TSP-BASED CURVE RECONSTRUCTION IN POLYNOMIAL TIME*

ERNST ALTHAUS† AND KURT MEHLHORN‡

Abstract. An instance of the curve reconstruction problem is a finite sample set V of an unknown collection of curves γ. 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 for single closed curves under otherwise weak assumptions on γ and V; γ must be a single closed curve. We extend his result along several directions.

- we weaken the assumptions on the sample,
- we show that Traveling Salesman based reconstruction also works for single open curves (with and without specified endpoints) and for collections of closed curves,
- we give alternative proofs, 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 set V of an unknown collection of curves γ. The task is to construct a graph G on V so that two points in V are connected by an edge of G iff the points are adjacent on γ. The curve reconstruction problem and the related surface reconstruction problem have 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 1.1 illustrates the curve reconstruction problem.

Many curve reconstruction algorithms have been proposed in the past; we restrict our discussion to algorithms that provably solve the reconstruction problem for a certain class of curves and under certain assumptions on the sample set. The algorithms differ with respect to the following aspects.

- Whether a collection of curves or just a single curve can be handled.
- Whether (collections of) open and closed curves can be handled or only (collections of) closed curves.
- Whether the sampling must be uniform or not. Uniform sampling with density ϵ requires that the sample set V contains at least one point from every curve segment of length ϵ. In non-uniform sampling, the sampling frequency may depend on local properties of the curve, e.g., can be lower in parts of low curvature.
- Whether non-smooth curves can be handled or not. A smooth curve has a tangent everywhere. Figure 2.1 shows non-smooth curves.

For uniformly sampled collections of closed, smooth curves 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 collections of closed smooth curves was first treated successfully by Amenta, Bern and Eppstein [ABE98] and subsequently improved algorithms such as [DK99, Gol99] appeared. Non-uniformly sampled collections of open and closed smooth curves were treated in [DMR00]. All papers mentioned so far require the curves to be 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 curve is less than π, see Figure 2.1. Giesen showed that the Traveling Salesman tour of the sample set V solves the curve reconstruction problem for uniformly sampled benign closed semi-regular curves. More precisely, he showed that for every benign semi-regular closed curve γ there exists a positive ϵ such that the optimal Traveling Salesman tour of V is a polygonal reconstruction of γ provided that for every $x \in \gamma$ there is a $p \in V$ with $\|xp\| \leq \epsilon$, where $\|xy\|$ is the Euclidean distance of the two points x and y. 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 several directions:

- We relate ϵ to local properties of the curve γ and show that the optimal Traveling Salesman tour solves the reconstruction problem even if sampling is non-uniform. For smooth curves our sampling condition is similar to the one used in [ABE98, DK99, Gol99, DMR00].

*This paper appeared in preliminary form in the Proceedings of the 11th ACM-SIAM Symposium Discrete Algorithms, January 2000. Research partially supported by EG-projects ALCOM-IT and GALIA.

†Max-Planck-Institute für Informatik, Im Stadtwald, 66123 Saarbrücken

‡Max-Planck-Institute für Informatik, Im Stadtwald, 66123 Saarbrücken

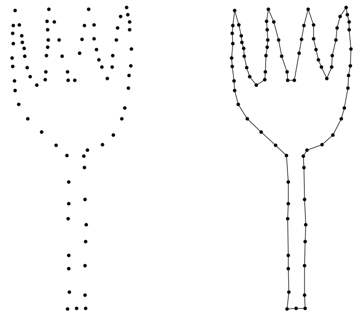

FIG. 1.1. *Part (a) shows a finite set V of points, part (b) shows the Traveling Salesman tour of the points.*

- We show that the Traveling Salesman path is able to reconstruct open curves for a suitable sampling condition. We treat the case of paths with and without specified endpoints.
- We show that the optimal Traveling Salesman tour (path) can be constructed in polynomial time if our sampling condition is satisfied.
- We give a simplified proof that the Traveling Salesman tour (path) solves the curve reconstruction problem.
- We show that an extension of the Traveling Salesman tour algorithm is able to reconstruct non-uniformly sampled collections of closed non-smooth curves.

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 [AMNS00].

A preliminary version of the present paper appeared in SODA 2000. It contained the result for open curves with specified endpoints and for closed curves. It did not contain the result for open curves without specified endpoints and for collections of closed curves. Also the proof of polynomiality was incomplete.

Since the publication of the conference version of this paper, Funke and Ramos [FR01] have presented an algorithm that work for collections of non-smooth curves. The algorithm is based on filtering the Delaunay diagram and can handle collections of open and closed curves; our algorithm can only handle collections of closed curves.

This paper is structured as follows. In Section 2 we give definitions and state our main results. In Sections 3, 4, and 5 we prove our main theorem for open curves with and without specified endpoints and for closed curves, respectively, assuming real arithmetic. The result is extended to finite precision arithmetic in Section 6. In Section 7 we discuss collections of closed curves. For computational purposes it is desirable to restrict the search for the reconstruction to a sparse graph defined on the sample set. In Section 8, we show that the edges of the polygonal reconstruction are in the Delaunay diagram for a slightly strengthened sampling condition. In Section 10 we relate our sampling condition to the sampling conditions used by other papers, and in Section 11 we relate our result to so-called necklace tours. Necklace tours are a polynomially solvable case of the Traveling Salesman problem. Curve reconstruction problems are allowed to "invent" curves, if the input does not satisfy the required sampling condition. This issue is discussed in Section 9. Finally, in Section 12 we offer conclusions and state open problems.

2. Definitions and Statements of Results. An open *curve* is given by an *embedding* $\gamma : [0,1] \to \mathbb{R}^2$ and a closed curve is given by an embedding $\gamma : S^2 \to \mathbb{R}^2$, where S^2 is the unit circle.

DEFINITION 2.1 ([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 $\mathrm{t}_l(\gamma(t_0))$ $(\mathrm{t}_r(\gamma(t_0)))$ 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 $\tau(\xi_n)$

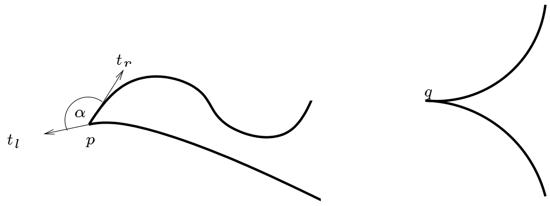

FIG. 2.1. *Two semi-regular curves, one benign and one not. In the left curve, the two tangents t_l and t_r at a point p of the curve are shown. The turning angle at p is α. In the curve on the right the turning angle at q is π as the left and right tangents at q point to opposite directions.*

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 2.1 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 to the same direction and the angle is π if the vectors point to opposite directions. 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.

A Traveling Salesman path (tour) for a set V of points is a path (cycle) passing through all points in V. An optimal Traveling Salesman path (tour) is a Traveling Salesman path (tour) of shortest length. An optimal Traveling Salesman path with specified endpoints a and b, where $a \in V$ and $b \in V$ is a shortest Traveling Salesman path among the paths with endpoints a and b. We can now state Giesen's result.

THEOREM 2.2 ([Gie99]). *For every benign semi-regular closed curve γ there exists an $\epsilon > 0$ with the following property: If V is a finite sample set of γ so that for every $x \in \gamma$ there is a $p \in V$ with $\|pv\| \leq \epsilon$, the optimal Traveling Salesman tour is a polygonal reconstruction of γ.*

The construction of optimal 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 specified 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:

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

We refer to this program as the *Subtour-ILP for the Traveling Salesman Path Problem with specified endpoints*. Slight variations of this ILP characterize paths with unspecified endpoints (see Section 4) and tours (see Section 5). 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 $0 \leq x_{uv} \leq 1$ gives the *Subtour-LP for the Traveling Salesman Problem with specified endpoints*. The objective function for both programs is $\sum_{uv} \|uv\| x_{uv}$, i.e., the total Euclidean length of the edges selected.

In general, the optimal solution of the Subtour-LP is fractional. Our main result states the optimal solution of the Subtour-LP is integral, whenever V is a sufficiently dense sample of a benign semi-regular curve. The reader might be interested to know that we discovered this fact in our experiments on curve reconstruction. We had implemented a branch-and-cut algorithm based on the Subtour-LP (the algorithm solved the Subtour-LP and was then supposed to branch on a fractional variable) and observed that the algorithm never branched. After seeing this behavior in a large number of examples, we formulated it as a conjecture and set out to prove it.

THEOREM 2.3 (Main Theorem). *Let γ be an closed benign semi-regular curve, let V be a finite set of samples of γ. If V satisfies the sampling condition given below then:*

- *The optimal Traveling Salesman tour of V is a polygonal reconstruction of γ.*
- *The Subtour-LP for Traveling Salesman tours has an optimal integral solution and this solution is unique.*

In the case of an open curve let a and b be the first and last sample point, respectively (in the order on γ). The statements above hold true for the optimal Traveling Salesman path with endpoints a and b and the Subtour-LP for Traveling Salesman paths with specified endpoints and also for the Traveling Salesman path with unspecified endpoints and for the Subtour-LP for Traveling Salesman paths with unspecified endpoints. The latter result required a strengthened sampling condition.

We will prove our main theorem for open curves with specified and unspecified endpoints in Sections 3 and 4, respectively. The proof for closed curves follows in Section 5.

Our main theorem suggests a reconstruction algorithm for benign semi-regular curves: *Solve the Subtour-LP. If the optimal solution is integral, output it.* We briefly discuss two strategies for solving the Subtour-LP.

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 X^* to the current LP satisfies all subtour elimination constraints and, if not, one adds a violated subtour elimination constraint to the LP. The check for a violated subtour elimination constraint is tantamount to a min-cut problem. One assigns capacity x_e^* to edge e for every edge e and computes a minimum (a,b)-cut. If the cut has value less than two, a violated inequality has been found. If the cut has value two, all cut constraints are satisfied. We use the simplex-based strategy in our in our experiments on curve reconstruction.

The Ellipsoid method, see [Sch86], solves the Subtour-LP in time polynomial in the size of the bit representations of the coefficients of the cost function. Distance values are, in general, non-rational numbers and hence the Ellipsoid method is not directly applicable in our setting. In Section 6, we extend our results to the situation where the position of the points and the distances between points are only approximately known and show how to obtain a polynomial time algorithm.

3. Open Curves. We assume that our open curve is oriented and write $p < q$ if p precedes q on γ. We use $B(p,r)$ and $B^0(p,r)$ to denote the closed and open ball with center p and radius r, respectively.

3.1. The Held-Karp Bound. Our proof of Theorem 2.3 exploits the connection between the Subtour-LP and what is known as the Held-Karp bound. The purpose of this section is to review the relevant facts about the Held-Karp bound.

Let $G = (V,E)$ be an undirected graph, let a and b be two designated vertices of G, and let c be an arbitrary cost function on the edges of G. A function $\mu : V \to \mathbb{R}$ is called a *potential function*. It gives rise to a modified distance function c_μ via $c_\mu(u,v) = c(u,v) - \mu(u) - \mu(v)$. Consider now any Traveling Salesman path T with endpoints a and b. Its costs under c and c_μ are related by $c_\mu(T) = c(T) - 2\sum_{v\in V}\mu(v) + \mu(a) + \mu(b)$ since the path uses two edges incident to every vertex except for a and b. Observe that $c_\mu(T) - c(T)$ does not depend on T and hence the optimal Traveling Salesman path for endpoints a and b is the same under both cost functions. Let T_0 be an optimal Traveling Salesman path for endpoints a and b.

Let MST_μ be a minimum spanning tree with respect to the cost function c_μ and let $C_\mu = c_\mu(\mathrm{MST}_\mu)$ be its cost. Then $C_\mu \le c_\mu(T_0)$, since a Traveling Salesman path is a spanning tree.

FACT 1. *Let μ be any potential function. If MST_μ is a Traveling Salesman path with endpoints a and b, it is an optimal Traveling Salesman path for endpoints a and b.*

Proof. From $C_\mu \le c_\mu(T_0)$, we conclude that MST_μ is an optimal Traveling Salesman path with respect to c_μ. Since the ranking of paths is the same under both cost functions, it is also optimal with respect to c. □

The inequality $C_\mu \le c_\mu(T_0) = c(T_0) - 2\sum_{v\in V}\mu(v) + \mu(a) + \mu(b)$ is valid for every potential function and hence

$$\max_\mu C_\mu + 2\sum_{v\in V}\mu(v) - \mu(a) - \mu(b) \le c(T_0)\ .$$

The quantity on the left is called the Held-Karp bound[1]. The following fact is crucial for our proof.

[1]A simple iterative algorithm can be used to approximate 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 $\mu(v) = 0$ for all $v \in V$ and compute MST_μ. As long as MST_μ is not a path with endpoints a and b, update μ as follows: $\mu(v) = \mu(v) + \alpha_t \cdot (d(v) - deg_\mu(v))$, where $d(a) = d(b) = 1$ and

FACT 2. *The Held-Karp bound is equal to the optimal objective value of the Subtour-LP.*

Proof. The fact follows by relaxing the degree constraints of the Subtour-LP in a Lagrangian fashion. For a short introduction to Lagrangian Relaxation see [CCPS98, page 259]. ☐

We remark (but will not use) that an optimal choice of μ in the Held-Karp bound is given by the optimal solution of the linear programming dual of the Subtour-LP; μ corresponds to the dual variables for the degree constraints. We next draw a simple consequence from the two facts above, which forms the basis for our proof.

LEMMA 3.1. *Let μ be any potential function. If MST_μ is the unique minimum spanning tree with respect to c_μ and moreover is a Traveling Salesman path with endpoints a and b, the Subtour-LP has a unique optimal solution and this solution is integral.*

Proof. If MST_μ is a Traveling Salesman path, it is optimal (Fact 1) and hence $c(\mathrm{MST}_\mu) = c(T_0)$. The Held-Karp bound is therefore equal to $c(T_0)$ and the same holds true for the optimal objective value of the Subtour-LP (Fact 2). The incidence vector of MST_μ is a feasible solution of the Subtour-LP of cost $c(\mathrm{MST}_\mu)$ and hence 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 $e \notin \mathrm{MST}_\mu$. Since MST_μ is unique there is a $\eta > 0$ so that decreasing the cost of e by η 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 can now describe our proof strategy for open curves. We define a potential function μ such that MST_μ is the unique minimum spanning tree in the complete network $G = (V, V \times V, c)$, where c is the Euclidean distance function, and moreover MST_μ coincides with the polygonal reconstruction (and hence is a Traveling Salesman path with endpoints a and b, where a and b are the first and last sample point, respectively). Then MST_μ and hence the polygonal reconstruction is the unique optimal solution of the Subtour-LP. We want to stress that the definition of μ is only needed for the proof of our main theorem. The reconstruction algorithm simply solves the Subtour-LP.

3.2. Intuition. When will the minimum spanning tree of the sample set be the correct reconstruction? Let $V = \{v_1, v_2, \ldots, v_n\}$ where we assume the points to be numbered according to their order on the curve. Kruskal's algorithm considers the edges $v_i v_j$ in increasing order of length and adds an edge to the spanning tree if it does not close a cycle. Kruskal's algorithm will therefore construct the path v_1–v_2–…–v_n if the potential function μ is such that

$$c_\mu(v_i, v_{i+1}) < c_\mu(v_h, v_j) \quad \text{whenever} \quad h \le i < j, j - h \ge 2 \; . \tag{3.1}$$

Let us consider two special situations: γ is essentially straight (any two left or right tangents to γ form an angle of less than $\pi/3$) or γ consists of a sharp corner (a points in which γ turns by at least $7\pi/24$) and two incident straight line segments. See Figure 3.1. We will show in Section 3.3 that any curve γ can be decomposed into subcurves which are either essentially straight or which consist of a sharp corner with two incident essentially straight legs.

