Connect the Dot: Computing Feed-Links with Minimum Dilation*

Boris Aronov[1], Kevin Buchin[2], Maike Buchin[3], Marc van Kreveld[3], Maarten Löffler[3], Jun Luo[4], Rodrigo I. Silveira[3], and Bettina Speckmann[2]

[1] Dep. Computer Science and Engineering, Polytechnic Institute of NYU, USA
aronov@poly.edu
[2] Dep. of Mathematics and Computer Science, TU Eindhoven, The Netherlands
{kbuchin,speckman}@win.tue.nl
[3] Dep. of Information and Computing Sciences, Utrecht University, The Netherlands
{maike,marc,loffler,rodrigo}@cs.uu.nl
[4] Shenzhen Institute of Advanced Technology, Chinese Academy of Sciences, China
jun.luo@sub.siat.ac.cn

Abstract. A *feed-link* is an artificial connection from a given location p to a real-world network. It is most commonly added to an incomplete network to improve the results of network analysis, by making p part of the network. The feed-link has to be "reasonable", hence we use the concept of dilation to determine the quality of a connection.

We consider the following abstract problem: Given a simple polygon P with n vertices and a point p inside, determine a point q on P such that adding a feedlink $\overline{pq}$ minimizes the maximum dilation of any point on P. Here the *dilation* of a point r on P is the ratio of the shortest route from r over P and $\overline{pq}$ to p, to the Euclidean distance from r to p. We solve this problem in $O(\lambda_7(n) \log n)$ time, where $\lambda_7(n)$ is the slightly superlinear maximum length of a Davenport-Schinzel sequence of order 7. We also show that for convex polygons, two feed-links are always sufficient and sometimes necessary to realize constant dilation, and that k feed-links lead to a dilation of $1 + O(1/k)$. For (α, β)-covered polygons, a constant number of feed-links suffices to realize constant dilation.

1 Introduction

Network analysis is a type of geographical analysis on real-world networks, such as road, subway, or river networks. Many facility location problems involve network analysis. For example, when a location for a new hospital needs to be chosen, a feasibility study typically includes values that state how many people

* This research has been supported by the Netherlands Organisation for Scientific Research (NWO) under BRICKS/FOCUS grant number 642.065.503, under the project GOGO, and under project no. 639.022.707. B. Aronov has been partially supported by a grant from the U.S.-Israel Binational Science Foundation, by NSA MSP Grant H98230-06-1-0016, and NSF Grant CCF-08-30691. M. Buchin is supported by the German Research Foundation (DFG) under grant number BU 2419/1-1.

F. Dehne et al. (Eds.): WADS 2009, LNCS 5664, pp. 49–60, 2009.

would have their travel time to the nearest hospital decreased to below 30 minutes due to the new hospital location. In a more global study of connectivity, one may analyze how many households are reachable within 45 minutes from a fire station. In this case, the households are typically aggregated by municipality or postal-code region, and the centroid of this region is taken as the representative point. This representative point might not lie on the road network. It might even be far removed from it, since nation-wide connectivity studies seldomly use detailed network data for their analysis. A similar situation occurs when the quality of the network data is not very high. In developing countries, data sets are often incomplete due to omissions in the digitization process, or due to lack of regular updates. In both cases a network study must be executed that involves a set of locations that are not connected to the network in the available data.

A workable solution in such cases is to connect the given locations to the known road network by *feed-links*. A feed-link is an artificial connection between a location and the known network that is "reasonable", that is, it is conceivable that such a connection exists in the real world [2,6]. A road network forms an embedded, mostly planar graph. Hence a location that does not lie on the network, lies inside some face of this graph. Such a face can be represented by a simple polygon. A feed-link is then a connection from the given location to the boundary of the simple polygon.

When computing feed-links we need to be able to judge their quality. That is, we have to assess if a particular connection could possibly exist in reality. To do this, we use the concept of *dilation*, also known as *stretch factor* or *crow flight conversion coefficient*. People in general do not like detours, so a connection that causes as little detour as possible, is more likely to be "real". Given an embedded plane graph, the dilation of two points p and q on the graph is the ratio of their distance within the graph to their Euclidean distance. The concept of dilation is commonly used in computational geometry for the construction of *spanners*: a t-spanner is a graph defined on a set of points such that the dilation between any two points is at most t, see [7,10,13,14,15].

