TSP-Based Curve Reconstruction in Polynomial Time

Ernst Althaus Kurt Mehlhorn*

Max-Planck-Institut für Informatik

Im Stadtwald, 66123 Saarbrücken, Germany

{althaus,mehlhorn}@mpi-sb.mpg.de

September 30, 1999

Abstract

An instance of the curve reconstruction problem is a finite sample V of an unknown curve γ and the task is to connect the points in V in the order in which they lie on γ. Giesen [Gie99] showed recently that the Traveling Salesman tour of V solves the reconstruction problem under fairly week assumptions on γ and V. We extend his result along three dimensions. We weaken the assumptions, we give an alternate proof, and we show that in the context of curve reconstruction the Traveling Salesman tour can be constructed in polynomial time.

1 Introduction

An instance of the curve reconstruction problem is a finite sample V of an unknown curve γ and the task is to construct a graph G on V such that two points in V are connected by an edge of G iff the points are adjacent on γ. The graph G is called a *polygonal reconstruction* of γ. The curve reconstruction problem and the related surface reconstruction problem received a lot of attention in the graphics and the computational geometry community. We are interested in *reconstruction algorithms with guaranteed performance*, i.e., algorithms which provably solve the reconstruction problem under certain assumptions on γ and V. Figure 1 illustrates the curve reconstruction problem.

If the curve is closed, smooth, and uniformly sampled, several methods are known to work ranging over minimum spanning trees [FG94], α-shapes [BB97, EKS83], β-skeletons [KR85], and r-regular shapes [Att97]. A survey of these techniques appears in [Ede98]. The case of non-uniformly sampled closed curves was first treated successfully by Amenta, Bern and Eppstein [ABE98] and subsequently improved algorithms such as [DK99, Gol99] appeared. Open non-uniformly sampled curves were treated in [DMR99]. All papers mentioned so far require that the underlying curve is smooth.

Giesen [Gie99] recently obtained the first result for non-smooth curves. He considered the class of *benign semi-regular curves.* An (open or closed) curve is *semi-regular* if a left and a right tangent exists in every point of the curve; the two tangents may however be different. A semi-regular curve is *benign* if the turning angle at every point of the

*research partially supported by EG-projects ALCOM-IT and GALIA.

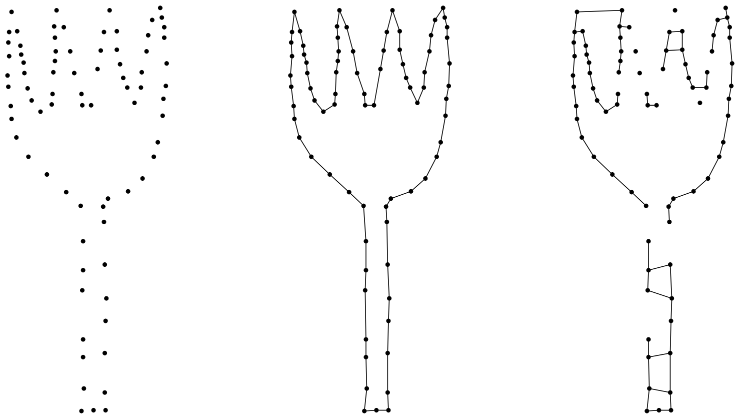

Figure 1: Part (a) shows a finite set V of points, part (b) shows the Traveling Salesman tour of the points, and part (c) shows the output of the reconstruction algorithm of Amenta, Bern, and Eppstein.

curve is less than π, see Figure 2. Giesen showed that the Traveling Salesman path, tour respectively, of the sample set V solves the curve reconstruction problem for uniformly sampled benign open, closed respectively, semi-regular curves. More precisely, he showed that for every benign semi-regular curve γ there exists a positive ϵ such that the Traveling Salesman path (tour) of V is a polygonal reconstruction provided that for every $x \in \gamma$ there is a $p \in V$ with $\|xp\| \leq \epsilon$. Giesen's result is an existence result; he did not quantify ϵ in terms of properties of the curve γ. We extend Giesen's result in three directions:

- We relate ϵ to local properties of the curve γ and show that the Traveling Salesman tour (path) solves the reconstruction problem even if the sampling is non-uniform. For smooth curves our sampling condition is similar to the one used in [ABE98, DK99, Gol99, DMR99].
- We show that the Traveling Salesman tour (path) can be constructed in polynomial time if the sampling condition is satisfied.
- We give a simplified proof that the Traveling Salesman tour (path) solves the curve reconstruction problem.

We have implemented a number of curve reconstruction algorithms. The JAVA-Applet http://review.mpi-sb.mpg.de:81/Curve-Reconstruction/ makes our implementations available. Our experiments show that the TSP-based curve reconstruction is able to solve the reconstruction problem for surprisingly small sampling density and that its speed is comparable to Delaunay diagram based reconstruction algorithms. The details are given in the companion paper [AMNS99].

This paper is structured as follows. In Section 2 we give definitions and state the main theorem. In Sections 3 and 4 we prove our main theorem for open and closed curves, respectively. In Section 5 we relate our sampling condition to the sampling conditions used by other papers, and in Section 6 we discuss the relationship to so-called necklace tours. Finally, in Section 7 we offer conclusions and state open problems.

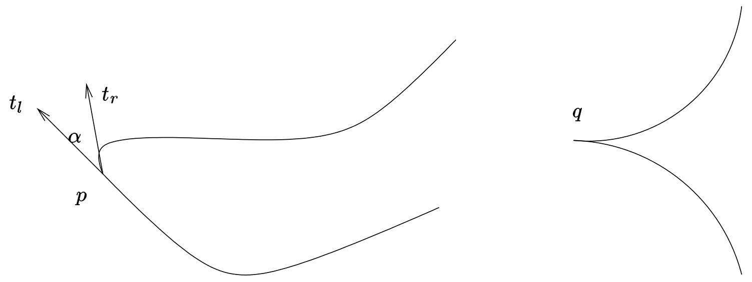

Figure 2: Two semi-regular curves, one benign and one not. In the left curve, the two tangents t_l and t_r at a point p of the curve are shown. The turning angle at p is α. In the curve on the right the turning angle at q is π as the left and right tangents at q point into opposite direction.

2 Definitions and Basic Properties

A open *curve* is given by an *embedding* $\gamma : [0,1] \to I\!R^2$ and a closed curve is given by an embedding $\gamma : S^2 \to I\!R^2$.

Definition 1 ([Gie99]) *Let*

$$T = \{(t_1, t_2)\ ;\ t_1 < t_2, t_1, t_2 \in [0,1]\}$$

and

$$\tau : T \to S^2, (t_1, t_2) \mapsto \frac{\gamma(t_2) - \gamma(t_1)}{|\gamma(t_2) - \gamma(t_1)|}.$$

The curve γ is called left (right) regular at $\gamma(t_0)$ with left (right) tangent $\mathrm{t}(\gamma(t_0))$ if for every sequence (ξ_n) in T which converges to (t_0, t_0) from the left (right) in the closure of T the sequence $\tau(\xi_n)$ converges to $\mathrm{t}(\gamma(t_0))$. We call γ semi-regular if it is left and right regular in all points $\gamma(t)$, $t \in [0,1]$. We call γ regular if it is semi-regular and the left and right tangent coincide in every point of the curve.

Figure 2 shows two semi-regular curves. Tangents are unit vectors. The angle between two vectors with the same source is the smaller of the two angles between the vectors. The angle is zero if the two vectors point into the same direction and the angle is π if the vectors point into the opposite direction. The angle between the left and right tangent at a point $p \in \gamma$ is called the *turning angle* at p. If the curve has a tangent at p, the turning angle at p is zero. If the turning angle at p is non-zero, we call p a *singularity* of the curve. A semi-regular curve is *benign* if the turning angle is less than π at every point of the curve. We can now state Giesen's result.

Theorem 1 ([Gie99]) *For every benign semi-regular curve γ there exists an $\epsilon > 0$ with the following property: If V is a finite sample of γ such that for every $x \in \gamma$ there is a $p \in V$ with $\|pv\| \leq \epsilon$, the Traveling Salesman path (tour if γ is closed) is a polygonal reconstruction of γ.*

The construction of Traveling Salesman paths or tours is an NP-hard problem. A successful method for solving the Traveling Salesman problem is to formulate the problem as an integer linear program (ILP) and to use a branch-and-cut algorithm based on the LP-relaxation of the problem. We give the formulation for Traveling Salesman paths with fixed endpoints a and b. We introduce a variable x_{uv} for every edge uv between two sample points and describe the set of all paths with endpoints a and b in the following way:

$$\begin{aligned}\sum_{v\in V} x_{uv} &= 2 \text{ for all } u \in V \setminus \{a,b\}\\ \sum_{v\in V} x_{uv} &= 1 \text{ for } u \in \{a,b\}\\ \sum_{u\in V', v\in V'} x_{uv} &\le |V'| - 1 \text{ for } V' \subset V,\ V' \neq \emptyset\\ x_{uv} &\in \{0,1\} \text{ for all } u,v \in V\end{aligned}$$

We refer to this program as the *Subtour-ILP*. The equality constraints in the Subtour-ILP are called *degree constraints*, the inequality constraints are called *subtour elimination constraints*, and the constraints $x_{uv} \in \{0,1\}$ are called the *integrality constraints*. Relaxing the integrality constraints to $0 \le x_{uv} \le 1$ gives the *Subtour-LP*. The objective function for both programs is $\sum_{uv} \|uv\| x_{uv}$, i.e., the total Euclidean length of the edges selected. It is well known that the Subtour-LP can be solved in polynomial time by use of the Ellipsoid method, see [Sch86]. A potentially exponential, but practically very efficient algorithm uses the simplex method and the cutting plane framework. One starts with the LP consisting only of the degree constraints and then solves a sequence of LPs. In each iteration one checks whether the solution to the current LP satisfies all subtour elimination constraints (= a min-cut problem) and, if not, one adds a violated cut constraint to the LP. We use the latter algorithm in our experimental framework for curve reconstruction. *We observed in our experiments that the optimal solution of the Subtour-LP is integral when the input set V is a dense sample of a semi-regular curve.* Our main theorem states that our experimental observation holds true generally.