For an essentially straight curve the minimum spanning tree will reconstruct for a large choice of potential functions. It will work without a potential function, i.e., $\mu(p) = 0$ for all p, and, more generally, it will work for any potential function that does not change too fast as a function of the position of its argument. For a point p which belongs to an essentially straight part of γ, we will essentially[2] define $\mu(p) = d(p)/3$, where $d(p)$ is maximal such that $B^0(p, d(p)) \cap \gamma$ is connected and essentially straight and $B^0(p, r)$ denotes the open ball with center p and radius r. This choice guarantees that $\mu(p)$ changes slowly with the position of its argument (with at most one third of the change in argument) and that $\mu(p)$ depends on local properties of the curve and is large in parts of γ that are intuitively simple to reconstruct. For sharp corners, the definition above leads to a potential value of zero.

Corners with a turning angle of more than $\pi/2$ will confuse the minimum spanning tree when used without a potential function as Figure 3.1 shows. One of our insights is that a simple potential function can be used to make the minimum spanning trees work. Assume that our curve consists of the two line segments $y = \pm m \cdot x$ for $0 \le x \le 1$ and let V be a finite set of samples. We define the potential as a function of the x-coordinates of the sample points. Fix $\pi(O)$ arbitrarily, let $p(x) = (x, mx)$ and $q(x) = (x, -mx)$

$d(v) = 2$ for all other nodes, t is the index of the iteration, and $deg_\mu(v)$ is the degree of v in MST_μ. The cost of MST_μ converges to the Held-Karp bound. We found in our experiments that the convergence is rather slow and that the cutting plane approach to solving the Subtour-LP is faster.

[2]The precise definition given in Section 3.4 is somewhat more involved.

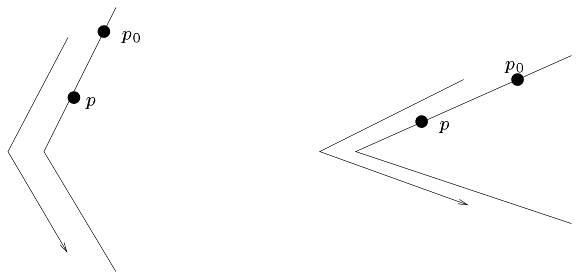

FIG. 3.1. *In the figure on the left, the Euclidean distance between p and p_0 grows as p moves away from p_0 along γ. In the figure on the right, the distance first grows and then shrinks again as p moves around the corner. The minimum spanning tree with $\mu(p) = 0$ for all p will reconstruct the curve on the left, but may fail on the curve on the right.*

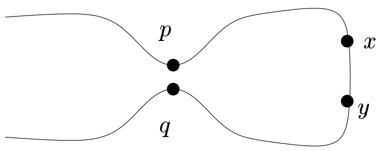

FIG. 3.2. *The edge pq does not belong to the polygonal reconstruction. Our definition of the potential function ensures that $c_\mu(p,q) > 0$. Our sampling condition that any edge in the polygonal reconstruction has non-positive reduced cost.*

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 conclude that the minimum spanning tree for the modified distance function reconstructs. In the argument above the choice of $\pi(O)$ is arbitrary. The actual choice of $\pi(O)$ will depend on local properties of the curve. For every sharp corner s we will consider the largest open disk $B^0(s, c_s)$ such that $\gamma \cap B^0(s, c_s)$ is connected and is essentially a sharp corner with two incident straight legs. We set $\pi(O) = c_s$.

We can now sketch our definition of μ. For every point $p \in \gamma$, the first definition $\mu(p) = d(p)/3$ is applicable. It assigns potential zero to sharp corners. Near sharp corners we use the second definition, namely $\mu(p) = c_s - \|sp_s\|$, where p is near the sharp corner s and p_s is the projection of p onto the angular bisector of the two tangents at s, see Figure 3.6.

The analysis above suggests that with this definition of μ, the minimum spanning tree solves the reconstruction task locally, i.e., if given the points in $V \cap \gamma'$, where γ' is a subcurve of γ that is either essentially straight or a sharp corner with two incident straight legs. In other words, inequality (3.1) holds if v_h,v_i, and v_j belong to the same γ'.

Does it also hold globally? Consider the situation shown in Figure 3.2. We have two points p and q that belong to distinct essentially straight parts of γ. We have $\max(d(p), d(q)) \le \|pq\|$ and hence $c_\mu(p,q) > 0$. More generally, we will show in Section 3.3 that any edge pq, where p and q do not belong to an either essentially straight subcurve or to a sharp corner with its incident legs, has positive modified cost.

The previous paragraph suggests our sampling condition. We require that any edge of the polygonal reconstruction has non-positive reduced cost. Then (3.1) holds certainly when $c_\mu(v_j, v_j) > 0$. When $c_\mu(v_h, v_j) \le 0$, v_h and v_j are guaranteed to lie in a common essentially straight subcurve or near a common sharp corner, and the local analysis applies.

This ends the informal description of our proof.

3.3. The Sampling Condition and the Global Reasoning for Open Curves. We will give the detailed definition of our potential function in Section 3.4; it assigns a positive real $\mu(x)$ to every point x of γ. Define the turning angle of a subcurve γ' of γ as the opening angle of the smallest double-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: $c_\mu(u, v) \le 0$.

(b) For any two adjacent samples u and v: $\gamma[u, v]$ turns by less than π. For two adjacent points p, q on the curve, $\gamma[p, q]$ denotes the subcurve of γ with endpoints p and q not containing an other

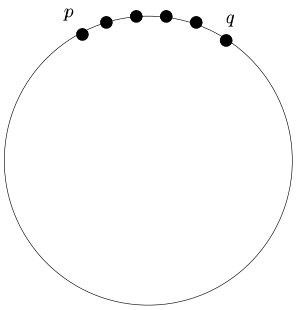

FIG. 3.3. *p and q are adjacent sample points and $c_\mu(p,q) \leq 0$ if p and q are sufficiently close. However, the sample set contradicts our intuition of what constitutes a dense sample set. Condition (b) excludes the case.*

sample point (In the case of closed curves we always have at least 3 sample points).

Condition (a) implies condition (b) in the case of open curves as we will see below. For closed curves, condition (a) does not suffices as the example in Figure 3.3 indicates. Condition (a) states that adjacent sample points must be sufficiently close in a metric sense and condition (b) states that the curve must not turn too much between adjacent sample points.

The sampling condition is easily satisfied. Let $\epsilon = \inf_{x \in \gamma} \mu(x)$. Then $\epsilon > 0$ since γ is compact and hence a sample set in which γ turns by less than π between adjacent samples and in which there is at least one sample point in every curve segment of length $\epsilon/2$ satisfies the sampling condition. We want to stress that the sampling condition can also be satisfied with non-uniform sampling. In regions of γ where μ is large, the sampling may be less dense than in regions where μ is small. In Section 10 we will relate our sampling condition to the conditions used in other papers on curve reconstruction.

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

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

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

We show that these conditions imply that MST_μ is the polygonal reconstruction of γ.

LEMMA 3.2. *Conditions (P1) and (P2) imply that MST_μ is the polygonal reconstruction of γ.*

Proof. Let MST_μ be any minimum spanning tree in (V, E, c_μ) and let $e = uv$ be any edge which does not belong to the polygonal reconstruction. We show that $e \notin \mathrm{MST}_\mu$. Observe first that any edge in the minimum spanning tree has non-positive modified cost since there is a spanning tree, namely the polygonal reconstruction, in which every edge has non-positive cost. This follows from the cycle rule for minimum spanning trees. So assume $c_\mu(e) \leq 0$. 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 $c_\mu(e') < c_\mu(e)$ for every edge e' on the part of the polygonal reconstruction between u and v. We conclude that $e \notin \mathrm{MST}_\mu$. □

3.4. The Definition of the Potential Function. In this section, we give the precise definition of our potential function. The definition depends on the parameters $\theta_{max\_sharp}, \theta_{turn}, f_{scale}, f_{wriggle}, \theta_{wriggle}$, and f_{shrink} whose choice is somewhat arbitrary but not completely independent. In Section 3.6 we summarize the conditions.

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 $\theta_{max\_sharp} = 7\pi/24 = 52.5°$.

A semi-regular curve may have an infinite number of singularities. For example, the convex hull of the points $(\cos(\pi/n), \sin(\pi/n))$, $n \geq 2$, has an infinite number of singularities, see Figure 3.4. 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 so that the tangents at q_i and r_i form an angle of at least $\pi/6$ and so that the sequence $q_1, r_1, q_2, r_2, \ldots$ increases 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.

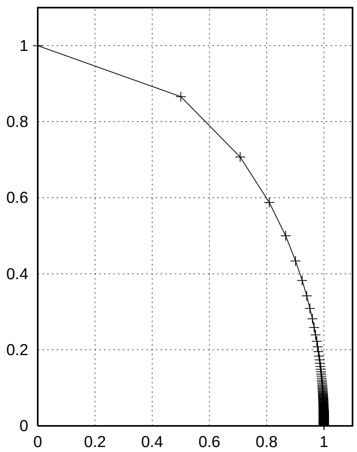

FIG. 3.4. *A semi-regular curve with an infinite number of corners*

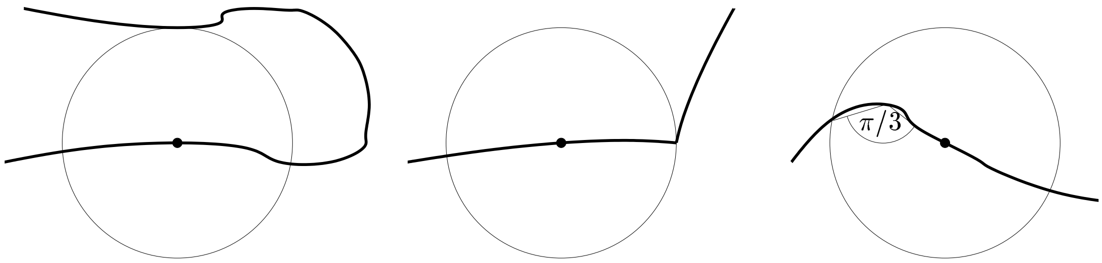

FIG. 3.5. *The closed ball intersects γ in more than one part, has a singularity on its boundary, or γ turns by $\pi/3$ im the ball.*

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.

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 $r \leq d(p)$.
- $B^0(p, d(p)) \cap \gamma$ turns by less than $\theta_{turn} = \pi/3$.

We will define our potential in parts that are far away from sharp corners as $f_{scale} d(p)$ and choose $f_{scale} = 1/3$. We define later what we mean exactly with far away.

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 (see Figure 3.5). For sharp corners $s \in S$ 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 10 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 $\alpha_s \geq 7\pi/24$).

For an angle α, let $\bar{\alpha} = \pi - \alpha$. For a point $p \in \gamma$, let p_s be the orthogonal projection of p onto b_s. See Figure 3.6 for an illustration of these definitions. For every $s \in S$ let δ_s be maximal so 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 angle between any segment with both endpoints on one leg and the tangent in s of the same leg is less than $\min\{f_{wriggle}\bar{\alpha}_s, \theta_{wriggle}\}$. We choose $f_{wriggle} = 1/4$ and $\theta_{wriggle} = \pi/9$. The second bound guarantees that the angle between any segment and the perpendicular bisector is less than $\pi/2$.
- For either of the two legs $d(p)$ increases as p moves away from s.

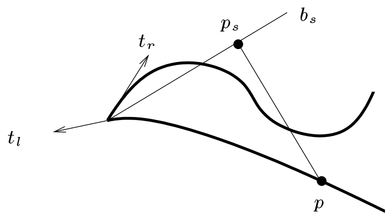

FIG. 3.6. *Illustration of the definition of b_s and p_s.*

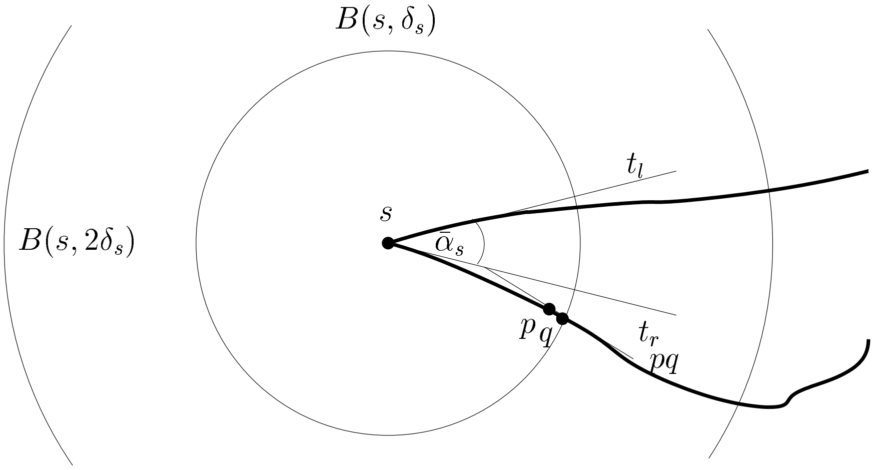

FIG. 3.7. *By the definition of δ_s, we know that γ is connected in $B^0(s,\delta_s)$, there is no sharp corner in $B^0(s,2\delta_s)$, the angle between a segment through two points of a leg in $B^0(s,\delta_s)$, say pq and the tangent of this leg, say t_r is less than $\min(\alpha_s/4, \pi/9)$, and the d values are increasing in $B(s,\delta_s)$.*

- $B^0(s, 2\delta_s)$ 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. For an illustration of the definition see Figure 3.7.Clearly, $\delta_s > 0$ for all $s \in S$.

Recall from Section 3.2 that we want to define the potential near a sharp corner as $c_s - \|sp_s\|$. The change form the potential for sharp corners and smooth areas is made by choosing the maximum of the two possibilities. We use the constant c_s to specify the exact point where we change from one definition to the other. We choose c_s maximal under the restriction that all points outside the $B(s, \delta_s f_{shrink})$ ball ($f_{shrink} = 1/5$) have potential $d(p)/3$. Thus let q^1 and q^2 be the points where the two legs intersect the boundary of the circle $B(s, \delta_s/5)$ and 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 3.5 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 5.

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

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

Observe that this definition "combines" the two cases discussed in Section 3.2. We use T to denote the set of points $p \in \gamma$ with $\mu(p) = d(p)/3$, i. e. the points that are not affected by the singularities. 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 μ is a continuous function. For a point $p \in B(s, \delta_s/5)$ we have $d(p) \leq \delta_s/5$ and hence $\mu(p) \leq c_s$.

We will frequently use the following simple observation.

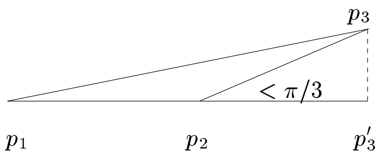

FIG. 3.8. *The situation in the proof of Lemma 3.8. We have* $\|p_1p_3\| \geq \|p_1p'_3\| > \|p_1p_2\| + \|p_2p_3\| \cos \pi/3$.

LEMMA 3.3. *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. ∎

3.5. Local Reasoning. 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 3.4. *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 3.5. *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 more than π according to our second sampling condition in the case of a closed curve. However γ' turns by less than π according to the definition of Γ. ∎

For open curves Lemma 3.5 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 3.6 (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 $\|pq\| \leq \mu(p) + \mu(q)$ by assumption. If $p \in T$ and $q \in T$, $\mu(p) = d(p)/3$ and $\mu(q) = d(q)/3$ 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 $\mu(p) \leq 2\delta_s/5$ and $\|pq\| \geq 4\delta_s/5$. If $q \in T$ then $\mu(q) = d(q)/3 \leq \|sq\|/3 \leq (\|pq\| + \delta_s/5)/3$ 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 $\mu(p) \leq 2\delta_s/5$ and $\mu(q) \leq 2\delta_t/5$, a contradiction. ∎

In order to show that the MST_μ coincides with the polygonal reconstruction, we show that the modified distance between two points p and r is either non-negative, or the modified distance from p to r is greater than the modified distance from p to any point q between p and r. Since we want to use these Lemmas also if we treat the problem with finite precision arithmetic, we quantify the change of the modified distance in the distance of q and r.

