, with the following two steps. First, all paths are combined into one cycle by using an auxiliary ILP. Secondly, the isolated points are inserted into this cycle by the CIF-idea. The details of the resulting algorithm, denoted by LPP, can be described as follows.

For the merging of paths into a single cycle we proceed similarly to the merging procedure described in Section 4.2 and transfer the quadratic costs depending on two incident edges in the original graph into linear weights depending on one edge in a new auxiliary graph. Let $P^1 = \left(v_1^1, v_2^1, \ldots, v_{|P^1|}^1\right)$, $P^2 = \left(v_1^2, v_2^2, \ldots, v_{|P^2|}^2\right)$, $\ldots$, $P^\tau = \left(v_1^\tau, v_2^\tau, \ldots, v_{|P^\tau|}^\tau\right)$, where $\cup_{i=1}^\tau P^i \subseteq V$, be the τ paths obtained from our rounding procedure (putting aside the isolated vertices for the time being). We now define an auxiliary graph $\widehat{G} = (\widehat{V}, F)$, where each path P^i, $1 \le i \le \tau$, is represented by two vertices $a^i, b^i \in \widehat{V}$. The edge set $F := F_P \cup F_C$ contains two types of edges: *Path edges* in F_P represent the original paths and connect the a-vertices with their b-counterparts. *Connecting edges* in F_C join vertices representing different paths. Formally we define

$$\widehat{V} := \left\{a^i, b^i \colon i = 1, \ldots, \tau\right\},$$
$$F_P := \left\{\{a^i, b^i\} \colon i = 1, \ldots, \tau\right\}$$
$$F_C := \left\{\{u, v\} \colon u \neq v \in \widehat{V}\right\} \setminus F_P$$

Example 5. *Consider the graph $G = (V, E)$ given in Figure 5 and let $P^1 = \left(v_1^1, v_2^1\right)$, $P^2 = \left(v_1^2, v_2^2, v_3^2\right)$ and $P^3 = \left(v_1^3, v_2^3, v_3^3\right)$ be three paths obtained from the rounding procedure. The auxiliary graph $\widehat{G} = (\widehat{V}, F)$ is depicted in Figure 6. Edges F_P, where $\{a^i, b^i\}$ corresponds to path P^i, $i = 1, 2, 3$, are depicted solid, edges F_C incident to a^1 dashed and the edges F_C incident to b^1 densely dotted (other edges of F_C are not depicted).*


Figure 5: LP-based approach: graph $G = (V, E)$ with paths P^1, P^2 and P^3

Figure 6: LP-based approach: graph $\widehat{G} = (\widehat{V}, F)$ (not all edges F_C are depicted)

Similarly to the auxiliary graph described in Section 4.2 $\widehat{G} = (\widehat{V}, F)$ transforms the *quadratic costs* depending on two edges into *linear weights* depending only on one

particular edge. Obviously, by choosing one edge $r \in F_C$, one can compute the cost of the new connection in the original graph $G = (V, E)$. Formally, we set $w_{uv} = 0$ for all $\{u, v\} \in F_P$ and

$$w_{a^i a^\ell} = c_{v_2^i v_1^i v_1^\ell} + c_{v_1^i v_1^\ell v_2^\ell},$$

$$w_{a^i b^\ell} = c_{v_2^i v_1^i v_{|P^\ell|}^\ell} + c_{v_1^i v_{|P^\ell|}^\ell v_{|P^\ell|-1}^\ell},$$

$$w_{b^i a^\ell} = c_{v_{|P^i|-1}^i v_{|P^i|}^i v_1^\ell} + c_{v_{|P^i|}^i v_1^\ell v_2^\ell},$$

$$w_{b^i b^\ell} = c_{v_{|P^i|-1}^i v_{|P^i|}^i v_{|P^\ell|}^\ell} + c_{v_{|P^i|}^i v_{|P^\ell|}^\ell v_{|P^\ell|-1}^\ell}$$

for the connecting edges F_C, $1 \leq i \neq \ell \leq \tau$. Note that in this way, the costs incurred along the paths are not explicitly included in the merging model, since they remain constant anyway.

Example 6. *Let us again consider the graphs depicted in Figures 5 and 6. Then*

$$w_{a^1 b^1} = 0,\ w_{a^2 b^2} = 0,\ w_{a^3 b^3} = 0,\ w_{a^1 a^2} = c_{v_2^1 v_1^1 v_1^2} + c_{v_1^1 v_1^2 v_2^2},\ w_{a^1, b^2} = c_{v_2^1 v_1^1 v_3^2} + c_{v_1^1 v_3^2 v_2^2},$$

$$w_{b^1 a^2} = c_{v_1^1 v_2^1 v_1^2} + c_{v_2^1 v_1^2 v_2^2},\ w_{b^1, b^2} = c_{v_1^1 v_2^1 v_3^2} + c_{v_2^1 v_3^2 v_2^2},\ etc.$$

Now, we want to merge the paths into one cycle $\widehat{T}$ in an optimal way. Assume that $\tau \geq 2$ (other cases will be considered below) and let $x_f \in \{0, 1\}$ be a binary variable indicating whether edge $f \in F$ is included in the cycle $\widehat{T}$ or not. Then we can model the merging problem as follows.

$$\min \sum_{f \in F} w_f x_f \tag{27}$$

$$\text{s.t.} \sum_{u \in \delta(v)} x_{uv} = 2 \qquad \forall\ v \in \widehat{V}, \tag{28}$$

$$\sum_{u \neq v \in S} x_{uv} \leq |S| - 1 \quad \forall\ S \subset \widehat{V},\ S \neq \emptyset, \tag{29}$$

$$x_f = 1 \qquad \forall\ f \in F_P, \tag{30}$$

$$x_f \in \{0, 1\} \quad \forall\ f \in F. \tag{31}$$

The objective function value (27) and the constraints (28), (29) and (31) are equivalent to the standard ILP model for the TSP. The additional constraints (30) guarantee that all paths in the original graph are included in the resulting cycle. Similarly to Section 4.2 we adapt the solution approach proposed by Pferschy and Staněk in [21] for this problem, i. e. we relax the subtour elimination constraints (29), solve the remaining ILP to integrality, add one subtour elimination constraint for every subtour and repeat the whole process until just one tour $\widehat{T}$ exists. This cycle can then be easily transformed into a subtour T in the original graph $G = (V, E)$.

At the end, the remaining isolated points are inserted into the tour T by the CIF-idea. Now let us briefly consider the special case that the rounding procedure yields only one path ($\tau = 1$): If $\left|P^1\right| \geq 3$, we just close the path to a cycle instead of solving the auxiliary ILP and then insert the remaining isolated points by CIF. If $\left|P^1\right| = 2$, we use this path as a starting edge of CIF similar to (6), choose the next vertex according to (7) and continue by running CIF. If no path is formed ($\tau = 0$), the default version of CIF is performed.

The choice of the rounding parameter ρ influences the number of (undesired) isolated points generated by the rounding procedure. However, there is a trade-off: Decreasing ρ, the number of isolated points decreases, but the generated paths tend to be of lower quality, and vice-versa. Therefore a careful setting of this parameter is required.

Finally, let us remark that in theory we only know that the size of the auxiliary graph $\widehat{G} = (\widehat{V}, F)$ is smaller or equal to the size of the original graph $G = (V, E)$, but in practice it is observed to be much smaller. Roughly spoken, this approach transforms the QTSP instances to TSP instances of at most equal size, i. e. in the described algorithm we solve the quadratic problem by using its linear version without increasing the instance size. Since the size of usual QTSP test instances is perceived as quite small if the data is used for TSP instances, the exact solution of the auxiliary TSP (with one additional set of constraints) does not increase the running time too much.

5.1 Enhanced LP-based strategies

In this section we introduce some enhanced strategies, which lead to better results for the LPP-algorithm. A summary of the different improvement strategies is given in Figure 7. LPP corresponds to the left-most possibility, i. e. after the *rounding to paths and isolated points, no cycles* step, *merging the paths to a cycle* yields one cycle, which is enlarged by the remaining isolated points in the last step.

The *rerun strategy* is the first possibility for enhancing LPP: After solving the LP relaxation and performing the rounding procedure, we have $x_e = 1$ for some $e \in E$. Now, we fix $x_e = 1$ for all such edges and resolve the LP with a smaller number of free variables and start the rounding procedure again in order to obtain $x_e = 1$ for more edges $e \in E$. This process is repeated until the number of isolated points remains unchanged. LPP combined with this *rerun strategy* will be denote as LPP$^\text{R}$.

