Edge-guarding Orthogonal Polyhedra

Nadia M. Benbernou* Erik D. Demaine† Martin L. Demaine† Anastasia Kurdia Joseph O'Rourke‡ Godfried Toussaint§ Jorge Urrutia¶ Giovanni Viglietta‖

Abstract

We address the question: How many edge guards are needed to guard an orthogonal polyhedron of e edges, r of which are reflex? It was previously established [3] that $e/12$ are sometimes necessary and $e/6$ always suffice. In contrast to the closed edge guards used for these bounds, we introduce a new model, open edge guards (excluding the endpoints of the edge), which we argue are in some sense more natural in this context. After quantifying the relationship between closed and open edge guards, we improve the upper bound to show that, asymptotically, $(11/72)e$ (open or closed) edge guards suffice, or, in terms of r, that $(7/12)r$ suffice. Along the way, we establish tight bounds relating e and r for orthogonal polyhedra of any genus.

1 Introduction

We consider a natural variation of the famous Art Gallery Problem: given an orthogonal polyhedron $\mathcal{P}$ (possibly with holes) in $\mathbb{R}^3$, select a minimum number of edges of $\mathcal{P}$ (called *edge guards*) so that the interior of $\mathcal{P}$ is fully guarded (i.e., each point of $\mathcal{P}$ is visible to at least one guard).

Although traditionally edge guards are *closed* in that they occupy the entire edge, we suggest that *open edge guards*, which exclude their endpoints, are a more natural model. We establish that at most three times as many open edge guards are needed to cover the same polyhedron as closed edge guards, a tight bound. Despite this apparent weakness, we improve the previous upper bound for closed edge guards to a better bound for open edge guards.

Open guards, open polyhedra. In line with our focus on open edge guards, we also focus on guarding *open polyhedra*, i.e., bounded polyhedra excluding their boundaries. Consider an orthogonal polyhedron to represent an empty room with solid walls, with the task to place guards who can detect unwelcome intruders. Because an intruder cannot hide *within* a wall but rather must be located inside the room, there is no need to guard the walls of the room, i.e., the boundary of the polyhedron. A guarding problem can alternatively be viewed as an illumination problem, with guards acting as light sources. Incandescent lights are modeled as point guards, and fluorescent lights are modeled as segment guards. In the latter case, it is more realistic to disregard the endpoints of the edge guards. The amount of light that a point interior to the polyhedron receives is proportional to the total length of the segments illuminating that point. Employing the open edge-guard model ensures that if a point is illuminated, it receives a strictly positive amount of light, and makes the model more realistic.

Notice that these two definitions of illuminated points (visible to an open edge guard or receiving a strictly positive amount of light from closed edge guards) cease to be equivalent if we consider polyhedra with boundary.

Previous work. Although guarding orthogonal polygons is a relatively well-studied topic, few positive results exist for orthogonal polyhedra. To the best of our knowledge, the only results relevant to the problem studied in this paper were given by Urrutia in his survey [3, Sec. 10]: for a polyhedron of e edges, $e/6$ closed edge guards always suffice, and $e/12$ guards are sometimes necessary (Figure 1). He also conjectured that the latter is the correct bound, i.e., that $e/12 + O(1)$ suffice. Because no proof of the upper bound was given in [3], another contribution here is that our proof for the $(11/72)e$ bound incorporates the essence of Urrutia's unpublished $e/6$ proof.

*Department of Mathematics, Massachusetts Institute of Technology, Cambridge, MA, USA, `nbenbern at mit.edu`

†MIT Computer Science and Artificial Intelligence Laboratory, 32 Vassar St., Cambridge, MA 02139, USA, `{edemaine, mdemaine}@mit.edu`

‡Department of Computer Science, Smith College, Northampton, MA, USA, `orourke@cs.smith.edu`

§Department of Music, Harvard University, Cambridge, MA, USA, Department of Computer Science, Tufts University, Medford, MA, USA, School of Computer Science, McGill University, Montreal, QC, Canada, `godfried at cs.mcgill.ca`

¶Instituto de Matemáticas, Universidad Nacional Autónoma de México. Partially supported by SEP-CONACYT of Mexico, Proyecto 80268 and by Spanish Government under Project MEC MTM2009-08652.