LEMMA 3.7. *Let p, q, and r be sample points on a regular subcurve with $p < q < r$. If $\{p, q, r\} \in B(t, d(t))$ for some point $t \in \gamma$ then $c_\mu(p, r) - c_\mu(p, q) \geq \|qr\|/6$.*

Proof. Let $t \in \gamma$ and $p < q < r \in B(t, d(t))$. Since the points are contained in a regular subcurve, we have (see Figure 3.8) $\angle(\vec{pq}, \vec{qr}) < \pi/3$ and hence $\|pr\| > \|pq\| + \|qr\| \cos \pi/3 = \|pq\| + \|qr\|/2$. Furthermore observe that $d(r) \leq d(q) + \|qr\|$. Thus

$$c_\mu(p, r) - c_\mu(p, q) = \|pr\| - \mu(p) - \mu(r) - \|pq\| + \mu(p) + \mu(r) \geq \|qr\|/2 - \|qr\|/3 = \|qr\|/6.$$

∎

LEMMA 3.8 (Prop. (P1) for regular regions). *MST_μ coincides with the polygonal reconstruction for regular subcurves.*

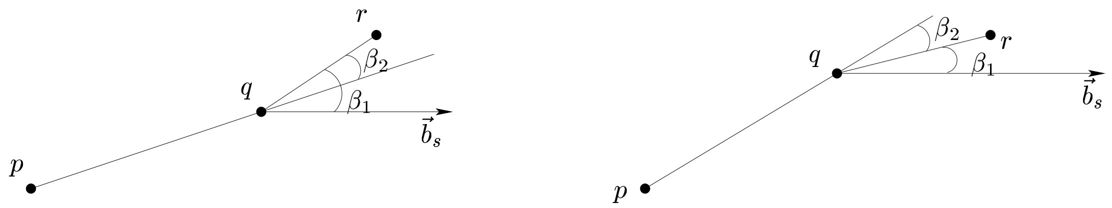

FIG. 3.9. *Illustration of the definition of β_1 and β_2, if p, q, and r are on the same leg.*

Proof. The Lemma above for both possible orientations of the curve implies directly that Prim's minimum spanning tree algorithm finds the polygonal reconstruction. ∎

LEMMA 3.9. *Let p, q, and r be sample points of γ with $p < q < r$. If $\{p, q, r\} \in B(s, \delta_s)$ for some sharp corner $s \in \gamma$ then $c_\mu(p, r) \geq 0$ or $c_\mu(p, r) - c_\mu(p, q) \geq \|qr\|(\sin \bar{\alpha}_s/4)^2/3$.*

Proof. We first argue that is suffices to prove the claim for the situations where either q and r both belong to T or neither of them does. Assume for the moment that those two cases have been dealt with. If exactly one of q and r belongs to T, there is a point u between q and r that belongs to the boundary of T. For this point, we have $d(u)/3 = c_s - \|su_s\|$ and hence u can be considered to be in T or outside T. The triples (p, q, u) and (p, u, r) are both in one of the special situations and hence we have $c_\mu(p, u) \geq 0$ or $c_\mu(p, u) - c_\mu(p, q) \geq \|qu\|(\sin \bar{\alpha}_s/4)^2/3$ and $c_\mu(p, r) \geq 0$ or $c_\mu(p, r) - c_\mu(p, u) \geq \|ur\|(\sin \bar{\alpha}_s/4)^2/3$. If $c_\mu(p, r) \geq 0$, we are done. Otherwise $c_\mu(p, r) < 0$ and hence $c_\mu(p, u) < 0$. Thus $c_\mu(p, r) - c_\mu(p, q) = c_\mu(p, r) - c_\mu(p, u) + c_\mu(p, u) - c_\mu(p, q) \geq (\|qu\| + \|ur\|)(\sin \bar{\alpha}_s/4)^2/3 \geq \|qr\|(\sin \bar{\alpha}_s/4)^2/3$, where the last inequality follows from the triangle inequality. We may from now on assume that q and r either both belong to T or neither of them does.

We need some further case distinctions. The first distinction is according to the sign of $c_\mu(p, q)$. The case $c_\mu(p, q) > 0$ is dealt with in the last paragraph of the proof.

We start with the assumption $c_\mu(p, q) \leq 0$. We make a further case distinction according to the position of s in the sequence $p < q < r$. In all four cases we employ a common strategy. We have $c_\mu(p, r) - c_\mu(p, q) = (\|pr\| - \|pq\|) - (\mu(r) - \mu(q))$. We bound $\|pr\| - \|pq\|$ from below and $\mu(r) - \mu(q)$ from above and estimate the difference of the bounds. In all cases we also use the estimates $\bar{\alpha}_s \leq 17\pi/24$, $\sin \bar{\alpha}_s/4 \leq 0.5281$, $(\sin \bar{\alpha}_s/4)^2/3 \leq 0.1$ and $\cos 2\pi/9 \geq 2/3$.

$s \leq p < q < r$: If $\{q, r\} \cap T = \emptyset$, $\mu(r) \leq \mu(q)$, and if $\{q, r\} \subseteq T$, $\mu(r) \leq \mu(q) + \|qr\|/3$. In either case[3], $\mu(r) \leq \mu(q) + \|qr\|/2$.
Let β_2 be the angle between the vectors $\vec{pq}$ and $\vec{qr}$. Then $\|pr\| - \|pq\| \geq \|qr\| \cos \beta_2$. By the sampling condition, we have $\beta_2 \leq 2\pi/9$, since the angle between pr, respectively qr, and the left tangent at s is at most $\pi/9$. Thus

$$c_\mu(p, r) - c_\mu(p, q) \geq \|qr\|(\cos(\beta_2) - 1/2) \geq \|qr\|/6 \geq \|qr\| \sin(\bar{\alpha}_s/4)^2/3 .$$

$p < q < r \leq s$: If $\{qr\} \subseteq T$, $\mu(r) \leq \mu(q) \leq \mu(q) + \|qr\|/2$. Let β_2 be the angle between the vectors $\vec{pq}$ and $\vec{qr}$. Then $\|pr\| - \|pq\| \geq \|qr\| \cos \beta_2$ and $\beta_2 \leq 2\pi/9$ and hence the argument used in the case $s \leq p < q < r$ applies.
Assume next that $\{q, r\} \cap T = \emptyset$. Let β_1 be the angle between the vectors $\vec{qr}$ and $\vec{b}_s$. Then $\mu(r) - \mu(q) = \|qr\| \cos \beta_1$. By the sampling condition, we have $\beta_1 \geq 3/4\bar{\alpha}_s$. Thus $\mu(r) - \mu(q) = \|qr\| \cos 3/4\bar{\alpha}_s$. Let β_2 as above be the angle between the vectors $\vec{pq}$ and $\vec{qr}$, see Figure 3.9 for an illustration. Then $\beta_2 \leq \bar{\alpha}_s/2$ (since the angle between pq, respectively qr, and the right tangent at s is at most $\bar{\alpha}_s/4$) and $\|pr\| - \|pq\| \geq \|qr\| \cos \beta_2 \geq \|qr\| \cos \bar{\alpha}_s/2$.
Combining bounds we obtain

$$\begin{aligned} c_\mu(p, r) - c_\mu(p, q) &\geq \|qr\|(\cos(\bar{\alpha}_s/2) - \cos(\bar{\alpha}_s/2 + \bar{\alpha}_s/4)) \\ &\geq \|qr\|(\cos(\bar{\alpha}_s/2) - \cos(\bar{\alpha}_s/2)\cos(\bar{\alpha}_s/4) + \sin(\bar{\alpha}_s/2)\sin(\bar{\alpha}_s/4)) \\ &\geq \|qr\| \sin(\bar{\alpha}_s/4)^2. \end{aligned}$$

$p < s \leq q < r$: Assume first that $\{q, r\} \cap T = \emptyset$. If q_s is as least as far from s as p_s, we have $\|pr\| - \|pq\| \geq 0$ and $\mu(q) - \mu(r) \geq \|qr\| \cos(17\pi/48 + \pi/9) = \|qr\| \cos(\pi/2 - 5\pi/144) = \|qr\| \sin(5\pi/144) > 0.1 \cdot \|qr\| \geq \|qr\|(\sin \bar{\alpha}_s/4)^2/3$. The claim follows.

[3]In Section 5.2 we will consider a modified potential function for which we only know $\mu(r) \leq \mu(q) + \|qr\|/2$. We want to reuse the proof there.

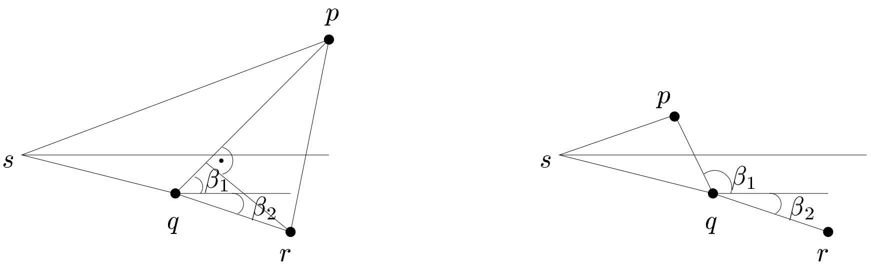

FIG. 3.10. *The case $p < s \le q < r$: In the left part q_s is closer to s than p_s and in the right part the converse is true. In both situations the definitions of β_1 and β_2 are illustrated.*

If q_s is closer to s than p_s, let β_1 be the angle between the vectors $\vec{qp}$ and $\vec{b}_s$ and β_2 be the angle between the vectors $\vec{qr}$ and $\vec{b}_s$; see the left part of Figure 3.10 for an illustration. We have $\mu(q) - \mu(r) = \|qr\| \cos \beta_2$ and $\|pq\| - \|pr\| \le \|qr\| \cos(\beta_1 + \beta_2)$. By the sampling condition, we have $\beta_2 \ge \bar{\alpha}_s/4$ and $\beta_1 \ge \angle(\vec{sp}, \vec{b}_s) \ge \bar{\alpha}_s/4$, since $\angle(\vec{qp}, \vec{b}_s) \ge \angle(\vec{sp}, \vec{b}_s)$ (moving along the line $\vec{sq}$ increases the angle). Combining bounds we obtain:

$$\begin{aligned} c_\mu(p,r) - c_\mu(p,q) &\ge \|qr\|(\cos\beta_2 - \cos(\beta_1 + \beta_2)) \\ &= \|qr\|(\cos\beta_2 - \cos\beta_1 \cos\beta_2 + \sin\beta_1 \sin\beta_2) \\ &\ge \|qr\| \sin(\bar{\alpha}_s/4)^2. \end{aligned}$$

We come to the case that $\{q, r\} \subseteq T$: If $p \in T$, we have $c_\mu(p,r) = \|pr\| - d(p)/3 - d(r)/3 > 0$, since $d(p) \le \|pr\|$ and $d(r) \le \|pr\|$.
So assume $p \notin T$. Since $q \in T$, we have $d(q)/3 \ge c_s - \|sq_s\|$ and hence $c_s \le d(q)/3 + \|sq_s\|$. Thus[4] $0 \ge c_\mu(p,q) = \|pq\| - (c_s - \|sp_s\|) - d(q)/3 \ge \|pq\| - (\|sq_s\| - \|sp_s\|) - 2d(q)/3 \ge \|pq\|/3 - (\|sq_s\| - \|sp_s\|)$ and hence $\|sq_s\| - \|sp_s\| \ge \|pq\|/3$. In particular, q_s lies further away from s than p_s.
Let β_1 be the angle between the vectors $\vec{qp}$ and $\vec{b}_s$ and β_2 be the angle between the vectors $\vec{qr}$ and $\vec{b}_s$; see the right part of Figure 3.10 for an illustration. Then $\cos\bar{\beta}_1 \ge 1/3$, $\beta_2 \ge \bar{\alpha}_s/4$, and $\|pr\| - \|pq\| \ge \|qr\| \cos(\pi - \beta_1 - \beta_2)$.
Combining bounds we obtain $c_\mu(p,r) - c_\mu(p,q) \ge \|qr\|(\cos(\bar{\beta}_1 - \beta_2) - 1/3)$. If $\cos(\bar{\beta}_1 - \beta_2) - 1/3 \ge 4/27$, we are done since $4/27 \ge (\sin\bar{\alpha}_s/4)^2/3$. So assume $\cos(\bar{\beta}_1 - \beta_2) \le 1/3 + 4/27 = 13/27 \le 1/2$. Then $\bar{\beta}_1 - \beta_2 \ge \pi/3$ and hence $\bar{\beta}_1 \ge \pi/3$ and hence $(1 - \cos\bar{\beta}_1)/\bar{\beta}_1 \ge 1/2$. Since $\cos(\bar{\beta}_1 - x)$ is convex in $[0, \bar{\beta}_1]$ and hence is above the line through the points $(0,1)$ and $(\bar{\beta}_1, \cos\bar{\beta}_1)$, we have

$$\cos(\bar{\beta}_1 - \beta_2) - 1/3 \ge \cos\bar{\beta}_1 + (1 - \cos\bar{\beta}_1)(\beta_2/\bar{\beta}_1) - 1/3 \ge \beta_2/2 \ge (\sin\bar{\alpha}_s/4)^2/3.$$

$p < q < s < r$: We have $\|qr\| \le \|qs\| + \|sr\|$. One of the previous cases applies to the triples (p,q,s) and (p,s,r) and hence $c_\mu(p,r) - c_\mu(p,q) = c_\mu(p,r) - c_\mu(p,s) + c_\mu(p,s) - c_\mu(p,q) \ge (\|qs\| + \|sr\|)(\sin\bar{\alpha}_s/4)^2/3 \ge \|qr\|(\sin\bar{\alpha}_s/4)^2/3$.

The discussion of the case $c_\mu(p,q) \le 0$ is now completed. So let us assume $c_\mu(p,q) > 0$. If $c_\mu(p,r) < 0$ there is a point q' between q and r with $c_\mu(p,q') = 0$. The first case applies to the triple (p,q',r) and hence $c_\mu(p,r) > c_\mu(p,q')$, a contradiction. Thus $c_\mu(p,r) \ge 0$. ☐

LEMMA 3.10 (Prop. (P1) for singular regions). *MST_μ coincides with the polygonal reconstruction for singular subcurves.*

Proof. The Lemma above for both possible orientations of the curve implies directly that Prim's minimum spanning tree algorithm finds the polygonal reconstruction. ☐

3.6. Conditions on the thresholds. In the preceding sections we showed that properties (P1) and (P2) hold, if one chooses the thresholds as in Section 3.4. There are other possible choices for the thresholds that make the arguments work. We now collect the conditions on the thresholds. Note that the Subtour-LP has an unique integral solution if there is a choice of the thresholds that make the MST$_\mu$ unique and equal to the polygonal reconstruction.

We introduced six thresholds:

θ_{max_sharp}: The minimum turning angle of a singularity, that we call sharp. We have chosen $\theta_{max\_sharp} = 7\pi/24$.

[4]Recall that we work under the assumption $c_\mu(p,q) \le 0$.

θ_{turn}: The maximal turning angle in a $B(p, d(p))$ ball. We have chosen $\theta_{turn} = \pi/3$.
f_{scale}: The factor by which we scaled the $d(p)$ value for the potential. We have chosen $f_{scale} = 1/3$.
$f_{wriggle}$: The factor by which we scaled the $\bar{\alpha}_s$ as maximal angle between a tangent and a segment in a sharp corner. We have chosen $f_{wriggle} = 1/4$.
$\theta_{wriggle}$: The maximal angle between a tangent and a segment in a sharp corner. We have chosen $\theta_{wriggle} = \pi/9$.
f_{shrink}: The factor by which we shrunk the δ_s ball to define the area where we use the potential function for sharp corners. We have chosen $f_{shrink} = 1/5$.

