



Project Number 288094

eCOMPASS

eCO-friendly urban Multi-modal route PIAnning Services for mobile uSers

STREP

Funded by EC, INFOS-G4(ICT for Transport) under FP7

eCOMPASS – TR – 046

Rectilinear Shortest Path and Rectilinear Minimum Spanning Tree with Neighborhoods

Yann Disser, Matus Mihalak, and Sandro Montanari

January 2014

Rectilinear Shortest Path and Rectilinear Minimum Spanning Tree with Neighborhoods

Yann Disser¹, Matús Mihalák², and Sandro Montanari²

¹ Institut für Mathematik, TU Berlin, Germany disser@math.tu-berlin.de

² Institute of Theoretical Computer Science, ETH Zurich, Switzerland
{matus.mihalak,sandro.montanari}@inf.ethz.ch

Abstract. We study the geometric shortest path and the minimum spanning tree problem with neighborhoods in the L_1 metric. In this setting, we are given a graph $\mathcal{G} = (\mathcal{R}, E)$, where $\mathcal{R} = \{R_1, \dots, R_n\}$ is a set of polygonal regions in the plane. Placing a point p_i inside each region R_i turns G into an edge-weighted graph $G_{\mathbf{p}}$, $\mathbf{p} = \{p_1, \dots, p_n\}$, where the cost of an edge is the distance between the points. The *Shortest Path Problem with Neighborhoods* asks, for given R_s and R_t , to find a placement \mathbf{p} such that the resulting shortest s - t path in $G_{\mathbf{p}}$ is smallest among all graphs $G_{\mathbf{p}}$. The *Minimum Spanning Tree Problem with Neighborhoods* asks for a placement \mathbf{p} such that the resulting minimum spanning tree of $G_{\mathbf{p}}$ has the smallest cost among all minimum spanning trees of all graphs $G_{\mathbf{p}}$. We study these problems in the L_1 metric, and show that the shortest path problem with neighborhoods is solvable in polynomial time, whereas the minimum spanning tree problem with neighborhoods is NP-hard, even if the neighborhood regions are segments.

1 Introduction

Given a sequence of n polygons in the plane, a start point s , and a target point t , TOURINGPOLYGONS is the problem to find a shortest path that starts at s , visits the n polygons in the given order, and ends at t . This problem is solvable in polynomial time, whenever the polygons are convex and disjoint [3]. In general, if the polygons are allowed to be non-convex and intersecting, the problem is NP-hard [3] (in the L_2 metric). For several years, the complexity of the problem for the case of non-convex, yet disjoint polygons has been open (even for the L_1 metric), which motivated the design of approximation algorithms [11] for the problem. Recently, Ahadi et al. [1] proved that TOURINGPOLYGONS is NP-hard even if the polygons are disjoint, for every L_p norm, $p \geq 1$, in the case where every polygon is formed by two joint line segments whose angles with the x -axis are in $\{0, \pm\pi/4, \pi/2\}$.

Apparently, the difficulty of TOURINGPOLYGONS lies in the shape of the polygons $P_i, i = 1, \dots, n$. It is a natural question to ask “How does the complexity of the problem changes when the sequence of the polygons is not given up front?” If we still require that all polygons are visited, the problem becomes to find a shortest path from s to t visiting all polygons (in an arbitrary order). The NP-hardness of this problem follows easily, as the related problem of finding a shortest tour starting at s , visiting all polygons, and

coming back to s , is NP-hard [12] (even if the polygons are points). Here, the difficulty obviously lies in finding the right order of the polygons to be visited.

In this paper we study a related question which relaxes the requirements “visit all polygons” and “follow the given order” in a natural way: Additionally to the set of n polygons P_1, \dots, P_n , we are given a set $E \subseteq 2^{\{P_1, \dots, P_n\}}$ of *allowed traversals*, where a pair $\{P_i, P_j\} \in E$ states that a tour may traverse from P_i to P_j ; Given two polygons P_s and P_t , we search for a shortest tour through the polygons such that for any two consecutively visited polygons P and P' along this tour we have $\{P, P'\} \in E$. We call this the *Shortest Path Problem with Neighborhoods*, SPN for short. An alternative view of it is the following: Given n polygons P_1, \dots, P_n and a graph $\mathcal{G} = (\{P_1, \dots, P_n\}, E)$, place a point inside each polygon P_i such that the shortest path in the resulting geometric graph is smallest possible (among all such placements). Obviously, in general, the problem remains NP-hard – TOURINGPOLYGONS is a special case of it (just set G to be the path induced by the order in which the polygons need to be visited). We are interested in polynomially solvable special cases of the problem.

