Wireless Localization within Orthogonal Polyhedra

Tobias Christ*   Michael Hoffmann*

Abstract

In the *wireless localization problem*, given a polygon $P \subset \mathbb{R}^2$, we have to place guards and fix their angular range such that P can be *described* using these guards. The guards describe P if for every point pair $p \in P$ and $q \notin P$, there is a guard that sees p but does not see q. We consider the analogous problem in 3D: given a polyhedron $P \subset \mathbb{R}^3$, place guards—which now are polyhedral cones—that collectively describe P. Generalizing a known result for 2-dimensional orthogonal polygons, we show that for any given 3-regular orthogonal polyhedron $P \subset \mathbb{R}^3$ with n vertices, it suffices to put a *natural vertex guard* onto every other vertex. (A natural vertex guard is a guard that is placed at a vertex v of P and the defining cone coincides with P in a sufficiently small neighborhood of v.) Furthermore, we show how to describe P with $3n/8$ (general) vertex guards.

1 Introduction

Art gallery problems are a classic topic in discrete and computational geometry. The wireless localization (or sculpture garden) problem, introduced by Eppstein et al. [5], differs from the classical setting in two respects: First, the guards are more powerful because they can "see through walls"; and second, their job is harder, because rather than asking them to collectively *see/cover* the entire polygon (which is very easy, not being impeded by walls), the guards must collectively *describe* the polygon instead.

The motivation for this model stems from communication in wireless networks. For illustration, suppose you run a café (modeled as a polygon P) and you want to provide wireless Internet access. 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, called 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.

---
*Institute of Theoretical Computer Science, ETH Zürich, `tobias.christ@inf.ethz.ch`, `hoffmann@inf.ethz.ch`

It is convenient to model a guard as a subset of the plane, namely the area where the broadcast from this guard can be received. This area can be described as an intersection or union of at most two halfplanes. The definition directly carries over to the 3-dimensional case, where a guard is a polyhedral cone, that is, an unbounded polyhedron with at most one vertex and a connected 1-skeleton.

We define a *guarding* of P to be a set of guards with the property that for each pair of points (p, q), where $p \in P$ and $q \notin P$, there is a guard g that *distinguishes* p and q, meaning $p \in g$ and $q \notin g$. It can be shown that this notion is equivalent to a description of P using a combination of the operations union and intersection over the guards or—in logical terminology—a monotone Boolean formula over the guards, that is, a formula using the operators AND and OR only, negation is not allowed. (See [2], Observation 1. If each pair (p, q) can be distinguished by a guard, we can find a formula in disjunctive normal form for P: For any point p, let G_p the intersection of all guards that contain p. Then $P = \bigcup_{p \in P} G_p$.)
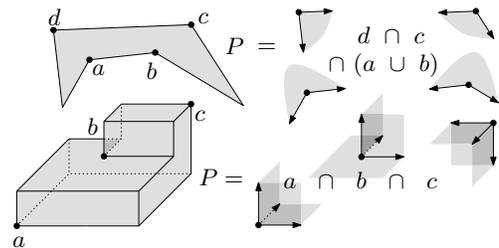
Natural locations for guards are points on the boundary of P. A guard that is placed at a vertex of P is called a *vertex guard*. A vertex guard on a vertex v is called *natural* if its shape is given by the shape of P at v. More precisely, if the intersection of an ε-ball around v with g equals the intersection with P. A guard placed anywhere on an edge e of P and the shape of which is given by the shape of P at e is called a *natural edge guard* (in 2D, edge guards are just the halfplanes defined by an edge, in 3D they are wedges whose only edge is the line through e). In 3D, there is a third class of natural guards. We call the closed halfspace defined by a face f of P a natural face guard on f. Note that both edge and face guards are cones without apex, so their exact position is undefined. We can think of them to be placed anywhere on their only edge (or anywhere on the bounding plane, respectively).