First of all, we must guarantee that $\mu > 0$, which is equivalent to

$$\boxed{\theta_{turn} < \theta_{max\_sharp}.}$$

We start our investigations by looking at property (P2), i.e. to guarantee that all edges which are not contained in a region have positive modified cost. Look at any edge pq with negative reduced cost. If p and q are in T, we have to show that there is a point x with $\|px\| \le d(x)$ and $\|qx\| \le d(x)$. For simplicity we have assumed that x is either p or q. This is reasonable, since the conditions we get are weaker than the conditions we need later. Thus we require $\max(d(p), d(q)) \ge \|pq\|$. Since $\|pq\| \le f_{scale}(d(p) + d(q))$ it suffices that

$$\boxed{f_{scale} < 1/2.}$$

If exactly one of p, q is not in T, we require that $2f_{shrink}\delta_s + f_{scale}\|sq\| \le \|pq\|$ (the first summand is an upper bound for the potential of p, the second for the potential of q). Using that $\|pq\| \ge 1 - f_{shrink}$ for all points outside the $B(s, \delta_s)$ ball, we conclude that it suffices that

$$\boxed{3f_{shrink} + f_{scale} < 1.}$$

Assume next that p and q are not in T. Then $p \in B(s_1, f_{shrink}\delta_{s_1})$ and $q \in B(s_2, f_{shrink}\delta_{s_2})$ for different sharp corners s_1 and s_2. Let w.l.o.g. $\delta_{s_1} \ge \delta_{s_2}$. It suffices that

$$\boxed{f_{shrink} < 1/3.}$$

Let us now turn to property (P1). In regular regions we require that "potential changes slower than the distance". Thus we need

$$\boxed{\cos(\theta_{turn}) > f_{scale}.}$$

For singular regions, we have to go through all cases of Lemma 3.9. If $s \le p < q < r$ we need similarly to the case of regular cures

$$\boxed{\cos(\theta_{wriggle}) > f_{scale}.}$$

Analogously, if $p < q < r < s$ and q and r are in T. If $p < q < r < s$ and q and r are not in T, we require that $\cos(1 - f_{wriggle}) < \cos(2f_{wriggle})$, which is equivalent to

$$\boxed{f_{wriggle} < 1/3.}$$

If $p < s \le q < r$ and $q, r \notin T$, we additional require that

$$\boxed{\theta_{max\_sharp}/2 + \theta_{wriggle} < \pi/2.}$$

If $q, r \in T$ we compute that $\cos(\bar{\beta}_1) \ge 1 - 2f_{scale}$ and require that $\cos(\bar{\beta}_1 - f_{wriggle}\alpha_s) > 1/3$. This is equivalent to

$$\boxed{f_{scale} \le 1/3.}$$

Note that the last requirement is the only one where the condition can also be satisfied with equality, since we subtract a positive number in the derivation. This ends our discussion.

4. Open Curves with Unspecified Endpoints. In the preceding sections we assumed that the first and last sample point (= endpoints of the Traveling Salesman path) are specified as part of the input. In this section we show that the Subtour-LP can also reconstruct when the endpoints are not specified. Of course, the requirements on the sample will be stronger. The argument used in this section is a variant of the argument used in Section 3.

We use the following formulation of the Subtour-ILP. The goal is to select a total of $n-1$ edges such that at most two of them are incident to any node and such that no subset V' with $V' \neq \emptyset$ is "over-full".

$$\begin{array}{llll} \min & \displaystyle\sum_{u,v\in V} c_{uv} x_{uv} & & \\ s.t. & \displaystyle\sum_{v \in V} x_{uv} \leq 2 & & \text{for all } u \in V \\ & \displaystyle\sum_{u,v\in V'} x_{uv} \leq |V'|-1 & & \text{for } V' \subset V,\, V' \neq \emptyset \\ & \displaystyle\sum_{u,v\in V'} x_{uv} = n-1 & & \\ & x_{uv} \in \{0,1\} & & \text{for all } u,v \in V. \end{array}$$

Selecting a total of $n-1$ edges such that at most two of them are incident to any node amounts to selecting a path and a set of cycles covering all nodes. The constraint that no set can be over-full implies that no cycles can be used and hence any solution must be a Traveling Salesman path. The Subtour-LP is obtained by replacing the integrality constraints $x_{uv} \in \{0,1\}$ by the linear constraints $0 \leq x_{uv} \leq 1$.

As in the case of open curves with specified endpoints, we have to show that the separation problem can be solved in polynomial time. The separation algorithm works as follows. Let x^* be the optimal solution. We assign a capacity of x^*_e to edge e for every edge e. Furthermore we introduce a artificial vertex s and edges us with capacity $2-\sum_{e\in\delta(u)} x^*_e$ for every node u of the graph. We compute a minimal cut in this graph and take as subset for the subtour elimination constraint the side of the cut that does not contain the artificial node s.

To see the correctness of this separation algorithm, we show that the subtour elimination constraint for S is violated iff the size of the cut is less than 2. Let $S \subset V$ be a subset of the nodes. First notice that the sum of the capacities of edges adjacent to a node is exactly 2 for every node in the above graph. Thus $2\sum_{e\in\gamma(S)} x^*_e + \sum_{e\in\delta(s)} x^*_e = 2\ S$ (for every node $u \in S$, we sum the capacities of the adjacent edges). We conclude that $\sum_{e\in\delta(s)} x^*_e < 2$ iff $\sum_{e\in\gamma(S)} x^*_e > \ S\ - 1$.

We consider only non-positive potential functions $\mu \leq 0$ in this section. Let a and b be fixed vertices and consider any Traveling Salesman path T with endpoints a and b. Its costs under c and c_μ are related by $c_\mu(T) = c(T) - 2\sum_{v\in V}\mu(v) + \mu(a) + \mu(b)$ since the path uses two edges incident to every vertex except for a and b. Observe that $c_\mu(T) - c(T)$ does not depend on T and hence the optimal Traveling Salesman path for endpoints a and b is the same under both cost functions. However, the relative order of Traveling Salesman path with distinct endpoints is changed.

Let MST_μ be a minimum spanning tree with respect to the cost function c_μ and let $C_\mu = c_\mu(\mathrm{MST}_\mu)$ be its cost. Then $C_\mu \leq c_\mu(T)$ for any Traveling Salesman path T.

FACT 3. *Let $\mu \leq 0$ be any potential function. If MST_μ is a Traveling Salesman path and $\mu(a) = \mu(b) = 0$ for the endpoints of this path, it is an optimal Traveling Salesman path.*

Proof. Let T_0 be an optimal Traveling Salesman path, say with endpoints u and v. Then $C_\mu \leq c_\mu(T_0)$, since T_0 is a spanning tree, $c(\mathrm{MST}_\mu) = c_\mu(\mathrm{MST}_\mu) + 2\sum_{v\in V}\mu(v) - \mu(a) - \mu(b) = c_\mu(\mathrm{MST}_\mu) + 2\sum_{v\in V}\mu(v)$, since MST_μ is a path with endpoints a and b and a and b have potential zero, and $c(T_0) = c_\mu(T_0) + 2\sum_{v\in V}\mu(v) - \mu(u) - \mu(v) \geq c_\mu(T_0) + 2\sum_{v\in V}\mu(v)$, since T_0 is a path with endpoints u and v, and since the potentials of u and v are non-positive. Thus

$$c(T_0) \geq c_\mu(T_0) + 2\sum_{v\in V}\mu(v) \geq c_\mu(\mathrm{MST}_\mu) + 2\sum_{v\in V}\mu(v) = c(\mathrm{MST}_\mu).$$

□

The inequality $C_\mu \leq c_\mu(T_0) = c(T_0) - 2\sum_{v\in V}\mu(v) + \mu(u) + \mu(v) \leq c(T_0) - 2\sum_{v\in V}\mu(v)$ is valid for every non-positive potential function (the last inequality uses non-positivity) and hence

$$\max_{\mu\leq 0}\big(C_\mu + 2\sum_{v\in V}\mu(v)\big) \leq c(T_0)\ .$$

The quantity on the left is called the Held-Karp bound. The following fact is crucial for our proof.

FACT 4. *The Held-Karp bound is equal to the optimal objective value of the Subtour-LP.*

Proof. The proof follows from [CCPS98, page 259]. Relaxing the degree constraints $\sum_{v\in V} x_{uv} \le 2$ in a Langrangean fashion, we obtain the problem

$$\begin{array}{rrcll} \max\limits_{\mu\le 0}\min\limits_{x\ge 0} & \sum\limits_{u,v\in V} c(uv)x_{uv} + \sum\limits_{u}\mu(u)(2-\sum\limits_{v\in V} x_{uv}) & & & \\ s.t. & \sum\limits_{u,v\in V'} x_{uv} & \le & |V'|-1 & \text{for } V'\subset V,\, V'\ne\emptyset \\ & \sum\limits_{u,v\in V} x_{uv} & = & n-1 & \\ & x_{uv} & \le & 1 & \text{for all } u,v\in V. \end{array}$$

Observe that we are only maximizing over non-positive potential functions μ. This stems from the fact that in contrast to Section 3, the degree constraints are now inequalities instead of equalities. The reformulation has the same objective value. The objective function of the LP can be reformulated as $\sum_{u,v\in V}(c(uv)-\mu(u)-\mu(v))x_{uv}+2\mu(V)$. We conclude that the LP is simply a minimum spanning tree problem for the cost function c_μ. □

We remark (but will not use) that the optimal choice of μ in the Held-Karp bound is given by the optimal solution of the linear programming dual of the Subtour-LP; μ corresponds to the dual variables for the degree constraints. We next draw a simple consequence from the two facts above.

LEMMA 4.1. *Let $\mu \le 0$ be any potential function. If MST_μ is the unique minimum spanning tree with respect to c_μ, is a Traveling Salesman path, and $\mu(a) = \mu(b) = 0$ for its endpoints a and b, the Subtour-LP has a unique optimal solution and this solution is integral.*

Proof. If MST_μ is a Traveling Salesman path, it is optimal (Fact 1) and hence $c(\mathrm{MST}_\mu) = c(T_0)$. The Held-Karp bound is therefore equal to $c(T_0)$ and the same holds true for the optimal objective value of the Subtour-LP (Fact 2). The incidence vector of MST_μ is a feasible solution of the Subtour-LP of cost $c(\mathrm{MST}_\mu)$ and hence 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 $e \notin \mathrm{MST}_\mu$. Since MST_μ is unique there is a $\eta > 0$ so that decreasing the cost of e by η 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. □

It remains to define the appropriate potential function. We obtain it as a modification of the potential function defined in the preceding section. We use $\bar\mu$ to denote it. Let V be a set of sample points and let a and b be the first and the last sample point.

Let $m = \min(\bar\mu(a), \bar\mu(b))$ and set $c_s = \min(c_s, m)$ for all sharp corners. This changes $\bar\mu$ (it makes $\bar\mu$ smaller for some points near sharp corners). Define new potential functions $\tilde\mu$ and μ by

$$\tilde\mu(p) = \min(\bar\mu(p), m) \quad \text{and} \quad \mu(p) = \tilde\mu(p) - m$$

for all $p \in \gamma$, i.e., first all potential values are capped at m and then m is subtracted. Then $\tilde\mu(p) \ge \mu(p)$ for all p and $\mu(p) \le 0$ for all p. Also $\mu(a) = \mu(b) = 0$.

We strengthen the sampling condition and require $c_{\tilde\mu}(pq) = \|pq\| - \tilde\mu(p) - \tilde\mu(q) \le 0$ for all edges in the reconstruction.

Since $c_s \le m$ for all sharp corners, we have $\tilde\mu(p) = \bar\mu(p)$ for all $p \notin T$. Here $\bar\mu$ denotes the original potential function, but with the capped c-values.

Suppose that V satisfies the strengthened sampling condition. Then the minimum spanning tree with respect to $c_{\tilde\mu}$ is equal to the polygonal reconstruction. This requires to check that Lemmas 3.6 to 3.10 stay true; we leave the straightforward but tedious check to the reader. The minimum spanning tree with respect to μ is the same as the minimum spanning tree with respect to $\tilde\mu$, since μ and $\tilde\mu$ differ only by a constant. We conclude that the minimum spanning tree with respect to μ is equal to the polygonal reconstruction and moreover unique. We finally observe that μ is non-positive and that $\mu(a) = \mu(b) = 0$. Thus MST_μ is the unique optimal solution of the Subtour-LP by Lemma 3.2.

5. Closed Curves.

We extend the result to closed curves in two steps.

- In Section 5.2 we alter the potential function to μ' so that the two longest edges of the polygonal reconstruction have the same modified cost and so that the minimum spanning trees with respect to the new modified cost are precisely the polygonal reconstruction minus one of the edges of maximal modified cost.

• In Section 5.1 we show that the preceding sentence implies our main theorem for closed curves. Observe that, in our write-up, the second step is dealt with first.

Readers familiar with the Held-Karp bound for Traveling Salesman tours may wonder why we are not arguing about 1-trees. We tried, but could not get the argument to work. A 1-tree is defined as follows. An arbitrary node $v \in V$ is fixed. A 1-tree consists of the two cheapest edges incident to v plus a minimum spanning tree of $V \setminus v$. We were unable to construct a potential function for which the optimal 1-tree coincides with the polygonal reconstruction. We were able to construct a potential function where the two cheapest edges incident to v were indeed the edges to the two neighbors in the polygonal reconstruction and were able to construct a potential function where the minimum spanning tree on $V \setminus v$ coincided with the polygonal reconstruction minus the two edges incident to v. We were unable to satisfy both conditions simultaneously.

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

$$\begin{array}{llcll}
\min & \sum\limits_{u,v\in V} c_{uv}x_{uv} & & & \\
s.t. & \sum\limits_{v\in V} x_{uv} & = & 2 & \text{for } u \in V \\
& \sum\limits_{u\in V', v\in V'} x_{uv} & \leq & |V'| - 1 & \text{for } V' \subseteq V, \emptyset \neq V' \neq V \\
& \sum\limits_{u,v\in V} x_{uv} & = & |V| & \\
& 0 \leq x_{uv} & \leq & 1 &
\end{array}$$

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 set of degree equalities to the objective function and obtain the following problem with the same objective value [CCPS98, pages 258–260]:

$$\begin{array}{llcll}
\max & 2\sum\limits_{u\in V} \mu(u) + \min \sum\limits_{u,v\in V} c_\mu(uv)x_{uv} & & & \\
s.t. & \sum\limits_{u\in V', v\in V'} x_{uv} & \leq & |V'| - 1 & \text{for } V' \subseteq V, \emptyset \neq V' \neq V \\
& \sum\limits_{u,v\in V} x_{uv} & = & |V| & \\
& 0 \leq x_{uv} & \leq & 1 &
\end{array}$$

In this formulation the maximization is over all choices of μ. For fixed μ the inner minimization is over the choices for the x_{uv}. We will show that for $\mu = \mu'$ 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 satisfies 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 optimal solution to the inner minimization problem.

It remains to prove that for $\mu = \mu'$, 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) with respect to the order of the points along the curve; see Figure 5.1. 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 tree 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 relaxed by:

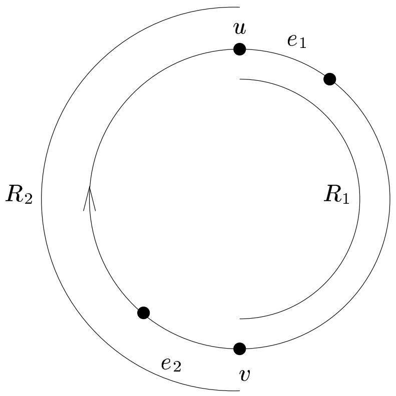

FIG. 5.1. *The notation used in the reformulations of the Subtour-LP. We have $e_i \in E_i$.*

$$
\begin{array}{llrcll}
\min\limits_{o_1,o_2 \geq 0} & \min & \sum\limits_{u,v \in V} c_{\mu'}(uv) x_{uv} & & & \\
 & s.t. & \sum\limits_{u \in R'_1, v \in R'_1} x_{uv} & \leq & |R'_1| - 1 & \text{for all } R'_1 \subseteq R_1, \text{ with } R'_1 \neq \emptyset \\
 & & \sum\limits_{u,v \in R_1} x_{uv} & = & |R'_1| - 1 - o_1 & \\
 & & \sum\limits_{u \in R'_2, v \in R'_2} x_{uv} & \leq & |R'_2| - 1 & \text{for all } R'_2 \subseteq R_2, \text{ with } R'_2 \neq \emptyset \\
 & & \sum\limits_{u,v \in R_2} x_{uv} & = & |R_2| - 1 - o_2 & \\
 & & \sum\limits_{uv \in C} x_{uv} & = & o_1 + o_2 & \\
 & & 0 \leq x_{uv} & \leq & 1 &
\end{array}
$$

Observe that we dropped some of the subtour elimination constraints.