In this paper we consider a single point p inside a simple polygon, whose boundary we denote by P. We solve the problem of placing one feed-link between p and P so that the maximum dilation over all points on P to p is minimized. We assume that a feed-link is a straight-line connection between p and exactly one point q on P. We allow the feed-link $\overline{pq}$ to intersect P in more points, see Fig. 1 (left), but assume that it is not possible to "hop on" the feed-link at any such point other than q (the white points in the figure provide no access to the feed-link). Fig. 1 (middle) shows that the feed-link yielding minimum dilation may intersect P in a point other than q. One could also choose to disallow feed-links that intersect the outside of P, or to use geodesic shortest paths inside P as feed-links, and measure the dilation of any point on P with respect to its geodesic distance to p. We also study the problem of connecting several feed-links to p to bound the dilation. Then any point on P uses exactly one of the feed-links to reach p over the network. Fig. 1 (right) shows that $n/2$ feed-links may be necessary to bound the dilation by a constant, if P has n vertices.

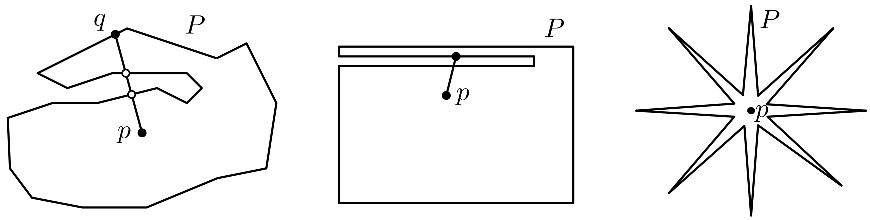

Fig. 1. A feed-link that intersects P gives no access to the feed-link other than q (left). A minimum dilation feed-link may intersect P in the interior of the feed-link (middle). Simple polygons may require many feed-links to achieve constant dilation (right).

In a recent paper [2] we showed how to compute the dilation of a polygon when a collection of feed-links to a point inside is given. We also gave heuristic algorithms to place one or more feed-links and compared them experimentally on generated polygons. The simple heuristic for one feed-link that connects p to the closest point on P is a factor-2 approximation for the optimal feed-link placement. We also studied the problem of placing as few feed-links as possible to realize a specified dilation. A simple incremental algorithm exists that uses at most one more feed-link than the minimum possible.

Results. In Section 2 we give an efficient algorithm to compute an optimal feed-link. For a simple polygon with n vertices, our algorithm runs in $O(\lambda_7(n) \log n)$ time, where $\lambda_7(n)$ is the maximum length of a Davenport-Schinzel sequence of order 7, which is only slightly superlinear [1,16]. If we are interested in the dilation with respect to only m fixed points on P, the running time reduces to $O(n + m \log m)$. Furthermore, we give a $(1 + \varepsilon)$-approximation algorithm for the general problem that runs in $O(n + (1/\varepsilon) \log(1/\varepsilon))$ time, for any $\varepsilon > 0$. The results in this section also hold with geodesic dilation and feed-links, or with feed-links that are not allowed to intersect the outside of P.

In Section 3.1 we show that for any convex polygon and any point inside, two feed-links are sufficient and sometimes necessary to achieve constant dilation. In this case the dilation is at most $3 + \sqrt{3}$. There are convex polygons where no two feed-links can realize a dilation better than $2 + \sqrt{3}$. We also show that we can realize a dilation of $1 + O(1/k)$ with k feed-links. Finally, in Section 3.2 we show that for (α, β)-covered polygons [8] (a class of realistic polygons), a constant number of feed-links suffices to obtain constant dilation. This result does not hold for most other classes of realistic polygons.

Notation. P denotes the boundary of a convex or simple polygon, and p is a point inside it. For two points a and b on P, $P[a, b]$ denotes the portion of P from a clockwise to b, its length is denoted by $\mu(a, b)$. Furthermore, $\mu(P)$ denotes the length (perimeter) of P. The Euclidean distance between two points p and q is denoted by $|pq|$. For two points q and r on P, the dilation of point r when the feed-link is $\overline{pq}$ is denoted by $\delta_q(r)$. For an edge e, $\delta_q(e)$ denotes the maximum dilation of any point on e when the feed-link is $\overline{pq}$.

2 Computing One Feed-Link with Minimum Dilation

Let $v_0, \ldots, v_{n-1}$ be the vertices of P and let p be a point inside P. We seek a point q on P such that the feed-link $\overline{pq}$ minimizes the maximum dilation to any point on P. We first consider the restricted case of minimizing the dilation only for m given points on P. Then we solve the general case. In both cases, the feed-link may connect to any point on P.

Let r be a point on P and let r' be the point opposite r, that is, the distance along P between r and r' is exactly $\mu(r, r') = \mu(r', r) = \mu(P)/2$. For any given location of q, r has a specific dilation. We study the change in dilation of r as q moves along P. If $q \in P[r', r]$, then the graph distance between p and r is $|pq| + \mu(q, r)$, otherwise it is $|pq| + \mu(r, q)$.

