

Natural Wireless Localization is NP-hard

Tobias Christ¹

Michael Hoffmann¹

Yoshio Okamoto²

¹ Institute for Theoretical Computer Science, ETH Zürich, {christt,hoffmann}@inf.ethz.ch


² Tokyo Institute of Technology, okamoto@is.titech.ac.jp

Abstract

We consider a special class of art gallery problems inspired by wireless localization. Given a polygonal region \mathcal{P} , place and orient guards each of which broadcasts a unique key within a fixed angular range. In contrast to the classical art gallery setting, broadcasts are not blocked by the boundary of \mathcal{P} . At any point in the plane one must be able to tell whether or not one is located inside \mathcal{P} only by looking at the set of keys received. We prove NP-hardness of one variant of the problem, namely the *natural* setting where guards may be placed aligned to a boundary edge or two consecutive boundary edges of \mathcal{P} only.

Introduction. Art gallery problems are a classic topic in discrete and computational geometry. A new direction has recently been introduced by Eppstein, Goodrich, and Sitchinava [1]. They propose to modify the concept of visibility by not considering the edges of the polygon/gallery as blocking. This changes the problem quite drastically because it breaks up a certain locality where the shape of the polygon dictates the possible placement of guards. A basic ingredient in hardness proofs for the classical setting is a small pocket/spike of the polygon which can only be guarded from a nearby point because the bounding polygon edges shield it away from the rest of the world. This argument breaks down if the edges do not block visibility. The motivation for this model stems from communication in wireless networks where the signals are not blocked by walls, either. For illustration, suppose you run a café (modeled, say, as a simple polygon P) and you want to provide wireless Internet access to your customers. But you do not want the whole neighborhood to use your infrastructure. Instead, Internet access should be limited to those people who are located within the café. To achieve this, you can install a certain number of devices, let us call them guards, each of which broadcasts a unique (secret) key in an arbitrary but fixed angular range. The goal is to place guards and adjust

their angles in such a way that everybody who is inside the café can prove this fact just by naming the keys received and nobody who is outside the café can provide such a proof. Formally this means that P can be described by a monotone Boolean formula over the keys, that is, a formula using the operators AND and OR only, negation is not allowed. (See [1, 2, 3].)



Notation. A *guard* g is a closed subset of the plane, whose boundary ∂g is described by a vertex v and two rays emanating from v . The ray that has the interior of the guard to its right is called the *left ray* ℓ_g , the other one is called the *right ray* r_g . The *angle* of a guard is the interior angle formed by its bounding rays. For a guard with angle π , the vertex is not unique. A *guarding* $\mathcal{G}(\mathcal{P})$ for a polygonal region \mathcal{P} is a set of guards such that there is a formula composed of this set and the operators union and intersection that defines \mathcal{P} . Natural locations for guards are the vertices and edges of the polygonal region. A guard which is placed at a vertex of \mathcal{P} is called a *vertex guard*. A vertex guard is *natural* if it covers exactly the interior angle of its vertex [1]. A guard placed anywhere on the line given by an edge of \mathcal{P} and broadcasting within an angle of π to the inner side of the edge is called a *natural edge guard*. A *natural guarding* is a guarding consisting of natural vertex and natural edge guards only [3]. From now on we consider natural guardings only.

The Natural Wireless Localization Problem. Given a finite union of simple polygons \mathcal{P} and an integer k , is there a natural guarding for \mathcal{P} using k guards?

Theorem 1 *The Natural Wireless Localization Problem is NP-complete.*

The problem is in NP. Given a set of natural guards we can check in polynomial time whether they can describe \mathcal{P} . To show the NP-completeness we reduce the Vertex Cover Problem to the Wireless Localization Problem. It is sufficient to consider 3-connected 3-regular planar graphs (see [4]). Let $G = (V, E)$ be such a graph of order n . Embed the vertices V into the plane, that is, think of G as a geometric graph. Denote the line defined by an edge $e \in E$ by \bar{e} . The set of edges incident to a vertex $v \in V$ is denoted by $E(v)$. Consider the line arrangement $L = \{ \bar{e} \mid e \in E \}$ defined by G . Add all perpendicular bisectors of edges, resulting in an *extended line arrangement* L' . We want L' to be in general position in the following *weak sense*: only the three lines defined by the edges incident to a vertex v meet in v and only \bar{e} and its perpendicular bisector meet in the midpoint of e denoted by $M(e)$. Furthermore, we want the embedding to be strictly convex.