Several other combinatorial optimization problems have been studied in this “with neighborhoods” spirit, such as the *Travelling Salesman Problem* [4, 5] the problems of finding (*minimum-length*) *convex hull* [9] or *enclosing circle* [10], or the *Minimum Spanning Tree Problem* [9, 13, 2]. In all these studies, there is a set of *neighborhood regions* R_1, \dots, R_n (not necessarily polygons), and the graph $\mathcal{G} = (\{R_1, \dots, R_n\}, E)$ is a complete graph.

In all these “with neighborhood” variants, one searches for a placement of points inside the region R_i such that the resulting geometric variant of the combinatorial optimization problem at hand has as cheap an optimal solution as possible. For example, the *Minimum Spanning Tree Problem with Neighborhoods* (MSTN for short) asks for a placement of n points, one in each region $R_i, i = 1, \dots, n$, such that the (euclidean) minimum spanning tree spanning these points is smallest among all such placements. Löffler and Kreveld showed [9] that this problem is NP-hard when the neighborhood regions are squares (not necessarily disjoint). Yang et al. showed that when the neighborhood regions are disjoint unit disks, the problem admits a PTAS [13] (i.e., it can be approximated arbitrary well). Dorrigiv et al. [2] later showed that MSTN in L_2 metric is APX-hard when the neighborhood regions are disjoint disks, not necessarily of unit diameter.

Our results. In Section 2, we show that SPN in L_1 metric can be solved in polynomial time if the neighborhood regions are axis-parallel polygons (convex or non-convex). In Section 3 we adapt the hardness result of MSTN in L_2 metric of Dorrigiv et al. [2] and show that MSTN in L_1 metric is APX-hard, even if the regions are line segments.

2 Shortest Path with Neighborhoods

In the Shortest Path Problem with Neighborhoods, or SPN, we are given a set of non-overlapping axis-parallel polygonal regions $\mathcal{R} = \{R_1, \dots, R_n\}$, a directed underlying graph $\mathcal{G} = (\mathcal{R}, E)$, and a pair $s, t \in \{1, \dots, n\}, s \neq t$. The problem asks for a placement p such that the cost of a shortest path between R_s and R_t in \mathcal{G}_p is smallest among all possible placements. We call such a placement an *optimal SPN placement*.

The SPN problem can be solved trivially if there is an edge in \mathcal{G} from R_s to R_t . In this case, a shortest st -path is the edge (R_s, R_t) , and an optimum placement minimizes the length of this edge. If $(R_s, R_t) \notin E$, the problem becomes more interesting, because it is not clear a priori what sequence of rectangles constitutes a shortest st -path, for an optimum placement p . Note that, if we know which regions constitute a shortest st -path, and in which order, the problem becomes the Touring Polygons Problem. Recall that this problem can be solved efficiently in case of convex polygonal regions, and it is NP-hard for any L_p metric (with $p \geq 1$) in case the regions are composed of at most two joint line segments whose angles with the x -axis are in $\{0, \pm\pi/4, \pi/2\}$. In the following we show that the above hardness of TOURINGPOLYGONS is due to the $\pm\pi/4$ segments, and that in fact for every input of rectilinear and even non-convex polygonal regions R_i , the more general SPN problem is solvable in polynomial time.

Given a finite set of points $P \subset \mathbb{R}^2$, the *Hanan grid* of P is induced by the vertical and horizontal lines passing through each point of P . The Hanan grid is a well-studied object, and it has been shown [6, 14] that the Hanan grid of a point set P contains a rectilinear minimum Steiner tree of P .

In the following, we show that to find an optimum SPN placement it is sufficient to consider only points of the Hanan grid induced by the corners of the regions in \mathcal{R} . Based on this, we provide an algorithm that computes an optimum SPN placement in time $\mathcal{O}(mn^2k^4)$, where $n = |\mathcal{R}|$, $m = |E|$, and k is the maximum number of corners of a region of \mathcal{R} . This implies that also TOURINGPOLYGONS with L_1 metric can be solved in polynomial time in case the regions are non-convex polygons whose segments form angles with the x -axis that are in $\{0, \pi/2\}$. To the best of our knowledge, this is the only known variant of TOURINGPOLYGONS that can be solved efficiently for non-convex regions.

2.1 Properties of optimum solutions

If at least one of the regions in \mathcal{R} has a non empty area, there exist infinitely many possible placements, and there may also exist infinitely many optimum SPN placements. It is crucial for an efficient algorithm to consider only a finite (polynomial) subset of points when looking for an optimum SPN solution.