We choose a fixed point v_0 on P and define two functions cw-dist(q) and ccw-dist(q) that measure the distance from p to v_0 via the feed-link $\overline{pq}$ and then from q either clockwise or counterclockwise along P, see Fig. 2. The dilation $\delta_q(r)$ of r can be expressed using either cw-dist(q) or ccw-dist(q), depending on the order in which v_0, q, r, and r' appear along P. In particular, we distinguish four cases that follow from the six possible clockwise orders of v_0, q, r, and r':

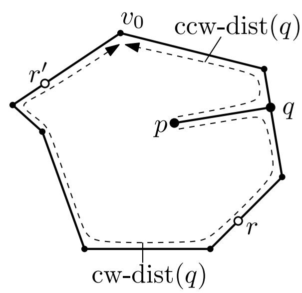

Fig. 2. cw-dist(q) and ccw-dist(q); shown is case 1 with order v_0qrr'

1. If the clockwise boundary order is v_0qrr' or $v_0r'qr$, then the dilation is $\delta_q(r) = (\text{cw-dist}(q) - \mu(r, v_0))\,/\,|rp|$.
2. If the clockwise boundary order is $v_0rr'q$, then the dilation is $\delta_q(r) = (\text{cw-dist}(q) + \mu(v_0, r))\,/\,|rp|$.
3. If the clockwise boundary order is $v_0qr'r$, then the dilation is $\delta_q(r) = (\text{ccw-dist}(q) + \mu(r, v_0))\,/\,|rp|$.
4. If the clockwise boundary order is v_0rqr' or $v_0r'rq$, then the dilation is $\delta_q(r) = (\text{ccw-dist}(q) - \mu(v_0, r))\,/\,|rp|$.

As q moves along P in clockwise direction, starting from v_0, three of the cases above apply consecutively. Either we have $v_0qrr' \to v_0rqr' \to v_0rr'q$, or $v_0qr'r \to v_0r'qr \to v_0r'rq$. We parameterize the location of q both by cw-dist(q) and ccw-dist(q). This has the useful effect that the dilation $\delta_q(r)$ of r is a linear function on the intervals where it is defined (see Fig. 3). In particular, for a fixed point r, $\delta_q(r)$ consists of three linear pieces. Note that we cannot combine the two graphs into one, because the parameterizations of the location of q by cw-dist(q) and ccw-dist(q) are not linearly related. This follows from the fact that cw-dist(q)+ ccw-dist(q) $= \mu(P) + 2 \cdot |pq|$.

We now solve the restricted case of minimizing the dilation only for m given points on P. For each point r we determine the line segments in the two graphs that give the dilation of r as a function of cw-dist(q) and ccw-dist(q). These line segments can be found in $O(n + m)$ time in total. Next, we compute the upper

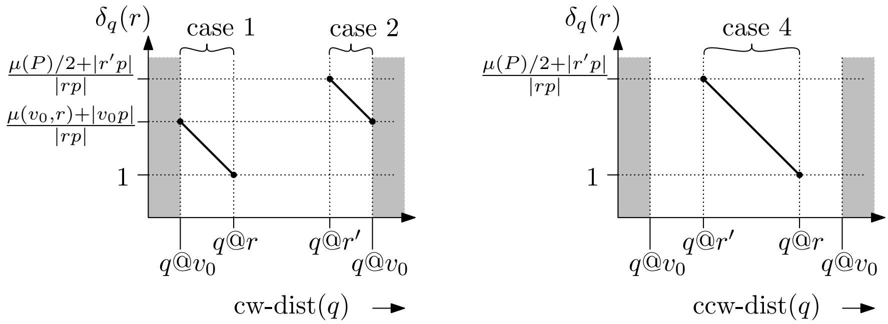

Fig. 3. Two graphs showing the dilation of a point r as a function of cw-dist(q) (left) and ccw-dist(q) (right); $q@r$ indicates "q is at position r"

envelope of the line segments in each of the two graphs. This takes $O(m \log m)$ time using the algorithm of Hershberger [12], and results in two upper envelopes with complexity $O(m \cdot \alpha(m))$. Finally, we scan the two envelopes simultaneously, one from left to right and the other from right to left, taking the maximum of the corresponding positions on the two upper envelopes, and recording the lowest value encountered. This is the optimal position of q.