‖Department of Computer Science, University of Pisa, Italy, `viglietta@di.unipi.it`. Partially supported by MIUR of Italy under project AlgoDEEP prot. 2008TFBWL4.

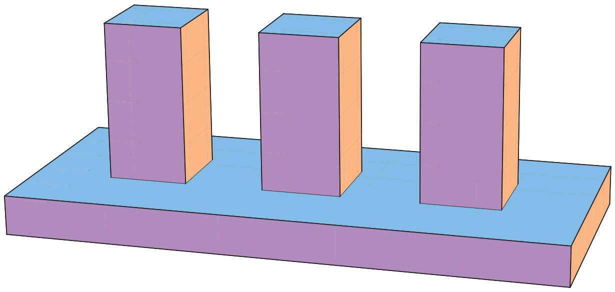

Figure 1: Lower bound example from [3]: k guards are required to guard a polyhedron with a total of $12k + 12$ edges.

2 Properties of orthogonal polyhedra

We start with precise definitions of necessary concepts. Given two points x and y, we denote by xy the (closed) straight line segment joining x and y, and by $\widetilde{xy}$ the corresponding *open segment*, i.e., the relative interior of xy.

Orthogonal polyhedra. A *cuboid* is defined as a compact subset of $\mathbb{R}^3$ bounded by 6 axis-orthogonal planes. The union of a finite non-empty set of cuboids is an *orthogonal polyhedron* if its boundary is a connected 2-manifold.

A *face* of an orthogonal polyhedron is a maximal planar subset of its boundary, whose interior is connected and non-empty. Faces are orthogonal polygons with holes, perhaps with degeneracies such as hole boundaries touching each other at single vertex, etc. A *vertex* of an orthogonal polyhedron is any vertex of any of its faces. An *edge* is a minimal positive-length straight line segment shared by two faces and connecting two vertices of the polyhedron. Each edge, with its two adjacent faces, determines a dihedral angle, internal to the polyhedron. Each such angle is 90° (at a *convex edge*) or 270° (at a *reflex edge*).

Visibility and guarding. *Visibility* with respect to a polyhedron $\mathcal{P}$ is a relation between points in $\mathbb{R}^3$: point x *sees* point y (equivalently, y is *visible* to x) if $xy \setminus \{x\}$ lies entirely in the interior of $\mathcal{P}$. Note that, according to the previous definition, the boundary of $\mathcal{P}$ occludes visibility; no portion of xy, except the endpoint x, can lie on the boundary of $\mathcal{P}$. Also, x is assumed to be invisible to itself when it belongs to the boundary. Given a point $x \in \mathcal{P}$, its *visibility region* $V(x)$ is the set of points that are visible to x. Similarly, the visibility region of a set $X \subseteq \mathbb{R}^3$, denoted by $V(X)$, is the set of points that are visible to at least one point in X.

The Art Gallery Problem we consider in this paper is: given an orthogonal polyhedron $\mathcal{P}$, efficiently select a (sub)set of its edges $e_1, e_2, \dots, e_k$, called the *guarding set*, so that the whole interior of $\mathcal{P}$ is guarded by the interiors of the selected edges. In other words, the interior of $\mathcal{P}$ must coincide with $V(\widetilde{e_1}) \cup V(\widetilde{e_2}) \cup \ldots \cup V(\widetilde{e_k})$. Our goal is also to minimize k, the number of selected edges. We bound k with respect to the total number e of edges of $\mathcal{P}$, or the number r of its reflex edges.

The notion of ε-guarding implies that each point is guarded by at least one positive-length segment. Guarding in our open polyhedra model is equivalent to ε-guarding in that, if a point is guarded, then it is also guarded by a positive-length segment, lying on some guard.