Theorem 2 *Let γ be an open benign semi-regular curve, let V be a finite set of samples of γ, and let a and b be the first and last sample point, respectively (in the order on γ). If V satisfies the sampling condition given below then:*

- *The Traveling Salesman path of V with endpoints a and b is a polygonal reconstruction of γ.*
- *The Subtour-LP has an optimal integral solution and this solution is unique.*

For a closed curve the statements above hold with the Traveling Salesman tour instead of the Traveling Salesman path.

For a given sample set V we can view the curve reconstruction problem as a problem in a complete network (V, E, c), where $c(u,v)$ is the Euclidean distance of u and v. We use T to denote the cost of a traveling salesman path with endpoints a and b. Our proof of Theorem 2 exploits the connection between the Subtour-LP and the so-called *Held-Karp bound*. The Held-Karp bound is based on the following lemma.

Lemma 1 (Held-Karp) *Let $\pi : V \to I\!R$ be an arbitrary function*[1]*. Define a* modified

[1] We follow common practice and call π a potential function.

cost function c_π *by* $c_\pi(u,v) = c(u,v) - \pi(u) - \pi(v)$*, let* MST_π *be a minimum spanning tree of* $N_\pi := (V, E, c_\pi)$*, let,* C_π *be its length with respect to* c_π*, and let* a *be the first and* b *be the last sample point. Then* $C_\pi + 2\sum_{v \in V} \pi(v) - \pi(a) - \pi(b)$ *is a lower bound on the length of the traveling salesman path with endpoints* a *and* b*.*

Proof: Let T_π denote the cost of the traveling salesman path with endpoints a and b with respect to the modified cost function. Then $\mathrm{MST}_\pi \le T_\pi$ since a path is a spanning tree and $T_\pi = T - 2\sum_{v \in V} \pi(v) + \pi(a) + \pi(b)$ by the definition of c_π. ∎

The Held-Karp bound[2] is $\max_\pi C_\pi + 2\sum_{v \in V} \pi(v) - \pi(a) - \pi(b)$. It is equal to the objective value of the Subtour-LP, see [CCPS98, page 259].

We can now outline the proof of the main result for open curves. The extension to closed curves is the subject of Section 4. We define a potential function π such that the minimum spanning tree MST_π is unique and coincides with the polygonal reconstruction of γ (and hence is a path with endpoints a and b). Then $T \le C$, since T is the length of the traveling salesman path, and $C \le T$, since a path is a spanning tree. Thus $T = C$, MST_π is a traveling salesman path, and the value of the Held-Karp bound is equal to the cost of the traveling salesman path. We conclude that the value of the Subtour-LP is equal to the cost of the traveling salesman path and that (the incidence vector of) MST_π is *an* optimal solution of the Subtour LP. We will next argue that it is the unique optimal solution. Assume that there is an optimal solution of the subtour LP with $x_e > 0$ for some $e \notin \mathrm{MST}_\pi$. Since MST_π is unique there is a $\mu > 0$ such that decreasing the cost of the edge e by μ will not change the minimum spanning tree and hence will not change the value of the Held-Karp bound. However, the objective value of the Subtour-LP will decrease. This is a contradiction to the equality of the two bounds.

We will define out potential function in Section 3.2; it assigns a positive real $\pi(x)$ to every point x of γ. Define the turning angle of a subcurve γ' of γ as the opening angle of the smallest cone that contains all (left and right) tangents to points $p \in \gamma'$. We require:

Sampling Condition:

(a) For any two adjacent (on γ) samples u and v: $c_\pi(u,v) \le 0$.

(b) For any two adjacent samples u and v: $\gamma[u,v]$ turns by less than π. Here $\gamma[u,v]$ denotes the subcurve of γ with endpoints u and v and containing no further sample point.

Condition (a) implies condition (b) in the case of open curves as we will see below. For closed curves, condition (b) is needed as the example in Figure 3 shows. The sampling condition is easily satisfied. Let $\epsilon = \inf_{x \in \gamma} \pi(x)$. Then $\epsilon > 0$ since γ is compact and hence a sample set such that γ turns by less than π between adjacent samples and such that there is at least one sample point in every curve segment of length $\epsilon/2$ satisfies the sampling condition.

[2] A simple iterative algorithm can be used to compute the Held-Karp bound. Let α_t, $t = 0, 1, \ldots$ be positive reals with $\alpha_t \to 0$ as $t \to \infty$ and $\sum_t \alpha_t = \infty$. Start with $\pi(v) = 0$ for all $v \in V$ and compute MST_π. As long as MST_π is not a path with endpoints a and b, update π as follows: $\pi(v) = \pi(v) + \alpha_t \cdot (d(v) - deg_\pi(v))$, where $d(a) = d(b) = 1$ and $d(v) = 2$ for all other nodes, t is the index of the iteration, and $deg_\pi(v)$ is the degree of v in MST_π. The cost of MST_π converges against the Held-Karp bound. We found in our experiments that the convergence is rather slow and that the cuttig plane approach to solving the Subtour-LP is faster.

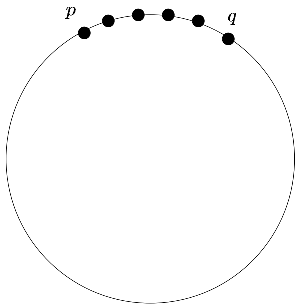

Figure 3: p and q are adjacent sample points and $c_\pi(p,q) \leq 0$ if p and q are sufficiently close. However, the sample contradicts our intuition of what constitutes a dense sample. Condition (b) excludes the case.

In order to show that MST_π is the polygonal reconstruction of γ from V we define a family Γ of (overlapping) subcurves γ' of γ such that:

(P1) Each subcurve γ' is connected and the minimum spanning tree (with respect to cost function c_π) of the points in $V \cap \gamma'$ is unique and coincides with the polygonal reconstruction of γ'.

(P2) For every edge e with $c_\pi(e) \leq 0$ there is a subcurve $\gamma' \in \Gamma$ containing both endpoints of e.

We show that (P1) and (P2) imply that MST_π is the polygonal reconstruction of γ. Let MST_π be any minimum spanning tree in (V, E, c_π) and let $e = uv$ be any edge which does not belong to the polygonal reconstruction. We show that $e \notin \mathrm{MST}_\pi$. Observe first that any edge in the minimum spanning tree has non-negative modified cost since there is a spanning tree, namely the polygonal reconstruction, in which every edge has non-negative cost. This follows from the cycle rule for minimum spanning trees. So assume $c_\pi(e) \leq 0$. Then there is a subcurve $\gamma' \in \Gamma$ containing both endpoints of e by (P2). We even have $\gamma[u,v] \subseteq \gamma'$ since γ' is connected by (P1). Moreover, the minimum spanning tree of $V \cap \gamma'$ is unique and coincides with the polygonal reconstruction of γ'. Thus $c_\pi(e') < c_\pi(e)$ for any edge e' on the part of the polygonal reconstruction between u and v. We conclude that $e \notin \mathrm{MST}_\pi$.

3 Open Curves

We assume that our open curve is oriented and write $p < q$ if p precedes q on γ.

3.1 Warm-Up

We start with two special cases to provide the reader with some intuition. Assume first that our curve consists of the two line segments $y = \pm m \cdot x$ for $0 \leq x \leq 1$ and let V be a finite set of samples. We call the two segments forming the curve the two legs of the curve, see Figure 4. We will show: (1) If $m \geq 1$ then the minimum spanning tree of V is the polygonal reconstruction. (2) If $m < 1$ then the minimum spanning tree of V does not necessarily give the polygonal reconstruction. (3) There is a (simple) potential

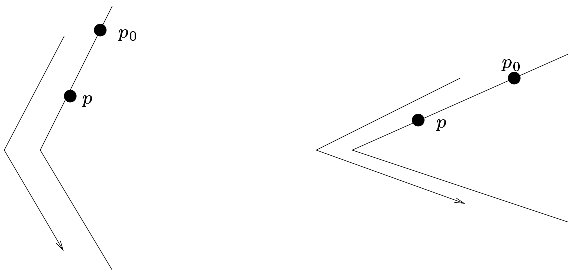

Figure 4: In the figure on the left, the Euclidean distance between p and p_0 grows as p moves away from p_0 along γ. In the figure on the right, the distance first grows and then shrinks again as p moves around the corner. The minimum spanning tree will reconstruct the curve on the left, but may file on the curve on the right.