Lemma 2 *Given a 3-regular and 3-connected planar graph $G = (V, E)$ it is possible to embed V into an $O(n^5) \times O(n^5)$ grid in time polynomial in n such that the extended line arrangement L' of G is in general position in the weak sense and every face is strictly convex.*

Proof. A drawing of G on an $O(n^2) \times O(n^2)$ grid in which all faces are strictly convex polygons can be obtained in linear time [5]. Let G be such a geometric graph on a $K \times K$ grid with $K \in O(n^2)$. The vertices and the midpoints of edges are called *points of interest*. Perturb the vertices locally one by one to finally get a new drawing which is in general position in the weak sense, that is, no lines of L' intersect the points of interest except those who have to. (We still denote the drawing after perturbation by G and the extended line arrangement by L' .)

An easy calculation reveals that on a $K \times K$ square grid the distance between any grid point v and any line l disjoint from v that is defined by two grid points is at least $1/(\sqrt{2}K)$. Define $\varepsilon := 1/(2\sqrt{2}K)$. If we perturb the vertices by less than ε , then the drawing remains strictly convex. (Destroying a convex face corresponds to moving a vertex v of the face and a line l defined by two other vertices of the face such that after the perturbation v lies on the other side of l . The distance of v and l is at least 2ε , therefore after having perturbed every vertex by less than ε , v cannot have changed sides.) In the same way one might argue that the drawing remains crossing-free, but this is not essential for our purposes.

Consider a small $k \times k$ grid around every vertex v inside an ε -ball around v and choose a new location for v in this grid as follows: Assume that we already have perturbed the first m vertices and that the extended line arrangement of the subgraph induced by the first m vertices is in general position in our weak sense. Perturbing the next vertex v we want to ensure that neither v nor one of the midpoints of its incident edges come to lie on a line of L' nor it induces a new line in L' that intersects a point of interest.

The extended line arrangement L' is built on $3n$ lines. (There are $3n/2$ edges giving rise to two lines each.) Therefore, we have $n + 3n/2$ points of interest. It is forbidden to put v

- on any of the lines from L' , except for those lines \bar{e} that are defined by an edge $e \in E(v)$ (less than $3n$ forbidden lines),
- on any line such that placing v there would result in one of the midpoints $M(e)$ of an edge $e \in E(v)$ lying on a forbidden line (less than $3 \cdot 3n = 9n$ forbidden lines),
- on any line defined by one of the neighbors of v and an arbitrary other point of interest (less than $(15/2)n$ forbidden lines),
- on any circle C centered at a point q of interest and passing through one of the neighbors of v (less than $(15/2)n$ forbidden circles), as placing v on C would cause the perpendicular bisector of the corresponding edge to pass through q .

All of these less than $27n$ forbidden loci are either lines or circles, each of which hits at most k of the k^2 grid points of the small grid around v . Choosing k larger than $27n$ we can always be sure to find an appropriate placement for v . Recall that this small grid is placed within an ε -ball around v . As $k = O(n)$ and $\varepsilon = \Omega(n^{-2})$, the distance between any two distinct points on a small grid is $\Omega(n^{-3})$. Overall, we refine the $K \times K$ grid by a factor of $O(n^3)$, which results in an $O(n^5) \times O(n^5)$ grid. \square

Let δ denote the minimum distance between any point q of interest and a line of L' disjoint from q . After perturbation the vertices lie on an $O(n^5) \times O(n^5)$ grid. Refining it by a factor of 2 ensures that all midpoints of edges are grid points, too. As on a $K \times K$ grid any grid point has distance at least $1/(\sqrt{2}K)$ to any disjoint line defined by two grid points, $\delta = \Omega(n^{-5})$.

