Beacon based routing and coverage

Michael Biro* Jie Gao† Justin Iwerks* Irina Kostitsyna†
Joseph S. B. Mitchell*

Abstract

We consider *beacon* based point-to-point routing and coverage problems motivated by sensor network applications. A beacon b is a point that can be *activated* to effect a 'gravitational pull' toward itself everywhere in a simply connected polygonal domain P.

We show that $\left\lfloor \frac{n}{2} \right\rfloor - 1$ beacons are sometimes necessary and always sufficient to route between any pair of points in P. We demonstrate that finding a minimum cardinality set of beacons to route from any source point $s \in P$ to a given destination $t \in P$ is NP-hard.

We show that it is NP-hard to find a minimum cardinality set of beacons to cover a simple polygon.

Keywords: path navigation, beacon, landmark, computational geometry, combinatorics, sensor networks

1 Introduction

The model of beacon based routing in this paper is an analog of geographical greedy routing in sensor networks in the continuous setting. In geographical routing [1, 4], each node is given a Euclidean coordinate and a message is delivered to the neighbor whose Euclidean distance to the destination is the smallest. When sensor distribution is very dense (i.e., close to infinity), geographical routing will always follow the straight line towards the destination, or, when the message hits the network boundary, may follow a boundary edge to greedily minimize the distance to the destination. This is precisely the model of beacon based routing in this paper, where the destination is a beacon.

Our model is also related to a family of routing schemes in sensor networks that use landmarks [2, 3, 6]. A subset of nodes, called landmarks, first flood the entire network such that each node records the distance to each landmark. For routing towards a destination, a function based on the distance vector to the landmarks is used to select the next hop. The one most similar to our model is the one adopted in [6]. In [6], the message is routed towards one landmark until the current node is equal distance away from the landmark as the destination. At this point another landmark is selected. The paper shows that by carefully choosing the landmarks the routing path is within a constant factor of the shortest path. In this paper we examine the combinatorial structures for landmark placement, to support this type of routing.

The first problem we consider is that of finding a minimum cardinality set of beacons to route between two points in a simple polygon P (minimum beacon routing set). In our model, a beacon can occupy a point location on the interior or the boundary of P, ∂P. When a beacon is *activated*, all points $p \in P$ move along straight lines toward b until they either reach b or make contact with ∂P. If contact is made with ∂P, p will follow along ∂P as long as its straight line distance to b decreases monotonically. p may alternate between moving in a straight line path toward b on the interior of P and following along ∂P. If p is unable to move so that its distance to b decreases monotonically, we say p is 'stuck' and has reached a local minimum or *dead point* on ∂P (see Figure 1). If p reaches b we say that b *attracts* p. Two points are *routed* if there is a sequence of beacons that can be activated and then deactivated, one at a time in order, such that a point beginning at a source s would visit each beacon in the sequence after it is activated and terminate at a destination t, which we always assume to be a beacon itself.

The traditional art gallery problem is concerned with finding a minimum cardinality set of guards to cover a simple polygon. Finding such a minimum sized guard set is known to be NP-hard [5]. We consider the analogous complexity problem with beacons serving as the guards. In this version, P is covered by a set of beacons B if for each $p \in P$, there exists some $b \in B$ such that b attracts p.

*Dept. Applied Mathematics and Statistics, Stony Brook University
†Dept. Computer Science, Stony Brook University

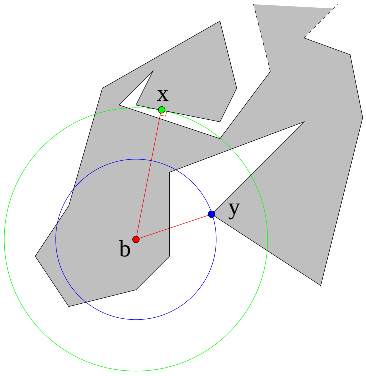

Figure 1: x and y are dead points with respect to the beacon, b.

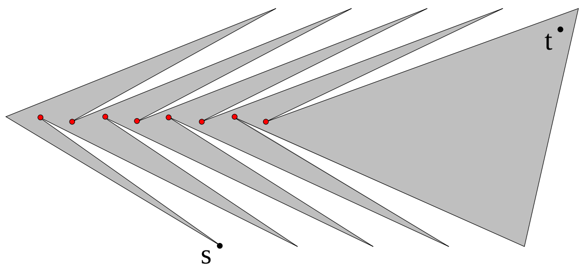

Figure 2: $\lfloor \frac{n}{2} \rfloor - 1$ beacons are sometimes necessary to route between all pairs of points in a simple polygon. Here, $n = 19$ and 8 beacons (light filled circles) are required to route from s to t.