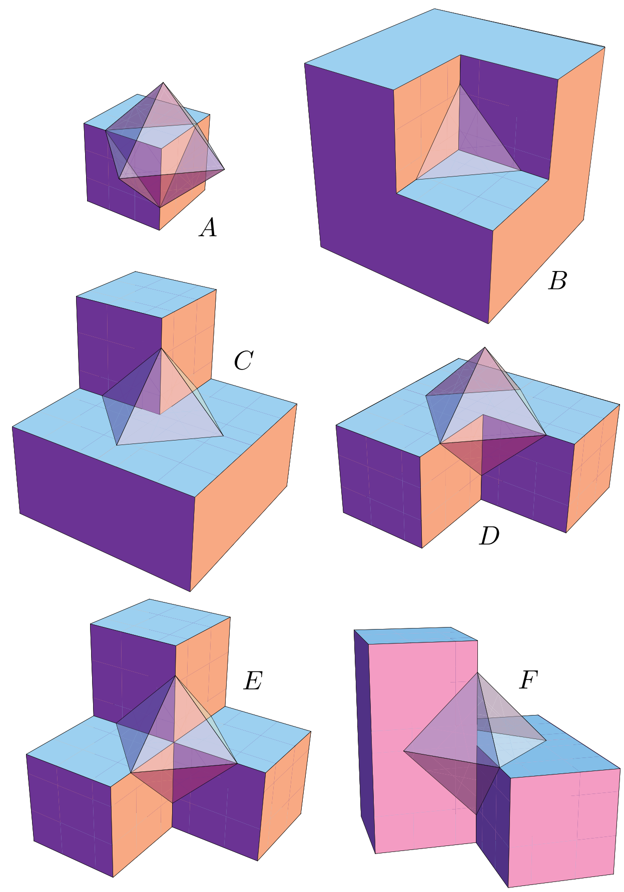

Figure 2: The six vertex types.

Vertex classification. Based on the number of incident reflex and convex edges, the vertices of orthogonal polyhedra form six distinct classes, denoted here by A, B, C, D, E and F, and are introduced as follows. Consider the eight octants determined by the coordinate axes intersecting at a given vertex, and place a sufficiently small regular octahedron around the vertex, such that each of its faces lies in a distinct octant. By definition, the set of faces that fall inside (or outside) the polyhedron is connected: recall that the boundary of a polyhedron is a 2-manifold. Consider all possible ways of partitioning the faces of the octahedron into two non-empty connected sets, up to isometry (refer to Figure 2):

- There is essentially a single way to select 1 face (resp. 7 faces). This corresponds to an A-vertex (resp. a B-vertex).
- There is a single way to select 2 faces (resp. 6 faces). This case does not correspond to a vertex of the orthogonal polyhedron: it implies that the considered point is not a vertex of any face on which it lies.
- There is a single way to select 3 faces (resp. 5 faces). This corresponds to a D-vertex (resp. a C-vertex).
- There are three ways to select 4 faces. One of them implies that the point lies in the middle of a face, hence it does not correspond to a vertex. The other two choices correspond to an E-vertex and an F-vertex, respectively.

Auxiliary results. We now present two useful properties of orthogonal polyhedra that will be employed to prove our main results. Let us denote by A the number of A-vertices in a given orthogonal polyhedron, and so on, for each vertex class.

Lemma 1 *In every orthogonal polyhedron with $r > 0$ reflex edges,* $3A + D \geqslant 28$.

Proof. Consider the bounding cuboid of the polyhedron and the set of orthogonal polygons (perhaps with holes, without degeneracies), formed by intersection of the faces of the cuboid and the polyhedron. The vertices of those polygons are either A-vertices or D-vertices of the polyhedron (convex vertices are A-vertices, and reflex vertices are D-vertices). Our strategy is to only look at the vertices belonging to the bounding faces and ensure that there is a sufficient number of them. Namely, we only need to show that there are at least

(a) 10 A-vertices, or

(b) 9 A-vertices and 1 D-vertex, or

(c) 8 A-vertices and 4 D-vertices.