Dobkin et al. [4] showed that n natural edge guards are sufficient for any simple polygon with n edges. Using both natural vertex guards and natural edge guards, $n - 2$ guards are sufficient and can be necessary [2]; using general vertex guards $8n/9$ are sufficient [3]. In the most general setting we do not have any restriction on the placement and the angles of guards. The best known upper bound is $\lfloor \frac{4n-2}{5} \rfloor$ and the best lower bound is $\lceil \frac{3n-4}{5} \rceil$ [2]. We believe that the upper bound can be improved to roughly $3n/4$ and the lower bound to $2n/3$, but this is still work in progress. The classical art gallery problem for orthogonal polygons was considered in [1].

For the $3D$ case, Dobkin et al. [4] observed that placing a face guard onto every face of a polyhedron suffices. To our knowledge, the 3-dimensional wireless localization problem has not been studied since then. In this work we focus on orthogonal polyhedra and prove that $n/2$ natural vertex guards suffice to guard a polyhedron with n vertices. If we allow general vertex guards, we can improve the bound to $3n/8$. We observe that there are orthogonal polyhedra that cannot be guarded by fewer than $n/4$ guards.

2 Notation and Basic Observations

An orthogonal polyhedron P is a polyhedron where all faces are orthogonal to one of the coordinate axes. Faces orthogonal to the x-axis (y-axis, z-axis) are called x-faces (y-faces, z-faces, respectively). Consequently, all edges of P are parallel to one of the coordinate axes and are called x-edges, y-edges and z-edges accordingly. Think of the x-axis as being oriented from left to right, the y-axis front to back, and the z-axis bottom up.

A polyhedron is a solid and closed subset of the space. We define its vertex set $V(P)$, its edge set $E(P)$ and its set of faces $F(P)$ in the usual way. Let $n = n(P) = |V(P)|$ be the number of vertices. Furthermore, we restrict our attention to bounded orthogonal polyhedra with the additional property that exactly three edges meet at every vertex. In other words, the graph of P has to be cubic. Eppstein and Mumford [6] use a similar definition and additionally require the polyhedral surface bounding a polyhedron to have the topology of a sphere. They call this class of polyhedra *simple orthogonal polyhedra*. We do not use this notation, as in this work, we do not need any topological conditions and allow the polyhedra to form handles (that is, their genus might be greater than 0) and to contain cavities (that is, their surface may be disconnected). Therefore, we use the term 3-*regular* instead. We define the *type* of a vertex v as follows. Assuming v to be the origin, the type of v is the set of octants P locally occupies around v. We call v *convex*, if only one octant is inside P. There are eight different possible types of convex vertices. We call v *reflex* if all but

Figure 1: Vertex types.
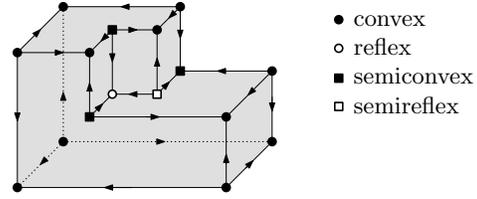
- • convex
- ○ reflex
- ■ semiconvex
- □ semireflex

one octant around v are in P. There are eight possible types of reflex vertices. Furthermore, there are vertex types where exactly three octants are occupied, so two of the adjacent edges are convex and one is reflex. We call such a vertex *semiconvex*. There are 24 different semiconvex types. Finally, there are vertex types where all but three octants are occupied, denoted as *semireflex*. See Genc [7], p. 38, for a classification of possible vertex types of general orthogonal polyhedra.

By definition guards are unbounded polyhedra with at most one vertex. In this context we restrict ourselves to 3-regular orthogonal guards. From now on, a guard is an unbounded orthogonal cone with at most one x-edge, at most one y-edge and at most one z-edge. Consequently, a guard has at most one x-face, at most one y-face and at most one z-face.