To implement the scan, we first add the vertices of P to the two envelopes. Since we need to compute the intersection points of the two envelopes we must unify their parameterizations. Consider the locations of q that fall within an interval I which is determined by two envelope edges e_1 and e_2. Since cw-dist$(q) = -$ccw-dist$(q) + 2 \cdot |pq| + \mu(P)$, the line segment of one envelope restricted to I becomes a hyperbolic arc in the parametrization of the other envelope. Hence e_1 and e_2 can intersect at most twice in a unified parametrization, and the scan takes time linear in the sum of the complexities of the two envelopes.

Theorem 1. *Given the boundary P of a simple polygon with n vertices, a point p inside P, and a set S of m points on P, we can compute the feed-link (which might connect to any point on P) that minimizes the maximum dilation from p to any point in S in $O(n + m \log m)$ time.*

Next we extend our algorithm to minimize the dilation over all points on P. Let $r_e(q)$ denote the point with the maximum dilation on a given edge e of P. Instead of considering the graphs of the dilation for a set of fixed points, we consider the graphs for the points $r_e(q)$ for all edges of P. The positions of $r_e(q)$ change with q. The graphs of the dilation do not consist of line segments anymore, but of more complex functions, which, however, intersect at most six times per pair, as we prove in the full paper. As a consequence, we can compute their upper envelope in $O(\lambda_7(n) \log n)$ time [12], where $\lambda_7(n)$ is the maximum length of a Davenport-Schinzel sequence of order 7, which is slightly superlinear [1,16].

Theorem 2. *Given the boundary P of a simple polygon with n vertices and a point p inside P, we can compute the feed-link that minimizes the maximum dilation from p to any point on P in $O(\lambda_7(n) \log n)$ time.*

Note that our algorithms ignore the degenerate case where p lies on a line supporting an edge e of P. In this case cw-dist(q) and ccw-dist(q) are both constant on e. This is in fact easy to handle, as we describe below when discussing geodesic dilation.

We can adapt our algorithms to not allow feed-links that intersect the exterior of P. We first compute the visibility polygon $V(p)$ of p with respect to P. The vertices of $V(p)$ partition the edges of P into parts that are allowed to contain q and parts that are not. The number of parts is $O(n)$ in total, and they can be computed in $O(n)$ time.

We compute the upper envelopes exactly as before. Before we start scanning the two envelopes, we add the vertices of P *and also* the vertices of the visibility polygon to the two envelopes. The envelopes now have the property that between two consecutive vertices, a feed-link is allowed everywhere or nowhere. During the scan, we keep the maximum of the dilation functions and record the lowest value that is allowed. The time complexity of our algorithms does not change.

We can also adapt our algorithms to use geodesic feed-links and geodesic shortest distances. In this case the feed-link is a geodesic shortest path between p and q, and the dilation of a point r on P is defined as the ratio of the graph distance between r and p (necessarily via q) and the geodesic shortest path between r and p.

By computing the shortest path tree of p inside P, we obtain the geodesic shortest distances of p to every vertex of P, and hence can partition P into $O(n)$ parts, such that the first vertex on a geodesic shortest path to p is the same (this first vertex can also be p itself) [11].

When we use cw-dist(q) and ccw-dist(q) to represent the location of q, we use the length of the geodesic from q to p instead of $|pq|$, plus the clockwise or counterclockwise distance to v_0. But now a value of cw-dist(q) or ccw-dist(q) does not necessarily represent a unique position of q anymore: when q traverses an edge of P and the geodesic from q to p is along this edge in the opposite direction, cw-dist(q) and ccw-dist(q) do not change in value. However, it is sufficient to consider only the location of q that gives the shortest feed-link (if any such feed-link is optimal, then the shortest one is optimal too). All other adaptations to the algorithms are straightforward, and we obtain the same time complexities.

3 Number of Feed-Links vs. Dilation

In this section we study how many feed-links are needed to achieve constant dilation. We immediately observe that there are simple polygons that need $n/2$ feed-links to achieve constant dilation, see Fig. 1. For convex polygons, we establish that two feed-links are necessary and sufficient to obtain constant dilation. For realistic ("fat") simple polygons, there are several definitions one can use to capture realism ([4,5,8,9,17] and others). Most of these definitions are not sufficient to guarantee constant dilation with a constant number of feed-links. However, for (α, β)-covered polygons [8] we can show that a constant number of feed-links suffices for constant dilation.

3.1 Convex Polygons

Let P be the boundary of a convex polygon and let p be a point inside P. We explore how many feed-links are necessary and sufficient to guarantee constant dilation for all points on P.