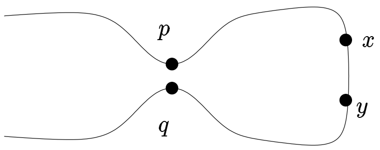

Figure 5: The minimum spanning tree MST_π will reconstruct the curve only if $c_\pi(x, y) < c_\pi(p, q)$. We therefore need to give x and y larger potential than p and q.

function π such that the minimum spanning tree M_π is the polygonal reconstruction even in the case $m < 1$.

Consider any point $p_0 = p(x_0) = (x_0, mx_0)$ on the "upper" leg. If $m \geq 1$, i.e., the angle between the two legs is at least $\pi/2$, any point on the upper leg between the origin and p_0 is closer to p_0 than any point on the lower leg and the first sample point (= closest to the origin) on the lower segment is the closest sample point to p_0 on the lower segment. This implies that Prim's minimum spanning tree algorithm when started with the first (= furthest from the origin) sample point on the upper leg will construct the polygonal reconstruction.

If $m < 1$, there is a point on the lower leg, say q, that is closer to p_0 than the origin and hence the minimum spanning tree is not the polygonal reconstruction if V consists of the points p_0, O, and q.

The crucial insight is that a potential function can be used to distort the metric so as to mimic the former case. We define the potential π as a function of x. Fix $\pi(O)$ arbitrarily, let $p(x) = (x, mx)$ and $q(x) = (x, -mx)$ and define $\pi(p(x)) = \pi(q(x)) = \pi(O) - x$. Then $c_\pi(p(x), p(x_0)) = \sqrt{1 + m^2}|x - x_0| - 2\pi(O) + x_0 + x$ and $c_\pi(q(x), p(x_0)) = \sqrt{(x - x_0)^2 + (mx + mx_0)^2} - 2\pi(O) + x_0 + x$. It is an easy exercise in calculus to show that $c_\pi(p(x), p(x_0))$ is an increasing function of $|x - x_0|$ and that $c_\pi(q(x), p(x_0))$ is an increasing function of x.

We will next have a look at smooth curves. Consider Figure 5. There are nearby sample points p and q which do not belong to the polygonal reconstruction. In order to exclude the edge pq from the minimum spanning tree we need to make its modified cost larger than the cost of all edges in the reconstruction. We will even make it positive. For a point $p \in \gamma$ let $d(p)$ be maximal such that the open ball $B^0(p, d(p))$ intersects γ in a

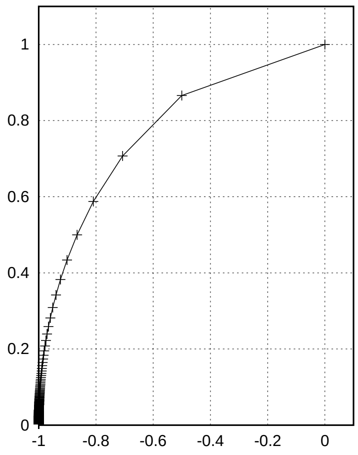

Figure 6: A semi-regular curve with an infinite number of corners

single component and such that $\gamma \cap B^0(p, d(p))$ turns by less than $\pi/3$. Let $\pi(p) = d(p)/3$. Then $c_\pi(p, q) > 0$. Moreover, π is larger in parts of the curve where the curvature is small and hence a smaller sampling density is required in parts of small curvature.

3.2 The Definition of the Potential Function

Singularities cause difficulties for most curve reconstruction algorithms; the difficulties grow with the turning angle. We call a singularity p a *sharp corner* if the turning angle at p is at least $7\pi/24 = 52.5°$; the choice of this threshold is somewhat arbitrary.

A semi-regular curve may have an infinite number of singularities. For example, the convex hull of the points $(\cos(\pi(1-1/n)), \sin(\pi(1-1/n)))$, $n \geq 2$, has an infinite number of singularities, see Figure 6. However, a semi-regular curve can have only a finite number of sharp corners. Assume otherwise. Then the sharp corners have an articulation point p. Let p_1, p_2, $\ldots$ be a increasing sequence of sharp corners converging to p. For any sharp corner p_i we can choose two points q_i and r_i in the vicinity of the corner such that the tangents at q_i and r_i form an angle of at least $\pi/6$ and such that the sequence q_1, r_1, q_2, r_2, $\ldots$ is increasing and converges to p. The sequence shows that γ has no left tangent at p.

We use S to denote the set of sharp corners of γ. We partition γ at the sharp corners and obtain a finite sequence γ_i, $0 \leq i \leq k$, of curves without sharp corners. We have $\gamma_i(1) = \gamma_{i+1}(0)$ for all i (if γ is a closed curve, interpret $k+1$ as zero) and $\gamma_i(x) \neq \gamma_j(y)$ otherwise.

We are now ready to define our potential function. The definition consists of two parts dealing with the neighborhoods of sharp corners and curve parts "far away" from all sharp corners, respectively. We start with the latter parts.

We use $B(p, r)$ and $B^0(p, r)$ to denote the closed and open ball with center p and radius r, respectively. For every point $p \in \gamma$ let $d(p)$ be maximal such that the open ball $B^0(p, d(p))$ has the following properties:

- $B^0(p, d(p)) \cap S = \emptyset$.
- $B^0(p, r) \cap \gamma$ is connected for all r with $r \leq d(p)$.

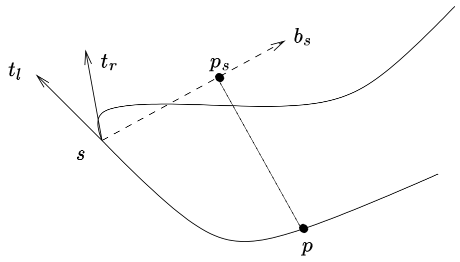

Figure 7: Illustration of the definition of b_s and p_s.

- $B^0(p, d(p)) \cap \gamma$ turns by less than $\pi/3$. The choice of this constant is again somewhat arbitrary.

Observe that the closed ball $B(p, d(p))$ has one of the following properties: it has a point in S on its boundary or it intersects γ in more than one component or γ turns by $\pi/3$ in the ball. For singularities $s \in S$ we have $d(s) = 0$ and for points $p \in \gamma \setminus S$ we have $d(p) > 0$. The function $p \mapsto d(p)$ is continuous. Thus $d(p)$ will be an increasing function of p as p moves away from a sharp corner for a neighborhood of any sharp corner. In Section 5 we will relate $d(p)$ to the distance of p to the medial axis of γ.

For sharp corners we define a quantity δ_s. For a sharp corner $s \in S$ let b_s be the bisector of the angle between the right tangent and the reversal of the left tangent, let α_s be the turning angle at s (we have $\alpha_s \geq 7\pi/24$), and let $m_s = \tan((\pi - \alpha_s)/2)$. Then $m_s \leq \tan 17\pi/48 \approx 2.03 < 9/4$. If one identifies s with the origin, b_s with the positive x-axis, the tangents have the line equations $y = \pm\, m_s \cdot x$. For a point $p \in \gamma$ let p_s be the orthogonal projection of p onto b_s. See Figure 7 for an illustration of these definitions. For every $s \in S$ let δ_s be maximal such that

- $\gamma \cap B^0(s, \delta_s)$ is connected. We call the two components of $(\gamma \setminus s) \cap B^0(s, \delta_s)$ the two legs of γ incident to s.
- The absolute value of the slope of the left or right tangent at any point on one of the legs is between $3/4m_s$ and $4/3m_s$.
- For any of the two legs $d(p)$ increases as p moves away from s.
- $B^0(s, 2\delta_s)$ contains no sharp corner different from s.

The last condition ensures that the balls $B^0(s, \delta_s)$, $s \in S$, are pairwise disjoint. Clearly, $\delta_s > 0$ for all $s \in S$. Let q^1 and q^2 be the points where the two legs intersect the boundary of the circle $B(s, \delta_s/5)$. Let

$$c_s = \min\left\{ d(q^i)/3 + \|sq_s^i\| \ ; \ i = 1, 2 \right\}.$$

Then $c_s \leq 2\delta_s/5$ since $d(q^i) \leq \|sq^i\| = \delta_s/5$ and $\|q_s^i\| \leq \|sq_i\| = \delta_s/5$.

Remark: In the proofs of Section 3.3 we will only use the fact that $0 < c_s \leq 2\delta_s/5$, the exact value of c_s does not matter. This will become important in Section 4. The definition of c_s made here results in the weakest sampling condition.