The *type* of an edge $e \in E(P)$, is given by its direction (parallel to the x-, y-, or z-axis) and by which quadrants around e are occupied by P in the plane orthogonal to e. Either one quadrant is occupied, in which case we call e a *convex edge* or three quadrants are occupied, in which case we call e a *reflex edge*. For example, we say an edge e is a convex z-$(++)$ edge if e is vertical and locally around e, the points in P are the points with both higher x- and y-coordinate. Or we say e is a reflex x-$(+-)$ edge if e is parallel to the x-axis and P occupies all but one quadrant around e, namely it leaves out the quadrant that lies behind e (higher y-coordinates) and below e (lower z-coordinates). There are 4 convex and 4 reflex edge types in any of the three directions, so totally, there are 24 different edge types. For each type, we fix a direction of the edge: We define the convex x-$(++)$ edges to be directed in negative x-direction (to the left). Similarly, we define the reflex x-$(++)$ edges to be oriented in positive x-direction (to the right). If we rotate P around the x-axis, an x-$(++)$ edge either becomes a x-$(+-)$ or a x-$(-+)$ edge. So rotating several times around all possible axes, each time by $\pi/2$, an edge can change from any type to any other type. We define the directions of all types in such a way that rotating by $\pi/2$ around any of the three coordinate axes flips the orientation of an edge.

Observation 1 *The edges of P can be oriented according to their type such that rotating by $\pi/2$ around any coordinate axis reverses the orientation.*
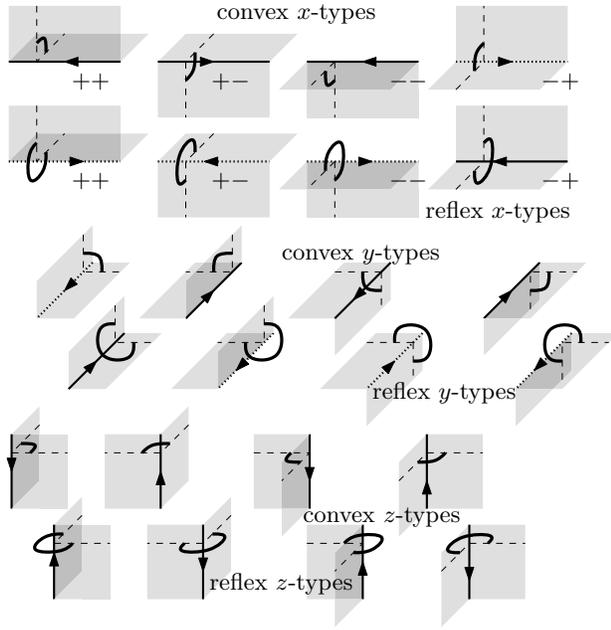
Figure 2: All possible edge types and their orientations.

See Figure 2 for a possible orientation of all edge types. There are exactly two ways to orient the convex types such that the observations holds and, independently, exactly two ways to orient the reflex types such that the observation holds. From now on, we think of every edge to be oriented as shown in Figure 2.

Observation 2 *At a vertex $v \in V(P)$ either all adjacent edges are pointing to v or all adjacent edges are pointing away from v.*

First consider just one convex vertex type. After observing the property for this type, it follows for all other convex vertex types directly. Repeatedly rotating the vertex by $\pi/2$ around any coordinate axis, we can go from one convex type to any other convex type. With each single rotation, the orientation of all three adjacent edges flip. So if the edges were pointing towards the vertex before, they are all pointing away after the rotation and vice versa, see Figure 3.

With this observation we have reproved that the graph of a 3-regular orthogonal polyhedron is bipartite, as observed for simple orthogonal polyhedra by Eppstein and Mumford [6]. (Their proof is somewhat easier, as it is a direct consequence of the fact that the graph is planar and the numbers of edges of every face is even.)