One feed-link is not sufficient to guarantee constant dilation. Consider a rectangle with width w and height $h < w$, and let p be its center. One of the points in the middle of the long sides will have dilation greater than $2w/h$, which can become arbitrarily large. Hence two feed-links may be necessary.

Two feed-links are also sufficient to guarantee constant dilation for all points on P. In fact we argue that we can always choose two feed-links such that the dilation is at most $3 + \sqrt{3} \approx 4.73$. This bound is not far from the optimum, since an equilateral triangle with p placed in the center has dilation at least $2 + \sqrt{3} \approx 3.73$ for any two feed-links. To see that, observe that one of the sides of the equilateral triangle does not have a feed-link attached to it (or only at a vertex), which causes the middle of that side to have dilation at least $2 + \sqrt{3}$.

Let q be the closest point to p on P. We choose $\overline{pq}$ as the first feed-link and argue that the dilation is now constant for all points in some part of P which includes q. Then we show how to place the second feed-link to guarantee constant dilation for the remaining part of P. Consider the smallest equilateral triangle Δ that contains P and which is oriented in such a way, that one of its edges contains q. Let e_0 be the edge of Δ containing q, and let e_1 and e_2 be the other edges, in clockwise order from e_0 (see Fig. 4). By construction, each edge of Δ is in contact with P. Let t_1 be a point of P in contact with e_1, and let t_2 be a point of P in contact with e_2.

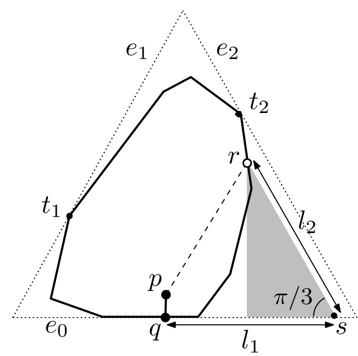

Fig. 4. The smallest equilateral triangle that contains P

Lemma 1. *For any point $r \in P[t_2, t_1]$, $\delta_q(r) \leq 3 + \sqrt{3}$.*

We prove Lemma 1 by arguing that $\mu(r, q) \leq l_1 + l_2$. The details can be found in the full paper. The second feed-link connects p to the point q' on $P[t_1, t_2]$ closest to p. Lemma 2 can be proven with similar arguments as Lemma 1.

Lemma 2. *For any point $r \in P[t_1, t_2]$, $\delta_{q'}(r) \leq 3 + \sqrt{3}$.*

These two lemmas jointly imply

Theorem 3. *Given the boundary P of a convex polygon and a point p inside it, two feed-links from p to P are sufficient to achieve a dilation of $3 + \sqrt{3}$.*

We now consider the general setting of placing k feed-links, where k is a constant. We prove that placing the feed-links at an equal angular distance of $\eta = 2\pi/k$ guarantees a dilation of $1 + O(1/k)$. To simplify the argument we choose $k \geq 6$ (the result for smaller k immediately follows from the result for two feed-links). Our proof uses the following lemma.

Lemma 3. *Let q_1 and q_2 be two points on the boundary P of a convex polygon such that the angle $\angle q_1 p q_2 = \eta \leq \pi/3$. Then for all points $r \in P[q_1, q_2]$, we have $\delta(r)/\max(\delta(q_1), \delta(q_2)) \leq 1 + \eta$.*

Note that $\overline{pq_1}$ and $\overline{pq_2}$ need not be feed-links in Lemma 3. The lemma implies that for $\eta = 2\pi/k$ we obtain the following result.

Theorem 4. *Given the boundary P of a convex polygon and a point p inside it, k feed-links from p to P are sufficient to achieve a dilation of $1 + O(1/k)$.*

Approximation algorithm for convex polygons. We can use Lemma 3 to obtain a linear-time $(1 + \varepsilon)$-approximation algorithm to place one feed-link optimally. We measure dilation only at $2\pi/\varepsilon$ points on P, and hence the running time of the approximation algorithm is $O(n + (1/\varepsilon)\log(1/\varepsilon))$ by Theorem 1. The points at which we measure the dilation are chosen on P such that the angle between two consecutive points measured at p is ε. Since Lemma 3 bounds the dilation between two consecutive points, the theorem follows.

Theorem 5. *For any $\varepsilon > 0$, given the boundary P of a convex polygon with n vertices and a point p inside it, we can compute a feed-link that approximately minimizes the maximum dilation from p to any point on P within a factor $1 + \varepsilon$ in $O(n + (1/\varepsilon)\log(1/\varepsilon))$ time.*

3.2 Realistic Polygons

A constant number of feed-links should guarantee constant dilation for *realistic* polygons. Therefore, we define a simple polygon to be *feed-link realistic* if there are two constants $\delta > 1$ and $c \geq 1$, such that there exist c feed-links that achieve a dilation of at most δ for any point on its boundary. Many different definitions of realistic polygons exist in the literature. We show that most of them do not imply feed-link realism. However, we also argue that polygons that are (α, β)-*covered* [8] are feed-link realistic.