We are now ready to define our potential-function π:

$$\pi(p) = \begin{cases} d(p)/3 & \text{if } p \text{ is in no } B(s, \delta_s). \\ \max\{c_s - \|sp_s\|, d(p)/3\} & \text{if } p \in B(s, \delta_s). \end{cases}$$

Observe that this definition "combines" the two cases discussed in Section 3.1. We use T to denote the set of points $p \in \gamma$ with $\pi(p) = d(p)/3$. Then $q^i \in T$ since $c_s \leq d(q^i)/3 + \|sq_s^i\|$ and hence $c_s - \|sq_s^i\| \leq d(q^i)/3$. Also since $d(p)$ and $\|sp_s\|$ increase as p moves away from s (at least as long as $p \in B^0(s, \delta_s)$), we have $p \in T$ for any curve point p that is not contained in $\cup_{s \in S} B^0(s, \delta_s/5)$. This also implies that π is a continuous function. For a point $p \in B(s, \delta_s/5)$ we have $d(p) \leq \delta_s/5$ and hence $\pi(p) \leq c_s$.

We will frequently use the following simple observation.

Lemma 2 *Let $u \in T \cap B(s, \delta_s)$ for some $s \in S$ and let v be a point on the other leg of s. Then $d(u) \leq \|uv\|$.*

Proof: Assume otherwise, i.e., $\|uv\| < d(u)$. We have $d(u) \leq \|us\|$ and hence $B^0(u, d(u)) \cap \gamma$ consists of at least two components, one containing u and one containing v. ∎

3.3 The Reconstruction

We consider the following family Γ of subcurves:

1. $B^0(p, d(p)) \cap T \cap \gamma$ for all $p \in T$.

2. $B^0(s, \delta_s) \cap \gamma$ for all $s \in S$.

We call the subcurves of the first kind regular subcurves and the subcurves of the second kind singular subcurves.

Lemma 3 *The subcurves $\gamma' \in \Gamma$ are connected.*

Proof: This is obvious for singular subcurves. So consider a subcurve $\gamma' = B^0(p, d(p)) \cap T \cap \gamma$ for some $p \in T$. The subcurve $B^0(p, d(p)) \cap \gamma$ is connected by definition. If γ' were not connected, $B^0(p, d(p)) \cap \gamma$ decomposes into three non-trivial segments γ_1, γ_2, and γ_3 with $\gamma_1 \cap T \neq \emptyset$, $\gamma_2 \cap T = \emptyset$ and $\gamma_3 \cap T \neq \emptyset$. This implies that γ_2 passes through a sharp corner, a contradiction to the definition of $d(p)$. ∎

Lemma 4 *Let u and v be adjacent sample points and let $\gamma' \in \Gamma$. If γ' contains u and v then $\gamma[u, v] \subseteq \gamma'$.*

Proof: γ' is connected and hence either $\gamma[u, v] \subseteq \gamma'$ or $\gamma \setminus \gamma[u, v] \subseteq \gamma'$. The latter case is impossible since $\gamma \setminus \gamma[u, v]$ is not connected in the case of an open curve and turns by

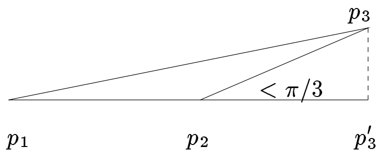

Figure 8: The situation in the proof of Lemma 6. We have $||p_1p_3|| \geq ||p_1p'_3|| > ||p_1p_2|| + ||p_2p_3|| \cos \pi/3$.

more than π by our second sampling condition in the case of an open curve. However γ' turns by less than π by the definition of Γ. ∎

For open curves Lemma4 holds true without the second sampling condition. Since curves $\gamma' \in \Gamma$ turn by less than π, the second sampling condition is implied by the first for open curves. We will next verify the properties (P1) and (P2).

Lemma 5 (Property (P2)) *Let $e = pq$ be an edge with non-positive modified cost. Then there is a subcurve $\gamma' \in \Gamma$ containing p and q.*

Proof: We have $||pq|| \leq \pi(p) + \pi(q)$ by assumption. If $p \in T$ and $q \in T$, $\pi(p) = d(p)/3$ and $\pi(q) = d(q)/3$ and hence $||pq|| \leq (d(p) + d(q))/3 \leq 2\max(d(p), d(q))/3$. Thus $\{p, q\} \subseteq B^0(x, d(x))$ for one of the endpoints x of e.

If one of the endpoints does not belong to T, say $p \notin T$, then $p \in B^0(s, \delta_s/5)$ for some sharp corner. If $q \in B^0(s, \delta_s)$ we are done. So assume otherwise. We have $\pi(p) \leq 2\delta_s/5$ and $||pq|| \geq 4\delta_s/5$. If $q \in T$ then $\pi(q) = d(q)/3 \leq ||sq||/3 \leq (||pq|| + \delta_s/5)/3$ and hence $||pq|| \leq (||pq|| + \delta_s/5)/3 + 2\delta_s/5$ or $2||pq||/3 \leq 7\delta_s/15$ or $||pq|| \leq 21\delta_s/30$, a contradiction to $||pq|| \geq 4\delta_s/5$. If $q \notin T$ then $q \in B^0(t, \delta_t/5)$ for some sharp corner t different from s and hence $||pq|| > 4(\delta_s + \delta_t)/5$. But $\pi(p) \leq 2\delta_s/5$ and $\pi(q) \leq 2\delta_t/5$, a contradiction. ∎

Lemma 6 (Property (P1) for regular regions) *MST_π coincides with the polygonal reconstruction for regular subcurves.*

Proof: Let p_1, p_2, and p_3 be three curve points on a regular subcurve with $p_1 < p_2 < p_3$. We claim that $c_\pi(p_1, p_2) < c_\pi(p_1, p_3)$. Observe that the claim implies directly that Prim's minimum spanning tree algorithm finds the polygonal reconstruction. The claim is equivalent to $||p_1p_2|| - d(p_2)/3 < ||p_1p_3|| - d(p_3)/3$. Since the points are contained in a regular subcurve, we have (see Figure 8) $\angle(\overrightarrow{p_1p_2}, \overrightarrow{p_2p_3}) < \pi/3$ and hence $||p_1p_3|| > ||p_1p_2|| + ||p_2p_3|| \cos \pi/3 = ||p_1p_2|| + ||p_2p_3||/2 \geq ||p_1p_2|| + ||p_2p_3||/3$. The claim now follows from the observation that $d(p_3) \leq d(p_2) + ||p_2p_3||$. ∎

Lemma 7 (Property (P1) for singular regions) *MST_π coincides with the polygonal reconstruction for singular subcurves.*

Proof: Let γ_1 and γ_2 be the two legs of the sharp corner. We identify s with the origin and the angular bisector of the two tangents with the positive x-axis. As a consequence we may identify the projection p_s of a curve point p on either leg onot the x-axis with a non-negative real, and we may view the two legs as monotone functions. We fix q on γ_1 and show:

1. For $p \in \gamma_1 : f_q(p_s) := \|pq\| - \pi(p) - \pi(q)$ is strictly decreasing for $p_s \in [0, q_s]$ and strictly increasing for $p_s \geq q_s$.
2. For $p \in \gamma_2$: Either $g_q(p_s) := \|pq\| - \pi(p) - \pi(q)$ is strictly increasing or $g_q(p_s)$ is positive.
3. For any $p \in \gamma_1$ with $p_s < q_s$ and any $r \in \gamma_2$ we have $f_q(p_s) < g_q(r_s)$ or $g_q(r_s)$ is positive.

The three items above together with the sampling condition that all edges in the polygonal reconstruction have non-positive modified cost imply that MST_π is the polygonal reconstruction. Consider Prim's algorithm starting at the first sample point (= furthest from s) of the first leg and assume inductively that it has constructed an initial segment of the polygonal reconstruction. Let q be the last node of the path P constructed so far and let p be the successor of q in the polygonal reconstruction. We need to show that p has smaller modified distance to P than any other unreached node and that q is the node in P which has smallest modified distance to p. We distinguish three cases.

If P does not yet contain all nodes of the first leg, p is the first unreached node on the first leg. The first item and the third item guarantee that p is the closest unreached node to any node in P and the first item guarantees that q is the closest reached node to p. Thus the edge pq will be added to the minimum spanning tree.

Consider next the case that P contains already all sample points in the first leg but not yet any sample point in the second leg. Then p is the first (= closest to s) sample point on the second leg. The modified distance between p and q is negative by the sampling condition. The second item (with the roles of the first and second leg interchanged) guarantees that for every point r on the second leg, q is the nearest point on the first leg (if there is any with negative modified distance). The second item guarantees also that p is the nearest unreached point to q. Thus the minimum spanning tree will be extended by the edge pq.

Finally assume that P contains all sample points in the first leg and at least one sample point in the second leg. Let p be the first (= nearest to s) unreached sample point on the second leg. The first item guarantees that the minimum spanning tree is extended by the edge pq.

It remains to prove the three claims. We do so in turn.