Corollary 1 *The graph of a 3-regular orthogonal polyhedron is bipartite.*

In 2D, placing a natural vertex guard onto every other vertex of an orthogonal polygon gives a valid guarding, see [5], Theorem 9, where this is proved for simple
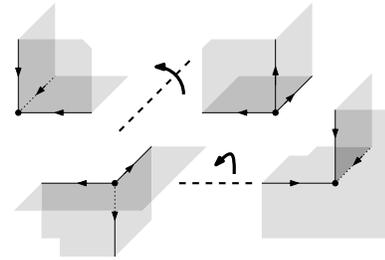
Figure 3: How the orientations of the edges adjacent to a convex and a semiconvex vertex flip when the polyhedron gets rotated by $\pi/2$ around a coordinate axis.

polygons. Because guarding a polygon and guarding its complement are equivalent problems (cf. [2], Observation 5), the same holds for polygons with holes.

Theorem 2 [5] *An orthogonal polygon with n vertices (possibly containing holes) can be guarded by $n/2$ guards placing a natural vertex guard onto every other vertex.*

For a face f of a polyhedron, let $\overline{f}$ denote the plane that contains f. A set G of guards *covers* a point r if there is an $\varepsilon > 0$ such that for any point pair $p, q \in B_\varepsilon(r)$ with $p \in P$ and $q \notin P$ there is some guard in G that contains p but does not contain q. We say that a set G of guards *covers* a face $f \in F(P)$, if G covers some point p in the interior of f. (p is in the *interior* of f if $p \in f$ and $p \notin e$ for any $e \in E(P)$.) and G *covers f completely*, if G covers all points in the interior of f. If $\{g\}$ covers f, then g has a face f_g with $\overline{f_g} = \overline{f}$.

Theorem 3 *For any integer $k \geq 2$, there are 3-regular orthogonal polyhedra with $4k$ vertices that cannot be guarded by fewer than k guards.*

Proof. Take a 2-dimensional orthogonal polygon Q with $2k$ pairwise non-collinear edges. Let P be a (right) prism with base Q. P has $2k + 2$ faces, k of which are x-faces. Consider a guarding $\mathcal{G}$ of P. Each face f of P has to be covered by at least one guard. No two x-faces are coplanar, so any guard can cover at most one x-face. Therefore, there are at least k guards in $\mathcal{G}$. □

3 Guarding with $n/2$ Natural Vertex Guards in 3D

Theorem 4 *A 3-regular orthogonal polyhedron P can be guarded with $n(P)/2$ natural vertex guards.*

Proof. Place a guard onto every vertex where all edges are pointing inwards. We show that for every inside/outside point pair (p, q) there is a guard g that distinguishes p and q. Denote the axis-parallel cube spanned by $p = (p_x, p_y, p_z)$ and $q = (q_x, q_y, q_z)$ by Q.
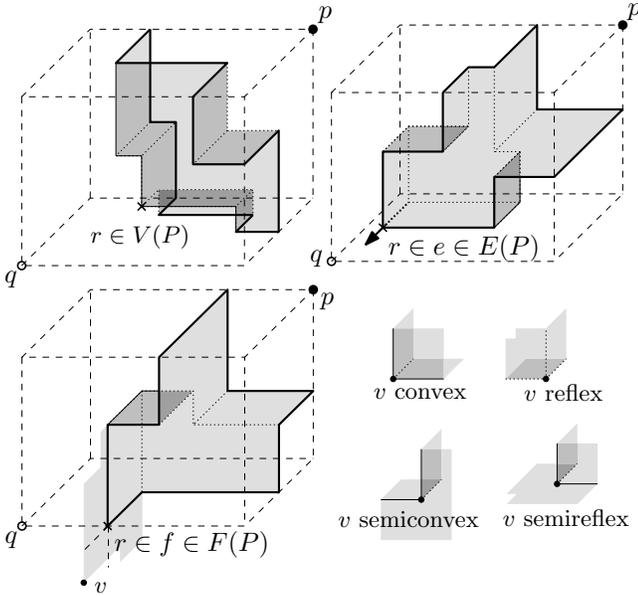
Figure 4: The three subcases of Case 1: Vertex v can have any of the types to the right.