The next idea focuses on the rounding procedure. In LPP (and in LPP$^\text{R}$) x_e is set to 1 only if the edge $e \in E$ does not close a cycle. Due to this condition we obtain only paths and isolated points in the next step, but we "lose" some information provided by the solution of the LP relaxation. Instead, we can also permit to close cycles in the rounding procedure, i. e. omit the condition C2, and stick only to the degree condition C1. Then we have to deal with a structure consisting of subtours, paths and isolated points. One obvious possibility is to merge the paths to one cycle in the same way as it is done in LPP and then merge all cycles (i. e. the ones directly obtained from the rounding procedure and the one created from the paths) to one subtour. This can be
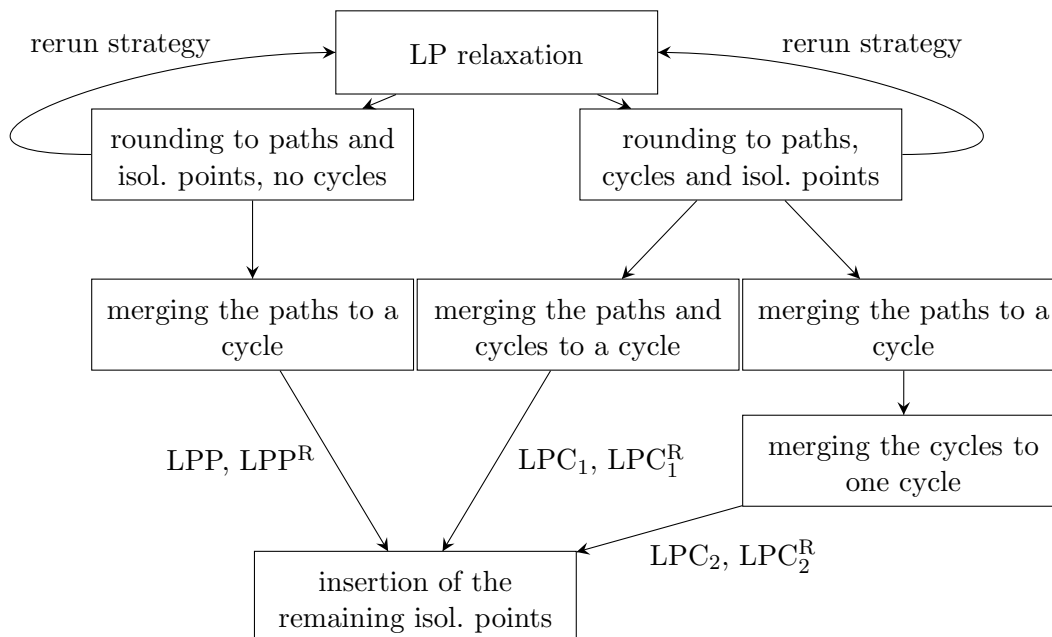
Figure 7: enhanced LP-based strategies

done by subsequently solving two auxiliary ILPs, namely (27)–(31) and then (11)–(19) as described in Section 4.2. As the last step, the isolated points are included by the CIF-idea. This approach corresponds to the right-most possibility in Figure 7 and can be used with or without the rerun strategy. We will denote the resulting algorithms by LPC_2^{R} and LPC_2, respectively.

LPC_2^{R} and LPC_2 use both auxiliary ILPs described in Sections 5 and 4.2 in succession. A natural idea is to combine these ILPs and to merge all cycles and paths at once. In particular, we deal with all cycles in the same way as in Section 4.2 and with all paths in the same way as in Section 5. The only difference is that we add *connecting edges* F_C also between all transformed cycles and all shrunk paths. The weights w_c of these new connections can be computed according to the same principle as it is done for other connecting edges. This variant corresponds to the middle possibility in Figure 7 and will be denoted by LPC_1^{R} (with rerun) and LPC_1 (without rerun).

Note that LPC_1 does not necessarily lead to the same or better results than LPC_2: if the paths and cycles are merged at once (LPC_1), no path edge is allowed to be removed. If, however, the paths are merged to a cycle first and then all cycles are merged together, the possibility of opening a path edge in the second step of LPC_2 is added.

Finally, we tested algorithms based on *randomised rounding* as introduced by Raghavan and Tompson [23]. Its main principle is to use the optimal solution values of the variables of a LP relaxation as probabilities for rounding up. We tested more variants of this approach applied to LPP and LPC_2 and although the obtained objective

function values are good for small values of n, the algorithms fail to produce good solutions in a reasonable computational time for larger instances. Thus we did not follow this idea any further.

6 Improvement approaches

The 2-*opt-heuristic* and the 3-*opt-heuristic* belong to the standard improvement algorithms used in the TSP context (for more details see [25]). Both algorithms work as follows: Basically, for a starting tour T, 2 or 3 edges are removed and the remaining segments of T are reconnected to obtain a new tour T'. If $z(T') < z(T)$, an improved tour was found and T' replaces T, otherwise T' is discarded. This process is repeated until no improvement is possible. We tested both algorithms and will denote them by +2O and +3O respectively. Note that depending on the precise definition, +3O is not always +2O-inclusive; we decided to use the +2O-inclusive variant.

In our computational experiments (see Section 7.3), we made two observations which are in line with the behavior observed for the analogous heuristics applied to the TSP.

1. +2O has relatively small running times in comparison to all stand-alone heuristics described in Sections 3, 4, and 5 and improves the solutions significantly.

2. +3O has running times comparable to the slower stand-alone heuristics, such as the LP-based ones, and yields significantly better solutions than +2O.

The naive extension of this approach to a 4-*opt-heuristic* can be expected to yield further improvements while the running times would increase dramatically. Therefore, we follow a different and more structured approach for removing parts of an existing solution while trying to find improvements.

6.1 Magnifying glass matheuristic

The magnifying glass heuristic can be understood as a *large neighborhood search*, where a (starting) solution is partially destroyed and repaired by corresponding operators; a process, which is iteratively repeated until a local optimum is reached (see [22]). The underlying idea of our approach works as follows: We consider a subset $S \subset V$ of vertices and remove all edges adjacent to S from the tour. Thus, we are left with several paths and isolated points. A new, possibly improved tour containing these paths is computed optimally by means of a quadratic programming model.

Formally, our algorithm works as follows. Let $T = (t_1, t_2, \ldots, t_n)$ be a tour in the graph $G = (V, E)$ and remember that $E(T) = \big\{\{t_1, t_2\}, \{t_2, t_3\}, \ldots, \{t_n, t_1\}\big\}$ is the set of all tour edges. Then for a set of vertices $S \subset V$, we remove all tour edges, which have at least one vertex in S, i.e. all $\{u, v\} \in E(T)$ with $\{u, v\} \cap S \neq \emptyset$. The remaining tour edges form a number of τ paths denoted by $P^1, P^2, \ldots, P^\tau$, where $P^i = (v_1^i, v_2^i, \ldots, v_{|P^i|}^i)$, $1 \leq i \leq \tau$. The idea is to reconnect these paths and the remaining isolated points

$I = V \setminus \{P^1 \cup P^2 \cup \ldots \cup P^\tau\}$ (mainly consisting of the vertices in S) in an optimal way to build a new tour T' as a possible improvement of T. This whole process is repeated for different choices of S, which constitutes a matheuristic type of improvement procedure.

Although the general procedure is defined for arbitrary sets S, the practical motivation is based on the consideration of vertex sets S corresponding to points in a geographical neighbourhood. If S contains all vertices of a certain region of the Euclidean plane, it can be expected that the removal of all connections within S yields only a small number of unconnected paths. Thus, the optimization should be mainly concerned with the improvement of the tour in a limited part of the given graph. This is the reason for using the expression "magnifying glass", meaning that we look at a small region of the graph and insert its vertices in a locally optimal way into the existing solution.

Following this intuition, we want to avoid short path fragments arising from the removal process which can be expected to lie close to the region of S. Therefore, in our implementation we enlarge S by all vertices of paths consisting of at most 3 vertices and reduce τ accordingly. Also all isolated points outside S arising from the removal of edges are added to S thus setting I equal to the enlarged set S.