Consider the left polygon in Fig. 5. At least c feed-links are required to obtain a dilation smaller than δ, if the number of prongs is c and their length is at least

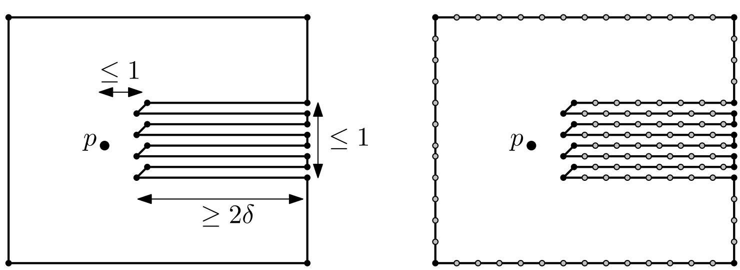

Fig. 5. A β-fat polygon (left) and an adaptation (right) that require many feed-links

δ times larger than the distance of their leftmost vertex to p. No feed-link can give a dilation at most δ for the leftmost vertex of more than one dent. However, the polygon is β-fat [5]. Definitions that depend on the spacing between the vertices or edge-vertex distances will also not give feed-link realism, because the left polygon in Fig. 5 can be turned into a realistic polygon according to such definitions. We simply add extra vertices on the edges to get the right polygon: it has edge lengths that differ by a factor of at most 2, it has no vertex close to an edge in relation to the length of that edge, and it has no sharp angles. The extra vertices obviously have no effect on the dilation. This shows that definitions like low density (of the edges) [17], unclutteredness (of the edges) [4,5], locality [9], and another fatness definition [18] cannot imply feed-link realism.

(α, β)-covered polygons. For an angle ϕ and a distance d, a (ϕ, d)*-triangle* is a triangle with all angles at least ϕ and all edge lengths at least d. Let P be the boundary of a simple polygon, let $\text{diam}(P)$ be the diameter of P, and let $0 < \alpha < \pi/3$ and $0 < \beta < 1$ be two constants. P is (α, β)*-covered* if for each point on P, an $(\alpha, \beta \cdot \text{diam}(P))$-triangle exists with a vertex at that point, whose interior is completely inside P [8]. Furthermore, P is (α, β)*-immersed* if for each point on P there is such a triangle completely inside P and one completely outside P. For ease of description, we assume that $\text{diam}(P) = 1$.

We use a result by Bose and Dujmović [3] that bounds the perimeter of P as a function of α and β.

Lemma 4. *The perimeter of an (α, β)-covered polygon P is at most $\frac{c}{\beta \sin \alpha}$, for some absolute constant $c > 0$.*

Also, we need a technical lemma that states that if the distance between two points on P is short enough, then it is proportional to the Euclidean distance.

Lemma 5. *If p and q can see each other on the inside of an (α, β)-covered polygon P and $\mu(p, q) < \beta$, then $\mu(p, q) < f(\alpha) \cdot |pq|$, where $f(\alpha) \leq \frac{2\pi}{\alpha \sin \frac{1}{4}\alpha}$.*

When P is (α, β)-immersed, each point on the boundary has an empty (α, β)-triangle outside P as well as inside P. This implies that the lemma also holds for two points p and q that can see each other on the *outside* of the polygon.

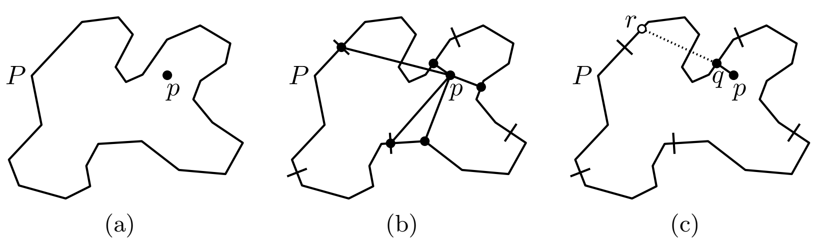

Fig. 6. (a) A polygon P that is (α, β)-immersed. (b) A feed-link to the closest point on each boundary piece of length β. (c) The dilation of r is constant, because the boundary distance between r and q is bounded by their Euclidean distance.