Edge orientations and hence guardings are symmetric under a rotation by an angle of π around a coordinate axis. Therefore we may suppose without loss of generality that **either** $q_x \leq p_x$, $q_y \leq p_y$ and $q_z \leq p_z$, **or** $p_x \leq q_x$, $p_y \leq q_y$ and $p_z \leq q_z$.

First consider the case that $q_x \leq p_x$, $q_y \leq p_y$ and $q_z \leq p_z$. Look at $Q \cap P$ and pick the point $r = (r_x, r_y, r_z) \in P \cap Q$ which minimizes $r_x + r_y + r_z$. The point r can arise in three different ways, as depicted in Figure 4: If $r \in V(P)$, r is a convex vertex such that P occupies the $(+++)$-octant. Therefore, there is a guard on it which distinguishes p and q. If $r = e \cap f$ is the intersection of an edge $e \in E(P)$ and a face f of Q, f must be adjacent to q and e must be a convex edge orthogonal to it. Therefore, e is pointing outward of Q to a guard g distinguishing p and q. If r is on an edge of Q and in the interior of a face f of P, look at f as a 2D-polygon: Drawing a horizontal and a vertical line through r divides f into four quadrants. f has a convex vertex in each quadrant, in particular in the quadrant opposite to $f \cap Q$. No matter which type v has as a vertex of P, it is going to distinguish p and q.

In the case where $p_x \leq q_x$, $p_y \leq q_y$ and $p_z \leq q_z$, we have to use a slightly different argument. Let A be the face of Q that is adjacent to p and orthogonal to the z-axis. Consider the 2-dimensional orthogonal polygon P' we get by intersecting P with the plane $\overline{A}$ and picking the connected component of the intersection that contains p. Let r be the point in $P' \cap A$ that maximizes $r_x + r_y$. As in the first case, there are three sub-cases to consider. If $r = A \cap e$, $e \in E(P)$, we observe that e is a convex z-edge of type $(--)$, so it is oriented
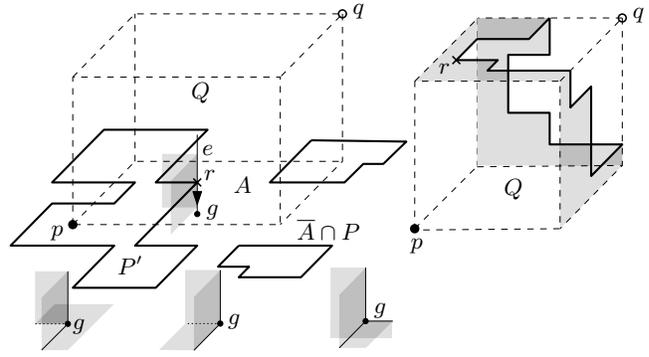
Figure 5: Case 2: Following e we find a guard g, which has one of several possible types.

downwards. So we can follow e to its end point outside Q where we find a guard g that distinguishes p and q, see Figure 5. If r is the intersection of an edge of A with a face f of P, then r divides f into four quadrants, in each of which we find a convex vertex of f (thought of as 2D-polygon). In particular, there is a convex vertex v of f in the quadrant opposite to the one containing $f \cap Q$ and there is a guard on v that distinguishes p and q. Finally, if r is a vertex of A, which means that A is completely contained in P, then pick another face B of Q adjacent to p, and repeat the argument. If all faces adjacent to p are completely inside P, then look at the top face C of Q and observe that $C \cap P$ must have at least one reflex vertex r. This vertex r lies on a reflex z-edge of P, which is pointing upward to a guard outside of Q that distinguishes p and q, see the example to the right in Figure 5. □

4 Improving the Bound to $3n/8$

Let P be an arbitrary 3-regular orthogonal polyhedron with n vertices. Every vertex is incident to one x-edge, one y-edge and one z-edge. Thus, there are exactly $n/2$ edges in each direction. In the guarding described above we used one guard per z-edge. In order to reduce the number of guards, we now place a natural edge guard onto roughly half of the z-edges only, namely onto those that are pointing downwards.