2 Our Results

In this section we present results pertaining to beacon based routing (Section 2.1) and coverage (Section 2.2). For the routing problem we show that $\lfloor \frac{n}{2} \rfloor - 1$ beacons are sometimes necessary and always sufficient to route between any pair of vertices in a simple polygon P. We show that it is NP-hard to find a minimum cardinality set of point beacons to route from all points $s \in P$ to a given destination point t. We also show that finding a minimum cardinality set of beacons B that covers a simple polygon P is NP-hard.

2.1 Beacon based routing

We begin with a combinatorial result for beacon based routing:

Theorem 1. $\lfloor \frac{n}{2} \rfloor - 1$ *beacons are sometimes necessary and always sufficient to route between any pair of vertices in a simple polygon.*

Proof. We can see from Figure 2 that $\lfloor \frac{n}{2} \rfloor - 1$ beacons are sometimes necessary.

To establish the upper bound, we first triangulate P and construct the dual graph G of the resulting triangulation. We then start from a leaf node of G and begin to peel off pairs of adjacent triangles. Let's call the two triangles σ_1 and σ_2 where σ_1 is the leaf triangle. We place a beacon at one vertex of the common edge of the two triangles and argue that one can navigate from any point p in either triangle to the beacon, or from the beacon to p, using *greedy routing*. We conduct case analysis on the number of triangles adjacent to σ_2 other than σ_1:

1. σ_2 has only one adjacent triangle σ_3. Suppose $\sigma_1 = \triangle ABC$, $\sigma_2 = \triangle BCD$. σ_3 is then either $\triangle BDE$ or $\triangle CDE$. If $\sigma_3 = \triangle BDE$, then we place a beacon b at the vertex B and otherwise we place b at C. In either case, since b is contained in each of the three triangles, any point p in these three triangles can navigate to b and vice-versa.

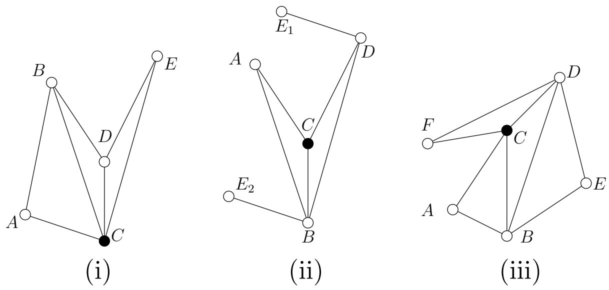

Figure 3: (i) The beacon is placed at a vertex common to three triangles; (ii) The beacon is placed at a vertex C common to σ_1 and σ_2, but not inside triangle σ_3; (iii)

2. σ_2 has two adjacent triangles σ_3, σ_4. Assume that $\sigma_1 = \triangle ABC$, $\sigma_2 = \triangle BCD$, $\sigma_3 = \triangle BDE$, $\sigma_4 = \triangle CDF$. We place a beacon b at C if $\angle FCB > 3\pi/2$, or at B if $\angle EBC > 3\pi/2$. We note that the two conditions cannot simultaneously be true. In particular, if $\angle FCB > 3\pi/2$, then $\angle BCD$ must be obtuse. If $\angle EBC > 3\pi/2$, then $\angle CBD$ must be obtuse. $\triangle BCD$ cannot have two obtuse angles. If neither of these conditions hold, then we place b arbitrarily at either B or C. Now, assume that we place b at vertex C. Then it must be the case that $\angle EBC \leq 3\pi/2$. Therefore, all points inside $\triangle BDE$ can reach b and vice versa. Also as C is a vertex of σ_1, σ_2 and σ_4, all points inside theses three triangles can reach b and vice versa. Hence, the claim is

true.

Given the basic steps as shown above, we will place beacons in a recursive manner: We take any leaf triangle σ_1 of the triangulation of P and place a beacon at one of the vertices of the shared edge of the pair σ_1 and its adjacent triangle, σ_2.