Theorem 6. *When P is (α, β)-immersed, we can place $\frac{c}{\beta^2 \sin\alpha}$ feed-links such that the dilation of every point on P is at most $1 + \frac{4\pi}{\alpha \sin\frac{1}{4}\alpha}$.*

Proof. We give a constructive proof. Given an (α, β)-immersed polygon and a point p inside it, we split P into pieces of length β. By Lemma 4 there are only $\frac{c}{\beta^2 \sin\alpha}$ pieces. On each piece, we place a feed-link to the closest point to p. Fig. 6(b) shows the resulting feed-links in an example.

For each point r on P, we show that the dilation is constant. Consider the piece of P containing r and the point q that is the closest point to p on that piece, as in Fig. 6(c). The segment $\overline{qr}$ may intersect P in a number of points. For each pair of consecutive intersection points, they can see each other either inside or outside P. Since P is (α, β)-immersed, Lemma 5 applies to each pair, and hence $\mu(q, r) \le f(\alpha) \cdot |qr|$. Also, we know that $|pq| \le |pr|$. We conclude that the dilation is bounded by

$$\begin{aligned}\delta_q(r) &= \frac{|pq| + \mu(q,r)}{|pr|} \le \frac{|pq| + f(\alpha)|qr|}{|pr|} \\ &\le \frac{|pq| + f(\alpha)(|pr| + |pq|)}{|pr|} \le \frac{|pr| + f(\alpha)(|pr| + |pr|)}{|pr|} = 1 + 2f(\alpha).\end{aligned}$$

□

When P is (α, β)-covered but not (α, β)-immersed, the proof no longer works since there can be two points that see each other outside the polygon, in which case Lemma 5 does not hold. However, we can still prove that (α, β)-covered polygons are feed-link realistic, albeit with a different dependence on α and β.

Let $C = \frac{4\pi c}{\beta^2 \alpha \sin\alpha \sin\frac{1}{2}\alpha}$ be a constant (depending on α and β). We incrementally place feed-links until the dilation is at most C everywhere. In particular, after placing the first i feed-links, consider the set of points on the boundary of P that have dilation worse than C. If q_{i+1} is the point of this set that is closest to p, then we let the next feed-link be $\overline{pq_{i+1}}$.

We now need to prove that this results in a constant number of feed-links. So, say we placed k feed-links this way, and let their points be $q_1 \dots q_k$. Obviously, we have $|pq_i| \le |pq_j|$ if $i < j$.

Lemma 6. *Unless $k = 1$, all points q_i are inside the circle D centered at p of radius $R = \frac{1}{2}\beta \sin\frac{1}{2}\alpha$.*

Inside the circle D, there cannot be edges of P of length β or longer. So, each point q_i has an empty (α, β)-triangle t_i with one corner at q_i and the other two corners outside D. Fig. 7(a) illustrates the situation, where the grey part is inside P. Let d_i be the direction of the bisector of t_i at q_i. In the full paper we prove that two directions d_i and d_j differ by at least $\frac{1}{2}\alpha$.

Lemma 7. *The angle between d_i and d_j is at least $\frac{1}{2}\alpha$.*

Theorem 7. *Given the boundary P of an (α, β)-covered polygon and a point p inside it, $\frac{4\pi}{\alpha}$ feed-links are sufficient to achieve a dilation of $\frac{4\pi c}{\beta^2 \alpha \sin\alpha \sin\frac{1}{2}\alpha}$.*

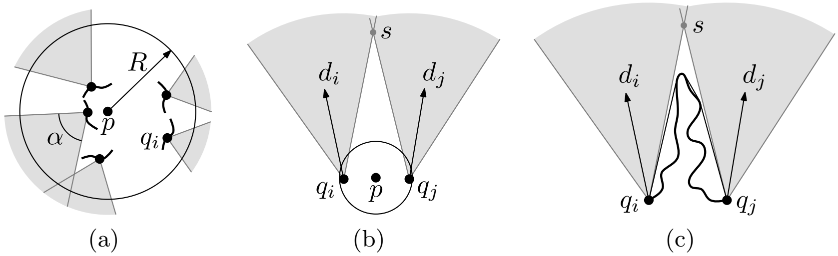

Fig. 7. (a) A circle around p of radius R contains points q_i such that $\overline{pq_i}$ is a feed-link. (b) If the angle between two bisecting directions d_i and d_j is small, the (α, β)-triangles intersect in s. (c) The boundary length between q_i and q_j cannot be too long.