Suppose each face of the bounding cuboid contains exactly one rectangle. If all the vertices of these rectangles coincide with the corners of the bounding cuboid, then the polyhedron is convex, contradicting the assumptions. Hence, there is a vertex x that is not a corner of the bounding cuboid. Let f denote a face containing x. At least one of the vertices, denoted by y, adjacent to x in the rectangle contained in f, is such that $\widetilde{xy}$ does not lie on an edge of the bounding cuboid. Let f' be the bounding face opposite to f, and f'' be the bounding face chosen as shown in Figure 3: out of

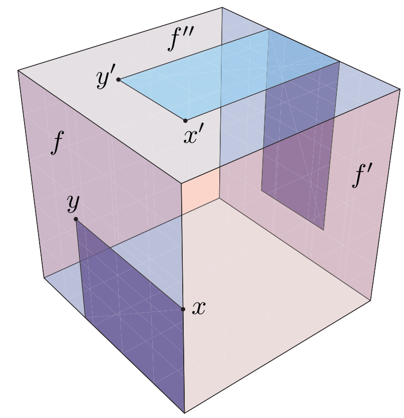

Figure 3: An illustration of the proof of Lemma 1.

the 4 faces surrounding f, f'' is the one that lies on the "side" of xy. f and f' contain two disjoint rectangles, and thus exactly 8 distinct A-vertices. Additionally, f'' has two extra A-vertices, lying on an edge $x'y'$ parallel to xy (refer to Figure 3). Collectively, f, f' and f'' contain at least 10 A-vertices, so (a) holds.

On the other hand, if there exists a bounding face f whose intersection with the polyhedron is not a single rectangle, then we need to analyze the following three cases. Let f' be the bounding face opposite to f.

- If f contains at least two polygons (those polygons' boundaries must be disjoint because f is a bounding face), then collectively f and f' contain at least 12 distinct A-vertices, so (a) holds. Indeed, every orthogonal polygon has at least 4 convex vertices.
- If f contains a polygon with at least one hole, then the polygon's external boundary contains at least 4 convex vertices (equiv. A-vertices), and the hole has at least 4 reflex vertices (equiv. D-vertices). f' also contains at least 4 convex vertices (A-vertices). Together f and f' contain at least 8 A-vertices and 4 D-vertices, so (c) holds.
- If f contains just one polygon, which is not convex, then such a polygon has at least 5 convex vertices and one reflex vertex. Together with f', there are at least 9 A-vertices and 1 D-vertex, so (b) holds.

□

Theorem 2 *For every orthogonal polyhedron with e edges in total, $r > 0$ reflex edges and genus $g \geqslant 0$,*

$$\frac{1}{6}e + 2g - 2 \leqslant r \leqslant \frac{5}{6}e - 2g - 12$$

holds. Both inequalities are tight for every g.

Proof. Let $c = e - r$ be the number of convex edges. Let A be the number of A-vertices, etc.. Double counting the pairs (edge, endpoint) yields (refer to Figure 2)

$$2c = 3A + C + 2D + 3E + 2F, \tag{1}$$

$$2r = 3B + 2C + D + 3E + 2F. \tag{2}$$

The angle deficit (with respect to 2π) of A- and B-vertices is $\pi/2$, the deficit of C- and D-vertices is $-\pi/2$, the defect of E- and F-vertices is $-\pi$. Hence, by the polyhedral version of Gauss-Bonnet theorem (see [2, Thm. 6.25]),

$$A + B - C - D - 2E - 2F = 8 - 8g. \tag{3}$$

Finally, since all the variables involved are non-negative,

$$9B + 3C + 3E + F \geqslant 0. \tag{4}$$

Subtracting 3 times (3) from 2 times (4) yields

$$-3A + 15B + 9C + 3D + 12E + 8F \geqslant 24g - 24.$$

Further subtracting (1) and adding 5 times (2) to the last inequality yields

$$2c - 10r + 24g - 24 \leqslant 0,$$

which is equivalent to $\frac{1}{6}e + 2g - 2 \leqslant r$.