Since distances between points are measured in L_1 metric and the neighborhood regions are axis-parallel, one of the most trivial approaches is to consider as possible placement points only the corners of regions in \mathcal{R} . It is however easy to construct instances where every optimum placement contains at least one point that is not a corner of a region in \mathcal{R} . We do not have to consider many more points other than the corners of the regions, though. Lemma 1 (below) shows that there always exists an optimum SPN placement where all points are points of the Hanan grid induced by the corners of the regions in \mathcal{R} , lying on the perimeters of those regions. To prove this fact, we use a particular property of the L_1 metric defined in terms of bounding boxes of points in \mathbb{R}^2 . Given $x, y \in \mathbb{R}^2$, the *bounding box* B_{xy} is the smallest axis-parallel rectangle containing x and y .

Proposition 1. For every $x, y, z \in \mathbb{R}^2$,

$$z \in B_{xy} \iff \|xy\| = \|xz\| + \|zy\|$$

$$z \notin B_{xy} \iff \|xy\| < \|xz\| + \|zy\|.$$

Lemma 1. *There exists an optimum placement \mathbf{p} such that every $p_i \in \mathbf{p}$ lies on the perimeter of R_i and is a grid point of the Hanan grid induced by the corners of the regions in \mathcal{R} .*

Proof. Let \mathbf{p} be an optimum placement and P a shortest st -path in $\mathcal{G}_{\mathbf{p}}$. We show how to move points in \mathbf{p} not satisfying the lemma to points of the Hanan grid on the perimeter of the regions in a way such that the resulting placement is still optimum. We distinguish between regions on (visited by) P and not on P .

A point in \mathbf{p} of a region not on P not satisfying the lemma can be trivially moved to an arbitrary corner of that region. Since the cost of P in the resulting placement is the same as in $\mathcal{G}_{\mathbf{p}}$, the resulting placement is still optimum.

We first show how to move points of regions on P not satisfying the lemma to the perimeter in a way such that the resulting placement is still optimum. Then, we show that every remaining point still not satisfying the lemma can be moved to a Hanan grid point on the perimeter of its region.

Note first that p_s of R_s must lie on its perimeter, otherwise we could obtain a better placement by moving it to a point on the perimeter of R_s closest to the point in the successor of R_s on P . The same argument holds by symmetry for p_t .

Let $R_j \notin \{R_s, R_t\}$ be a region on P , and consider $p_i, p_k \in \mathbf{p}$, where R_i is the predecessor of R_j on P and R_k is its successor. Consider the bounding box $B_{p_i p_j}$, and let p_c be a point on the perimeter of R_j contained in B_{ij} . By Proposition 1 and triangle inequality, we have

$$\|p_i p_j\| + \|p_j p_k\| = \|p_i p_c\| + \|p_c p_j\| + \|p_j p_k\| \geq \|p_i p_c\| + \|p_c p_k\|.$$

Thus, moving p_j to p_c does not increase the cost of P . The resulting placement is still optimum, and p_j now lies on the perimeter of R_j . We can apply this operation to every point in the interior of its region.

We now show how to move points in \mathbf{p} to Hanan grid points on the perimeters of their regions in such a way that the resulting placement is still optimum. By the above, we can assume each point of \mathbf{p} to be lying on the perimeter its region, and that only points of regions on P may not be grid points.

Let $p_j = (x_j, y_j) \in \mathbf{p}$ be a point on the perimeter of R_j not on the Hanan grid. Since R_j is axis-parallel, p_j lies on a line of the grid. Thus, either x_j is the x -coordinate of a grid point, or y_j is the y -coordinate of a grid point. We consider only the latter case; the former is symmetric.

Let x_l be the largest x -coordinate of a grid point lying to the left of p_j , and x_r be the smallest x -coordinate of a grid point lying to the right of p_j . We define the set $\{(x, y) \in \mathbb{R}^2 \mid x_l < x < x_r\}$ as the *vertical stripe* of p_j .

Consider a sequence R_i, \dots, R_k of consecutive regions on P of maximal length such that R_j is in the sequence, and every point in \mathbf{p} of a region in the sequence lies in the vertical stripe of p_j . None of the points in the sequence is a grid point. We first consider the case where $R_i \neq R_s$ and $R_k \neq R_t$.