1. If P has at most one more triangle besides σ_1 and σ_2 or P has only two more triangles but both adjacent to σ_2, then we are done. By the above arguments we can reach from any source point to any destination point by using the single beacon: First route from the source to the beacon and then route from the beacon to the destination (which is always a beacon).

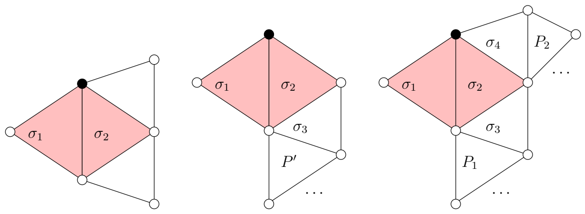

Figure 4: (i)

2. Otherwise, we peel off σ_1 and σ_2. There are two subcases to consider.
   (a) σ_2 is only adjacent to one more triangle σ_3 (i.e., σ_2 has degree 2 in the dual graph). In this case peeling off σ_1 and σ_2 will still leave a simple polygon P'. We can recursively 'beaconize' P'. Now we argue that one can navigate with the union of these beacons. In particular, if the source and destination pair are both in $\sigma_1 \cup \sigma_2$ or both in P', then we can navigate by induction hypothesis. If the source and destination pair are separated in $\sigma_1 \cup \sigma_2$ and P', we can use the beacon x of σ_2 and the beacon y of σ_3 (if it is a different beacon) to help guide the message across the two pieces. By the analysis of the basic case, any point inside σ_3, in particular, the beacon y, is reachable to and from the beacon x in σ_2. Thus navigation works in this case.
   (b) σ_2 is adjacent to two more triangles σ_3 and σ_4. Thus peeling off σ_1 and σ_2 will partition the triangulation to two pieces called P_1 and P_2. Suppose that P_1 contains σ_3 and P_2 contains σ_4. By the same argument as above, by using beacon x in σ_2 we can navigate between the three pieces $\sigma_1 \cup \sigma_2$, P_1 and P_2.

With the algorithm, we can see that each time we peel off two or three triangles at a time and place one beacon. Thus the total number of beacons we place would be at most $\lfloor \frac{n-2}{2} \rfloor = \lfloor \frac{n}{2} \rfloor - 1$.

□

In the subsequent theorem, we establish the hardness of all source routing in a simple polygon.

Theorem 2. *It is NP-hard to find a minimum cardinality set of beacons to route from all source points s to a given destination point t in a simple polygon.*

Proof. We prove hardness by reducing from the LINE HITTING problem: given an arrangement of n lines in the plane, place a minimum cardinality set of points S so that each line intersects ('hits') at least one point of S. Given an instance of the LINE HITTING problem, we construct a 'spike box' large enough for its rectangular body to contain a positive length segment of each line and all intersection points formed by the arrangement. The spike box contains a protruding zigzag spike gadget for each line in the arrangement. We let the destination point t be any point in the body of the spike box that is not on any line and let there be a source point s at the end of each spike. A zigzag spike gadget (see Figure 5) is constructed at either of the two locations a line exits the body of the spike box so that activating a beacon anywhere in the dark grey region $G(s)$ attracts s at the top of the zigzag spike to the beacon. Note that the light grey regions in the spike are regions in which an activated beacon could attract a point down from the above left and right hand horizontal edges, but not down the slanted edges. Since each beacon must be activated once in a sequence, the most efficient way to route from s in a zigzag spike to the body of the spike box is clearly by placing one beacon in $G(s)$ (else we would need to use multiple beacons to accomplish this task). We can make each region $G(s)$ as narrow as we please so we will treat these regions as line segments.

Since t is in the body of the spike box, t can attract any other point in the body, with no other beacons required. Since t is not on any line of the arrangement, t is not in $G(s)$. Therefore, we need a beacon in each $G(s)$ to route from s to the body of the spike box.

Given an arrangement of lines that can be covered with k points, we can route to t with k beacons by placing the beacons on their corresponding lines.

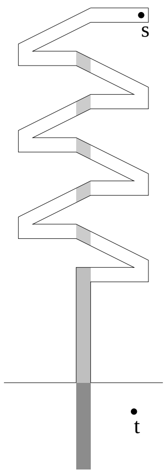

Figure 5: A zigzag spike gadget

Since each line is covered by a point, there is a point in each $G(s)$ and so every point can be routed to t. If we can route to t using $k < n$ beacons, then we can cover the lines with k points. Since k is small, there must be a beacon in each $G(s)$. Place the points on the lines corresponding to the beacons, and there will be a point hitting every line. □

2.2 Beacon coverage

Finally, we establish the hardness of covering a simple polygon with a minimum cardinality set of becaons. The proof is again based on a reduction from the LINE HITTING problem.