We start with the first item. Assume that we move p away from q such that p_s changes by some $\delta > 0$, i.e., if $p_s \geq q_s$ then p_s increases by δ and if $p_s < q_s$ then p_s decreases by δ. The Euclidean distance of p and q grows by more than δ since each leg is a monotone function. On the other hand, $\pi(p)$ changes by at most δ. This can be seen as follows. View π as a function of p_s. It has a derivative except when p lies on the boundary of T. For p in the boundary of T, we have a left and right derivative. The derivate is 1 when $p \notin T$ and the derivative is less than 1 when $p \in T$ (observe that p moves by at most

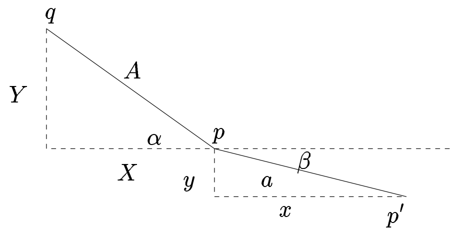

Figure 9: The notation used in the proof of the third item.

$(4/3)m_s x < 3x$ when p_s changes by x and hence the derivative of $d(p)/3$ with respect to p_s is less than one).

We turn to the second item. We consider first the case that $p \notin T$. Let $r \notin T$ with $\delta = r_s - p_s > 0$. We have $\pi(r) = \pi(p) - \delta$ and $\|qr\| > \|qp\| - \delta$ (since the distance along b_s decreases by δ and the distance orthogonal to b_s increases). Thus $g_q(p_s) < g_q(r_s)$.

Assume next that $p \in T$. Then $\pi(p) = d(p)/3 \geq c_s - p_s$. We have $d(p) \leq \|pq\|$ and $d(q) \leq \|pq\|$ by Lemma 2. We next distinguish cases according to whether $q \in T$ or $q \notin T$. If $q \in T$ then $g_q(p_s) = \|pq\| - d(q)/3 - d(p)/3 \geq \|pq\|/3 > 0$. If $q \notin T$ and $p_s \leq q_s$ then $\pi(q) = c_s - q_s \leq d(p)/3 + p_s - q_s \leq d(p)/3$. And hence $g_q(p_s) = \|pq\| - d(p)/3 - \pi(q) \geq \|pq\|/3 > 0$.

This leaves the case that $p \in T$, $q \notin T$ and $p_s > q_s$. If $g_q(p_s) > 0$ there is nothing to show. So assume $0 \geq g_q(p_s) = \|pq\| - (c_s - q_s) - d(p)/3$. Since $p \in T$ we have $d(p)/3 \geq c_s - p_s$ and hence $c_s \leq d(p)/3 + p_s$. Thus $0 \geq \|pq\| - (p_s - q_s) - 2d(p)/3 \geq \|pq\|/3 - (p_s - q_s)$ and hence $p_s - q_s \geq \|pq\|/3$. Let $A = \|pq\|$, let $Y = \|qq_s\| + \|pp_s\|$ (the distance of q and p orthogonal to b_s), and let $X = p_s - q_s$. Then $A^2 = X^2 + Y^2$, $X \geq A/3$, and $Y \geq (3m_s/4)X$. Consider now a point $p' \in \gamma_2 \cap B^0(s, \delta_s)$ and further from s than p. Let $a = \|pp'\|$, $x = p'_s - p_s$, and $y = \|p'p'_s\| - \|pp_s\|$, see Figure 9. We claim that $\|qp'\| - \|qp\| > a/3$. For this it suffices to show that the cosine of the angle φ formed by the vectors $\overrightarrow{qp}$ and $\overrightarrow{pp'}$ is more than 1/3. Let α and β be the angles between a horizontal line and the vectors $\overrightarrow{qp}$ and $\overrightarrow{pp'}$, respectively. Then $\varphi = |\beta - \alpha|$, $3m_s/4 \leq \tan\beta \leq 4m_s/3$, since $(3m_s/4)x \leq y \leq (4m_s/3)x$, $\tan\alpha \geq 3m_s/4$, since $Y \geq (3m_s/4)X$, and $\cos\alpha \geq 1/3$ since $X \geq A/3$. If $\alpha \geq \beta$ then $\cos\varphi = \cos(\alpha - \beta) > \cos\alpha \geq 1/3$ and if $\alpha < \beta$ then[3] $\varphi = \beta - \alpha \leq \arctan 4m_s/3 - \arctan 3m_s/4 = \arctan 7m_s/12(1+m_s^2) < \arctan m_s/(1+m_s^2) < \arctan 1/2$. Thus $\alpha - \beta = \arctan 1/2 < 45°$ (since $\tan 45° = 1$) and hence $\cos(\beta - \alpha) > 1/3$ certainly.

The third item follows from the other two items and the fact that $f_q(0) = g_q(0)$. ∎

4 Closed Curves

We extend the result to closed curves in two steps. In Section 4.2 we we alter the potential function to π' such that the two longest edges of the polygonal reconstruction have the same modified length and such that the minimum spanning trees are precisely the

[3]We use the identity $\arctan x - \arctan y = \arctan((x-y)/(1+xy))$, the fact that $m/(1+m^2) \leq 1/2$ for all m and the fact that the arctan function is strictly increasing.

polygonal reconstruction minus one of the edges of maximal modified cost. In Section 4.1 we show that the preceding sentence implies our main theorem for closed curves.

4.1 The Subtour LP

Assume that the potential function π' has been constructed. The Subtour-LP for the traveling salesman problem can be formulated as follows:

$$\begin{array}{rcll} & & \text{minimize } \sum_{u,v\in V} c_{uv}x_{uv} & \\ & & \text{subject to} & \\ \sum_{v\in V} x_{uv} & = & 2 & \text{for } u \in V \\ \sum_{u\in V',v\in V'} x_{uv} & \leq & |V'|-1 & \text{for } V' \subseteq V,\, \emptyset \neq V' \neq V \\ \sum_{u,v\in V} x_{uv} & = & |V| & \\ & & 0 \leq x_{uv} \leq 1 & \end{array}$$

The last equality is redundant but helpful for our Lagrangian-relaxation. The length of the polygonal reconstruction is an upper bound for the objective value of the Subtour LP. We relax the first set of equalities to the objective function and obtain the following problem with the same objective value [CCPS98, pages 258–260]:

$$\begin{array}{rcll} & & \text{maximize } 2\sum_{u\in V} \pi(u) + \min \sum_{u,v\in V} c_\pi(uv)x_{uv} & \\ & & \text{subject to} & \\ \sum_{u\in V',v\in V'} x_{uv} & \leq & |V'|-1 & \text{for } V' \subseteq V,\, \emptyset \neq V' \neq V \\ \sum_{u,v\in V} x_{uv} & = & |V| & \\ & & 0 \leq x_{uv} \leq 1 & \end{array}$$

In this formulation the maximization is over all choices of π and for fixed π the inner minimization is over the choices for the x_{uv}. We will show that for $\pi = \pi'$ the polygonal reconstruction is the unique optimal solution for the minimization problem and hence the objective value of the maximization problem is at least the length of the polygonal reconstruction. It cannot be larger and hence the objective value of the maximization problem is equal to the length of the reconstruction. This proves that the polygonal reconstruction is an optimal solution to the Subtour LP. We still need to argue uniqueness. Assume that there is another optimal solution for the subtour LP. Since the solution satiesfies the degree constraints it will give the same value to the inner minimization problem as the polygonal reconstruction and hence the polygonal reconstruction is not the unique optimum solution to the inner minimization problem.

It remains to prove that for $\pi = \pi'$, the polygonal reconstruction is the unique optimal solution for the inner minimization problem. Orient γ arbitrarily, let e_1 and e_2 be edges in the polygonal reconstruction which have maximal modified cost, and let u and v be the starting nodes of e_1 and e_2, respectively, and let R_1 and R_2 be the sample points from u to v (both inclusive) resp. from v to u (both inclusive) wrt to the order of the points along the curve; see Figure 10. Then $R_1 \cup R_2 = V$ and $|R_1 \cap R_2| = 2$. Let E_i, $i = 1, 2$, be the set of edges having both endpoints in R_i and let C be the remaining set of edges. Then $e_i \in E_i$ and any edge $e \in C$ has a modified cost larger than e_1 (and hence e_2). Otherwise there would be a minimum spanning trees that is not contained in the polygonal reconstruction. In an optimal solution to the inner LP the total weight of the edges in E_i is $|R_i| - 1 - o_i$ for some $o_i \geq 0$, $i = 1, 2$, and the total weight of the edges in C is $o_1 + o_2$. Thus the inner LP is equivalent to:

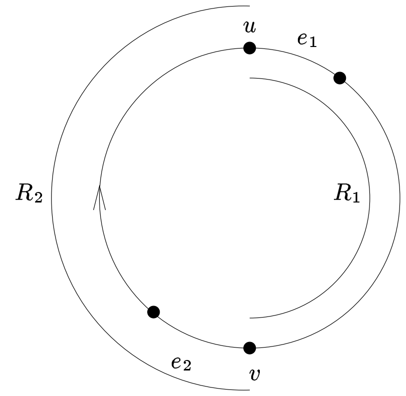

Figure 10: The notation used in the reformulations of the Subtour LP. We have $e \in R_i$.