Let $R_{i'}$ be the predecessor of R_i on P , and $R_{k'}$ be the successor of R_k on P . If $p_{i'}$ lies to the left of the vertical stripe of p_j , we move every point p_i, \dots, p_k horizontally



Fig. 1. Moving points in a vertical stripe.

to the x -coordinate x_l . Otherwise, we move them horizontally to x_r . Figure 1 illustrates an example of such a moving. The cost difference of P before and after moving the points can be expressed as

$$\sum_{(R_a, R_b) \in P'} \|p_a, p_b\| - \|p'_a, p'_b\|, \quad (1)$$

where P' is the sub-path between $R_{i'}$ and $R_{k'}$, and p'_a (resp. p'_b) is the new location of p_a (p_b). Since points are only moved horizontally, their y -differences do not change. Thus, (1) can be rewritten as

$$\sum_{(R_a, R_b) \in P'} |x_a - x_b| - |x'_a - x'_b|. \quad (2)$$

Before moving them, all points p_i, \dots, p_k are contained in the vertical stripe of p_j ; therefore the cost of P' before the moving is at least $|x_{i'} - x_i| + |x_{k'} - x_k|$. After the moving, the x -coordinates of all p_i, \dots, p_k become $x' \in \{x_l, x_r\}$. Thus, (2) is at least

$$|x_{i'} - x_i| + |x_{k'} - x_k| - |x_{i'} - x'| - |x_{k'} - x'|. \quad (3)$$

If $p_{i'}$ and $p_{k'}$ lie on the same side of the vertical stripe of p_j , the new coordinate x' is closer to both $x_{i'}$ and $x_{k'}$. If $p_{i'}$ and $p_{k'}$ lie on different sides of the vertical stripe, then $|x_{i'} - x'| + |x_{k'} - x'| = |x_{i'} - x_{k'}|$. In both cases, (3) is positive; that is, the cost of P does not increase and the new placement is still optimum.

The case where $R_i = R_s$ follows trivially from above, because we can define $p_{i'}$ to be the point p_i itself. In this way, the distance between p_i and $p_{i'}$ is always 0, no matter where the point is moved. The same holds also for the remaining cases. \square

2.2 Algorithm

We now present an algorithm that computes an optimum SPN placement by exploiting the structural properties of optimal placements established in Lemma 1. To do so, we create an auxiliary graph from \mathcal{R} and E with the property that a shortest path between two designated vertices of this graph yields a minimum SPN placement. Such a path can be found using standard shortest path techniques, such as Dijkstra's algorithm. The

auxiliary graph $\mathcal{D} = (V_{\mathcal{D}}, E_{\mathcal{D}})$ is defined as follows. There is a vertex in $V_{\mathcal{D}}$ for every point on the perimeter of a region that is also a point of the Hanan grid induced by the corners of the regions in \mathcal{R} , and two additional vertices v_s and v_t . In the following, we say “a vertex v of region R_i ” to indicate a vertex corresponding to a point of R_i . There is an edge in $E_{\mathcal{D}}$ from v_s to every vertex of R_s , and from every vertex of R_t to v_t . Also, let u be a vertex of R_i and R_j be a region such that $(R_i, R_j) \in E$. For every segment composing the perimeter of R_j , there is an edge in $E_{\mathcal{D}}$ from u to its closest vertex on that segment. Furthermore, there is an edge in $E_{\mathcal{D}}$ from u to the next vertex along the perimeter of R_i , in both directions. We assign a cost to an edge $(u, v) \in E_{\mathcal{D}}$ equal to 0 if either $u = v_s$ or $v = v_t$, and equal to $\|uv\|$ otherwise. The following theorem shows that a shortest path between v_s and v_t in \mathcal{D} yields an optimum SPN placement.

Theorem 1. *Given a shortest path $P_{\mathcal{D}}$ from v_s to v_t in \mathcal{D} , let \mathbf{p} be a placement as follows. For each region $R_i \in \mathcal{R}$, if R_i has vertices on $P_{\mathcal{D}}$, \mathbf{p} contains the first of them. Otherwise, \mathbf{p} contains one of its corners chosen arbitrarily. The placement \mathbf{p} is an optimum SPN placement.*

Proof. To see that \mathbf{p} is optimum, consider the vertices on $P_{\mathcal{D}}$ chosen as points of \mathbf{p} in the order as they appear on $P_{\mathcal{D}}$. Since the regions of these points are connected in \mathcal{G} , the path $P_{\mathcal{D}}$ corresponds to an st -path P in \mathcal{G} . By triangle inequality, the cost of P in $\mathcal{G}_{\mathbf{p}}$ is at most the cost of $P_{\mathcal{D}}$. For the sake of contradiction, suppose there exists an optimum placement \mathbf{q} and a shortest st -path Q in $\mathcal{G}_{\mathbf{q}}$ whose cost is smaller than the cost of P in $\mathcal{G}_{\mathbf{p}}$. Without loss of generality, we can assume the points in \mathbf{q} to satisfy Lemma 1. Thus, every point of \mathbf{q} corresponds to a vertex of \mathcal{D} . We construct a path $Q_{\mathcal{D}}$ from v_s to v_t in \mathcal{D} as follows. The first edge is (v_s, q_s) ; after that, for every edge (R_i, R_j) on Q , consider the points $q_i, q_j \in \mathbf{q}$, the bounding box B_{q_i, q_j} , and the at most two segments on the perimeter of R_j on which q_j lies. By construction, q_i is connected in \mathcal{D} to a vertex on both segments; let v be one of them chosen arbitrarily such that $v \in B_{q_i, q_j}$. We add to $Q_{\mathcal{D}}$ the path that from q_i goes to v , and follows the perimeter of R_j to q_j . By Proposition 1, the cost of this path is equal to the distance $\|q_i q_j\|$. The last edge on $Q_{\mathcal{D}}$ is (q_t, v_t) . To see that the cost of $Q_{\mathcal{D}}$ is equal to the cost of Q in $\mathcal{G}_{\mathbf{q}}$ it is sufficient to notice that the first and the last edge of $Q_{\mathcal{D}}$ have cost 0 and, for every edge (R_i, R_j) on Q , the sub-path from q_i to q_j in $Q_{\mathcal{D}}$ has cost equal to $\|q_i q_j\|$. This results in a contradiction, because we have then found a path from v_s to v_t with cost smaller than $P_{\mathcal{D}}$. \square