Such a guarding for sure covers all x– and y-faces, but we have not done anything about the z-faces yet. An easy solution would be to place a natural face guard g_f onto every z-face f. This would yield a valid guarding, but the number of z-faces could be as large as $n/4$ (even after permuting the coordinate axes). So instead, we replace every natural edge guard g_e on a z-edge e pointing down to a vertex v of f by a vertex guard g_v: $g_v := g_e \cap g_f$ if f has the interior of P above and $g_v := g_e \cup g_f$ if f has the interior of P below. So— in some sense—we combine the natural edge guards
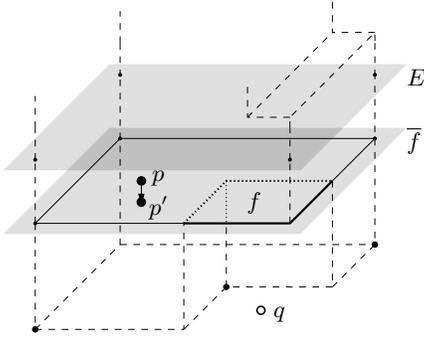
Figure 6: An increasing event.

and the natural face guards to vertex guards. (Note that these new guards are not necessarily natural vertex guards.) However, some z-faces may not be covered still. We call a z-face f *good*, if the guards we place on vertices of f cover f completely, and we call f *bad*, otherwise. If a z-face f is bad, then we put a natural face guard g_f onto f. So we have to make sure that this does not happen too often and that we will not use more than roughly $n/8$ face guards in this way. Let $\mathcal{G}$ be the set of vertex guards as described above together with the face guards on bad z-faces.

Lemma 5 $\mathcal{G}$ *is a valid guarding of P.*

Proof. We use a sweep argument. For simplicity, we assume that P has no coplanar z-faces. Imagine sweeping a plane E orthogonal to the z-axis upwards and look at the *intersection polygon* $Q = P \cap E$. Whenever E is coplanar with a z-face f of P (called the *event face*), the intersection polygon Q changes: the new intersection polygon Q' is either bigger or smaller, $Q' = Q \cup f$ or Q' is (the closure of) $Q \setminus f$. We call the first case an *increasing event*, the second case a *decreasing event*.

The claim is that using guards encountered so far only, we are able to guard P as far as we have seen it: At any moment the set $\tilde{\mathcal{G}}$ of guards that lie below the sweep plane E, together with an imaginary face guard g_E that is the closed halfspace below E, is a guarding of the part $\tilde{P} = P \cap g_E$ of P below E. We prove this claim by induction on the number of event faces processed.

At the beginning $\tilde{P}$ is empty. At some point, we hit the first z-face f of P. The vertices of f correspond to z-edges of P starting at f and going upwards. According to our rule, there is a guard on every second vertex of f, which we can think of as a 2-dimensional guarding for f extending to the region orthogonally above f.

Increasing Event. Let f be an increasing event face, that is, the interior of P lies above f (Figure 6). We claim that after E has passed f (but no other event face yet), we still have a valid guarding for $\tilde{P}$. Consider a point pair $p \in \tilde{P}$ and $q \notin \tilde{P}$. We may suppose without

loss of generality that both p and q lie in the closed halfspace below $\overline{f}$: If one of the points, say, p lies above $\overline{f}$, then consider the orthogonal projection p' of p onto $\overline{f}$ instead. As all guards in $\tilde{\mathcal{G}}$ are located in the closed halfspace below $\overline{f}$, none of them distinguishes p and p'.

If p and q are both below $\overline{f}$, then by induction there is a guard that distinguishes them. If both p and q lie in $\overline{f}$, then we are in a 2-dimensional situation and find a guard that distinguishes them because our guarding contains a 2-dimensional guarding of Q'. (The vertices of Q' correspond to z-edges. There is a guard on every other z-edge that—intersected with $\overline{Q'}$—is a natural 2D vertex guard of Q'. See Theorem 2.)