$$
\begin{array}{c}
\text{minimize over all non-negative values of } o_1 \text{ and } o_2\\
\min \sum_{u,v\in V} c_{\pi'}(uv)x_{uv}\\
\text{subject to}\\
\begin{array}{rcll}
\sum_{u\in R_1', v\in R_1'} x_{uv} & \leq & |R_1'|-1 & \text{for all } R_1' \subseteq R_1, \text{ with } R' \neq \emptyset\\
\sum_{u,v\in R_1} x_{uv} & = & |R_1'|-1-o_1 & \\
\sum_{u\in R_2', v\in R_2'} x_{uv} & \leq & |R_2'|-1 & \text{for all } R_2' \subseteq R_2, \text{ with } R' \neq \emptyset\\
\sum_{u,v\in R_2} x_{uv} & = & |R_2|-1-o_2 & \\
\sum_{uv\in C} x_{uv} & = & o_1+o_2 &
\end{array}\\
0 \leq x_{uv} \leq 1
\end{array}
$$

Consider the inner minimization problem for fixed values of o_1 and o_2. The first two lines describe a partial minimum spanning tree for R_1 and the next two lines a partial minimum spanning tree for R_2. More precisely, the system is minimized, if one chooses the $\lfloor |R_1|-1-o_1 \rfloor$ shortest edges of the minimum spanning tree and fills the fractional part with the next edge[4]. The same is true for R_2. The LP takes its minimum for $o_1 = o_2 = 0$

[4]For $o_1 = 0$ this is the characterization of minimum spanning trees by a linear program, see [CCPS98, Theorem2.8]. The proof extends to non-zero values of o_1. We sketch the proof for the sake of completeness. The dual linear program has an unconstrained variable y_{R_1} and a non-positive variable $y_{R'}$ for every proper subset of R_1. It reads $\max(|R_1|-1)y_{R_1} + \sum_{R';\emptyset\neq R'\neq R_1}(|R'|-1)y_{R'}$ subject to $y_{R_1} + \sum_{R';e\in\gamma(R')} y_{R'} \leq c_e$ for every edge e and $y_{R'} \leq 0$ for every R', $\emptyset \neq R' \neq R_1$. In this footnote we use $\gamma(R)$ to denote all edges having both endpoints in R. We use Kruskal's algorithm to construct a primal and dual solution simultaneously. The dual solution will always be feasible, the primal solution will always satisfy $x(\gamma(R')) \leq |R'|-1$ for all non-empty R' (but will satisfy $x(\gamma(R_1)) = n-1-o_1$ only in the last step), and we will always have complementary slackness.

We start with the empty spanning forest, i.e., $x_e = 0$ for all e, and all dual variables equal to zero. We declare all singleton sets R' active. We then increase y_{R_1} and decrease $y_{R'}$ for all active R' at the same rate until the first constraint $y_{R_1} + \sum_{e;e\in\gamma(R)} y_R \leq c_e$ becomes tight. This will be for the minimum weight edge (let us assume for simplicity that all weights are distinct); observe that the $y_{R'}$ for singletons sets R' are irrelevant. We add e to the spanning forest, i.e., we change x_e to one. Observe that this preserves complementary slackness. The addition of e to the spanning forest combines two active sets R_1' and R_2' into a set R'. We declare R_1' and R_2' inactive and R' active. Observe that R' corresponds to a subtree of the spanning tree and hence the primal constraint $x(\gamma(R')) = |R'|-1$ is tight. Thus $y_{R'}$ is a dual variable which complementary slackness allows to become non-zero. We proceed, i.e., we increase y_{R_1} and decrease $y_{R'}$ for every active set R' at the same rate. Observe what happens. For every edge f contained in an active set R' the left hand side of $y_{R_1} + \sum_{e;e\in\gamma(R)} y_R \leq c_f$ will not change, because the increase of y_{R_1} is balanced by the decrease of $y_{R'}$. For any edge f connecting two active sets the left since any edge in C has larger cost than any edge in the minimum spanning tree. For $o_1 = o_2 = 0$, the system describes the minimum spanning trees for R_1 and R_2. Thus the polygonal reconstruction is the unique optimal solution of the inner minimization problem and hence of the Subtour LP.

4.2 The Modified Potential Function

We show how to alter the potential such that the two longest edges of the polygonal reconstruction have the same length and such that all minimum spanning trees for V remain part of the polygonal reconstruction.

Let e_{max} be the edge of the polygonal reconstruction with largest modified cost. We claim that one of the following cases arises:

(1) There is a sharp corner such that both endpoints of e_{max} are outside the ball $B(s, \delta_s/5)$.

(2) There is a point $v \in \gamma$ such that $\|vu\|/2 \geq \pi(u)$ for both endpoints u of e_{max} and such that v does not lie in the $B^0(s, \delta_s)$ ball of any sharp corner s.

The first case certainly arises when there are at least three sharp corners. Assume that the first case does not hold. We make a further case distinction: Either both endpoints of e_{max} lie in T (this will be certainly the case when there is no sharp corner) or some endpoint of e_{max} lies in $B(s, \delta_s/5) \setminus T$ for some sharp corner s. In the former case the curve must leave the union of the $B(u, d(u))$ balls of the two endpoints u of e_{max} (because the curve can turn by at most $2\pi/3$ within the union of these balls) and any curve point v outside the two balls (since $\|uv\| \geq d(u) = 3\pi(u)$ for any of the endpoints) and outside $\cup_{s\in S} B^0(s, \delta_s)$ will work; v exists since the regions $(B^0(s, \delta_s))_{s\in S}$ are pairwise disjoint and since γ turns by less than π in any such region. So assume that some endpoint of e_{max} lies in $B(s, \delta_s/5) \setminus T$ of some sharp corner s. Consider the leg of s that does not contain the other endpoint of e_{max}. It contains no endpoint of e_{max} that lies in T. Let v be the point on this leg of s with distance δ_s from s and let u be an endpoint of e_{max}. If $u \in B(s, \delta_s/5)$ then $\|uv\| \geq 4\delta_s/5$ and $\pi(u) \leq \max(c_s, d(u)/3) \leq 2\delta_s/5$ and if $u \notin B(s, \delta_s/5)$ then $u \in T$ and hence u does not lie on the same leg as v does. Then $d(u) \leq \|uv\|$ by Lemma 2 and we are done. In either case we have shown that one of items above holds. We also need the following Lemma.

Lemma 8 *Let s be a sharp corner and let x and y be adjacent sample points such that $s \in \gamma[x, y]$. Then $x, y \in B(s, 3\delta_s/3)$ and either x or y lies in $B(s, \delta_s/5) \setminus T$.*

Proof: Since the regions $B^0(p, d(p)) \cap T$ do not contain any sharp corner and since the regions $(B^0(t, \delta_t))_{t\in S}$ are pairwise disjoint, we have $x, y \in B^0(s, \delta_s)$ by Lemmas 4 and 5. Assume $x \notin B(s, 3\delta_s/5)$. Then

$$c_\pi(xy) \geq c_\pi(sx) \geq \|sx\| - 2\delta_s/5 - \|sx\|/3 \geq 2\|sx\|/3 - 2\delta_s/5 > 2\delta_s/5 - 2\delta_s/5 = 0,$$

---

hand side will increase. When the next edge becomes tight, we add it to the spanning tree, ... We finish when we adding a partial edge makes the constraint $x(\gamma(R_1)) = n - 1 - o_1$ tight.

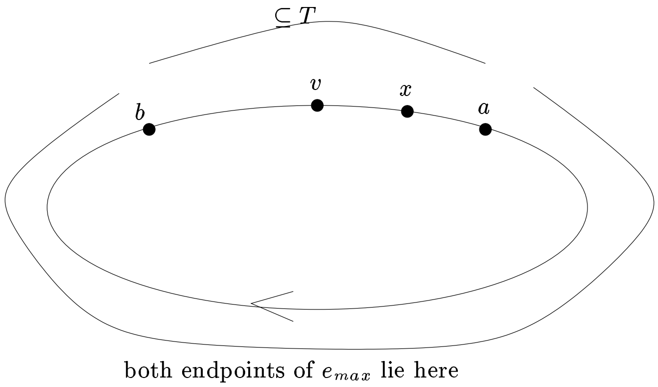

Figure 11: The construction of the potential function π'.

where the first inequality follows from the third claim in the proof of Lemma 7. Assume next that $x, y \in T$. Lemma 2 implies $d(x), d(y) \leq \|xy\|$ and hence $c_\pi(x, y) \geq \|xy\|/3 > 0$. ∎

Assume first that there is a sharp corner s such that both endpoints of e_{max} lie outside the $B(s, \delta_s/5)$ ball. We decrease c_s continuously down to zero. For $c_s = 0$, we have $\pi(x) = d(p)/3$ for all $x \in B(s, \delta_s)$ and hence the edge xy in the reconstruction that connects the two legs of s (it exists by Lemma 8) has positive modified cost: $c_{\pi'}(x, y) = \|xy\| - \pi'(x) - \pi'(y) \geq \|xy\| - \|xy\|/3 - \|xy\|/3 > 0$ since $d(x), d(y) \leq \|xy\|$ and hence $c_\pi(x, y) \geq \|xy\|/3 > 0$ by Lemma 2. The modified cost of the edge e_{max} is not affected by the change of c_s. Thus there must be a value of c_s for which the two largest modified costs in the reconstruction are the same. All edges in the reconstruction still have non-positive modified cost (because e_{max} has) and the minimum spanning tree with respect to $c_{\pi'}$ remains the polygonal reconstruction by the remark in Section 3.2 that only the fact that $0 < c_s \leq 2\delta_s/5$ is used in our proofs.