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[5]. The same is true for R_2. The LP takes its minimum for $o_1 = o_2 = 0$ since any edge in C has higher 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.

[5] For $o_1 = 0$ this is the characterization of minimum spanning trees by a linear program, see [CCPS98, Theorem 2.8]. 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 cost edge (let us assume for simplicity that all costs 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 hand side will increase. When the next edge becomes tight, we add it to the spanning tree, ... We finish when adding a partial edge makes the constraint $x(\gamma(R_1)) = n - 1 - o_1$ tight.

5.2. The Modified Potential Function and the Local Reasoning. We show how to alter the potential so that the two longest edges of the polygonal reconstruction have the same modified cost and so that all minimum spanning trees for V remain part of the polygonal reconstruction. Note that the new potential is defined according to a given sample set, whereas the original potential only depends on the curve.

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

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

(2) There is a point $v \in \gamma$ so that $\|vu\|/2 \geq \mu(u)$ for both endpoints u of e_{max} and so 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 certainly be 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 $\|uv\| \geq d(u) = 3\mu(u)$ 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 $\mu(u) \leq \max(c_s, d(u)/3) \leq 2\delta_s/5$ 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 3.3 and we are done. In either case we have shown that one of items above holds. We also need the following Lemma.

LEMMA 5.1. *Let s be a sharp corner and let x and y be adjacent sample points so that $s \in \gamma[x, y]$. Then $x, y \in B(s, 3\delta_s/5)$ 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 3.5 and 3.6. Assume $x \notin B(s, 3\delta_s/5)$. Then

$$c_\mu(x, y) \geq c_\mu(s, x) \geq \|sx\| - 2\delta_s/5 - \|sx\|/3 \geq 2\|sx\|/3 - 2\delta_s/5 > 2\delta_s/5 - 2\delta_s/5 = 0,$$

where the first inequality follows from the third claim in the proof of Lemma 3.10. Assume next that $x, y \in T$. Lemma 3.3 implies $d(x), d(y) \leq \|xy\|$ and hence $c_\mu(x, y) \geq \|xy\|/3 > 0$. □

We now turn to the definition of the modified potential function and the proof that all minimum spanning trees for V with respect to the modified potential function are subsets of the polygonal reconstruction.

Assume first that there is a sharp corner s so 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 $\mu(x) = d(p)/3$ 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 5.1) has positive modified cost: $c_{\mu'}(x, y) = \|xy\| - \mu'(x) - \mu'(y) \geq \|xy\| - \|xy\|/3 - \|xy\|/3 > 0$ since $d(x), d(y) \leq \|xy\|$ and hence $c_\mu(x, y) \geq \|xy\|/3 > 0$ by Lemma 3.3. 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 $c_{\mu'}$ remains the polygonal reconstruction by the remark in Section 3.4 that only the fact that $0 < c_s \leq 2\delta_s/5$ is used in our proofs. This completes the discussion of the first case.

We come to the case that there is a point $v \in \gamma$ so that $\mu(u) \leq \|vu\|/2$ for both endpoints u of e_{max} and so 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 $l \in \mathbb{R}$, $0 \leq l \leq \mu(v)$, we define a as the first point with $\mu(a) = \|va\|/2 + l$ and b as the last point with $\mu(b) = \|vb\|/2 + l$ and define a potential μ' by

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

Observe that a, b, and μ' depend on l. For simplicity of notation we do not make this dependence explicit in the notation. Figure 5.2 illustrates the definition of μ'.

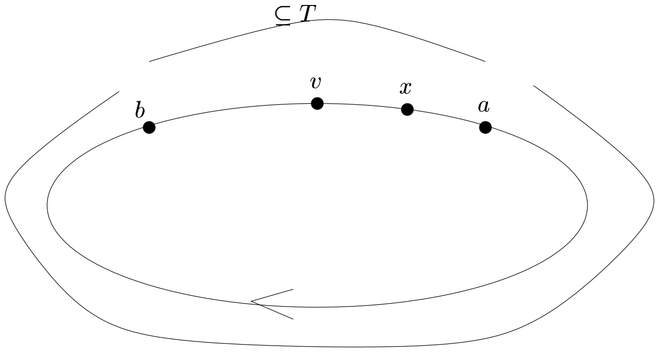

FIG. 5.2. *The construction of the potential function μ'.*

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 $c_{\mu'}$ uses only edges of the polygonal reconstruction.

We first show the existence of a and b for all choices of l. For $l = \mu(v)$ we have $a = b = v$. For $0 \leq l < \mu(v)$ we have $\mu(v) > l + \|vv\|/2$ and $\mu(u) \leq \|uv\|/2 \leq l + \|uv\|/2$ for either endpoint of e_{max}. From the continuity of μ we conclude that there is a point a between v and u for which $\mu(a) = \|va\|/2 + l$ and a point b between u and v for which $\mu(b) = \|vb\|/2 + l$. We have shown that a and b exist and that e_{max} does not depend on l. Furthermore we know that $\mu'(p) \leq \mu(p)$ for all p by 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 $\mu(r) = d(r)/3 \leq \|sr\|/3 = \delta_s/5 \leq \|rv\|/2 \leq \|rv\|/2 + l$. Continuity of μ 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 $l = \mu(v)$, we have $a = b = v$ and hence $\mu = \mu'$. Thus e_{max} has the maximal modified cost $c_{\mu'}$ 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 $c_{\mu'}(xy)$. 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 first show that $a < x$ and $y < b$ is impossible. Assume otherwise. Then $\mu(x) \leq \mu(a) + \|ax\|/3 = \|va\|/2 + \|ax\|/3$ and $\mu(y) \leq \mu(b) + \|by\|/3 = \|vb\|/2 + \|by\|/3$ and $\|xy\| > \|xa\|/2 + \|ab\| + \|by\|/2$ since γ' turns by less than $\pi/3$. Thus $c_\mu(x, y) > 0$, a contradiction to our sampling condition. Thus either $v \leq x \leq a$ or $b \leq y \leq v$ or both. We may assume without loss of generality that $v \leq x \leq a$. If $y < b$, we have $\mu(y) \leq \mu(b) + \|by\|/3 = \|vb\|/2 + \|by\|/3$, $\|xy\| > \|xv\|/2 + \|vb\| + \|by\|/2$, and $\mu'(x) = \|vx\|/2$ and hence $c_{\mu'}(x, y) > 0$. If $b \leq y$, we have $\|xy\| > \|xv\|/2 + \|vy\|/2$, $\mu'(y) = \|vy\|/2$ and $\mu'(x) = \|vx\|/2$ and hence $c_{\mu'}(x, y) > 0$. 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. This completes the definition of the modified potential function in the second case.

In order to verify that all minimum spanning trees are subsets of the polygonal reconstruction is suffices to show that Lemma 3.6 holds for the new potential function and that $c_{\mu'}(p, r) > c_{\mu'}(p, q)$ for any three points with $p < q < r$ and $\{p, q, r\} \subseteq \gamma'$ for some $\gamma' \in \Gamma$.

From $\mu'(x) \leq \mu(x)$ for all x, we conclude $c_{\mu'}(pq) \leq c_\mu(pq)$ for all edges pq. Thus Lemma 3.6 stays true.

If $\{p, q, r\} \subseteq B(u, d(u)) \cap T$ for some $r \in \gamma$, we have $\angle(\vec{pq}, \vec{qr}) < \pi/3$ and hence $\|pr\| - \|pq\| > \|qr\|/2$ and $\mu'(r) \leq \mu'(q) + \|qr\|/2$. Thus $c_{\mu'}(p, r) > c_{\mu'}(p, q)$.

Assume next that $\{p, q, r\} \subseteq B(s, \delta_s)$ for some sharp corner s. If $c_{\mu'}(p, r) \geq 0$ we are done. Otherwise $c_\mu(p, r) < 0$, since the modified distance wrt. μ is at most the modified distance wrt. μ'. We make the same case distinction as in the proof of Lemma 3.9. In this proof we bounded $\|pr\| - \|pq\|$ from below and $\mu(r) - \mu(q)$ from above. We now need to bound $\mu'(r) - \mu'(q)$. Since $\mu'(x) = \mu(x)$ for $x \notin T$, we only need to reconsider the case that r and q are in T. We have $\mu'(r) \leq \mu'(q) + \|qr\|/2$ and hence the arguments used in the cases $p < q < r \leq s$, $s \leq p < q < r$, and $p < q < s < r$ stay valid. We only need to reconsider the case $p < s \le q < r$.

If $\mu(p) \ne \mu'(p)$, we have $\mu(q) = \mu'(q)$ and $\mu(r) = \mu'(r)$ and are done. So assume $\mu(p) = \mu'(p)$. If $\mu(q) = \mu'(q)$ we have $c_{\mu'}(p,r) - c_{\mu'}(p,q) \ge c_\mu(p,r) - c_\mu(p,q) \ge 0$, since $\mu'(r) \le \mu(r)$. Otherwise $\mu'(r) - \mu'(q) \le 0 \le \mu(r) - \mu(q)$, since r is closer to v than q and both q and r are in $T \cap B(s, \delta_s)$. Thus $c_{\mu'}(p,r) - c_{\mu'}(p,q) \ge c_\mu(p,r) - c_\mu(p,q) \ge 0$.

6. Solving the Subtour-LP. The Subtour-LP has an exponential number of constraints. The Ellipsoid method [Sch86] allows one to solve LPs with an exponential number of constrains in time polynomial in the number of variables if the following conditions are satisfied:

- The coefficients of the variables in the constraints are polynomially bounded. This is the case for the Subtour-LP.
- The separation problem can be solved in polynomial time, i.e., given a vector x^*_{uv} of polynomially bounded values for the variables, one can decide in polynomial time, whether the vector satisfies all constraints and, if not, exhibit a violated constraint. This is the case. It is trivial to check the degree constraints and the constraint that all values lie between zero and one. In order to check the subtour elimination constraints (we discuss the case of tours, having already discussed the cas of path in Sections 3 and 4). Consider the complete network on V and assign capacity x^*_{uv} to edge uv. Consider any subset V' of V with $\emptyset \ne V' \ne V$ and observe that $2|\,V'\,| = \sum_{u \in V'} \sum_{v \in V} x^*_{uv} = 2\sum_{u \in V', v \in V'} x^*_{uv} + \sum_{u \in V'} \sum_{v \notin V'} x^*_{uv}$. We conclude that the subtour elimination constraint for V' is satisfied iff the capacity of the cut $(V', V \setminus V')$ is at least two. Some subtour elimination constraint is satisfied if the minimum cut is less than two. A minimum cut can be computed in polynomial time.
- The coefficients in the objective function are polynomially bounded. This is not the case since the bit-representation of an Euclidean length is in principal infinite. Furthermore there are curves where the bit-representation of any point is infinite.

We show that it suffices to know the sample points and their Euclidean distances only approximately. More precisely, we show the following: Let S be a set of points. Let m be the minimal distance between any two points in S and for a point $p \in S$ let p^* be a closest point of γ. If

1. $\|pp^*\| \le \rho m/10$ for all $p \in S$, where ρ is a constant depending on γ, and
2. the set $S^* = \{p^* \mid p \in S\}$ satisfies a strengthened sampling condition

then the Subtour-LP has a unique optimal integral solution even when the distances between sample points are only known up to an error of $m\rho/10$. Moreover, the Subtour-LP can be solved in polynomial time and the optimal solution is a tour connecting the points in S in the order in which the points S^* lie on γ. We also show how to estimate ρ from the sample set and without knowledge of γ.

Now we make precise what we mean by a approximate sample set of a curve γ. Recall that for an angle α we have defined $\bar\alpha$ by $\bar\alpha = \pi - \alpha$. Let $\alpha_0 = \max_{s \in S} \alpha_s$. We define $\bar\alpha_0 = 17\pi/24$, if there is no sharp corner.

DEFINITION 6.1. *Let γ be a benign semi-regular curve. We call a set S of points $p \in \mathbb{Q}^2$ an* approximate sample set *of a curve γ, if for all $p \in S$*

$$\|p\gamma\| \le (\sin \bar\alpha_0/4)^2 m/30.$$

For the following let $\rho = (\sin \bar\alpha_0/4)^2 m/3$. Then $m > 7\rho$ (since $\bar\alpha_0 < 17\pi/24$) and $\|pp^*\| \le \rho/10$ for the points in an approximate sample. We need the following sampling condition.

Sampling Condition for Approximate Sample Sets:

(a) For any two adjacent (on γ) samples u^* and v^*: $\|u^*v^*\| \le 9/10(\mu(u^*) + \mu(v^*))$.

(b) For any two adjacent samples u^* and v^*: $\gamma[u^*, v^*]$ turns by less than π. For two adjacent points p, q on the curve, $\gamma[p,q]$ denotes the subcurve of γ with endpoints p and q not containing an other sample point.

Let γ be a benign semi-regular curve and S an approximate sample set satisfying the sampling condition. For any two sample points p and q let $||pq||_\approx$ be a rational number, so that $\|pq\| - ||pq||_\approx \le \rho/10$. Then $\|p^*q^*\| - ||pq||_\approx \le 3\rho/10$. Note that for all $p, q \in S$ there exists a choice of $||pq||_\approx$ which has a bit representation of polynomial size in m and $\bar\alpha_0$. We consider the approximate Subtour-LP of the approximate sample set S.

THEOREM 6.2. *Let γ be an open (closed) benign semi-regular curve and S be an approximate sample set of γ satisfying the sampling condition above.*

- *The approximate Subtour-LP for S has an unique optimal integral solution.*
- *The approximate Subtour-LP can be solved in polynomial time.*

Proof. We define the potential function of an approximate sample point p as $\mu(p) := \mu(p^*)$. We call the new modified cost function $c_\mu^{\approx}$. We show:

1. If p and q are adjacent sample points then $c_\mu^{\approx}(p,q) < -3\rho/10$ and hence there is a minimum spanning tree in which each edge has length less than $-3\rho/10$.
2. If p and q are sample points which are not contained in some $\gamma' \in \Gamma$ then $c_\mu^{\approx}(p,q) \geq -3\rho/10$ and hence no such edge belongs to a minimum spanning tree.
3. If $p < q < r$ are three sample points contained in some $\gamma' \in \Gamma$ and $c_\mu^{\approx}(p,r) < -3\rho/10$ then $c_\mu^{\approx}(p,r) > c_\mu^{\approx}(pq)$, and hence the minimum spanning tree reconstructs locally.