The above theorem shows how to construct an optimum SPN placement once a shortest path between v_s and v_t in \mathcal{D} is known. Since edge costs in \mathcal{D} are greater or equal than 0, we can find such a path with Dijkstra’s algorithm in time $\mathcal{O}(|V_{\mathcal{D}}| \log |V_{\mathcal{D}}| + |E_{\mathcal{D}}|)$. The sizes of $V_{\mathcal{D}}$ and $E_{\mathcal{D}}$ depend on the number of points on the perimeters of the regions that are the grid points of the Hanan grid induced by the corners of \mathcal{R} . To evaluate this number, consider a line of the Hanan grid. Each time this line intersects (cut in two nonempty parts) an orthogonal segment on the perimeter of a region, an additional vertex is introduced. Conversely, each segment of the perimeter of a region can in the worst case be intersected by every grid line orthogonal to it. If k is the maximum number of corners of a region in \mathcal{R} (and therefore on the number of segments of its perimeter), and $|\mathcal{R}| = n$, the number of grid lines is $\mathcal{O}(nk)$. Thus, the number




Fig. 2. An example of a planar 3-SAT instance on 5 variables. Dashed line are parts of the spinal path, solid lines are clauses.

of grid points lying on the perimeter of one region is $\mathcal{O}(nk^2)$, and the size of $V_{\mathcal{D}}$ is $\mathcal{O}(n^2k^2)$. To evaluate the size of $E_{\mathcal{D}}$, consider an edge $(R_i, R_j) \in E$ and a vertex v of R_i . By construction, there is an edge from v to a vertex on each of the at most k segments on the perimeter of R_j . Furthermore, v is connected to at most two vertices on the perimeter of R_i . If we have $|E| = m$ edges and $\mathcal{O}(nk^2)$ vertices in each region, the size of $E_{\mathcal{D}}$ is therefore $\mathcal{O}(mnk^3)$. Therefore, the running time of Dijkstra's algorithm for computing a shortest path from v_s to v_t in \mathcal{D} is $\mathcal{O}(n^2k^2 \log nk + mnk^3)$.

3 Minimum Spanning Tree with Neighborhoods

In the Minimum Spanning Tree Problem with Neighborhoods, or MSTN, we are given a set of regions $\mathcal{R} = \{R_1, \dots, R_n\}$ and an underlying graph $\mathcal{G} = (\mathcal{R}, E)$. The problem asks for a placement \mathbf{p} such that the cost of a minimum spanning tree in $\mathcal{G}_{\mathbf{p}}$ is smallest among all possible placements.

It is known [2] that, if distances are measured in L_2 norm and the neighborhood regions are disks, the MSTN problem does not admit an FPTAS unless $\text{P} = \text{NP}$. We will adapt their proof and show that MSTN does not admit an FPTAS also in the case where distances are measured in L_1 metric, the neighborhood regions are non-overlapping axis-parallel segments, and the underlying graph is complete. In the following, the term MSTN refers to this variant.

The reduction is from the planar 3-SAT problem. Planar 3-SAT is a variant of 3-SAT where the graph associated with the formula is planar. The graph contains a vertex for each variable and each clause, and there is an edge from a variable to a clause if the clause contains a literal of that variable. Planar 3-SAT was shown to be NP-hard by a reduction from the standard 3-SAT problem [8]. Furthermore, it was shown that in the plane embedding used in the reduction there always exists a so-called *spinal path* passing through every vertex corresponding to a variable without crossing any edge of the graph. Knuth and Raghunathan [7] observed that there always is a simple embedding where the variables are arranged on a straight line (the spinal path), and the clauses are drawn as three legged segments completely above or below them, in a way such that none of the legs cross each other. Figure 2 shows an example of such an embedding.