Theorem 3. *It is NP-hard to find a minimum cardinality set of beacons whose union covers a simple polygon.*

Proof. Given an instance of the LINE HITTING problem, we construct a 'spike box' large enough for its rectangular body to contain a positive length segment of each line and all intersection points formed by the arrangement. The spike box contains an arrow shaped spike gadget for each line in the arrangement that protrudes from the body. We place a spike gadget at either location where a line exits the body of the spike box. Each gadget has an arbitrarily small entrance (see Figure 6. The minimum number of beacons to cover the tip of each spike is exactly the number of beacons needed to cover the whole spike box. This is because any beacon placed in the convex body of the spike box covers the body of the spike box. Therefore, if any of the spike-covering beacons are in the body, each spike is covered as well as the body, so the entire polygon is covered. If all the spike-covering beacons are in the spikes, then they are each in the region of points $R(p)$ that attract exactly one tip, say p. Since all intersections of regions $R(p)$ for each p are contained in the body of the spike-box, each beacon is responsible for only one tip, and so we can move it, as long as it remains in $R(p)$. Specifically, we can move it to the body of the spike-box, and cover the entire polygon with the same number of beacons.

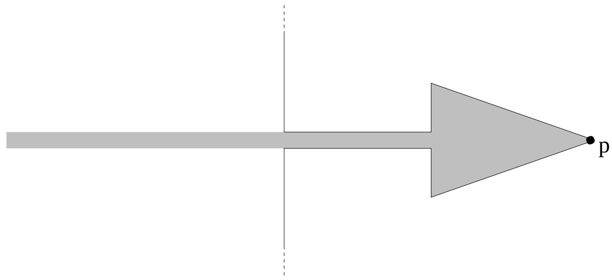

Figure 6: An 'arrow spike'. p is attracted by all points in the grey shaded region, which extends into the body of the spike box.

If we can cover each line in the arrangement with k points, then we can cover the polygon with k beacons. Given an instance of k points covering the lines, place k beacons in their corresponding places in the spike box. Since each line is covered, there is a beacon in $R(p)$ for each spike tip p, and so we can cover the polygon with k beacons. If we can cover the polygon with k beacons, then we can cover the lines with k points. Given the beacon placement, place the points on the corresponding points of the lines. Since each spike-tip is covered, there is a beacon in each $R(p)$ and so there is a point on each line, covering the lines with k points. □

Acknowledgements. This work was partially supported by various sources.

References

[1] P. Bose, P. Morin, I. Stojmenovic, J. Urritia Routing with Guaranteed Delivery in Ad Hoc Wireless Networks. *Wireless Networks*, 7(6):609–616, 2001.

[2] Q. Fang, J. Gao, L. Guibas, V. de Silva, L. Zhang. GLIDER: Gradient landmark-based distributed routing for sensor networks. *Proc. of the 24th Conference of the IEEE Communication Society (INFOCOM'05)*, 339–350, March, 2005

[3] R. Fonseca, S. Ratnasamy, J. Zhao, C. T. Ee, D. Culler, S. Shenker, I. Stoica Beacon vector

routing: scalable point-to-point routing in wireless sensornets. *Proc. of the 2^{nd} Symposium on Networked Systems Design and Implementation, (NSDI'05)*, 329–342, May, 2005.

[4] B. Karp and H. Kung GPSR: Greedy perimeter stateless routing for wireless networks. *Proc. of the ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom)*, 243–254, 2000.

[5] D. T. Lee and A. K. Lin. Computational complexity of art gallery problems. *IEEE Transactions on Information Theory*, 32(2):276–282, 1986.

[6] A. Nguyen, N. Milosavljevic, Q. Fang, J. Gao, L. J. Guibas. Landmark selection and greedy landmark-descent routing for sensor networks. *Proc. of the 26^{th} Annual IEEE Conference on Computer Communications (INFOCOM'07)*, 661–669, May, 2007.

[7] J. O'Rourke. *Art gallery theorems and algorithms.* Oxford University Press, Inc., 1987.