Let us now consider the reconnecting procedure. We have a set of paths P^i, $1 \leq i \leq \tau$, and set of isolated points I. In order to connect them into a new tour, we define a complete auxiliary graph $\widehat{G} = (\widehat{V}, F)$ with vertex set $\widehat{V}$ representing all isolated points I and containing two vertices a^i, b^i for each path P^i, $1 \leq i \leq \tau$. We model the reconnecting problem as an instance of QTSP with a smaller number of vertices and the following coefficients:

$$
\begin{aligned}
w_{uvt} &= c_{uvt} && \text{for all } u \neq v \neq t \in I, \\
\left. \begin{aligned}
w_{a^i b^i t} &= c_{v^i_{|P^i|-1} v^i_{|P^i|} t} \\
w_{t a^i b^i} &= c_{t v^i_1 v^i_2}
\end{aligned} \right\} && \text{for all } 1 \leq i \leq \tau. \text{ and for all } t \in I
\end{aligned}
$$

Now, an optimal tour can be obviously found by using the standard QTSP integer program, i.e. by (20)–(23) or, as we do, by the linearisation (24), (21)–(23), (25), (26) described in Section 5, with the additional constraints

$$x_{a^i b^i} = 1 \quad \forall\, 1 \leq i \leq \tau. \tag{32}$$

These constraints ensure that the edges $\{a^i, b^i\}$, $1 \leq i \leq \tau$, representing the paths P^i are contained in the tour. To solve this ILP to optimality, we adopt the *integral approach* proposed by Aichholzer et al. in [2], i.e. we relax the subtour elimination constraints (22), solve the remaining ILP to integrality, add one subtour elimination constraint for every generated subtour and repeat the whole process until the solution consists of just one tour $\widehat{T'}$. This cycle can then be easily transformed into a new tour T' in the original graph with $z(T') \leq z(T)$. Again note that in this way, all costs incurred along the paths are not explicitly included in the model, since they remain constant anyway.

Example 7. *Consider the graph $G = (V, E)$ with the tour T and with the set S consisting of all vertices inside the square area bounded by the dotted line in Figure 8. After removing all tour edges having at least one vertex in S (dashed in Figure 8), we obviously obtain two paths P^1 and P^2 (solid in Figure 8) and 8 isolated vertices.*


Figure 8: graph $G = (V, E)$ with subset S and paths P^1 and P^2

Figure 9: graph $\widehat{G} = (\widehat{V}, F)$ (depicted are only edges fixed by (32))

The auxiliary graph $\widehat{G} = (\widehat{V}, F)$ can be seen in Figure 9. It consists of all isolated vertices I and of the vertices a^1, b^1 and a^2, b^2 representing the paths P^1 and P^2, respectively. Although $\widehat{G}$ is complete, for a better representation only the edges $\{a^1, b^1\}$, $\{a^2, b^2\}$ fixed by (32) are depicted.

This procedure could be performed for any subset S. Since we deal with the AngleTSP and the AngleDistanceTSP and all our test instances are based on points in the Euclidean plane, we generate a sequence of sets S systematically by moving a square (our magnifying glass) over the set of points. Given a fixed parameter s, we place a square of size $s \times s$ in the left upper corner of the point set of the instance. All vertices inside form the first set S. After obtaining a new tour T', we shift the square to the right by $2/3 \cdot s$ (rounded to the nearest integer) and find a new tour T''. This process is repeated until we reach the right upper corner of the instance. Then we proceed with the next row, placed $2/3 \cdot s$ (rounded to the nearest integer) below the first one. In this way, row after row, we continue to sweep over the point set until we reach the bottom right corner of the instance. Clearly, this approach guarantees that every vertex $v \in V$ is contained in at least one set $S \subset V$. To stimulate a consideration of point connections in multiple directions, we use overlapping sets S.

Concerning the choice of s we note that all our test instances are based on points in the Euclidean plane, which are uniformly distributed on $\{0, \ldots, 500\}^2$ (for further details see Section 7.1). Thus, by setting $s = 500\sqrt{\frac{k}{n}}$ (rounded to the nearest integer), we have k vertices in every square of size $s \times s$ in expectation. This improvement heuristic will be denoted $+\mathrm{M}^k$, where k is the expected number of vertices in each square.

6.2 Lens metaheuristic

A natural improvement strategy is given by the use of *metaheuristics*, which offer possibilities of escaping local optima, as they may occur e.g. in +2O and +3O. While the design options for metaheuristics are unlimited, we decided to restrict ourselves to approaches based on the *lens neighbourhood*. This neighbourhood can be evaluated quickly and still allows a substantial diversification of the current solution.

6.2.1 Lens neighbourhood

Let $T = (t_1, t_2, \ldots, t_n)$ be a tour in $G = (V, E)$. Then, for a fixed thickness parameter γ, we position a lens on all tour edges $e_t \in E(T)$ in the same way as described for the *lens procedure* in Section 2. Let t_l be a vertex in the lens, which is positioned on the tour edge $\{t_i, t_{i+1}\}$ (the wrapping around of indices at the beginning and at the end of the tour T follows in a natural way). Then we can remove the tour edges $\{t_{l-1}, t_l\}$, $\{t_l, t_{l+1}\}$ and $\{t_i, t_{i+1}\}$ from the tour T and add the edges $\{t_i, t_l\}$, $\{t_l, t_{i+1}\}$ and $\{t_{l-1}, t_{l+1}\}$ in order to obtain a new tour T'.

Example 8. *Let us consider Figure 10 depicting the graph $G = (V, E)$ and the tour $T = \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11\}$. Assuming that we are solving the AngleTSP, $z(T) \approx 13.0661$. The lens positioned on the edge $\{1, 2\}$ with $\gamma = 30°$ (≈ 0.5236) is depicted dashed in Figure 10 and contains only the vertex 8. By removing the edges $\{7, 8\}$, $\{8, 9\}$ and $\{1, 2\}$ and by adding the edges $\{1, 8\}$, $\{8, 2\}$ and $\{7, 9\}$ we obtain the new improved tour $T' = \{1, 8, 2, 3, 4, 5, 6, 7, 9, 10, 11\}$ with $z(T') \approx 12.5664$ depicted in Figure 11.*


Figure 10: graph $G = (V, E)$ with tour T and with a lens on edge $\{1, 2\}$ Figure 11: graph $G = (V, E)$ with tour T'

The *lens neighbourhood* $L(T)$ of a tour T is defined as the set of all tours T' obtained by considering all edges $e_t \in E(T)$, positioning a lens on e_t and creating a new tour T' for every vertex t_l contained in such a lens as described above. Obviously, the lens neighbourhood is a subset of the 3-*opt-neighbourhood*. Since the metaheuristics we design below are based on the evaluation of a large number of iterations, the 3-opt-neighbourhood

would be impractical, while our lens neighbourhood reduces the 3-opt-neighbourhood, allows a parametrization of this reduction and keeps "good" neighbours in the tour.

Of course, one could use the lens neighbourhood as a simple local search heuristic most probably yielding worse solutions than +3O with lower running times. However, our goal is to escape from local optima. Thus we allow a worsening of the solution by using the concept of *simulated annealing*. This metaheuristic is very well-known and has been widely used during the last decades, thus we only refer to [24] for more details.

Given a tour T in $G = (V, E)$ let $L(T) = \{T'_1, T'_2, \ldots, T'_{|L(T)|}\}$ be the lens neighbourhood of this tour. Then we define the *positive* and the *negative neighbourhood* as

$$L^+(T) := \big\{T' \in L(T)\colon z(T') < z(T)\big\} \text{ and } L^-(T) := L(T) \setminus L^+(T),$$

respectively. We tested three variants of applying the acceptance rule of simulated annealing with the lens neighbourhood:

Random choice: Choose a random tour $T' \in L(T)$ among all neighbours, accept it if $T' \in L^+(T)$ and use the acceptance rule if $T' \in L^-(T)$.

Least-cost choice: Take $T' = \operatorname{argmin}_{t' \in L(T)} \big\{z(t')\big\}$ and deal with it in the same way as in the previous case, i. e. T' is taken if $T' \in L^+(T)$ and the acceptance rule is applied if $T' \in L^-(T)$.

Combined choice: If $L^+(T) \neq \emptyset$, choose and accept a random $T' \in L^+(T)$. Otherwise take $T' = \operatorname{argmin}_{t' \in L^-(T)} \big\{z(t')\big\}$ and apply the acceptance rule.

The acceptance rule accepts the new tour $T' \in L^-(T)$ with probability

$$P = e^{\frac{\frac{z(T) - z(T')}{z(T)} n}{t(\ell)}}, \tag{33}$$

where $t(\ell)$ is a so-called temperature parameter depending on the *iteration number* $\ell \in \{0, 1, \ldots, I\}$ with I denoting the maximum number of iterations. In our implementation, one *iteration* corresponds to one choice of $T' \in L(T)$. We use a linearly decreasing temperature function $t(\ell)$ with $t(0) = 1$ and $t(I) = 0$, i. e. we set

$$t(\ell) := \frac{I - \ell}{I}. \tag{34}$$

Note that the algorithm can stop before reaching I iterations, e. g. if the lens neighbourhood is empty.