Proof. We place feed-links incrementally as described, until all points on P have dilation at most C. By Lemma 7 there cannot be more than $\frac{4\pi}{\alpha}$ feed-links, because otherwise some pair q_i and q_j would have (α, β)-triangles with directions d_i and d_j whose angle is smaller than $\frac{1}{2}\alpha$. □

4 Conclusions

We presented an efficient algorithm to compute an optimal feed-link for the boundary of a simple polygon and a point inside it. Furthermore, we showed that two feed-links are sometimes necessary and always sufficient to guarantee constant dilation for convex polygons; by placing k feed-links, we can even guarantee a dilation of at most $1 + O(1/k)$. Finally, we considered the number of feed-links necessary for realistic polygons, and proved that (α, β)-covered polygons require only a constant number of feed-links for constant dilation. For other definitions of realistic polygons such a result does provably not hold.

It is open whether the optimal feed-link can be placed in $O(n \log n)$ time or even faster. It is also open whether a linear-time, $(1+\varepsilon)$-approximation algorithm exists for computing an optimal feed-link in a simple polygon (we proved this only for convex polygons).

A number of interesting and challenging extensions of our work are possible. Firstly, the optimal placement of more than one feed-link seems difficult. Secondly, we did not consider the situation where several points lie inside P and need to be connected via feed-links. Here we may or may not want to allow one feed-link to connect to another feed-link. Thirdly, assume we are given an incomplete road network N and several locations, which might fall into different faces of the graph induced by N. How should we place feed-links optimally?

Acknowledgements. We thank Tom de Jong for introducing us to this problem and Mark de Berg for pointing us to the results in [3].

References

1. Agarwal, P., Sharir, M., Shor, P.: Sharp upper and lower bounds on the length of general Davenport-Schinzel sequences. Journal of Combinatorial Theory, Series A 52, 228–274 (1989)
2. Aronov, B., Buchin, K., Buchin, M., Jansen, B., de Jong, T., van Kreveld, M., Löffler, M., Luo, J., Silveira, R.I., Speckmann, B.: Feed-links for network extensions. In: Proc. 16th International Symposium on Advances in Geographic Information Systems, pp. 308–316 (2008)
3. Bose, P., Dujmović, V.: A note on the perimeter of (α, β)-covered objects (manuscript, 2009)
4. de Berg, M.: Linear size binary space partitions for uncluttered scenes. Algorithmica 28, 353–366 (2000)
5. de Berg, M., van der Stappen, A., Vleugels, J., Katz, M.: Realistic input models for geometric algorithms. Algorithmica 34, 81–97 (2002)
6. de Jong, T., Tillema, T.: Transport network extensions for accessibility analysis in geographic information systems. In: Proc. AfricaGIS (2005)
7. Ebbers-Baumann, A., Grüne, A., Klein, R.: The geometric dilation of finite point sets. Algorithmica 44, 137–149 (2006)
8. Efrat, A.: The complexity of the union of (α, β)-covered objects. SIAM Journal on Computing 34, 775–787 (2005)
9. Erickson, J.: Local polyhedra and geometric graphs. Computational Geometry: Theory and Applications 31, 101–125 (2005)
10. Farshi, M., Giannopoulos, P., Gudmundsson, J.: Finding the best shortcut in a geometric network. In: Proc. 21st Symposium on Computational Geometry, pp. 327–335 (2005)
11. Guibas, L.J., Hershberger, J., Leven, D., Sharir, M., Tarjan, R.E.: Linear-time algorithms for visibility and shortest path problems inside triangulated simple polygons. Algorithmica 2, 209–233 (1987)
12. Hershberger, J.: Finding the upper envelope of n line segments in $O(n \log n)$ time. Information Processing Letters 33, 169–174 (1989)
13. Langerman, S., Morin, P., Soss, M.: Computing the maximum detour and spanning ratio of planar paths, trees, and cycles. In: Alt, H., Ferreira, A. (eds.) STACS 2002. LNCS, vol. 2285, pp. 250–261. Springer, Heidelberg (2002)
14. Narasimhan, G., Smid, M.: Approximating the stretch factor of Euclidean graphs. SIAM Journal on Computing 30, 978–989 (2000)
15. Narasimhan, G., Smid, M.: Geometric Spanner Networks. Cambridge University Press, Cambridge (2007)
16. Sharir, M., Agarwal, P.: Davenport-Schinzel Sequences and Their Geometric Applications. Cambridge University Press, Cambridge (1995)
17. van der Stappen, A.F., Overmars, M.H., de Berg, M., Vleugels, J.: Motion planning in environments with low obstacle density. Discrete & Computational Geometry 20, 561–587 (1998)
18. van Kreveld, M.: On fat partitioning, fat covering and the union size of polygons. Computational Geometry: Theory and Applications 9, 197–210 (1998)