If $p \in \overline{f}$ and q lies below, then either there is a face guard g_f that does the job or f is a good face. In the latter case, the new vertex guards (i.e., those placed when handling the event face f) collectively cover f. If $p \in f$, then one of these new guards distinguishes p and q. Otherwise, a point p' slightly below p lies within P as well. By induction there is some old guard (i.e., a guard placed before handling the event face f) to distinguish p' from q. As such a guard cannot distinguish between p and p', it also distinguishes p and q. Symmetrically, if $q \in \overline{f}$ and p is below, consider a point q' located slightly below q and note that $q' \notin P$ because the event is increasing. By induction, there is an old guard that distinguishes p and q' but cannot distinguish q and q'. Hence this guard also distinguishes p and q.

Decreasing Event. Consider a point pair $p \in \tilde{P}$ and $q \notin \tilde{P}$. As above, we may suppose without loss of generality that p lies in the closed halfspace below $\overline{f}$. However, if q lies above the event face f—that is, the orthogonal projection q' of q onto $\overline{f}$ lies in f—we cannot simply replace q by q', because $q' \in P$. But we know that some guard that was placed when handling the event face f distinguishes q and q', and every guard placed when handling f contains the closed halfspace below $\overline{f}$. Therefore, if p lies below $\overline{f}$, then this guard distinguishes p and q. If $p \in \overline{f}$, then recall that we have a 2-dimensional guarding for Q', which must contain a guard that can distinguish p and q'. This guard classifies q' as outside and so it does with q. Hence it distinguishes p and q. We may thus assume that q lies in the closed halfspace below $\overline{f}$ as well. It follows inductively that there exists an old guard that distinguishes p and q. $\qquad\square$

Lemma 6 *Under a random rotation around the z-axis by a multiple of $\pi/2$ and independently, a reflection with respect to the plane $z = 0$ with probability $1/2$, a z-face with 4 or 6 vertices is good with probability at least $1/2$.*

Proof. Each vertex v of f corresponds to a z-edge e_v of P. e_v either starts at v going upwards or it ends at v. We call v a *starting vertex* or an *ending vertex*, respectively. If v is a starting vertex and e_v is pointing

to v, there is a guard g_v on v. Else, if e_v is oriented upwards or v is an ending vertex, there is no guard on v. Under a reflection in the xy-plane, starting vertices turn into ending vertices and vice versa, see Figure 7.

Consider f as a 2D-polygon. If it has 6 vertices, exactly 5 are convex and one is reflex. One of the convex 2D-vertex-types appears twice, the other three appear exactly once and are referred to as *unique*, therefore. If f has 4 vertices, they are all unique. So in any case we have at least three unique (convex) vertices. Moreover, these vertices appear consecutively along the boundary of f, which implies that for at least one of them the incident z-edge is directed towards the vertex.

If there is a guard on some unique vertex of f, then f is good because this guard covers f completely. (Note that f may be bad if there is a guard at some non-unique vertex only, see Figure 7 (B).) If two adjacent unique vertices of f are starting, then—by the remark above—at least one of them has a guard. When reflecting P at the plane $z = 0$, all ending vertices turn into starting vertices and vice versa. Hence, if two adjacent unique vertices of f are ending, then at least of them has a guard after this reflection and so f is good with probability at least $1/2$. It remains to consider the case that the three unique vertices of f follow the pattern starting–ending–starting or ending–starting–ending and neither of the starting vertices, with and without reflection, has a guard (Figure 7 (C)). When rotating around the z-axis by $\pi/2$ or $3\pi/2$, starting vertices remain starting and ending vertices remain ending. But edge orientations flip and so a starting vertex without guard turns into a starting vertex with guard. As a result, both the original face and the reflected variant turn good after such a rotation. So again with probability at least $1/2$, the face f appears as a good face. $\square$

Theorem 7 *Let P be a 3-regular orthogonal polyhedron. Then P can be guarded with $3n(P)/8$ guards.*