Our preliminary computational experiments have shown that among these three variants, *least-cost choice* leads to slightly better results than the *random* and the *combined choice*. Therefore we perform only *least-cost choice* in the reminder of the paper and denote it by +L.

7 Computational results

7.1 Test set-up

We describe the setup of our computational tests in this section.

7.1.1 Benchmark instances

Our benchmark instances are based on the instance specification developed in [9] and [12] and also used in [2] from where we take over the integer point sets. As in [2] we do not round the costs to integers. Furthermore, we add point sets for new sizes $n = 5$ and $85 \leq n \leq 200$ generated by the same algorithm as used by the previous papers.

We test two types of quadratic instances[1]:

AngleTSP-instances are based on points in the Euclidean plane: First, we randomly choose n points independently from the discrete uniform distribution $\{0, \ldots, 500\} \times \{0, \ldots, 500\}$. Then we compute the *turning angles* α, multiply by 1000, in the same way as [9] and [2], and then we round them to 12 decimal places.

AngleDistanceTSP-instances extend the above AngleTSP-instances by combining the turning angles with the Euclidean distances between the points in a weighted sum. Taking the same point sets in the plane we denote the Euclidean distances between vertices i and j as d_{ij}. Following the literature, we construct QTSP-distances for a weighting parameter $\rho \in \mathbb{R}_0^+$, setting $\lambda_1 = 100\rho$ and $\lambda_2 = 100$, the cost coefficients (4) reduce to:

$$c_{ijk} := 100 \left(\rho \cdot \alpha_{ijk} + \frac{d_{ij} + d_{jk}}{2} \right) \tag{35}$$

Again, all distances are rounded to 12 decimal places.

These instances were introduced in [29] for an approximate solution of the *TSP for Dubins vehicle*. They lie between AngleTSP-instances (for $\rho \to \infty$) and standard TSP instances (for $\rho = 0$). For compatibility with the literature we set $\rho = 40$ for all our tests as done in [9] and [2].

For each size $n = 5, 10, \ldots, 200$ we took/generated 10 point sets used for 10 AngleTSP- and 10 AngleDistanceTSP-instances.

7.1.2 Test environment

All tests were run on an *Intel(R) Core(TM) i7-4790 CPU @ 3.60GHz with 32 GB RAM* under *Linux*[2] and all programs, including the reference algorithm from the literature,

---

[1]All instances, the lower bound values, the optimal objective function values (whenever they are available) as well as the complete test results for all combinations of algorithms are available on `arXiv`.

[2]Precise version: *Linux 4.4.0-75-generic #96-Ubuntu SMP x86_64 x86_64 x86_64 GNU/Linux.*

were implemented in *Python*[3]. ILPs were set up using *PuLP*[4] and solved with *Gurobi*[5]. Moreover, in order to guarantee the relative reproducibility of our computational results, we (i) allowed no additional swap memory and (ii) ran all tests separately without other user processes in background.

7.1.3 Evaluation layout

As a basis for comparisons for smaller instances with $n \leq 75$ (AngleTSP-instances) resp. $n \leq 100$ (AngleDistanceTSP-instances) we computed optimal objective function values by using the *integral ILP based approach* described in [2]. For larger instances (i. e. $n \geq 80$ resp. $n \geq 105$) running times become prohibitive and we settled for lower bounds. These were computed by using the linearisation (24), (21)–(23), (25), (26) described in Section 5, relaxing the subtour elimination constraints (22) and the integrality constraints (23) and (26) and by solving the remaining LP using an LP solver. I. e. we used the same LP relaxation as in Section 5 for our LP-based approaches.

For every instance X and every heuristic H we define the **objective function value ratio**

$$r_H^X := \frac{z(H^X)}{\underline{z}(X)}, \tag{36}$$

where $z(H^X)$ denotes the objective function value of the solution given by the heuristic H for instance X and $\underline{z}(X)$ is the lower bound for instance X. Obviously, $r_H^X \geq 1$ for all H and X; $r_H^X - 1$ expresses the relative gap between the objective function value of the heuristic solution and the lower bound. We always report the *geometric mean ratio* values for the heuristic H and all instances X of the same type (AngleTSP or AngleDistanceTSP) and the same size n. Such *means of objective function value ratios* will be denoted by $\overline{r}^g$. Recall the jump implied by the change from optimal values to lower bounds as the basis of the comparison when moving from smaller to larger instances.

For the running times, we report the *arithmetic means* $\overline{t}^a$ again for every heuristic over all instances of the same type (AngleTSP or AngleDistanceTSP) and size n.

To estimate the meaning of gaps to the lower bound for larger instances, we try to give an estimation for the unkown differences between lower bounds and optimal solution values by computing these values for the smaller instances. A certain extrapolation of these values to larger instances seems justified. It turns out that (see Table 1):

1. The objective function value ratio means $\overline{r}^g$ are significantly smaller for the AngleDistanceTSP-instances.

---

[3]Precise version: *Python 3.5.2.*
[4]Precise version: *PuLP 1.6.0.*
[5]Precise version: *Gurobi 6.5.1 (Linux x86_64)* in the standard setting (i. e. the solver is allowed to use multiple threads if solving an ILP).

2. We can expect them to be over 10% for larger AngleTSP-instances and over 6.5% for larger AngleDistanceTSP.

n	5	10	15	20	25	30	35	40	45	50
AngleTSP	1.0000	1.0109	1.0296	1.0653	1.0956	1.0924	1.0958	1.1036	1.1122	1.1042
AngleDistanceTSP	1.0000	1.0000	1.0000	1.0029	1.0135	1.0211	1.0359	1.0306	1.0430	1.0514

n	55	60	65	70	75	80	85	90	95	100
AngleTSP	1.1029	1.1094	1.0866	1.0997	1.1035	–	–	–	–	–
AngleDistanceTSP	1.0500	1.0510	1.0498	1.0548	1.0537	1.0549	1.0578	1.0620	1.0632	1.0662

Table 1: optimal objective function value ratio means $\overline{r}^g$

7.2 Stand-alone heuristics

The proposed stand-alone heuristics can be divided into three groups: the *simple construction and related heuristics* (see Section 3), the *heuristics based on convex hulls* (see Section 4) and the *LP-based heuristics* (see Section 5). In the following, we first discuss the results with respect to the objective function values and later we put them into the context of running times.

Consider first the simple construction and related heuristics. This group contains three parameter-free heuristics, NN_S, NN_S^2 and CIF, and one approach, which uses LENS (see Section 2) as sub-procedure, namely NN_S^L. We tested NN_S^L with the lens thickness parameter $\gamma = 10°, 20°, \ldots, 90°$ and obtained the best results for $\gamma = 40°$ (≈ 0.6981); thus we used this value of γ for all further experiments. The objective function value ratio means for all four heuristics can be seen in Figure 12.

If we consider CIF, which is our "reference method" from [13]), we can see that it is outperformed by all other heuristics of this group for AngleTSP-instances and that it yields results similar to NN_S and NN_S^L for the AngleDistanceTSP-instances. A comparison of the nearest-neighbour-based heuristics (NN_S, NN_S^2 and NN_S^L) gives us NN_S^2 as a clear winner. Furthermore, comparing to NN_S, we can see that NN_S^L performs slightly worse for AngleTSP- and slightly better for AngleDistanceTSP-instances.

Figure 13 sums up results for CH, CH_C and CH_C^L, which were tested only for the AngleTSP-instances. Again, after making preliminary tests, we fixed the lens thickness parameter $\gamma = 40°$ (≈ 0.6981) for CH_C^L. The choice of parameter C for CH_C and CH_C^L always entails a trade-off: for larger values of C, we obviously obtain better objective function values but larger computation times. Our preliminary tests have shown that there are relatively large objective function value differences between $C = 2$, $C = 10$ and $C = 20$; the step to $C = 30$, however, only slightly improves the objective function values while the running times significantly increase. Thus we set $C = 20$ for all our further tests.

Figure 13 gives us a clear ranking: CH_C^L yields the best objective function values followed by CH_C and CH. This effect, however, seems to be less significant for larger values of n.

Figure 12: simple construction and related heuristics:

×·×·×·×	NN_S	+·+·+·+	NN_S^2	⅄·⅄·⅄·⅄	NN_S^L
⅄·⅄·⅄·⅄	CIF				

The third group of stand-alone heuristics consists of LPP, LPP^R, LPC_1^R and LPC_2^R. Again, preliminary tests gave us clear evidence for the value of the rounding parameter ρ, which we fixed to $\rho = 0.5$ in all four heuristics.