We turn to the first item. Let p and q be adjacent sample points. We have $6\rho < m - \rho \leq \|pq\| - \rho \leq \|p^*q^*\| + 2\rho/10 - \rho \leq \|p^*q^*\| \leq \mu(p^*) + \mu(q^*)$. Thus

$$\begin{aligned}
c_\mu^{\approx}(p,q) &= \|pq\|_{\approx} - \mu(p) - \mu(q) \\
&\leq \|p^*q^*\| + 3\rho/10 - 1/10(\mu(p^*) + \mu(q^*)) - 9/10(\mu(p^*) + \mu(q^*)) \\
&< \|p^*q^*\| - 3\rho/10 - 9/10(\mu(p^*) + \mu(q^*)) \\
&\leq -3\rho/10.
\end{aligned}$$

For the second item, consider points p and q that are not contained in any common subcurve $\gamma' \in \Gamma$. We have

$$\begin{aligned}
c_\mu^{\approx}(p,q) &\geq \|p^*q^*\| - 3\rho/10 - \mu(p^*) - \mu(q^*) \\
&= c_\mu(p^*,q^*) - 3\rho/10 \\
&\geq -3\rho/10.
\end{aligned}$$

We come to the third item. Consider three sample points $p < q < r$ that are contained in some $\gamma' \in \Gamma$ and for which $c_\mu^{\approx}(pr) < -3\rho/10$. Then $c_\mu(p^*,r^*) < 0$ and hence (using Lemmas 3.7 and 3.9 and the fact that $(\sin\bar{\alpha}_0/4)^2/3 \leq 1/6$):

$$\begin{aligned}
c_\mu^{\approx}(p,r) - c_\mu^{\approx}(p,q) &\geq c_\mu(p^*,r^*) - c_\mu(p^*,q^*) - 6\rho/10 \\
&\geq \|qr\|(\sin\bar{\alpha}_0/4)^2/3 - 6\rho/10 \\
&\geq \rho - 6\rho/10 > 0.
\end{aligned}$$

□

What have we achieved at this point. We have shown that the Subtour-LP reconstructs provided our sample set S satisfies the sampling condition for approximate sample sets *and* we are given approximate distances of polynomial size that differ by at most $\rho/10$ from the true distances. We could compute approximate distances with the required property, if we were given ρ or alternatively α_0 as an additional input. We now show how to compute a lower bound bound on α_0, which leads to a polynomial precision in the input size, without any additional knowledge of the curve.

LEMMA 6.3. *Let m and M be the minimal, respectively maximal, distance between two sample points. Then $\sin(\bar{\alpha}_0/4) \geq m/(15M)$.*

Proof. If γ has no sharp corners, $\alpha_0 = 17\pi/24$ and there is nothing to show. So assume otherwise and let s be any sharp corner. We prove that there is a sample point p in $B(s,\delta_s) \cap T$ on each leg of the sharp corner and the use this fact to bound $\bar{\alpha}_s$ from below.

We look at an arbitrary of the two orders obtained by splitting the curve at s and prove that there is a sample point behind s in $B(s,\delta_s) \cap T$. Assume otherwise. Let x be the first sample point behind s outside $B(s,\delta_s)$ and let y be the sample point preceding x. Since every edge of the polygonal reconstruction must lie in at least one subcurve $\gamma' \in \Gamma$, y must lie behind s. By assumption y does not lie in T. Assume first that $x \in T$. Then $\|xy\| \geq 4\delta_s/5$ and $\mu(y) \leq 2\delta_s/5$. Thus $c_\mu(x,y) = \|xy\| - \mu(x) - \mu(y) \geq \|xy\| - 2\delta_s/5 - \|xy\|/3 \geq 2\|xy\|/3 - 2\delta_s/5 \leq 8\delta_s/15 - 2\delta_s/5 > 0$, a contradiction. Assume now $x \notin T$. Let s' be the corner so that $x \in B(s',\delta_{s'})$ and assume w.l.o.g. $\delta_s \geq \delta_{s'}$. Then $\mu(x) \leq 2\delta_s/5$, $\mu(y) \leq 2\delta_s/5$ and $\|xy\| \geq 2\delta_s - \delta_s/5 - \delta_{s'}/5 \geq \delta_s$. Thus $c_\mu(x,y) > 0$, a contradiction.

Let p be the first sample point behind s in $B(s,\delta_s) \cap T$ and let q be the adjacent sample point behind p. Then $q \in T$; if $q \notin T$ then $q \in B(s',\delta_{s'})$ for some sharp corner s' and hence s' would have no sample point in $B(s',\delta_{s'}) \cap T$. The distance between p and q is at least m. Also $\|pq\| \leq d(p)/3 + d(q)/3 \leq$

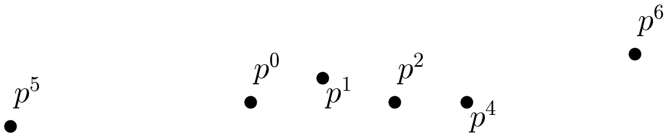

FIG. 7.1. *The points p^1 up to p^5 are joined with p^0, but p^6 is not joined.*

$d(p)/3 + (d(p) + \|pq\|)/3$ and so $d(p) \geq \|pq\| \geq m$. Consider the intersections of the two legs of s with the boundary of the δ_s-ball centered at s. The intersections have distance at least m (since d-values grow along each leg) and s sees the intersections under an angle of at least $\bar{\alpha}_0/2$. Thus $\sin \bar{\alpha}_0/4 \geq m/(2\delta_s)$. Since there is at least one sample outside the ball $B(s, \delta_s)$ and at least one sample inside the ball $B(s, \delta_s/5)$, we have $M \geq 4\delta_s/5$. Thus $\sin \bar{\alpha}_0/4 \geq 4m/(10M)$. □

7. Collections of Closed Curves. In the preceding sections we showed that the Subtour-LP formulation of the Traveling Salesman problem is able to reconstruct *single* closed and open curves. In this section we extend the algorithm so that is can handle collections of closed curves. We do not know how to handle collections of open and closed curves. Please note that the algorithms [DMR00, FR01] can handle open and closed curves.

The algorithm works in rounds. The first round constructs an initial partition of the sample points and subsequent rounds merge blocks of the partition. The construction of the initial partition and the merging is done conservatively, i.e., all points in the same block provably belong to the same curve. In the first round, every point is joined to points close to it; Section 7.1 gives the details. In later rounds, see Section 7.2 we solve the Subtour-LP for each block and then analyze the solution. If the Subtour-LP fails on a block or if curves constructed for different blocks interfere, some blocks are merged.

Throughout this section we assume our set of sample points to satisfy a somewhat strengthened sampling condition. The strengthened sampling condition leads to denser sampling near sharp corners. We change μ to μ' by decreasing the δ- and c-values of sharp corners. We set $\delta'_s \leq \delta_s$ such that for any two points p and q in $B(s, \delta'_s/7)$ the angle between the segment pq and the corresponding tangent in s is at most $\pi/40$. Furthermore we decrease the value c_s to c'_s such that $c'_s - \|sp_s\| \leq d(p)/3$ for every $p \notin B(s, \delta'_s/60)$. This enlarges the region where the potential is defined by $d(p)/3$ from T to T'. Recall that the choice of c_s only guaranteed that the points outside $B(s, \delta_s/5)$ belonged to T, whereas T' contains all points outside the balls $B(s, \delta'_s/60)$.

7.1. The Initial Partition. We define a graph on our set of sample points. The connected components of this graph form the initial partition. For a sample point $p = p^0$ let p^1, p^2, ... be the other sample points in order of increasing distance (ties are broken arbitrary). We always join p^0 with p^1 and p^2. We join p^0 and p^i, $i \geq 3$, if $\angle \overrightarrow{p^{k-1}p^{k-2}}, \overrightarrow{p^{k-1}p^k} \geq 2\pi/3$ for all k with $2 \leq k \leq i-1$. Observe, that the decision whether p^i is joined to p^0 depends only on the points p^0 up to p^{i-1}, but not on the point p^i itself. This is essential for making connections between the points on different legs of a sharp corner but also hinders the extension to open curves. Figure 7.1 illustrates the definition.

LEMMA 7.1. *If p and its two adjacent sample points are in T', then p is only joined with points in $B(p, d(p))$ and is joined with both adjacent sample points.*

Proof. Since p is in T', $\gamma \cap B(p, d(p))$ consists of a single component and turns by less than $\pi/3$. Also the two sample points q, r adjacent to p on γ lie in $B(p, d(p))$. We show this for q. We have $d(q) \leq d(p) + \|pq\|$, $\mu(q) = d(q)/3$ since $q \in T'$, and, by our sample condition, $\|pq\| \leq d(p)/3 + d(q)/3$. Thus $\|pq\| \leq d(p)/3 + (d(p) + \|pq\|)/3$ and hence $\|pq\| \leq d(p)$.

Assume w.l.o.g. that q is considered before $r = p^i$. Orient $\gamma \cap B(p, d(p))$ such that $r < p < q$. Then $q = p^1$ and $p^1 < p^2 < \ldots < p^{i-1}$. Since $\gamma \cap B(p, d(p))$ turns by less than $\pi/3$, we have $\angle \overrightarrow{p^{k-1}p^{k-2}}, \overrightarrow{p^{k-1}p^k} \geq 2\pi/3$ for all k with $2 \leq k \leq i-1$ and $\angle \overrightarrow{p^{i-1}p^{i-2}}, \overrightarrow{p^{i-1}p^i} \leq \pi/3$. Thus p is joined with p^1 up to p^i, but not with p^{i+1}. □

We turn to the non-smooth parts of the curve. We first show that our sampling condition implies that each leg of a sharp corner must contain several sample points.

LEMMA 7.2. *a) Let p be a sample point in $B(s, \delta'_s) \setminus B(s, 2\delta'_s/60)$. Then both adjacent sample points lie on the same leg as p, the one closer to s has distance at most $\|ps\|/2$ from p and the one further from s has distance at most $\|ps\|$ from p.*

b) For every leg ℓ of a sharp corner s we have at least one sample point in $B(s, 2 \cdot 2^j \delta'_s/60) \setminus B(s, 2^j \delta'_s/60)$ for $j = 1, 2, 3, 4$.

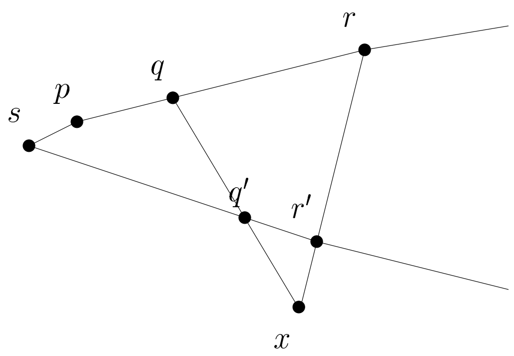

FIG. 7.2. *The angle between $q'q$ and $r'r$ must be large that the region is grown further. This contradicts the fact that the distance between q and r is short.*

Proof. a) We have $\mu'(p) \leq \|sp\|/3$. The point q on the same leg as p with distance exactly $\|sp\|/2$ to s lies in T' and hence $\mu'(q) \leq \|sp\|/6$. Thus $c_{\mu'}(p,q) \geq \|sp\|/2 - \|sp\|/6 - \|sp\|/3 = 0$. Lemma 3.9 implies that $c_{\mu'}(p,x) \geq 0$ for any point x between s and q and for any point x on the other leg. Thus there must be a sample point between p and q. We conclude that both adjacent sample points lie on the same leg as p and that the one closer to s satisfies the distance constraint stated. For the one further from p we consider the point q on the same leg as p, further away from s than p, and having distance $\|sp\|$ from p. Then $\mu'(q) \leq 2\|sp\|/3$ and hence $c_{\mu'}(p,q) \geq \|sp\| - 2\|sp\|/3 - \|sp\|/3 = 0$. We now argue as above.

We turn to part b). Part a) implies that if one of the annuli contains a point, the adjacent annuli do also and hence all annuli do. We conclude that either all annuli contain a sample point or none does. Assume the latter and let p be the first sample point on ℓ outside $B(s,\delta'_s)$. The sample point q preceding p must lie in $B(s, 2\delta'_s/60)$. Thus $\|pq\| \geq \|ps\| - 2\delta'_s/60$, $\mu'(p) \leq \|ps\|/3$, $\|ps\| \geq \delta'_s$, and $\mu'(q) \leq \max(2\delta'_s/60, c_s) \leq 2\delta'_s/5$, a contradiction. □

LEMMA 7.3. *A sample point p is only joined with points of the same curve.*

Proof. Consider a sample point p. If p and both adjacent sample points of p are in T', the claim follows from Lemma 7.1.

Otherwise $p \in B(s, 2\delta'_s/60)$ by Lemma 7.2, part a). Let $p = p^0, p^1, \ldots, p^j$ be the sample points in $B(s,\delta'_s)$ ordered according to their distance from p. Since both legs of the sharp corner contain sample points in the annuli $B(s, 2\cdot 2^j\delta'_s/60) \setminus B(s, 2^j\delta'_s/60)$ for $j = 2,3,4$, we must have a subsequence of length at least three such that the first and the last element of the subsequence, call them q and r respectively, lie on the same leg as p and further away from s than p, and the points in between (there is at least one) lie on the other leg. Figure 7.2 illustrates the situation. We show that no point after r is joined to p. Assume otherwise. Let q' be the point added directly after q and r' be the point added directly before r.

We want to bound the angle between the segments $q'q$ and $r'r$. If q' is equal to r' this angle is at least $2\pi/3$. Otherwise the angle between $q'q$ and the segment between q' an the sample q'' added after q' is at least $2\pi/3$ and the angle between $r'r$ and the segment between r' an the sample r'' added before r' is at least $2\pi/3$ (since r is not the last point joined to p). Also q', q'', r'', and r' lie all on the same leg and hence the angle between $q'q$ and $r'r$ is at least $\pi/3 - 4\pi/40$. Here we use the strengthened sampling condition.

Let x be the intersection between the lines supporting qq' and rr'. We have $d(q) < d(r) \leq d(q) + \|qr\|$, $\|rr'\| \geq d(r)$ and $\|qq'\| \geq d(q)$ by Lemma 3.3, $\|qr\| \leq (d(q)+d(r))/3$ by our sampling condition, and hence $\|rx\| \geq \|rr'\| \geq d(r) \geq (d(r)+d(q))/2 \geq 3\|qr\|/2$ and $\|qx\| \geq \|qq'\| \geq d(q) \geq (d(q)+d(r)-\|pq\|)/2 \geq \|qr\|$. Application of the "theorem of cosines[6]" with $D = \|qr\|$ yields $\cos(\angle(q'q, r'r)) \geq [(3D/2)^2 + D^2 - D^2]/(2 \cdot D \cdot 3D/2) = 3/4$ and hence $\angle(q'q, r'r) < \pi/3 - \pi/10$, a contradiction. □

LEMMA 7.4. *For any curve γ_j all sample points in $T' \cap \gamma_j$ belong to the same component.*

Proof. Breaking γ_j at its sharp corners gives us a collection of subcurves γ_{jk}. Consider any subcurve γ_{jk}. The sample points on γ_{jk} come in three groups: first a group of points outside T', then a group of points in T', and finally a group of points outside T'. Lemma 7.1 implies that all sample points in $\gamma_{jk} \cap T'$ belong to a single component. Consider now two adjacent subcurves incident to a sharp corner s. In both subcurves all points in T' belong to the same component. Let p and q be the points in the components containing the points T' which are closest to s and do not connect to both neighbors. We

[6] $\|qr\|^2 = \|qx\|^2 + \|rx\|^2 - 2\|qx\| \cdot \|rx\| \cdot \cos(\angle(q'q, r'r))$

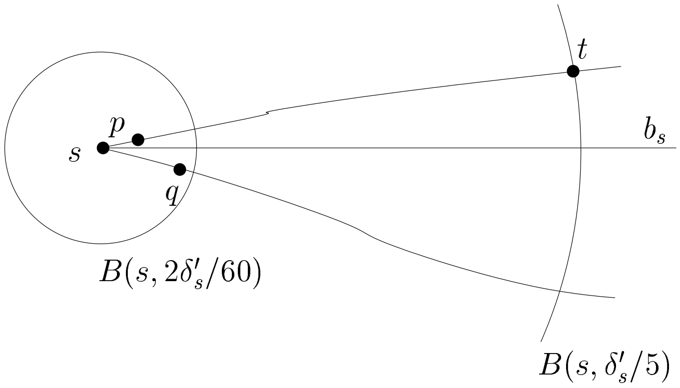

FIG. 7.3. *The definition of* t.

claim that p and q are connected to points on the other leg. Assume otherwise, say p is not connected to a point on the other leg. Since p is not connected to both neighbors we have $p \in B(s, \delta'_s/60)$. Let $p = p^0, p^1, \dots, p^i$ be the sample points connected to p and ordered according to their distance from p. Then $s < p < p^1 < \dots < p^i$ since p is not connected to both neighbors and since p is not connected to a sample point on the other side. Also we have shown in the proof of Lemma 7.3 that $p^i \in B(s, \delta'_s)$ Thus $\angle(p^{i-1}p^{i-2}, p^{i-1}p^i) \geq 2\pi/3$, a contradiction to the fact that p^i is the last point joined with p.