We come to the case that there is a point $v \in \gamma$ such that $\pi(u) \leq \|vu\|/2$ for both endpoints u of e_{max} and such that v lies outside the $B^0(s, \delta_s)$ balls of all sharp corners. We split the curve at v and orient the resulting open curve arbitrarily. For any $l \in \mathbb{R}$, $0 \leq l \leq \pi(v)$, we define a as the first point with $\pi(a) = \|va\|/2 + l$ and b as the last point with $\pi(b) = \|vb\|/2 + l$ and define a potential π' by

$$\pi'(p) = \begin{cases} \pi(p) & \text{if } a \leq p \leq b \\ \|vp\|/2 + l & \text{otherwise} \end{cases}$$

Observe that a, b, and π' depend on l. For simplicity of notation we do not make this dependence explicit in the notation. Figure 11 illustrates the definition of π'.

We need to argue that a and b exist for all choices of l, that the cost of e_{max} does not depend on l, that there is a choice of l for which the two largest modified costs are the same, and that every minimum spanning tree with respect to $c_{\pi'}$ uses only edges of the polygonal reconstruction.

For $l = \pi(v)$ we have $a = b = v$. For $0 \leq l < \pi(v)$ we have $\pi(v) > l + \|vv\|/2$ and $\pi(u) \leq \|uv\|/2 \leq l + \|uv\|/2$ for either endpoint of e_{max}. The existence of a and b follows from the continuity of π. Also $a \leq u \leq b$ for either endpoint u of e_{max} and $\pi'(p) \leq \pi(p)$ for all p. The latter claim follows from $\pi'(v) = l \leq \pi(v)$ and the definition of a and b.

We next show that $a \leq r$, where r is the first point on γ with $\|sr\| = 3\delta_s/5$ for some sharp corner s, and hence $a \in T$. We have $\|rv\| \geq 2\delta_s/5$ and $\pi(r) = d(r)/3 \leq \|sr\|/3 = \delta_s/5 \leq \|rv\|/2 \leq \|rv\|/2+l$. Continuity of π implies that a lies between v and r. Similarly, b lies in T and after the last point r on γ with $\|sr\| = 3\delta_s/5$ for some sharp corner.

We next argue that there is a choice of l for which the two largest modified costs in the reconstruction are the same. From $a \leq u \leq b$ for any endpoint u of e_{max} we conclude that the modified cost of e_{max} does not depend on l. For $l = \pi(v)$, we have $a = b = v$ and hence $\pi = \pi'$. Thus e_{max} has the maximal modified cost $c_{\pi'}$ among all edges in the reconstruction. For $l = 0$, we consider the reconstruction edge between the last and the first sample points x and y, respectively, and show that it has positive modified cost $c_{\pi'}(xy)$. There must be a subcurve γ' containing $\gamma[x, y]$ and hence v. Since $v \notin B(s, \delta_s)$ for any sharp corner, γ' is a regular subcurve and hence $\gamma[x, y] \subseteq T$. We show first that $a < x$ and $y < b$ is impossible. Assume otherwise. Then $\pi(x) \leq \pi(a) + \|ax\|/3 = \|va\|/2 + \|ax\|/3$ and $\pi(y) \leq \pi(b) + \|by\|/3 = \|vb\|/2 + \|by\|/3$ and $\|xy\| > \|xa\|/2 + \|ab\| + \|by\|/2$ since γ' turns be less than $\pi/3$. Thus $c_\pi(x, y) > 0$, a contradiction to our sampling condition. Thus either $v \leq x \leq a$ or $b \leq y \leq v$ or both. We may assume w.l.o.g. that $v \leq x \leq a$. If $y < b$, we have $\pi(y) \leq \pi(b) + \|by\|/3 = \|vb\|/2 + \|by\|/3$, $\|xy\| > \|xv\|/2 + \|vb\| + \|by\|/2$, and $\pi'(x) = \|vx\|/2$ and hence $c_{\pi'}(x, y) > 0$. If $b \leq y$, we have $\|xy\| > \|xv\|/2 + \|vy\|/2$, $\pi'(y) = \|vy\|/2$ and $\pi'(x) = \|vx\|/2$ and hence $c_{\pi'}(x, y) > 0$. In either case we have shown that for $l = 0$ there is an edge in the reconstruction with positive modified cost. Continuity implies that there is a value of l for which the the two largest modified costs in the reconstruction are the same.

We call an edge unaffected if the change of potential does not change the modified cost of the edge and call it affected otherwise. The modified cost of an effected edge is increased. All unaffected edges of the minimum spanning tree MST_π remain in the minimum spanning tree $\mathrm{MST}_{\pi'}$. Assume that there is an edge $e = pq \in \mathrm{MST}_{\pi'}$ which does not belong to the polygonal reconstruction. Then e must be an affected edge and hence the potential of one its endpoints, say p, was decreased. Also since there is a spanning tree of negative cost edges, e.g., the polygonal reconstruction, e has negative modified cost and hence there is a subcurve γ' containing both of its endpoints. At least one of p and q must lie in $\gamma[b, v] \cup \gamma[v, a]$, say p. Thus $p \in T$. Since pq is not an edge of the polygonal reconstruction there must be a sample point u in $\gamma[p, q]$.

If $\gamma[p, q] \subseteq B(r, d(r)) \cap T$ for some $r \in \gamma$, we have $\angle(\overrightarrow{pu}, \overrightarrow{uq}) < \pi/3$ and hence $\|pq\| > \|pu\| + \|uq\|/2$ and $\pi' q \leq \pi' u + \|uq\|/2$. Thus $c_{\pi'}(pq) > c_{\pi'}(pu)$. Similarly, $c_{\pi'}(pq) > c_{\pi'}(qu)$. Thus pq cannot be in the minimum spanning tree.

If $p, q \in B(s, \delta_s)$ for some sharp corner, we distinguish two cases. Assume first that p and q are on the same leg of s. Recall the proof of the first claim in the proof of Lemma 7. It uses only the fact that the derivative of the potential as a function of the projection on the bisector of the sharp corner is at most one. This is still true and hence the proof stays valid. So assume that p and q lie on different legs. We claim, that u lies on the same leg as p. Let r be a sample-point in $B(s, 3\delta_s/5)$ on the same leg as p; r exists by Lemma 8. Then $\pi'(r) = \pi(r)$ and $\pi(q) = \pi'(q)$. Therefore $c_{\pi'}(pq) > c_{\pi'}(pr)$, since this inequality holds for π and since $c_{\pi'}(pq) - c_{\pi'}(pr) = c_\pi(pq) - c_\pi(pr)$. For any reconstruction edge xy with $r \leq x < y \leq q$, we have $c_\pi(xy) \leq c_\pi(rq) < c_\pi(pq) < c_{\pi'}(pq)$ and hence $c_{\pi'}(xy) < c_{\pi'}(pq)$ since $c_\pi(rq) = c_{\pi'}(pq)$.

5 Our Sample Condition and the Local Feature Size

The papers [ABE98, DK99, Gol99, DMR99] investigated the reconstruction problem for smooth curves. A curve is smooth if it twice-differentiable. They expressed the sampling condition in terms of the so-called *local feature size*. The local feature size $f(p)$ at a curve point p is the distance of p from the medial axis of γ. The *medial axis* of a curve is the closure of the set of points in the plane which have at least two nearest (with respect to the Euclidean metric) points on the curve. They required a sampling condition of the form: For any $p \in \gamma$ there must be a sample point $v \in V$ with $\|pv\| \le \epsilon \cdot f(p)$; here ϵ is a parameter which depends on the algorithm. All algorithms require $\epsilon \le 1/3$.

The experimental results in [AMNS99] suggest that the TSP-algorithms works for sparser sample sets than the algorithms mentioned above. We do not know whether this observation is a fact and can prove only a much weaker result.

Lemma 9 *Let $\epsilon = 1/20$. If for any $p \in \gamma$ there must be a sample point $v \in V$ with $\|pv\| \le \epsilon \cdot f(p)$ then V satisfies our sampling condition.*

Proof: Let x and y be two adjacent sample points and let p be the intersection of $\gamma[x, y]$ with the perpendicular bisector of x and y. Then $\|px\| = \|py\| \le \epsilon \cdot f(p)$ and hence $\|xy\| \le 2\epsilon f(p)$. It therefore suffices to show that $\pi(x), \pi(y) \ge \epsilon f(p)$ or $d(x), d(y) \ge 3\epsilon f(p)$.