To see that the left-hand side inequality is tight for every r and g, consider the staircase-like polyhedron with holes depicted in Figure 4. If the staircase has k "segments" and g holes, then it has a total of $6k+12g+6$ edges and $k + 4g - 1$ reflex edges.

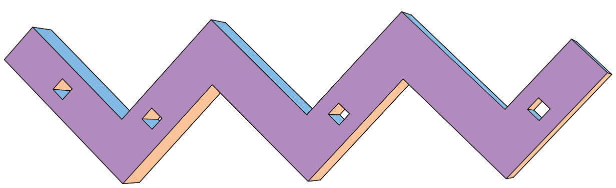

Figure 4: A polyhedron that achieves the tight left-hand side bound in Theorem 2.

According to Lemma 1, $3A + D \geqslant 28$, unless the polyhedron is a cuboid. Then

$$9A + 3D + 3E + F \geqslant 84. \tag{5}$$

Subtract 3 times (3) from 2 times (5):

$$15A - 3B + 3C + 9D + 12E + 8F \geqslant 24g + 144.$$

Subtract (2) and add 5 times (1):

$$2r - 10c + 24g + 144 \geqslant 0,$$

which is equivalent to $r \leqslant \frac{5}{6}e - 2g - 12$.

To see that the right-hand side inequality is also tight, consider a cuboid with a staircase-like well carved in it, and a number of cuboidal "poles" carved out from the surface of the well (i.e., the *negative* version of Figure 4). If the staircase has k "segments" and g poles, then the polyhedron has a total of $6k + 12g + 18$ edges and $5k + 8g + 3$ reflex edges. □

Notice that the statement of the previous theorem does not hold if we change the definition of orthogonal polyhedron by dropping the condition of connectedness of the boundary. Indeed, consider a cube and remove several smaller disjoint cubic regions from its interior. The resulting shape has unboundedly many reflex edges and just 12 convex edges.

Finally, the next proposition characterizes visibility regions of points belonging to polyhedra.

Proposition 3 *The visibility region of any point in a polyhedron or on its boundary is an open set.*

Proof. Let x be a point in a polyhedron $\mathcal{P}$. Let f be a face of $\mathcal{P}$, not containing x. The region of space "occluded" by f is a closed set $O(x, f)$, shaped like a truncated unbounded pyramid with apex x and base f. The number of faces is finite. Forming the union of all $O(x, f)$, for every face f not containing x, we obtain a closed set $O(x)$.

The region occluded by the faces containing x is the corresponding (unbounded) solid angle, external with respect to $\mathcal{P}$, which is a closed set. Its union with $O(x)$ is again a closed set, and therefore the complement of $O(x)$ is an open set, which by definition is $V(x)$. □

Observe that the visibility regions of open and closed edges are also open sets, since they are unions of open sets.

3 Open vs. closed edge guards

We now establish the relationship between the number of open and closed edge guards required to guard the interior of an orthogonal polyhedron.

Theorem 4 *Any orthogonal polyhedron guardable by k closed edge guards is guardable by at most $3k$ open edge guards, and this bound is tight.*

Proof. Given a set of k closed edges that guard the entire polyhedron, we first construct a guarding set of open edges of size at most $3k$ and then show that this set also guards the entire polyhedron. The construction is simple: for each closed edge uv from the original guarding set, place the open edge $\widetilde{uv}$ into the new guarding set. For the endpoint u, also add a reflex edge $\widetilde{uw}$ with $w \neq v$, if such edge exists, or any other edge incident to u otherwise. Similarly, an incident edge is selected for the other endpoint v. Hence, for each edge of the original guarding set, at most 3 open edges are placed in the new guarding set.