Thus p connects to a point u on the other leg and q connects to a point v on the other leg. If either u or v belong to the component containing the points in T' we are done. So assume otherwise. Then u is closer to s than q and v is closer to s than p and hence $pvuq$ builds a convex quadrangle. The two segments pu and qv are crossing. Thus either $\|pv\| < \|pu\|$ or $\|qu\| < \|qv\|$ (since $\|pv\| + \|qu\| < \|pu\| + \|qv\|$ and hence either p or q will be joined with a sample closer to s, a contradiction. $\square$

We call the component containing all sample points in $\gamma_j \cap T'$ the main component of γ_j.

LEMMA 7.5. *The Subtour-LP applied to the main component of* γ_j *reconstructs* γ_j.

Proof. We show that the sample points in $\gamma_j \cap T'$ satisfy our original sampling condition. Consider two points p and q in $\gamma_j \cap T'$ that are adjacent along γ_j. If they are also adjacent in the full sample, we are done. Assume otherwise. Then $\{p, q\} \subseteq B(s, 2\delta'_s/60)$ for some sharp corner s. Let t be the point such $c_s = d(t)/3 + \|st_s\|$, see Figure 7.3. We have $\|pq\| \leq 2 \cdot (2\delta'_s/60) \cdot \sin 3\overline{\alpha}_s/4 + \|sp_s\| - \|sq_s\|$, since $\angle(sp, sp_s) \leq 3\overline{\alpha}_s/4$ and similarly for q, $d(t) \geq \delta'_s/5 \sin \overline{\alpha}_s/2$, since the ball around t with radius $\delta_s/5 \sin \overline{\alpha}_s/2$ does not intersect the other leg, $\mu(p) \geq c_s - \|sp_s\| = d(t)/3 + \|st_s\| - \|sp_s\| = d(t)/3 + \|t_s p_s\| \geq \frac{\delta'_s}{15} \sin \overline{\alpha}_s/2 + \|t_s p_s\|$ and, by the same argument, $\mu(q) \geq \frac{\delta'_s}{15} \sin \overline{\alpha}_s/2 + \|t_s q_s\|$. Thus $c_\mu(p, q) \leq \frac{4\delta'_s}{60} \sin 3\overline{\alpha}_s/4 + \|sp_s\| - \|sq_s\| - 2\frac{\delta'_s}{15}\frac{2}{3} \sin 3\overline{\alpha}_s/4 - \|t_s p_s\| - \|t_s q_s\| \leq 0$. $\square$

7.2. Merging Components. The initial partition may contain too many components. For every curve γ_j, there is a main component which contains all sample points in $\gamma_j \cap T'$ and maybe other components. We call them minor components. Each minor components is contained in $B(s, 2\delta'_s/60)$ for some sharp corner s. The reconstruction based on the Subtour-LP is guaranteed to succeed for the main component of every curve. For the minor components it may or may not produce a tour. In this section we describe a strategy for merging components.

We define a region R_{pq} for every edge pq of the computed reconstruction. If the region R_{pq} for an edge pq of the computed reconstruction contains a sample point from another component, we join the component containing p and q with the component of the point closest to pq and lying in another component. We continue until the components stabilize.

Before we define the regions R_{pq} we draw an important consequence of the merging rule. For an edge in a minor component the point closest to it and in another component is guaranteed to lie on the same curve. This follows from the fact that a minor component is contained in $B(s, 2\delta'_s/60)$ for some sharp corner, that the corresponding major component has a point in $B(s, 4\delta'_s/60)$, and that any point within $B(s, \delta'_s)$ belongs to the same curve.

In the arguments to follow we can therefore concentrate on edges in the reconstruction of the main component. In particular, we can use the fact that the Subtour-LP correctly reconstructs the main component. We come to the definition of the regions R_{pq}. We define R_{pq} as the union of a region R'_{pq} and the circumcircle of the segment pq. For every sample point p we define β_p to be the angle between

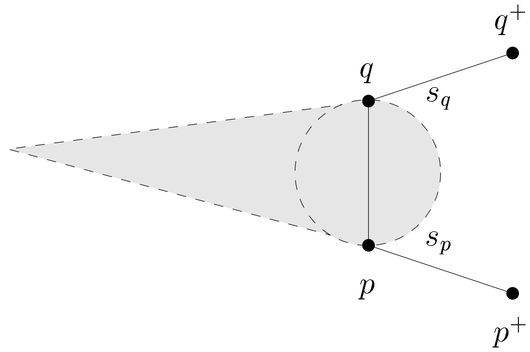

FIG. 7.4. *The definition of R_{pq}.*

the two segments incident to p. The following paragraph motivates our definition of the region R'_{pq}.

Assume pq is the segment in the main component connecting the two legs of a sharp corner s (we say that the edge straddles the sharp corner) and let $\bar{\alpha}_s$ be the angle between the two tangents at s. Let s_p be a segment defined by two adjacent sample points on the leg of p, let s_q be a segment defined by two adjacent sample points on the leg of q, both lying in $B(s, \delta'_s/7)$. Let θ_s be the angle formed by the segments. Since either segment forms an angle less than $\bar{\alpha}_s/4$ with the corresponding tangent at s, we have $\bar{\alpha}_s/2 \leq \theta_s \leq 3\bar{\alpha}_s/2$ or $2\theta_s/3 \leq \bar{\alpha}_s \leq 2\theta_s$, i.e., θ_s is a good estimator for $\bar{\alpha}_s$. Since again the angle between any tangent on a leg and the appropriate tangent in the corner is at most $\alpha_s/4$, we know that the angle between the corner point and the points p and q is between $\bar{\alpha}_s/2$ and $3\bar{\alpha}_s/2$, thus between $\theta_s/3$ and $3\theta_s$.

We come to the definition of R'_{pq}. For an edge pq let p^+ and q^+ be the other neighbors of p and q, respectively, and let θ_s be the angle between the segments pp^+ and qq^+, i.e.$\theta_s = \beta_p + \beta_q - \pi$. We define R'_{pq} as the set of all points r with

- $\angle(\overrightarrow{rp}, \overrightarrow{rq}) \geq \theta_s/3$,
- $\angle(\overrightarrow{pr}, \overrightarrow{pq}) \leq \pi - \beta_p + \min(\pi/20, \theta_s)$ and r lies on the opposite halfspace with respect to the line pq as p^+, and
- $\angle(\overrightarrow{qr}, \overrightarrow{pq}) \geq \pi - \beta_q + \min(\pi/20, \theta_s)$ and r lies on the opposite halfspace with respect to the line qp as q^+.

For an illustration of this definition see Figure 7.4. Note that $\pi - \beta_q + \min(\pi/20, \theta_s) \leq \pi - \beta_q + \theta_s = \beta_p \leq \pi$.

LEMMA 7.6. *If the main component of a curve does not yet contain all sample points from the curve, it will grow.*

Proof. The main component contains all points in T'. Consider a sample point on γ_j which does not belong to the main component and let p and q be its adjacent sample points in the main component. Then $\{p, q\} \subseteq B(s, 2\delta'_s/60)$ for some sharp corner. If p and q lie on the same leg of s the centerball of pq contains the subcurve between p and q and if pq straddles the sharp corner, the region R'_{pq} contains the subcurve. ☐

To prove that we do not merge components that do not belong to the same curve, we need the following three Lemmas.

LEMMA 7.7. *Every point r in R'_{pq} has distance at most $\|pq\|/sin(\theta_s/3)$ from p and q.*

Proof. Let m_1 and m_2 be the points on the perpendicular bisector of pq with distance $\|pq\|/(2\sin(\theta_s/3))$ from p and q (see Figure 7.5).

By the Inscribed Angle Theorem, every point which sees pq under an angle of at least $\theta_s/3$ lies inside the union of the the balls with center m_1 or m_2 through p and q. Thus any point in the region R_{pq} that distance at most $\|pq\|/\sin(\theta_s/3)$ from p and q. ☐

LEMMA 7.8. $\theta_s \geq \alpha_s/2$.

Proof. If the segment pq straddles the corner, $\theta_s \geq \alpha_s/2$, since pp^+ lies completely on one leg and qq^+ lies completely on the other leg. If p^+ and q^+ are on the same leg, this follows directly from the sampling condition. If the segment pp^+ straddles the corner, the angle formed by the segments pp^+ and qq^+ is smaller than the angle formed by the segments sp^+ and qq^+, which is at least $\alpha_s/2$. ☐

LEMMA 7.9. *Let pq be a segment of the polygonal reconstruction, with $p, q \in T$ and $d(p) \geq d(q)$. Then $d(p) \geq 3\|pq\|/2$.*

Proof. Assume otherwise. Than $c_\mu(p, q) = \|pq\| - d(p)/3 - d(q)/3 \geq \|pq\| - 2d(p)/3 > \|pq\| + \|pq\| \geq 0$.

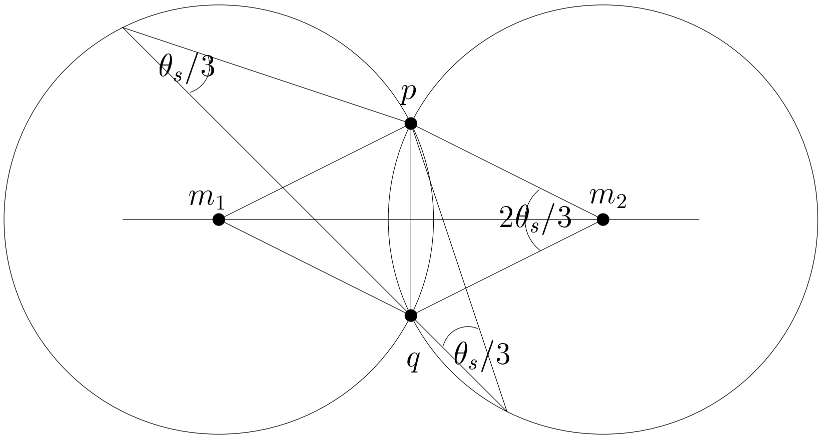

FIG. 7.5. *The Inscribed Angle Theorem: The central angle is equal to twice the inscribed angle*

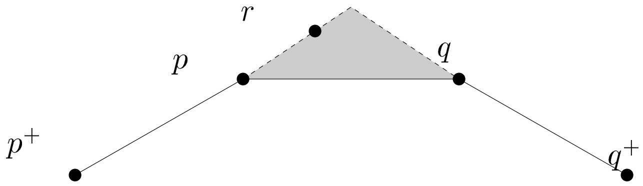

FIG. 7.6. *The segments rp and rq form a small angle with the segment pq. Thus $\|pr\|$ can not be much larger as $\|pq\|$.*

□

We now turn to the proof that we do not merge a curve with points from an other component.

LEMMA 7.10. *The region R_{pq} of an edge pq in the polygonal reconstruction of the main component contains no sample point of another component.*

Proof. This is obvious for the center ball of pq. We turn to the region R'_{pq}.

Assume first p and q are in T' and w.l.o.g. $d(p) \geq d(q)$. Assume there is a point r outside the the $B(p, d(p))$ ball in R'_{pq}. Then $\|pr\| \geq d(p) \geq 3\|pq\|/2$ by Lemma 7.9.

We know $\angle(\overrightarrow{pq}, \overrightarrow{pp^+}) \geq 2\pi/3$ and $\angle(\overrightarrow{pp^+}, \overrightarrow{pr}) \geq 19\pi/20$. Hence $\angle(\overrightarrow{pq}, \overrightarrow{pr}) \leq (\pi - 2\pi/3) + (\pi - 19\pi/20) = 23\pi/60$, see Figure 7.6. Analogously $\angle(\overrightarrow{qp}, \overrightarrow{qr}) \leq 23\pi/60$, thus $\angle(\overrightarrow{rp}, \overrightarrow{rq}) \geq 7\pi/30$. So[7] $\|pr\| \leq \sin(23\pi/60)/\sin(7\pi/3)\|pq\| < 3\|pq\|/2$.

Assume next that one of p and q does not belong to T'.

We will show that $R'_{pq} \subseteq B(s, \delta'_s)$. Let r be any point in R'_{pq} and let θ be the angle under which r sees the segment pq. Then $\theta \geq \theta_s/3$ by the definition of R'_{pq}.

Assume first $\theta_s \geq \pi/6$. We know $\|pr\| \leq \|pq\|/\sin(\theta_s/3) \leq 6\|pq\|$. Thus $\|sr\| \leq 6(8\delta'_s/60)+2\delta'_s/60 < \delta'_s$.

Assume now $\theta_s < \pi/6$. Look at the triangle $\triangle pqr$ and assume w. l. o. g. $\pi - \beta_p > \pi/2$ (see Figure 7.7). We show that the corner point s is almost as far away from p as r. The angle at r is at least $\theta_s/3$, the angle at p is at most $\pi + \theta_s - \beta_p$. Thus $\|pr\| \leq \|pq\| \sin(\pi + \theta_s - \beta_p)/\sin(\theta_s/3)$.

The sharp corner s form an angle of at most $3\theta_s$ with points p and q, since $\bar{\alpha}_s \leq 2\theta_s$ by Lemma 7.8. The angle at p of the triangle $\triangle pqs$ is between $\pi + \theta_s - \beta_p$ and $\pi - \theta_s + \beta_p$.

We have to distinguish two cases according to β_p. Assume first $\beta_p > \pi - 2\theta_s$. We conclude $\|rs\| \leq \|ps\| + \|pr\| \leq 2\delta'_s/60 + \|pq\| \sin(\pi + \theta_s - \beta_p)/\sin(\theta_s/3) \leq 2\delta'_s/60 + \|pq\| \sin(3\theta_s)/\sin(\theta_s/3) \leq 2\delta'_s/60 + 10\|pq\| \leq 2\delta'_s/60 + 40\delta'_s/60 \leq \delta'_s$.

Assume now $\beta_p < \pi - 2\theta$. We conclude $\|rs\| \leq \|ps\| + \|pr\| \leq 2\delta'_s/60 + \|pq\| \sin(\pi + \theta_s - \beta_p)/\sin(\theta_s/3) \leq 2\delta'_s/60 + \|ps\|(\sin(3\theta)/\sin(\pi - \theta_s + \beta_p))(\sin(\pi + \theta_s - \beta_p)/\sin(\theta_s/3)) \leq 2\delta'_s/60 + \|ps\|(\sin(3\theta_s)/\sin(\theta_s))(\sin(3\theta_s)/\sin(\theta_s/3)) \leq 2\delta'_s/60 + 27\|ps\| \leq 2\delta'_s/60 + 52\delta'_s/60 < \delta'_s$. □

8. Curve Reconstruction and the Delaunay Diagram.

Most previous curve reconstruction algorithms use sampling conditions that guarantee that the polygonal reconstruction is a subset of the

[7] In this proof we make frequent use of the fact that $a/\sin\alpha = b/\sin\beta = c/\sin\gamma$ for a triangle with sides a, b, c and corresponding angles α, β, and γ.