The reduction starts from a plane embedding of an instance of planar 3-SAT and constructs an instance of MSTN such that a solution to the latter indicates whether the former is satisfiable. First, we define three types of gadgets: a gadget for each variable, a gadget for each clause, and a gadget for the spinal path. Then, we show how to replace each variable, clause and the spinal path with a corresponding gadget resulting in an

instance of MSTN. From our construction, it will be easy to see that the size of the resulting MSTN instance is polynomially bounded. Finally, we provide two threshold values t_1 and t_2 , with $t_1 < t_2$, and we prove that an optimum solution of the constructed MSTN instance has a cost smaller than t_1 if and only if the initial 3-SAT formula is satisfiable. If the formula is not satisfiable, the cost of an optimum solution of the MSTN instance is at least t_2 . This proves that MSTN does not admit an FPTAS unless $P = NP$.

An important tool in the definitions of the gadgets is a so-called wire. A *wire* is a set of points (i.e., regions) placed in close succession, so that any minimum spanning tree (for any placement) will contain the edges connecting the points. To ensure this, it is sufficient to place two consecutive points in a wire at a suitably small distance. Since the edges between consecutive points in wires do not form a cycle, any minimum spanning tree in any placement will contain the edge connecting them. However, this suitably small distance must still be large enough to guarantee that a wire can be realized with a polynomial number of points. Since in the following construction the smallest non-zero distance between any two regions (other than those for the wires) is at least $d/2$, for a constant $d := 0.25$, a suitably small value for the points of a wire is, for example, $d/4$.

3.1 Reduction Gadgets

Variable gadget. For each variable there are $k = 6c + 6$ segments, or rectangles with empty area, of length α , where c is the maximum number of clauses in which the variable appears as a literal that are completely above or below the variable vertex in the embedding. Note that $k \geq 12$, because a variable appears at least once in a clause. In the following we specify the value of the parameter α more precisely, and we show it to be polynomial in the number of clauses and variables.

As illustrated in Figure 3, the segments are placed along the perimeter of a rectangle with sides of length $3c\alpha + d$ and $3\alpha + d$. In its interior we place a wire for every two segments consecutive in clockwise order. Each of these wires ends on the line bisecting the angle formed by the corresponding segments; for parallel segments, the endpoint is at distance d from their common point. For perpendicular segments, the endpoint is at distance d from the intersection of the lines passing through the segments. We connect these wires in the bounded region in a tree-like structure as in Figure 3. We call this arrangement of wires in the internal region a *k-tree*.

A placement of points inside a variable gadget is called a *configuration* if, for every two consecutive segments in clockwise order, the placement contains either their two closest points or their two farthest points. For a variable gadget there exist exactly two different configurations. To see this, consider two consecutive segments in a variable gadget and a configuration placement. If the placement contains their two closest points, we can place points in the remaining segments in exactly one way in order to obtain a configuration. Similarly, if the placement contains their two farthest points, we have exactly one way to place points in the remaining segments. We associate these two possible configurations with the two assignments to the variable.

Clause gadget. Clause gadgets are composed of at most three wires meeting at a single point following the embedding. As in Figure 3, each wire of a clause gadget approaches the common point of two adjacent horizontal segments of a variable gadget. Clauses that




Fig. 3. A variable gadget with $k = 18$. The variable appears in A with negative sign and in B, C with positive sign. Thicker lines are the segments, the rest wires.

are located above the spinal path in the rectilinear embedding approach variable gadgets from above, while clauses that in the rectilinear embedding are located below the spinal path approach variable gadgets from below. Furthermore, clause wires approach a variable gadget in the same clockwise order as the edges connecting the variable vertex to the corresponding clauses in the rectilinear embedding.

A clause wire terminates at distance $1 + 2d$ from the common point of the approached segments along the vertical line passing through it. The approached segments are chosen such that their common point is contained in a configuration satisfying the clause. That is, an edge with cost $1 + 2d$ connects the clause wire to the segments in a configuration placement satisfying the clause.

Spinal path gadget. The spinal path gadget consists of wires following the embedding of the planar 3-SAT instance. As in Figure 3, the spinal path gadget approaches every variable gadget twice, once from the left and once from the right. For each side, the spinal path wire is split in two parts, each approaching two adjacent vertical segments. The point at which a part terminates is located at distance 1 from the common point of the approached segments along the horizontal line passing through it.

3.2 The reduction

Given an instance of a planar 3-SAT and its rectilinear embedding on the plane [7], we create an instance of MSTN and provide two threshold values t_1, t_2 , with $t_1 < t_2$. We show that if the 3-SAT instance is satisfiable, then there is a placement with a minimum spanning tree of cost at most t_1 , and if the 3-SAT instance is unsatisfiable, then the cost of a minimum spanning tree for any placement is at least t_2 .