The ball $B = B(p, f(p))$ contains no medial axis point in its interior and hence γ intersects B^0 in a single component (by Lemma 1 of [ABE98]) and γ turns by less than $\pi/3$ in $B^0(p, f(p)/5)$. Assume otherwise. Let u and v be the intersections of γ with the boundary of $B^0(p, f(p)/5)$. The tangents at u and v form an angle of $\pi/3$. Consider a disk D centered at a point c of the perpendicular bisector of u and v and having u and v in its boundary and let m be the midpoint of u and v, see Figure 12. The perpendiculars to the tangents in u and v intersect the bisector in the points a and b with a being closer to m than b. For c between m and a, γ enters D a x and leaves it at y. For c beyond b, γ leaves D at x and enters it at y. For c between a and b, γ leaves D at x and at y and hence intersects the boundary of D at least three times and thus (again by Lemma 1 of [ABE98]) there is a medial axis point in the disc centered at a and passing through u and v.

We have $\alpha + \beta = \pi/3$ and $\alpha \ge \beta$. Thus $\alpha \ge \pi/6$ and hence $(\|uv\|/2)/\|ua\| = \sin\alpha \ge 1/2$ or $\|ua\| \le \|uv\| \le 2f(p)/5$. Thus $\|pa\| \le \|pu\| + \|ua\| \le 3f(p)/5$. We conclude that $B(a, \|ua\|)$ is contained in $B(p, f(p))$ and hence the latter ball contains a point of the medial axis.

We have now shown that $d(p) \ge f(p)/5$ and hence $d(x) \ge f(p)/5 - \|xp\| = f(p)/5 - \epsilon f(p) = 3\epsilon f(p)$ for $\epsilon \le 1/20$. ∎

6 Necklace Tours

We have shown that curve reconstruction gives rise to a polynomially solvable case of the Euclidian traveling salesman problem. In this section we relate our results to known solvable case, the so-called *necklace tours*. Let V be a set of points in the plane and assume that there is a set of disks centered at the points in V such that each disk intersects with

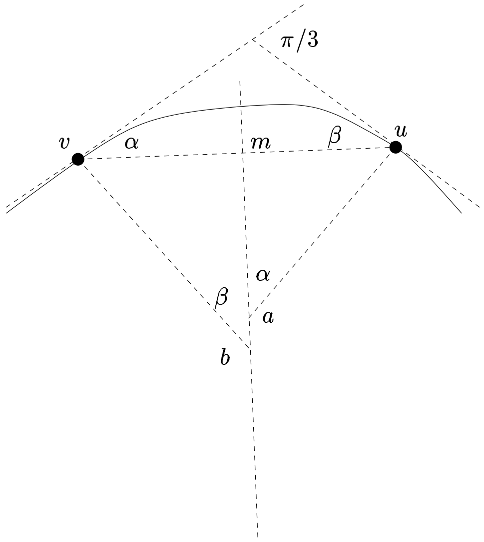

Figure 12: The bisector of u and v is shown dashed and m is the midpoint of u and v. The perpendiculars to the tangents in u and v intersect the bisector in the points a and b with a being closer to m than b.

exactly two other disks and such that the intersection graph of the disks is connected. The intersection graph of the disks defines a tour on V; the two neighbors of a point v correspond to the two disks that intersect the disk associated with v. The tour is called a necklace tour and is known to be an optimal traveling salesman tour of V; see [BDvD$^+$98].

We cannot claim the necklace tours are a special case of our result, since there is no curve underlying a necklace tour. The optimality proof for necklace tours is a special case of our argument. We simply define the potential of any point v as the radius of the disk associated with v. Then exactly the edges in the tour have non-positive cost. Any tour which uses an edge outside the necklace tour must include edges of positive cost and can include only a subset of the edges in the necklace tour. This implies that the necklace tour is optimal.

Figure 13 shows an example of a TSP problem which is covered by Theorem 2 and whose solution is not a necklace tour.

7 Conclusions

We have shown that TSP-based curve reconstruction permits a polynomial time algorithm. In the companion paper [AMNS99] we show that TSP-based curve reconstruction is reasonably efficient and works for smaller sampling density than other reconstruction algorithms. TSP-based curve reconstruction reconstructs a single curve. The generalization to collections of open and closed curves is open. Surface reconstruction is the three-dimensional analogue of curve reconstruction. There are no surface reconstruction algorithms that can handle non-smooth surfaces. It is open whether global minimization can also be applied to surface reconstruction.

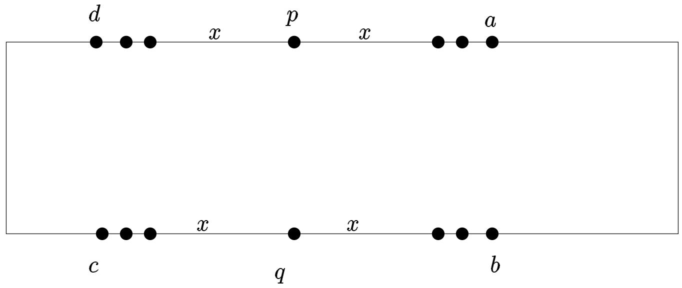

Figure 13: The underlying curve is a rectangle and the parts of the curve between a and b, c and d, respectively, are densely sampled. Let D be the distance between the horizontal lines and let x be the distance from p and q to their neighbors in the tour. In a necklace tour the disks centered at p and q must have radius $(1-\epsilon)x$ for some ϵ and hence the tour is not a necklace tour if $2(1-\epsilon)x > D$. In our terminology we have $d(p) = d(q) = d(z) = D$ for p, q and their neighbors z on the tour provided the horizontal edges of the rectangle are long enough. Thus our sampling condition is satisfied if $x \leq 2D/3$ and hence the tour will be constructed by the subtour LP.

References

[ABE98] N. Amenta, M. Bern, and D. Eppstein. The crust and the β-skeleton: Combinatorial curve reconstruction. *Graphical Models and Image Processing*, pages 125–135, 1998.

[AMNS99] E. Althaus, K. Mehlhorn, S. Näher, and S. Schirra. Experiments on curve reconstruction. unpublished, 1999.

[Att97] D. Attali. r-regular shape reconstruction from unorganized points. In *Proceedings of the 13th Annual ACM Symposium on Computational Geometry (SCG'97)*, pages 248–253, 1997.

[BB97] F. Bernardini and C.L. Bajaj. Sampling and reconstructing manifolds using α-shapes. In *Proceedings of the 9th Canadian Conference on Computational Geometry (CCCG'9 7)*, pages 193–198, 1997.

[BDvD+98] R. E. Burkard, V. G. Deineko, R. van Dal, J. A. A. van der Veen, and G. J. Woeginger. Well-solvable special cases of the traveling salesman problem: A survey. *SIAM Review*, 40:496–546, 1998.

[CCPS98] W.J. Cook, W.H. Cunningham, W.R. Pulleyblank, and A. Schrijver. *Combinatorial Optimization.* John Wiley & Sons, Inc, 1998.

[DK99] T.K. Dey and P. Kumar. A simple provable algorithm for curve reconstruction. In *Proceedings of the 10th ACM-SIAM Symposium on Discrete Algorithms (SODA'99)*, pages 893–894, 1999.

[DMR99] T.K. Dey, K. Mehlhorn, and E.A. Ramos. Curve reconstruction: Connecting dots with good reason. In *Proceedings of the 15th Annual ACM Symposium on Computational Geometry (SCG'99)*, pages 197–206, 1999. www.mpi-sb.mpg.de/ mehlhorn/ftp/cure.ps.

[Ede98] H. Edelsbrunner. Shape reconstruction with the Delaunay complex. In *LATIN'98: Theoretical Informatics*, volume 1380 of *Lecture Notes in Computer Science*, pages 119–132, 1998.

[EKS83] H. Edelsbrunner, D.G. Kirkpatrick, and R. Seidel. On the shape of a set of points in the plane. *IEEE Transactions on Information Theory*, 29(4):71–78, 1983.

[FG94] L.H. Figueiredo and J.M. Gomes. Computational morphology of curves. *Visual Computer*, 11:105–112, 1994.

[Gie99] J. Giesen. Curve reconstruction, the TSP, and Menger's theorem on length. In *Proceedings of the 15th Annual ACM Symposium on Computational Geometry (SCG'99)*, pages 207–216, 1999.

[Gol99] C. Gold. Crust and anti-crust: A one-step boundary and skeleton extraction algorithm. In *Proceedings of the 15th Annual ACM Symposium on Computational Geometry (SCG'99)*, pages 189–196, 1999.

[KR85] D.G. Kirkpatrick and J.D. Radke. A framework for computational morphology. *Computational Geometry*, pages 217–248, 1985.

[Sch86] A. Schrijver. *Theory of Linear and Integer Programming.* Wiley, 1986.