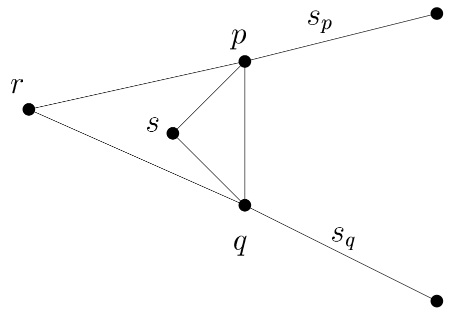

FIG. 7.7. *The two triangles pqr and pqs.*

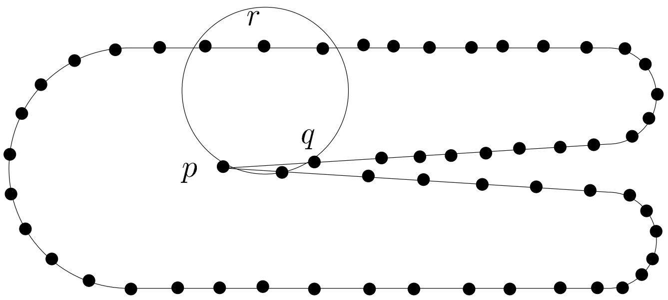

FIG. 8.1. *The edge pq does not belong to the Delaunay Triangulation; p is a sharp corner with* $\delta_p = \|pr\|/2$. *Also* $c_p \approx 2\delta_p/5 = \|pr\|/5$. *If* $\|pq\| < \|pr\|/5$, *the edge pq has negative reduced cost. Thus our sampling condition is satisfied.*

Delaunay diagram. Our sampling condition does not imply that the Traveling Salesman Tour is a subgraph of the Delaunay triangulation, see Figure 8.1. This fact can be interpreted positively and negatively: positively, as an indication of the strength of the TSP-reconstruction, and negatively, since the optimal Traveling Salesman tour must be searched for in the complete graph on the sample set. In this section we show that a slight strengthening of our sample condition implies that the polygonal reconstruction is contained in the Delaunay Diagram.

Additional condition on the sample set: An edge pq of the polygonal reconstruction with an endpoint not in T has length at most $4\delta_s \sin(\alpha_s/2)/5$, where s is the sharp corner with $\{p,q\} \subseteq B(s,\delta_s)$.

LEMMA 8.1. *If the sample set V satisfies the strengthened sampling condition, the polygonal reconstruction is contained in the Delaunay diagram of V.*

Proof. Let pq be an edge of the polygonal reconstruction. We construct a Delaunay ball B for it.

Assume first that $\{p,q\} \subseteq T$. The center ball of pq is contained in $B(p,d(p))$ and hence empty of sample points; otherwise, γ would either intersect the ball in more than one component or turn by more than $\pi/3$ within the ball.

Assume next that one of the endpoint of the edge, say p, does not lie in T. Let s be the sharp corner with $p \in B(s,\delta_s)$. We distinguish the cases whether p and q are on the same leg of s or not.

Assume first that p and q lie on different legs, see Figure 8.2. We put the center m of B into the halfspace containing s and having p and q in its boundary and set the radius r of B to $(\|pq\|/2)/\sin(\bar{\alpha}_s/2) \le 2\delta_s/5$; the upper bound on the radius follows from the strengthened sampling condition. Since one of p or q is contained in $B(s,\delta_s/5)$, we conclude $B(m,r) \subseteq B(s,\delta_s)$.

Assume now that there is a sample point $x \in B(m,r)$. We discuss the case that x lies on the same leg as q and leave the other case to the reader. Since p and q are adjacent samples, x cannot lie on the segment sq and hence $\angle(\vec{sp},\vec{qx}) = \angle(\vec{sp},\vec{sq}) + \angle(\vec{sq},\vec{qx}) \le \angle(\vec{qp},\vec{sq}) + \angle(\vec{sq},\vec{qx}) = \angle(\vec{qp},\vec{qx})$ (We have $\angle(\vec{sp},\vec{sq}) < \angle(\vec{qp},\vec{sq})$, since moving along the segment $\vec{sq}$ increase the angle). Since x lies in the lune of $B(m,r)$ defined by pq, we have $\angle(\vec{qx},\vec{qp}) \le \bar{\alpha}_s/2$. Thus $\angle(\vec{sp},\vec{qx}) \le \bar{\alpha}_s/2$. On the other hand, the angle between the two tangents at s is $\bar{\alpha}_s$ and hence $\angle(\vec{sp},\vec{qx}) > \bar{\alpha}_s - 2\bar{\alpha}_s/4 = \bar{\alpha}_s/2$, a contradiction.

Assume next that p and q lie on the same leg, see Figure 8.3. The center m of B lies on the same side of the angular bisector of the cone defined by the tangents $\mathrm{t}_r(s), -\mathrm{t}_l(s)$ as p and q and sees the segment pq under an angle of $\bar{\alpha}_s$. As above, we conclude that $B(m,r) \subseteq B(s,\delta_s)$. Then x must be contained in the lune of $B(m,r)$ defined by the segment pq. Since p and q are adjacent sample points, x must lie

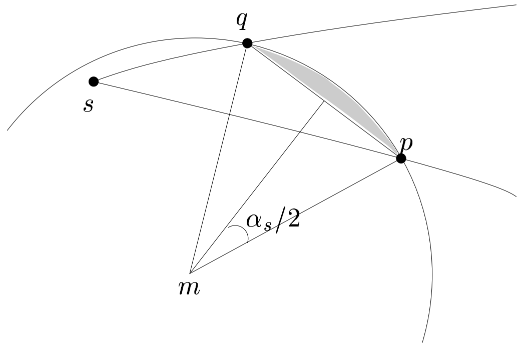

FIG. 8.2. *The sample x must lie on the upper leg behind q and in the shaded lune.*

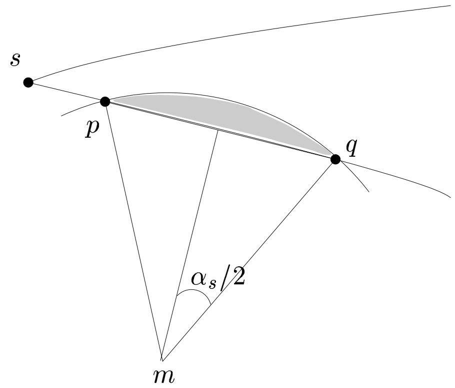

FIG. 8.3. *The sample x must lie on the upper leg and in the shaded lune.*

on the other leg. Assume w.l.o.g. that p is closer to s than q. We have $\angle(\vec{pq}, \vec{px}) \leq \bar{\alpha}_s/2$ and hence $\bar{\alpha}_s = \angle(\mathrm{t}_r(s), -\mathrm{t}_l(s)) < \angle(\vec{ps}, \vec{px}) + 2(\bar{\alpha}_s/4) \leq \bar{\alpha}_s$, a contradiction. ☐

9. Monotonicity. Intuitively, a larger sample set makes the reconstruction task simpler. We discuss how various sampling conditions and reconstruction algorithms behave with respect to larger sample sets.

A sampling condition is called *monotone*[8] if any superset of a set satisfying the sampling condition also satisfies the sampling condition.

For closed curves and open curves with specified endpoints, our sampling condition is monotone. For open curves with unspecified endpoints the superset must satisfy the additional constraint that the additional points must lie in $\gamma[a, b]$, where a and b are the extreme sample points in the old sample. The sampling conditions used in the papers [ABE98, DK99, Gol99, DMR00] are also monotone; again the additional constraint is needed for open curves. The conditions in [FR01] are not monotone.

All algorithms mentioned in the present paper come with a guarantee of the form: If the curve γ is from a certain class of curves and the sample set V is sufficiently dense, the algorithm will reconstruct γ. It is not specified what the algorithm does, if the hypothesis of the theorem is not satisfied. The algorithm may either fail, i.e., indicate that it could not find a curve, or "invent" a curve. From a practical point of view this situation is unsatisfactory as a user has in general no way of distinguishing reconstruction from invention. The situation is aggravated by the fact that the sampling densities required by the theorems are quite high and that the algorithms tend to work for smaller densities and hence are likely to be used in situations not covered by the theorems. *It would be nice to have algorithms that never invent curves.*

A reconstruction algorithm is called *self-consistent* if is has the following property. On an input V it either outputs `FAILURE` or `SUCCESS`. In the latter case it also outputs a curve Γ passing through V such that for any sample V' from Γ with $V \subseteq V'$, it will also output Γ. A reconstruction algorithm that is not self-consistent can change its mind if given additional sample points that seem to confirm the output of the algorithm.

THEOREM 9.1. *The algorithm in [DMR00] and the TSP-algorithm are self-consistent, the algorithms in [ABE98, DK99, Gol99, FR01] are not self-consistent.*

[8]Be aware, this choice of name reflects a bias on behalf of the authors.

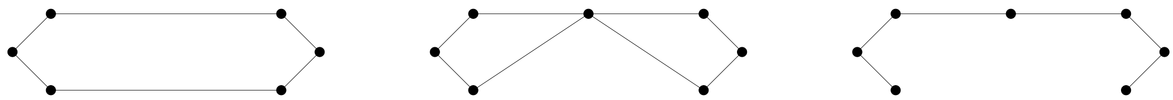

FIG. 9.1. *For the sample set in the left figure, the algorithms in [ABE98, DK99, Gol99] produce the hexagon, as shown. If one adds a sample point in the middle of one of the long segments, the algorithm in [DK99] produces the output of the middle figure, the algorithms in [ABE98, Gol99] produce the output of the right figure. Thus none of these algorithms is self-consistent.*

Proof. The algorithm in [DMR00] is constructed to be self-consistent. For the algorithms in [ABE98, DK99, Gol99, FR01] it is easy to come up with examples that show non-self-consistency (see Figure 9.1).

It remains to show self-consistency for the TSP-algorithm. We show that if the solution for the Subtour-LP for a set V is unique and integral then the same is true for the Subtour-LP for $V \cup \{z\}$ for any z on an edge of the integral solution of the Subtour-LP of V (and different from all points in V).

The claim is a simple consequence of the so-called splitting-off lemma (see [Lov79], Problem 6.53). Consider an optimal solution of the Subtour-LP for $V \cup \{z\}$. The value C of this solution is at most the length C_0 of the optimal tour of V (since the tour for V is also a tour for $V \cup \{z\}$). The splitting-off lemma allows to construct a solution for the Subtour-LP for V from the solution for $V \cup \{z\}$. It implies the existence of a set of triples (e, f, r), where e and f are edges incident to z and r is a non-negative real, such that

- for each edge e incident to z the sum of the third components of all triples containing e is equal to x_e (the value of the edge e in the Subtour-LP for $V \cup \{z\}$) and
- a solution for the Subtour-LP for V can be obtained by modifying the solution for the Subtour-LP for $V \cup \{z\}$ as follows: for each edge uv with $u \neq z$ and $v \neq z$ increase x_{uv} by the sum of the third components of all triples (uz, zv, r). Delete all edges incident to z.

For any edge $e = uv$ let y_e be the increase of x_e in this construction. The cost of the obtained solution is $C + \sum_{uv} y_{uv}(c_{uv} - c_{uz} - c_{zv})$. This cost is at most C (and hence at most C_0) since $c_{uv} \leq c_{uz} + c_{zv}$ by the triangle inequality. Since the solution for the Subtour-LP for V is unique and is equal to the tour for V, the cost of the solution cannot be smaller than C_0 and hence for any edge $e = uv$ with $y_{uv} > 0$ we must have $c_{uv} = c_{uz} + c_{zv}$, i.e., z lies on the line segment $\overline{uv}$. Moreover, $y_e + x_e$ must be integral for every edge $e = uv$.

In the tour for V there is only one edge passing through z (since optimal tours are non-self-intersecting) and hence there can be only one edge uv with $y_{uv} > 0$. Thus our set of triples consists of a single triple (uz, zv, r) and since the degree constraint at z must be satisfied for the optimal solution of the Subtour-LP for $V \cup \{z\}$ we conclude that $r = 1$. We conclude that the optimal solution of the Subtour-LP for $V \cup \{z\}$ is unique and integral. ☐

10. Our Sample Condition and the Local Feature Size. The papers [ABE98, DK99, Gol99, DMR00] 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 $\|pv\| \leq \epsilon \cdot f(p)$; here ϵ is a parameter which depends on the algorithm. All algorithms require $\epsilon \leq 1/2$.

The experimental results in [AMNS00] 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 10.1. *Let γ be a smooth curve and $\epsilon < 1/10$. If for any $p \in \gamma$ there is a sample point $v \in V$ with $\|pv\| \leq \epsilon \cdot f(p)$, then V satisfies our sampling condition.*

Before we prove Lemma 10.1, we show the following Lemma.

LEMMA 10.2.

$$f(p) < 3d(p)$$

Proof. The following fact was shown in [ABE98]

FACT 5. *Let r be a point of a smooth curve γ. Let furthermore q and s be points on γ with distance less than $f(r)$ from r with $q < r < s$. Then $\angle(\vec{rq}, \vec{rs}) > \pi/3$.*

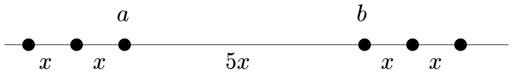

FIG. 11.1. *The underlying part of the curve is a line segment. Hence, there is no condition on the sample set. In a necklace tour the disks centered at a and b have radius at most $2x$. Thus the two disks will not intersect.*

By the definition of $d(p)$, either $B(p, d(p)) \cap \gamma$ is not connected or $B(p, d(p))$ contains three points turning by $\pi/3$.

In the first case, there is a medial axis point in $B(p, d(p))$ by Lemma 1 of [ABE98] and hence $f(p) \leq d(p)$.

We turn to the second case. Let $q < r < s \in B(p, d(p)) \cap \gamma$ forming an angle of $\pi/3$. By Fact 5 we conclude $f(r) \leq \max(\|qr\|, \|rs\|)$. Thus $f(p) \leq \|pr\| + \max(\|qr\|, \|rs\|) < 3d(p)$. □

Proof. [of Lemma 10.1]

We have to show that for an ϵ–sampled curve, with $\epsilon < 1/10$, the modified cost of the edge between two adjacent sample points is less then 0.

Let p and q be adjacent sample points. Let w.l.o.g. $d(p) \leq d(q)$. Lemma 3.4. of [DK99] states $\|pq\| \leq 2\epsilon f(p)/(1-\epsilon) = 1/5 f(p)/(1-1/10)$. Thus $c_\mu(p,q) = \|pq\| - d(p)/3 - d(q)/3 < 1/5 f(p)/(1-1/10) - f(p)/9 - f(p)/9 \leq 0$ □

Our sample condition depends on several parameters and we have set these parameters to particular values in Section 3. In Section 3.6 we discussed the dependency between the parameters. For smooth curves we can set f_{scale} to $1/2$ and then the argument above works for $\epsilon = 1/7$

11. Necklace Tours.

We have shown that curve reconstruction gives rise to a polynomially solvable case of the Euclidean traveling salesman problem. In this section we relate our results to a 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 so that each disk intersects with exactly two other disks and so 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 that 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 11.1 shows an example of a TSP problem which is covered by Theorem 2.3 and whose solution is not a necklace tour.

12. Conclusions.

We have shown that TSP-based curve reconstruction permits a polynomial time algorithm. In the companion paper [AMNS00] we show that TSP-based curve reconstruction is reasonably efficient and works for smaller sampling density than other reconstruction algorithms. We have also shown how to extend TSP-based reconstruction to collections of closed curves. The generalization to collections of open and closed curves stays open.

Surface reconstruction is the three-dimensional analogue of curve reconstruction. There are no surface reconstruction algorithms that can handle non-smooth surfaces. Whether global minimization can also be applied to surface reconstruction remains an open question. A first step in this direction was recently taken by [AGJ00]

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.

[AGJ00] U. Adamy, J. Giesen, and M. John. New techniques for topologically correct surface reconstruction. In *IEEE Visualization 2000*, 2000. Accepted for publication.

[AMNS00] E. Althaus, K. Mehlhorn, S. Näher, and S. Schirra. Experiments on curve reconstruction. In *ALENEX*, pages 103–114, 2000. www.mpi-sb.mpg.de/ mehlhorn/ftp/exp-curve.ps.

[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'97)*, 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.

[DMR00] T.K. Dey, K. Mehlhorn, and E.A. Ramos. Curve reconstruction: Connecting dots with good reason. *Computational Geometry: Theory and Applications*, 15(4):229–244, 2000. www.mpi-sb.mpg.de/ mehlhorn/ftp/cure.ps.gz.

[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.

[FR01] S. Funke and E.A. Ramos. Reconstructing collections of curves with corners and endpoints. In *Proceedings of the 8th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'01)*, 2001. Accepted for publication.

[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.

[Lov79] L. Lovász. *Combinatorial problems and exercises*. North-Holland Publishing Co., Amsterdam, 1979.

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