Theorem 2. *MSTN with L_1 metric and neighborhood regions as segments is APX-hard.*

Proof. To create an instance of MSTN, replace in the given embedding every variable, clause, and the spinal path with a gadget as explained above. The wires forming the spinal path, the m clause gadgets and the k -trees in the internal region of each variable gadget have a fixed cost in every MST, denoted as c_{wires} . The remaining cost of the spanning tree is given by connecting the segments of the variable gadgets to the k -trees and the spinal path and clause wires.

Suppose there exists a satisfying assignment. Then, we place points in each variable gadget in a configuration according to its value in the assignment. We provide an upper bound t_1 on the cost of a minimum spanning tree in this placement by constructing a spanning tree and evaluating its cost. For each pair of consecutive segments having their closest points in the placement, the spanning tree connects them to the k -tree of the corresponding variable with cost d . If there is a total of K segments among all variable gadgets, the spanning tree requires a cost of $(K/2)d$ to connect all of them to the k -trees (note that K is even). For each clause gadget, consider a variable satisfying it in the assignment. We connect the corresponding endpoint of the clause wire to one of the segments it approaches with an edge with cost $1 + 2d$. Overall, the cost for connecting all the clause wires to the tree is $m(1 + 2d)$. For each part of the spinal path gadget approaching a variable gadget, exactly one of its endpoint approaches a point of the placement. The spanning tree contains the $2n$ edges of cost 1 connecting them. Overall, the cost of an optimum MSTN solution in case a satisfying assignment exists is therefore at most

$$t_1 := c_{wires} + (K/2)d + (1 + 2d)m + 2n.$$

If there is no satisfying assignment, we show that the cost of an optimum MSTN solution is at least $t_2 := t_1 + d$. To see this, consider an optimum placement where every point in a segment is one of its extreme points. The existence of such an optimum placement is guaranteed by the fact that wires approaching variable gadgets and wires of the k -trees terminate either to the left or to the right of horizontal segments, and above or below vertical segments.

First, we provide an upper bound on the minimum spanning tree cost for such a placement. To provide this upper bound, we construct a spanning tree and we evaluate its cost in the placement. Then, we use this upper bound to show that, in every minimum spanning tree, clause wires are connected to the tree either with an edge from one of the endpoints to one of the approached segments, or with an edge from one of its endpoints to the approached k -tree endpoint. Finally, we show that the cost of any optimum MSTN solution is at least t_2 .

The spanning tree contains the wires composing the spinal path, the clauses, and all k -trees. Every segment in a variable gadget is connected to an endpoint of the k -tree with an edge of cost d . Every part of a wire of the spinal path approaching a variable gadget is connected to one of the approached segments by one of its endpoints with an edge with cost 1. Similarly, an endpoint of each clause wire chosen arbitrarily is connected to a k -tree of a variable appearing in that clause with an edge with cost $1 + 3d$. The cost of such a spanning tree is

$$c_{wires} + 2n + m(1 + 3d) + Kd. \tag{4}$$

We now prove that, in every minimum spanning tree, each clause is connected to it either with an edge from one of its endpoint to an approached segment, or with an edge from one of its endpoints to the corresponding k -tree endpoint. Suppose this is not the case, and there is a clause whose endpoint in the MST is connected neither to an approached segment, nor to the corresponding k -tree endpoint. By construction, the next closest object is located at distance at least α , where α is the above defined length of the segments of the variable gadgets. Since the wires composing the spinal path, the clauses and the k -trees are part of every MST, by setting

$$\alpha := 2n + m(1 + 3d) + Kd + 1$$

we get a contradiction, because the cost of a minimum spanning tree would then be greater than (4), the cost of the spanning tree shown above. Thus, in every minimum spanning tree every clause is connected via one of its endpoints to one of the approached regions.

Finally, we show that if the formula is not satisfiable, any optimum MSTN solution has cost greater than t_2 . Clearly, we cannot provide a configuration for each variable gadget such that every clause where that variable appears can be connected to it with an edge with cost $1 + 2d$, otherwise the formula would be satisfiable. Therefore, in an optimum solution, either at least one variable gadget is not set in a configuration, or every variable gadget is in a configuration and for at least one clause no wire endpoint approaches a point of the placement.