Proof. Guard P as described at the beginning of the section. Whenever we have to use a face guard g_f, we charge it to the edges of f. The edges of a z-face f are also edges of P and each x- or y-edge appears exactly once as an edge of a z-face. So the edges of a z-face f of degree $d \geq 8$ each get charged at most $1/d \leq 1/8$. The edges of a face f with degree $d = 4$ or $d = 6$ get charged $1/d \leq 1/4$ if we have to use a face guard g_f. If we pick a random rotation around the z-axis and independently decide to reflect P with respect to the plane $z = 0$ with probability $1/2$, we have shown that f ends up as a bad face with probability at most $1/2$. So the x- or y-edges gets charged at most $1/4$ with probability at most $1/2$ and 0 otherwise. So the expected charge of an $x-$ or y-edge is at most $1/8$. The vertex guards get charged to their corresponding z-edge. Under a random
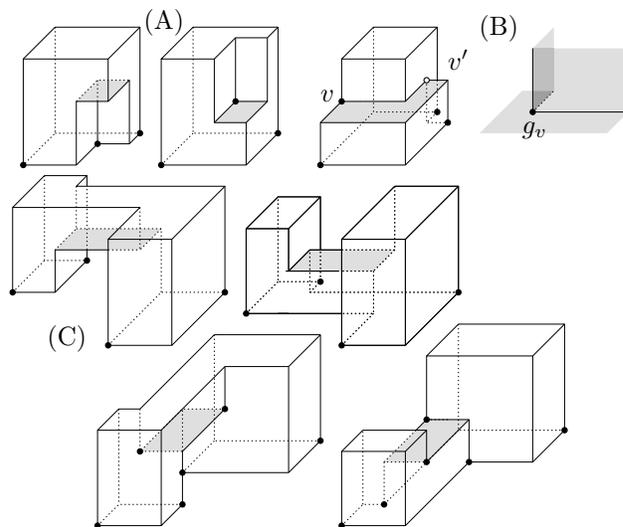
Figure 7: (A) a bad face that turns good after reflection; (B) a face that is bad even though there is a vertex v with a guard g_v on it: g_v does not cover the face completely; (C) a face that stays bad after reflecting, but both itself and the reflected version turn good after rotating around the z-axis by $\pi/2$.

rotation, a z-edge gets charged 1 with probability $1/2$ and 0 otherwise. Therefore, the expected total charge is going to be at most $\frac{1}{8}n + \frac{1}{2}n/2 = 3n/8$, so there is a rotation (possibly combined with a reflection) such that at most $3n/8$ guards are used. $\square$

References

[1] J. Abello, V. Estivill-Castro, T. Shermer, and J. Urrutia. Illumination of orthogonal polygons with orthogonal floodlights. *Internat. J. Comput. Geom. Appl.*, 8(1):25–38, 1998.

[2] T. Christ, M. Hoffmann, Y. Okamoto, and T. Uno. Improved bounds for wireless localization. *Algorithmica*, 57:499–516, July 2010.

[3] T. Christ and A. Mishra. Wireless localization with vertex guards. In *Abstracts of the 27th European Workshop on Computational Geometry (EuroCG '11)*, Morschach, Switzerland, March 2011.

[4] D. P. Dobkin, L. Guibas, J. Hershberger, and J. Snoeyink. An efficient algorithm for finding the CSG representation of a simple polygon. *Algorithmica*, 10:1–23, 1993.

[5] D. Eppstein, M. T. Goodrich, and N. Sitchinava. Guard placement for efficient point-in-polygon proofs. In *Proc. 23rd Annu. Sympos. Comput. Geom.*, pages 27–36, 2007.

[6] D. Eppstein and E. Mumford. Steinitz theorems for orthogonal polyhedra. In *Proc. 26th Annu. Sympos. Comput. Geom.*, pages 429–438, 2010.

[7] B. Genc. *Reconstruction of Orthogonal Polyhedra*. PhD thesis, University of Waterloo, 2008.