The ranking, shown in Figure 14, is the same for AngleTSP- and AngleDistanceTSP-instances: ordered from the worst to the best heuristic we get LPP, LPP^R, LPC_1^R and LPC_2^R, where the last two approaches yield very similar results. The step between LPP and LPP^R (i. e. between not using vs. using the rerun strategy) is significant, the differences between LPC_1^R and LPC_2^R are rather marginal.

Let us now consider the efficiency of the particular algorithms shown as trade-off between the running time and the solution quality. Therefore we plot for every heuristic
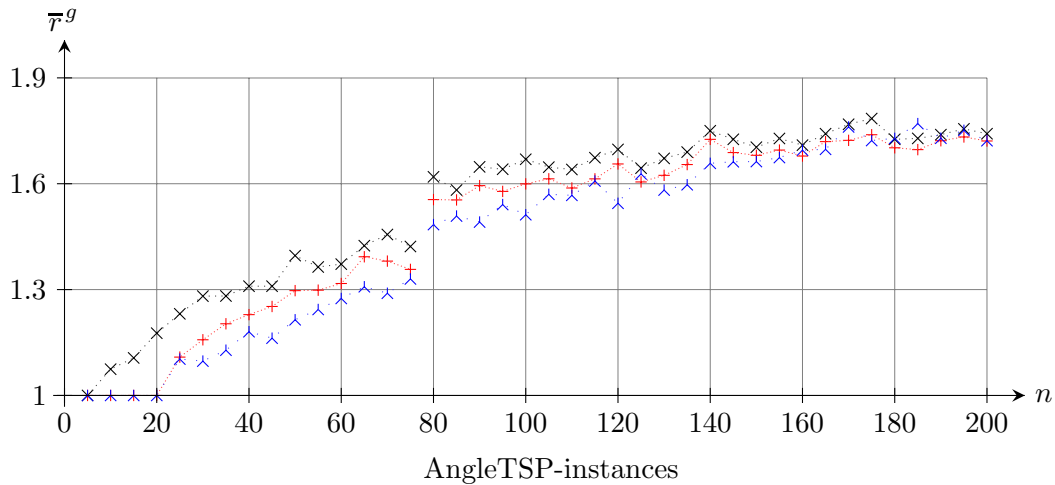
Figure 13: heuristics based on convex hulls:

×·×·×·×   CH   +·+·+·+ CH_C ⅄·⅄·⅄·⅄ CH_C^L

the objective function value ratio $\overline{r}^g$ against the running time $\overline{t}^a$ taken as means over all instances with $105 \leq n \leq 200$ in Figure 15 (the exact values can be found in Tables 2 and 3 in the Appendix). Thus, each heuristic is represented by a single point in the (running time, objective function ratio)-space. We concentrated on the larger instance sizes $105 \leq n \leq 200$, because

1. then, as mentioned, all objective function values are compared with the lower bounds (and not with the optima) and

2. larger values can be expected to give a better indication for the general behavior (note that e. g. CH_C always yields an optimum if $n \leq C$).

Consider now the trade-offs for the AngleTSP-instances. Among the "fast" heuristics (CIF, CH, CH_C and CH_C^L), CH_C^L clearly Pareto-dominates the other ones. In this context it is interesting that it is superior to CH in both the objective function values and the running times. NN_S^2 can be seen as a good trade-off between the tour quality and the time investment. If better objective function values are sought, LPP^R, LPC_1^R or LPC_2^R would be the obvious choice, LPC_2^R, however, brings only a marginal improvement in comparison to LPC_1^R. The results for the AngleDistanceTSP-instances are in many aspects similar to the results for the AngleTSP-instances, there is, however, one main difference: NN_S^2 significantly outperforms all other heuristics, although its running time advantage decreases.
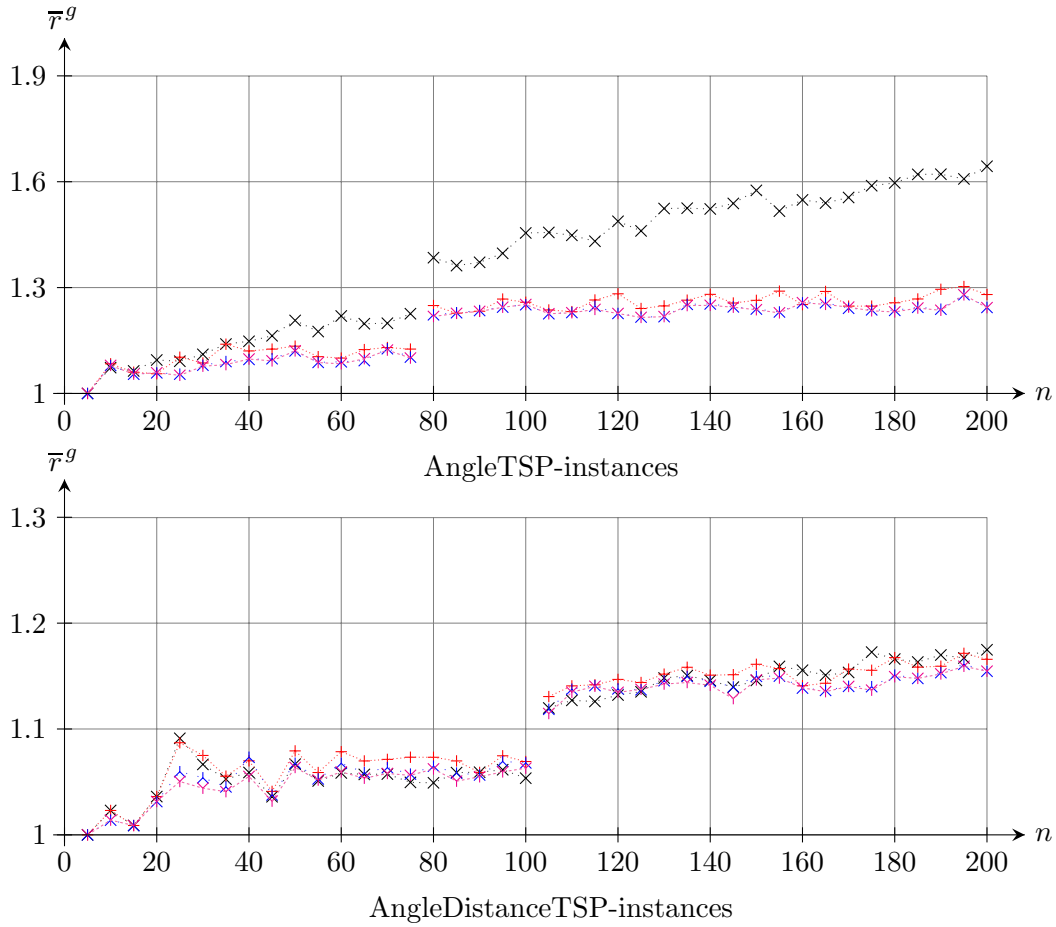
29

Figure 14: LP-based heuristics:

×–×–×–×   LPP   +–+–+–+   LPP^R ⅄–⅄–⅄–⅄   LPC$_1^R$

Y–Y–Y–Y   LPC$_2^R$

7.3 Improvement heuristics

Concerning improvement heuristics, it would not make sense to present results for all possible combinations of stand-alone and improvement heuristics in detail. Instead, we discuss the results for the best stand-alone heuristic combined with various improvement approaches; in particular, we took LPC$_2^R$ for the AngleTSP-instances and NN$_S^2$ for the AngleDistanceTSP-instances as the stand-alone heuristic (results for all other combinations are summed up in Tables 2 and 3 in the Appendix).