In the former case, the cost of a minimum spanning tree is at least $c_{wires} + 2n + (K/2)d + j(1 + 2d) + (m - j)\delta$, where j is the number of clauses that can be satisfied by the assignment corresponding to the configuration and δ is the minimum cost necessary to connect a clause that is not satisfied by the assignment. By the above, we know that in any minimum spanning tree a clause wire is connected by one of its endpoint to one of the approached segments or the corresponding k -tree endpoint. Since every variable gadget is in a configuration, the smallest distance δ between an endpoint of a non satisfied clause wire and a point in the placement is at least $1 + 3d$. Thus, any optimum MSTN solution where every variable gadget is set in a configuration results in a minimum spanning tree with cost at least t_2 .

In the latter case, there exists at least one variable gadget that is not in a configuration. Let then a be the overall number of segments for which the point in the placement is not the closest or the farthest to the point in one of the consecutive segments. Note that a is even, therefore $a \geq 2$, and the cost of a spanning tree is at least

$$c_{wires} + 2n + \frac{(K + a)}{2}d + m(1 + 2d) \geq t_2.$$

Suppose now that there exists an FPTAS for MSTN. Given an instance of planar 3-SAT, we construct the gadget presented above and calculate t_1 . We then set a parameter $\epsilon < d/t_1$, so a $(1 + \epsilon)$ -approximate solution to the MSTN problem would tell us whether the cost of the corresponding optimum solution is smaller than t_1 or greater than t_2 , and thus, whether there exists a satisfying assignment for the planar 3-SAT instance. \square

4 Conclusions

We considered the Shortest Path Problem and the Minimum Spanning Tree Problem with Neighborhoods in L_1 metric. We showed that the first problem can be solved efficiently if the neighborhood regions are rectilinear non-convex polygons, while the second one does not admit a PTAS unless $P = NP$ even in the case where the regions are segments.

An interesting open problem is to consider a variant of SPN where the goal is to find a placement *maximizing* the cost of a shortest path between two given regions (rectilinear polygons). In this case, it is easy to adapt the proof of Lemma 1 to show that the Hanan grid induced by the corners of the regions still contains the points of an optimum placement. However, it is not clear how to design a polynomial time algorithm computing it.

References

1. Ahadi, A., Mozafari, A., Zarei, A.: Touring disjoint polygons problem is NP-hard. In: Proc. 7th International Conference on Combinatorial Optimization and Applications (COCOA). pp. 351–360 (2013)
2. Dorigiv, R., Fraser, R., He, M., Kamali, S., Kawamura, A., López-Ortiz, A., Seco, D.: On minimum-and maximum-weight minimum spanning trees with neighborhoods. In: Proc. 10th International Workshop on Approximation and Online Algorithms (WAOA). pp. 93–106 (2012)
3. Dror, M., Efrat, A., Lubiw, A., Mitchell, J.S.B.: Touring a sequence of polygons. In: Proc. 35th Annual ACM Symposium on Theory of Computing (STOC). pp. 473–482 (2003)
4. Dumitrescu, A., Mitchell, J.S.B.: Approximation algorithms for TSP with neighborhoods in the plane. In: Proc. Twelfth Annual ACM-SIAM Symposium on Discrete algorithms (SODA). pp. 38–46 (2001)
5. Elbassioni, K.M., Fishkin, A.V., Mustafa, N.H., Sitters, R.: Approximation algorithms for euclidean group TSP. In: Proc. 32nd International Colloquium on Automata, Languages and Programming (ICALP). pp. 1115–1126 (2005)
6. Hanan, M.: On Steiner’s problem with rectilinear distance. *SIAM Journal on Applied Mathematics* 14(2), 255–265 (1966)
7. Knuth, D., Raghunathan, A.: The problem of compatible representatives. *SIAM Journal on Discrete Mathematics* 5(3), 422–427 (1992)
8. Lichtenstein, D.: Planar formulae and their uses. *SIAM J. Comput.* 11(2), 329–343 (1982)
9. Löffler, M., Kreveld, M.: Largest and smallest convex hulls for imprecise points. *Algorithmica* 56(2), 235–269 (2010)
10. Löffler, M., van Kreveld, M.J.: Largest bounding box, smallest diameter, and related problems on imprecise points. *Comput. Geom.* 43(4), 419–433 (2010)
11. Pan, X., Li, F., Klette, R.: Approximate shortest path algorithms for sequences of pairwise disjoint simple polygons. In: Proc. 22nd Canadian Conference on Computational Geometry (CCCG). pp. 175–178 (2010)
12. Papadimitriou, C.H.: The euclidean travelling salesman problem is NP-complete. *Theoretical Computer Science* 4(3), 237–244 (1977)
13. Yang, Y., Lin, M., Xu, J., Xie, Y.: Minimum spanning tree with neighborhoods. In: Proc. Third International Conference on Algorithmic Aspects in Information and Management (AAIM). pp. 306–316 (2007)
14. Zachariasen, M.: A catalog of Hanan grid problems. *Networks* 38(2), 76–83 (2001)