The main steps of the reduction are the following. We successively replace every vertex by a small convex hexagon. There are exactly two

ways to guard such a polygon optimally: if the vertices are numbered in order, either put guards on all even vertices or put guards on all odd vertices. The idea is that one way of guarding this vertex gadget corresponds to the vertex being in the vertex cover and the other way of guarding it corresponds to the vertex not being in the cover.

Then, successively replace every edge by an edge gadget, that is, a simple polygon having the property that it needs at least 4 natural guards unless there is a ray of some other guard passing through the gadget, in which case it needs 3 guards only. Thus, the edge gadget “asks” one of its incident vertex gadgets to help out with such a ray, which corresponds to the edge asking for one of its incident vertices to be in the cover.

Finally, we will add a “punishment gadget” to every vertex, to ensure that every vertex put into the cover is paid for.

Vertex gadget. Every vertex $v \in V$ is replaced by a convex hexagon centered at v . Let $E(v) = \{e_1, e_2, e_3\}$ be the edges incident to v . Let $P(v)$ be a polygon that has two edges collinear to e_1 , two edges collinear to e_2 and two edges collinear to e_3 , all of them having length λ . We can avoid creating edges that are on the same line as edges already defined in other vertex gadgets by varying the edge length $0 < \lambda < \delta$. Observe that every line defined by the edges of $P(v)$ is parallel to one of the lines in the extended line arrangement L' with distance at most δ .




Figure 1: The two ways to guard $P(v)$ for a vertex v and $P(w)$ for a pointed vertex w .

Observation 1 *There are exactly two ways to guard the vertex gadget $P(v)$ of an vertex v using 3 guards only: Either the rays of the guards go into the same direction as the edges of v , which we call the first way to guard it, or they go into the other direction, which we call the second way.*

In the special case of a *pointed* vertex, that is, a vertex for which all incident edges lie within an angle of less than π , the rays parallel to the middle edge go into the other direction. See Figure 1.




Figure 2: An edge gadget $Q(e)$.

Edge gadget. For every edge $e \in E$ we insert a polygon $Q(e)$ such that no line through any edge defined in some other gadget intersects $Q(e)$, except for the lines parallel to e that are defined by the two vertex gadgets corresponding to the endpoints of e . Furthermore, we have to ensure that no line defined by an edge of $Q(e)$ intersects any other edge gadget already created. If λ is chosen to be sufficiently small, the lines defined by $Q(e)$ remain inside a δ -strip around the perpendicular bisector of e within our area of interest (up to outermost intersection of the perpendicular bisector with other lines of the extended line arrangement). Hence the lines defined by $Q(e)$ do not intersect any other edge gadget, as these gadgets are located at an intersection of other lines of L' . Note that choosing a very small λ and therefore a “very flat” $Q(e)$ does not change the properties of $Q(e)$ as far as guardings are concerned.

For an edge e that is the middle edge of a pointed vertex on the outer face, do not place $Q(e)$ on the midpoint $M(e)$ but somewhere within the outer face along \bar{e} . Note that at this point we use that G forms a strictly convex embedding: As all pointed vertices lie on the outer face, two pointed vertices can only be connected by an edge of the outer face, which cannot be a middle edge. In other words, at most one endpoint of e can be pointed.

If $Q(e)$ is hit by a ray r_g of a guard g or a left ray $\ell_{g'}$ of a guard g' from the incident vertex gadgets, then there $Q(e)$ can be guarded using 3 guards only. (See Figure 2, $Q(e) = a \cap b \cap (c \cup g)$ or $Q(e) = a \cap b \cap (c \cup g')$, respectively.) If there are no such rays crossing $Q(e)$, then at least 4 guards are needed to guard $Q(e)$, see Figure 3. (If there was a guarding using 3 guards only, then there would have to be a natural vertex guard on every second vertex. In either case there would be a pair of points (depicted by crosses in Figure 3), one inside $Q(e)$ the other outside of $Q(e)$, that could not be distinguished by the guards.)




Figure 3: An edge gadget hit by no ray.