Again, we must fix some parameters first. Similar to CH$_C$ and CH$_C^L$ we have a trade-off in +M: k expresses the expected number of vertices in our magnifying glass, thus larger values of k lead to better objective function values, but worse running times. After
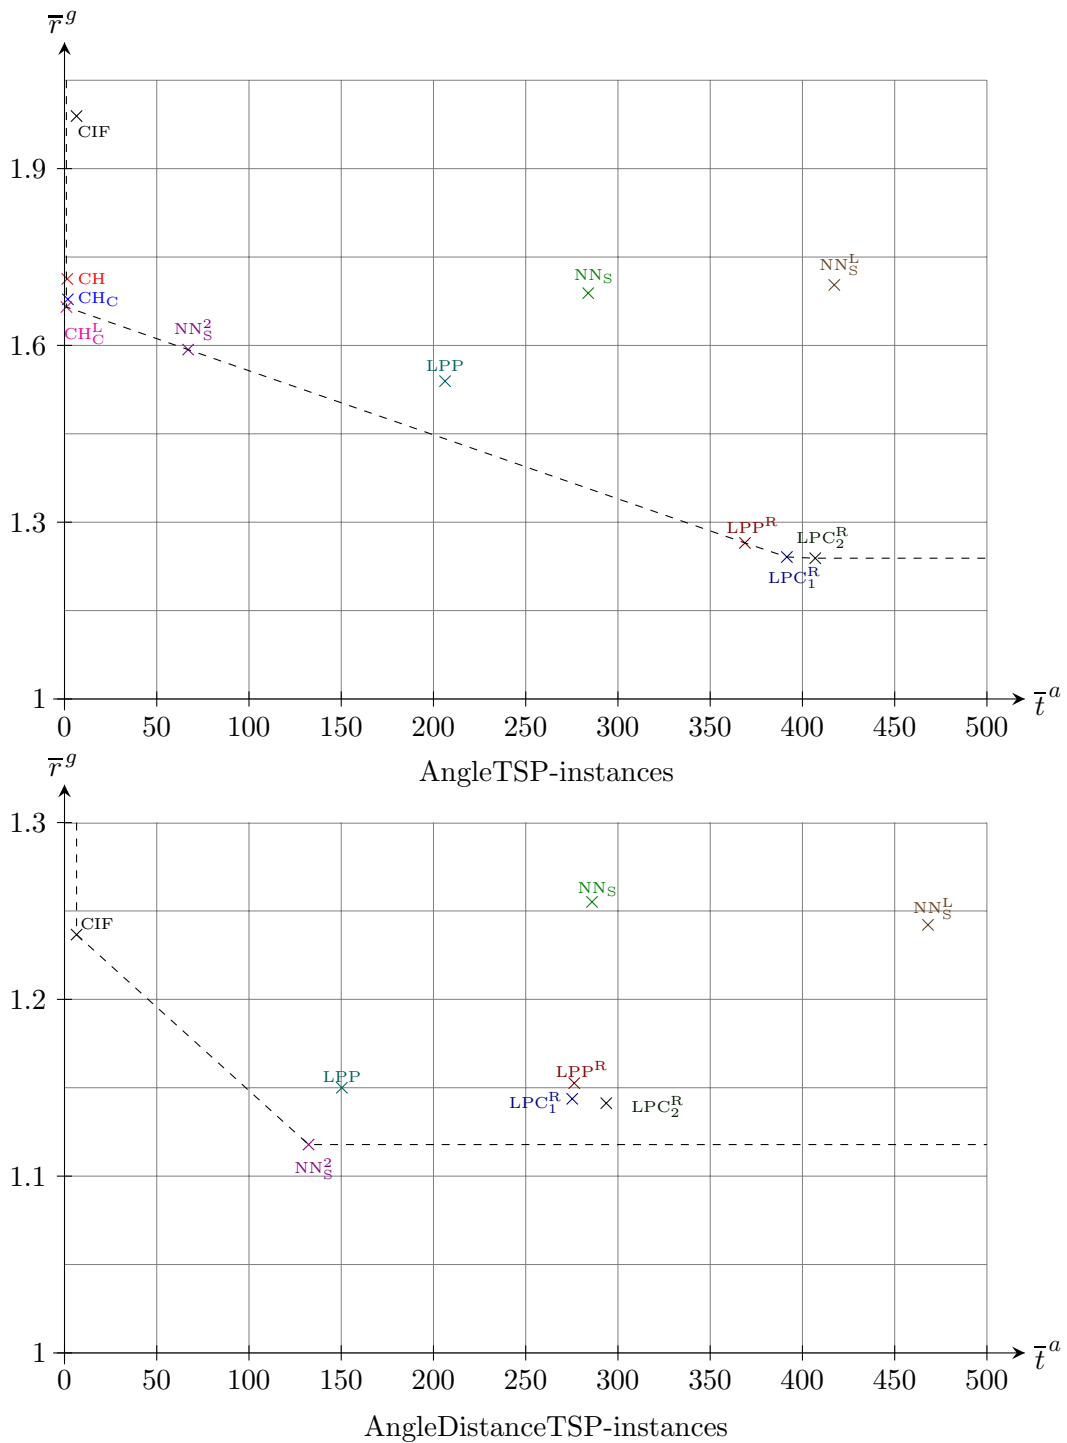
Figure 15: stand-alone heuristic trade-offs for instances with $105 \le n \le 200$; Pareto-frontier is plotted dashed

31

various tests we decided to use two variants, namely $+\mathrm{M}^{15}$ and $+\mathrm{M}^{20}$; the first one is roughly half-way between $+2\mathrm{O}$ and $+3\mathrm{O}$ with respect to the computation times and the latter one yields the overall best results with respect to the objective function values. In our metaheuristic $+\mathrm{L}$, we have just two parameters: after preliminary tests we could set the lens thickness to $\gamma = 20°$ (≈ 0.3491) and the maximum number of iterations to $I = 2000$.

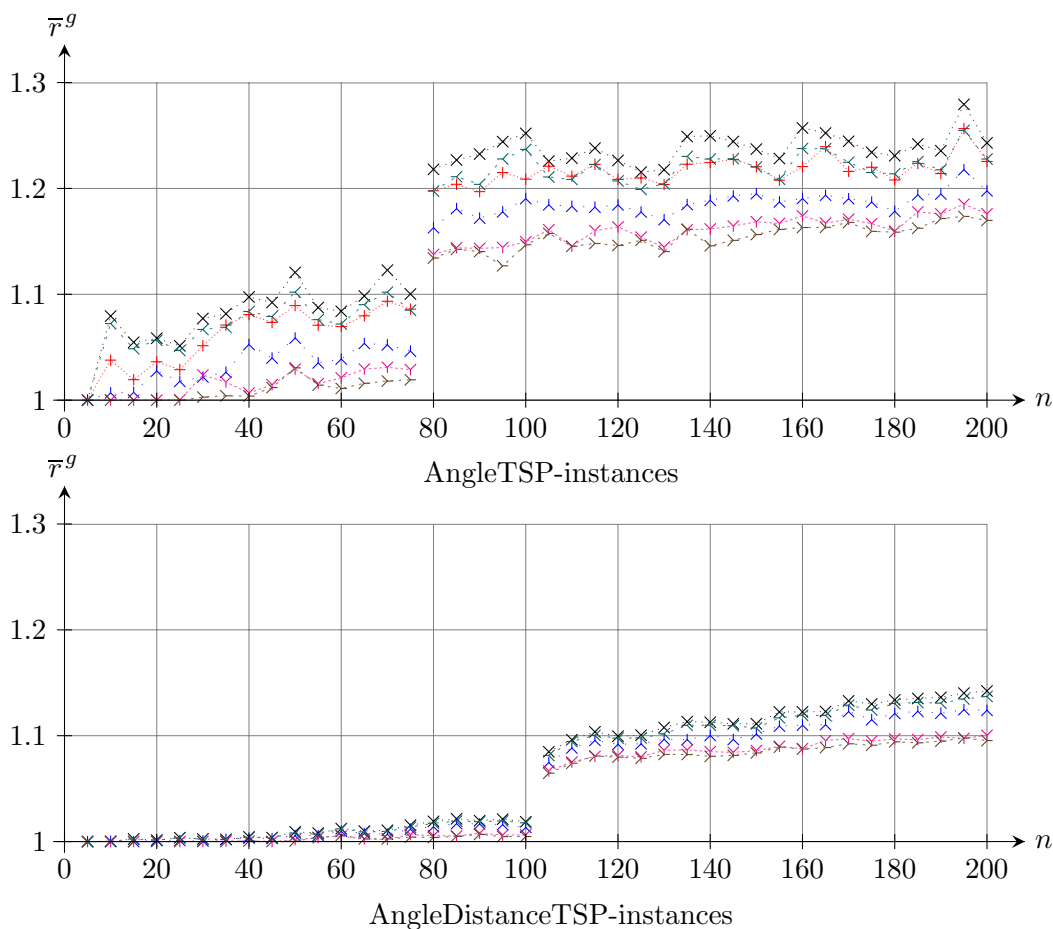Figure 16: improvement heuristics:

×·×·×·×	$\mathrm{LPC}_2^{\mathrm{R}}$	+·+·+·+	$\mathrm{LPC}_2^{\mathrm{R}}+2\mathrm{O}$	⅄·⅄·⅄·⅄	$\mathrm{LPC}_2^{\mathrm{R}}+3\mathrm{O}$
⅄·⅄·⅄·⅄	$\mathrm{LPC}_2^{\mathrm{R}}+\mathrm{M}^{15}$	⋗·⋗·⋗·⋗	$\mathrm{LPC}_2^{\mathrm{R}}+\mathrm{M}^{20}$	⋖·⋖·⋖·⋖	$\mathrm{LPC}_2^{\mathrm{R}}+\mathrm{L}$

Next, consider Figure 16. We can see that the ranking is the same for the AngleTSP- and AngleDistanceTSP-instances: ordered from the worst to the best improvement heuristic we get $+\mathrm{L}$, $+2\mathrm{O}$, $+3\mathrm{O}$, $+\mathrm{M}^{15}$ and $+\mathrm{M}^{20}$, where $+\mathrm{L}$ and $+2\mathrm{O}$ yield similar objective function values (note that $+2\mathrm{O}$ makes no sense for $\mathrm{NN}_{\mathrm{S}}^2$; see Section 3.1 for more details). By considering concrete relations between the particular improvement heuristics, $+3\mathrm{O}$

performs better for AngleTSP-instances and can be placed between +2O and +M^{15} as well as +M^{20} for this group, for AngleDistanceTSP-instances, however, +M^{15} and +M^{20} clearly beat all other approaches.

Figure 17 sums up the trade-offs of the particular improvement heuristics. Obviously, +2O performs much better than +L, because it reaches similar objective function values faster. Very interesting is the performance of +M^{15} and +M^{20}: +M^{15} beats +3O both in solution qualities and running times. The step from +M^{15} to +M^{20} costs a significant amount of computation time while the objective function value improvement is rather marginal. Summing up, +2O and +M^{15} offer the best quality/cost ratio.

Finally, by considering Table 2 in the Appendix, we can see that the overall best mean objective function value ratio $\overline{r}^g$ could be obtained by the combinations LPC$_1^R$+M^{20} for the AngleTSP- and LPC$_2^R$+M^{20} for the AngleDistanceTSP-instances, i. e. not by improving the outcome of the best stand-alone heuristics respectively. Of course, we cannot estimate how far the solutions are above the optimum in these cases. However, Table 1, allows a rough intuition that the difference between the lower bound gap and the optimal objective function gap is around 10 percentage points for the larger AngleTSP- and 6.5 percentage points for the larger AngleDistanceTSP-instances. Thus the best obtained objective function value ratios $\overline{r}^g$ of 15.73% and 8.26% can be roughly translated into 5% and 1.5% gaps to the unknown optimal solution values for the AngleTSP- and AngleDistanceTSP-instances, respectively.

8 Conclusions

Quadratic variants of classical combinatorial optimization problems have received increasing attention in recent years. They permit the inclusion of interdependency effects in the objective function while the standard problems only add up the costs of elements in a feasible solution subset selected from a given ground set. In this paper we consider two quadratic variants of the famous TSP corresponding to the following geometric cases: Given points in the Euclidean plane, which should be connected by a Hamiltonian cycle, the cost of such a tour is either the sum of *turning angles* in every point (AngleTSP) or a linear combination of turning angles and Euclidean distances (AngleDistanceTSP). Both problems have applications in robotics and other planning problems for moving objects of high inertia.

One main contribution of the paper is a number of constructive heuristics and tour improvement algorithms, which are based on geometric properties of good solutions, particularly encouraging the exploitation of the convex hull for the given point set and subsets thereof. A second contribution utilizes the impressive performance of todays ILP solvers. We use ILP models to patch together paths and subtours arising either from constructive procedures or from relaxations of more general ILP models. Finally, we also solve ILP models as an improvement step by iteratively selecting a subregion of the point set and removing all solution edges from this area. Then the resulting

Figure 17: improvement heuristic trade-offs for instances with $105 \leq n \leq 200$; Pareto-frontier is plotted dashed

subproblem is solved to optimality while keeping fixed all tour connections outside this area. This matheuristic can be thought of as a magnifying glass which is successively placed on certain parts of the Euclidean plane and rewires the tour in this area in an optimal way. The described principle has similarities to other well-known meta- and matheuristic ideas, such as large neighbourhood search.

Extensive computational experiments on test instances previously used in the literature, but also on larger sized instances, were performed and reported in a condensed way. We managed to take a major step forward compared to previously reported heuristics. The obtained gaps to the lower bound of roughly 15% for AngleTSP (8% for AngleDistanceTSP) for instances up to $n = 200$ points could be translated in a cursory way to an expected gap of 5% (1.5%) to the unknown optimal solution values. Note for comparison that the previously best performing heuristic reported in the literature (corresponding to CIF+3O) yields average gaps to the lower bound of 31% (12%) and requires the running times for the 3-opt heuristic, which rapidly increases for larger instances. Running times to reach our best solution values obtained on a standard PC average from 5 to 10 minutes, but with a moderate growth rate for even larger instances.

For future research the still existing optimality gaps of our approaches suggest opportunities for further progress. In particular, there seems to be plenty of potential for developing sophisticated metaheuristics. Moreover, our approaches are motivated by the geometry of the point set and the implied cost coefficients in the objective function. Moving away from this structure to more general cost values, possibly still preserving some weaker geometric relationship, could give rise to challenging new variants of the QTSP. In this context, note that our LP-based approaches could be used for the general QTSP as well. Finally, the underlying idea of our magnifying glass matheuristic, which performed very well in our QTSP setting, can be rather easily applied to other combinatorial optimization problems with a spatial structure of the given objects and decomposable solution sets.

Acknowledgements

References

[1] A. Aggarwal, D. Coppersmith, S. Khanna, R. Motwani, and B. Schieber, "The angular-metric traveling salesman problem," *SIAM Journal on Computing*, vol. 29, no. 3, pp. 697–711, 2000.

[2] O. Aichholzer, A. Fischer, F. Fischer, F. J. Meier, U. Pferschy, A. Pilz, and R. Staněk, "Minimization and maximization versions of the quadratic travelling salesman problem," *Optimization*, vol. 66, no. 4, pp. 521–546, 2017.

[3] E. Amaldi, G. Galbiati, and F. Maffioli, "On minimum reload cost paths, tours, and flows," *Networks*, vol. 57, no. 3, pp. 254–260, 2011.

[4] D. L. Applegate, R. E. Bixby, V. Chvátal, and W. J. Cook, *The Traveling Salesman Problem: A Computational Study*. Princeton University Press, 2006.

[5] A. Assad and W. Xu, "The quadratic minimium spanning tree problem," *Naval Research Logistics*, vol. 39, pp. 399–417, 1992.

[6] A. Ćustić, R. Zhang, and A. P. Punnen, "The quadratic minimum spanning tree problem and its variations," *Discrete Optimization*, vol. 27, no. 4, pp. 73–87, 2018.

[7] G. Dantzig, R. Fulkerson, and S. Johnson, "Solution of a large-scale traveling-salesman problem," *Journal of the Operations Research Society of America*, vol. 2, no. 4, pp. 393–410, 1954.

[8] C. Duin and S. Voß, "The pilot method: a strategy for heuristic repetition with application to the Steiner problem in graphs," *Networks*, vol. 34, no. 3, pp. 181–191, 1999.

[9] A. Fischer, "A polyhedral study of quadratic traveling salesman problems," Ph.D. dissertation, Chemnitz University of Technology, Department of Mathematics, 2013, available at: http://www.qucosa.de/fileadmin/data/qucosa/documents/11834/Dissertation_Anja_Fischer.pdf.

[10] ——, "An analysis of the asymmetric quadratic traveling salesman polytope," *SIAM Journal on Discrete Mathematics*, vol. 28, no. 1, pp. 240–276, 2014.

[11] A. Fischer and F. Fischer, "An extended approach for lifting clique tree inequalities," *Journal of Combinatorial Optimization*, vol. 30, no. 3, pp. 489–519, 2015.

[12] A. Fischer and C. Helmberg, "The symmetric quadratic traveling salesman problem," *Mathematical Programming A*, vol. 142, no. 1-2, pp. 205–254, 2013.

[13] A. Fischer, F. Fischer, G. Jäger, J. Keilwagen, P. Molitor, and I. Grosse, "Exact algorithms and heuristics for the quadratic traveling salesman problem with an application in bioinformatics," *Discrete Applied Mathematics*, vol. 166, pp. 97–114, 2014.

[14] ——, "Computational recognition of RNA splice sites by exact algorithms for the quadratic traveling salesman problem," *Computation*, vol. 3, no. 2, pp. 285–298, 2015.

[15] T. F. Gonzalez, *Handbook of Approximation Algorithms and Metaheuristics*. Chapman & Hall/CRC, 2007.

[16] G. Gutin and A. P. Punnen, Eds., *The Traveling Salesman Problem and Its Variations*. Springer US, 2007.

[17] H. Hu and R. Sotirov, "Special cases of the quadratic shortest path problem," *Journal of Combinatorial Optimization*, 2017, to appear, doi.org/10.1007/s10878-017-0219-9.