Observation 2 *At least 4 guards are needed for $Q(e)$ unless one of the incident vertex gadgets is guarded in the first way, in which case 3 guards suffice.*

The punishment gadget. For every vertex v we add a punishment gadget similar to the edge gadgets. Choose an edge $e \in E(v)$ and place a polygon $R(v)$ as shown in Figure 2 somewhere along \bar{e} , (seen from v) on the side opposite to where the edge gadget is. (For a pointed vertex, pick one of the two edges incident to the reflex angle.) Again make sure not to create any kind of interference with other gadgets. We can always find an appropriate placement for the gadgets by putting them somewhere far away from everything else and choosing λ small enough. For lack of space, we have to defer details to the full paper.

Observation 3 *Every punishment gadget needs at least 4 guards if there are no rays of other guards intersecting it. If its corresponding vertex gadget is guarded in the second way, then there is a guarding using only 3 guards.*

All in all we get a finite union of simple polygons $\mathcal{P}(G)$ composed of all vertex gadgets, all edge gadgets and all punishment gadgets. It is possible to construct $\mathcal{P}(G)$ in polynomial time. Let $n = |V|$ be the number of vertices and $m = |E| = 3n/2$ the number of edges. The following lemma completes the proof of Theorem 1.

Lemma 3 *There is a vertex cover of G using k vertices if and only if there is a natural guarding for $\mathcal{P}(G)$ using at most $3m + 6n + k$ guards.*

Proof. If there is a vertex cover of G of size k we can find a guarding of $\mathcal{P}(G)$ as follows. Guard every vertex that is in the cover in the first way, every vertex that is not in the cover in the second way. For every vertex guarded the first way we have to put 4 guards on the corresponding punishment gadget, for every vertex we covered the second way only 3. For all edge gadgets 3 guards are sufficient, because since we started from a vertex cover there will be the necessary rays around to do it with 3 for all of them. So all in all we need

$3m + 3n + 4k + 3(n - k) = 3m + 6n + k$ guards. This proves the first direction of the lemma.

Suppose we are given a natural guarding of $\mathcal{P}(G)$ using at most $3m + 6n + k$ guards. We say a guard *belongs* to a gadget, if it is a natural vertex guard or a natural edge guard on this gadget. We can find a vertex cover of G as follows. Take all vertices whose punishment gadget has more than 3 guards belonging to it and put them into the cover. Furthermore, add one of the endpoints of all edges whose gadget has more than 3 guards belonging to it. So far we have not put more than k vertices into the cover since there cannot be more than k edge or punishment gadgets with more than 3 guards. (See the observations above. At least 3 guards belong to every vertex gadget and at least 3 guards belong to each of the other $m + n$ gadgets, so there remain at most k guards that can create such edge- or punishment gadgets with 4 or more guards.) Indeed, this set forms a cover of G because every edge that has no endpoint whose gadget is guarded in the first way needs at least 4 guards. Therefore one of its endpoints got added to the cover. \square

Outlook. The reduction sketched above shows that it is hard to find an optimal guarding for a collection of polygons or equivalently—looking at the complement—for a polygon with holes. Currently we are working to extend the result to the case where the input consists of a single simple polygon. We also plan to investigate the complexity of finding general, not necessarily natural, guardings.

References

- [1] D. Eppstein, M. Goodrich, N. Sitchinava. Guard placement for efficient point-in-polygon proofs. Proc. 23rd SoCG (2007), pp. 27–36.
- [2] M. Damian, R. Flatland, J. O’Rourke, S. Ramaswami. A new lower bound on guard placement for wireless localization. 17th FWCG (2007). <http://arxiv.org/abs/0709.3554>
- [3] T. Christ, M. Hoffmann, Y. Okamoto, T. Uno. Improved bounds for the wireless localization problem. Proc. 11th SWAT (2008), pp. 77–89.
- [4] R. Uehara. NP-complete problems on a 3-connected cubic planar graph and their applications. Technical Report TWCU-M-0004, Tokyo Woman’s Christian University, 1996. <http://www.jaist.ac.jp/~uehara/ps/triangle.ps.gz>
- [5] I. Bárány, G. Rote. Strictly convex drawings of planar graphs. Documenta Mathematica 11 (2006), pp. 369–391.