[18] G. Jäger and P. Molitor, "Algorithms and experimental study for the traveling salesman problem of second order," in *Proceedings of COCOA*, ser. Lecture Notes in Computer Science, vol. 5165. Springer, 2008, pp. 211–224.

[19] E. L. Lawler, J. K. Lenstra, A. H. G. Rinnooy Kan, and D. B. Shmoys, Eds., *The Traveling Salesman Problem: A Guided Tour of Combinatorial Optimization*. Wiley-Interscience Series in Discrete Mathematics and Optimization, 1985.

[20] A. Medeiros and S. Urrutia, "Discrete optimization methods to determine trajectories for Dubins' vehicles," *Electronic Notes in Discrete Mathematics*, vol. 36, pp. 17–24, 2010.

[21] U. Pferschy and R. Staněk, "Generating subtour elimination constraints for the TSP from pure integer solutions," *Central European Journal of Operations Research*, vol. 25, no. 1, pp. 231–260, 2017.

[22] D. Pisinger and S. Ropke, "Large neighborhood search," in *Handbook of Metaheuristics*, M. Gendreau and J.-Y. Potvin, Eds. Springer, 2010, pp. 399–419.

[23] P. Raghavan and C. D. Tompson, "Randomized rounding: A technique for provably good algorithms and algorithmic proofs," *Combinatorica*, vol. 7, no. 4, pp. 365–374, 1987.

[24] C. R. Reeves, *Modern heuristic techniques for combinatorial problems*. John Wiley & Sons, Inc., 1993.

[25] G. Reinelt, *The Traveling Salesman: Computational Solutions for TSP Applications*. Springer, 1994.

[26] B. Rostami and F. Malucelli, "Lower bounds for the quadratic minimum spanning tree problem based on reduced cost computation," *Computers and Operations Researc*, vol. 64, pp. 178–188, 2015.

[27] B. Rostami, F. Malucelli, D. Frey, and C. Buchheim, "On the quadratic shortest path problem," in *Experimental Algorithms*, ser. Lecture Notes in Computer Science, vol. 9125.   Springer, 2015, pp. 379–390.

[28] B. Rostami, F. Malucelli, P. Belotti, and S. Gualandi, "Lower bounding procedure for the asymmetric quadratic traveling salesman problem," *European Journal of Operational Research*, vol. 253, pp. 584–592, 2016.

[29] K. Savla, E. Frazzoli, and F. Bullo, "Traveling salesperson problems for the Dubins vehicle," *IEEE Transactions on Automatic Control*, vol. 53, no. 6, pp. 1378–1391, 2008.

[30] D. P. Williamson and D. B. Shmoys, *The Design of Approximation Algorithms*. Cambridge University Press, 2011.

[31] B. Woods, A. Punnen, and T. Stephen, "A linear time algorithm for the 3-neighbour traveling salesman problem on Halin graphs and extensions," 2015, available as: `arXiv:1504.02151`.

Appendix

In the following two tables the best value in every row is marked by an asterisk and the best value in every column is highlighted by underlying.

	no improv.	+2O	+3O	$+M^{15}$	$+M^{20}$	+L	min
NN_S	1.6887	1.6560	1.5466	1.3637	1.3026*	1.6811	1.3026*
NN_S^2	1.5928	1.5928	1.5254	1.3446	1.2867*	1.5824	1.2867*
NN_S^L	1.7026	1.4266	1.2726	1.2077	1.1886*	1.6406	1.1886*
CIF	1.9892	1.4630	1.3121	1.2697	1.2336*	1.8689	1.2336*
CH	1.7128	1.6999	1.5798	1.3645	1.3069*	1.7057	1.3069*
CH_C	1.6783	1.6741	1.6026	1.3736	1.3098*	1.6748	1.3098*
CH_C^L	1.6652	1.4703	1.2757	1.2096	1.1882*	1.6083	1.1882*
LPP	1.5393	1.3339	1.2376	1.2118	1.1887*	1.4771	1.1887*
LPP^R	1.2648	1.2313	1.1975	1.1718	1.1615*	1.2426	1.1615*
LPC_1^R	1.2410	1.2210	<u>1.1884</u>	1.1655	<u>1.1573*</u>	1.2231	<u>1.1573*</u>
LPC_2^R	<u>1.2389</u>	<u>1.2202</u>	1.1888	<u>1.1652</u>	1.1577*	<u>1.2210</u>	1.1577*
min	<u>1.2389</u>	<u>1.2202</u>	<u>1.1884</u>	<u>1.1652</u>	<u>1.1573*</u>	<u>1.2210</u>	<u>1.1573*</u>

AngleTSP-instances

NN_S	1.2551	1.1570	1.1254	1.0993	1.0934*	1.2387	1.0934*
NN_S^2	<u>1.1179</u>	1.1179	1.1057	1.0883	1.0856*	<u>1.1140</u>	1.0856*
NN_S^L	1.2422	1.1484	1.1201	1.0975	1.0920*	1.2244	1.0920*
CIF	1.2366	1.1504	1.1209	1.1035	1.0953*	1.2303	1.0953*
LPP	1.1500	1.1134	1.0990	1.0880	1.0857*	1.1432	1.0857*
LPP^R	1.1526	1.1077	1.0944	1.0860	1.0836*	1.1396	1.0836*
LPC_1^R	1.1437	1.1064	1.0936	1.0858	1.0828*	1.1333	1.0828*
LPC_2^R	1.1413	<u>1.1058</u>	<u>1.0931</u>	<u>1.0854</u>	1.0826*	1.1310	1.0826*
min	<u>1.1179</u>	<u>1.1058</u>	<u>1.0931</u>	<u>1.0854</u>	<u>1.0826*</u>	<u>1.1140</u>	1.0826*

AngleDistanceTSP-instances

Table 2: objective function value ratio means for all test instances with $105 \leq n \leq 200$

	no improv.	+2O	+3O	$+M^{15}$	$+M^{20}$	+L	min
NN_S	283.8885*	284.1263	413.1111	465.3341	569.1563	356.7633	283.8885*
NN_S^2	67.0159*	67.1282	193.0359	240.8241	345.4880	139.7709	67.0159*
NN_S^L	417.2671*	417.8017	571.5596	517.4951	605.8565	487.5652	417.2671*
CIF	6.5365*	7.2117	168.7839	110.9418	<u>172.3227</u>	77.9529	6.5365*
CH	1.5028*	1.7001	137.9146	179.3441	298.8040	74.5875	1.5028*
CH_C	1.8191*	1.9635	<u>111.6876</u>	187.8074	303.5915	74.9145	1.8191*
CH_C^L	<u>1.0236*</u>	<u>1.4739</u>	165.6817	<u>108.0078</u>	201.2742	<u>71.9514</u>	<u>1.0236*</u>
LPP	206.2802*	206.7912	340.4795	295.4575	380.8795	277.6503	206.2802*
LPP^R	368.8422*	369.1380	483.0896	455.4174	532.4050	440.6636	368.8422*
LPC_1^R	391.5621*	391.8200	501.6059	475.4146	548.8180	463.3042	391.5621*
LPC_2^R	406.9456*	407.1874	518.1517	491.3909	560.5582	478.6038	406.9456*
min	<u>1.0236*</u>	<u>1.4739</u>	<u>111.6876</u>	<u>108.0078</u>	<u>172.3227</u>	<u>71.9514</u>	<u>1.0236*</u>

AngleTSP-instances

NN_S	285.9435*	286.5523	412.0495	332.5851	367.3309	357.0333	285.9435*
NN_S^2	132.3110*	132.4236	230.7380	168.1561	195.0742	202.4256	132.3110*
NN_S^L	468.0509*	468.6615	592.9431	513.0854	548.2855	539.1046	468.0509*
CIF	<u>6.5468*</u>	<u>7.0853</u>	<u>127.9475</u>	<u>41.8982</u>	<u>71.4263</u>	<u>73.8265</u>	<u>6.5468*</u>
LPP	150.3579*	150.7859	252.0784	183.6876	211.1502	220.0040	150.3579*
LPP^R	276.3159*	276.7537	376.4522	310.7960	338.4974	347.1968	276.3159*
LPC_1^R	275.2136*	275.6351	375.8346	309.4856	335.4789	345.7596	275.2136*
LPC_2^R	293.6624*	294.0777	392.0432	327.9660	354.4542	364.1617	293.6624*
min	<u>6.5468*</u>	<u>7.0853</u>	<u>127.9475</u>	<u>41.8982</u>	<u>71.4263</u>	<u>73.8265</u>	<u>6.5468*</u>

AngleDistanceTSP-instances

Table 3: running time means for all test instances with $105 \leq n \leq 200$