To prove the equivalence of the two guarding sets, we need to show that the volume that was guarded by an endpoint u of the closed edge from the original guarding set, is guarded by some point belonging to the interior of uv or the interior of uw, as chosen above, i.e., $V(u) \subseteq V(\widetilde{uv}) \cup V(\widetilde{uw})$.

Let x be any point previously guarded by u, $x \in V(u)$. By Proposition 3, a ball $\mathcal{B}$ centered at x belongs to $V(x)$. Then we create a right circular cone $\mathcal{C}$ with apex u, whose base is centered at x and is contained in $\mathcal{B}$. Clearly, $\mathcal{C} \subset V(u)$. Let $\mathcal{D}$ be a small-enough ball centered at u that does not intersect any face of the polyhedron except those containing u (refer to Figure 5). We prove that $\mathcal{D} \cap \mathcal{P} \subseteq V(\widetilde{uv} \cap \mathcal{D}) \cup V(\widetilde{uw} \cap \mathcal{D})$.

If u is an A-vertex, then $\mathcal{D} \cap \mathcal{P} \subseteq V(\widetilde{uv} \cap \mathcal{D})$. If u is a B-vertex (as illustrated), then of the 8 octants determined by orthogonal planes crossing at u, one is external to $\mathcal{P}$. Out of the 7 octants that need to be guarded, 6 are guarded by $\widetilde{uv} \cap \mathcal{D}$. The same holds for $\widetilde{uw}$, and together they guard all 7 octants (two of the octants guarded by $\widetilde{uv}$ are missing a face, but those two faces are guarded by $\widetilde{uw}$).

In all other cases (u is a C-, D-, E- or F-vertex), either uv or uw is a reflex edge. Assume without loss of generality that uv is reflex. Then, $\widetilde{uv} \cap \mathcal{D}$ sees all of $\mathcal{D} \cap \mathcal{P}$ (refer to Figure 2).

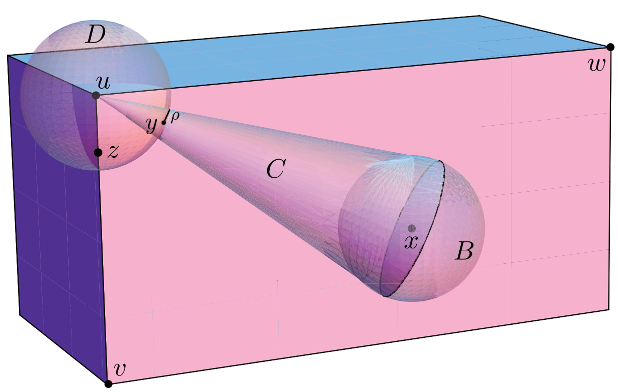

Figure 5: Construction from the proof of Theorem 4.

The boundaries of $\mathcal{D}$ and $\mathcal{C}$ intersect at a circle of radius $\rho > 0$. Let y be the center of that circle. There is a point z on $\widetilde{uv} \cap \mathcal{D}$ or on $\widetilde{uw} \cap \mathcal{D}$ that sees y, and hence the entire open segment $\widetilde{uz}$ sees y. Pick a point t on $\widetilde{uz}$ such that $||ut|| < \rho$. Then t sees x.

A similar argument holds for the visibility region of the other endpoint, v, of uv.

To see that 3 is the best achievable ratio between the number of open and closed edge guards, consider the polygon in Figure 6 and extrude it to an orthogonal prism. Each large dot in that figure represents the projection of some distinguished point located in the interior of the prism. The only (closed) edges that can see more than two selected points are the highlighted edges (on the lower or upper base of the prism). Picking those edges as guards yields the minimum set of guards, and together they guard the entire polyhedron. On the other hand, the relative interior of any edge can see at most one point of interest. At least as many open edge guards as there are distinguished points are necessary. $\square$

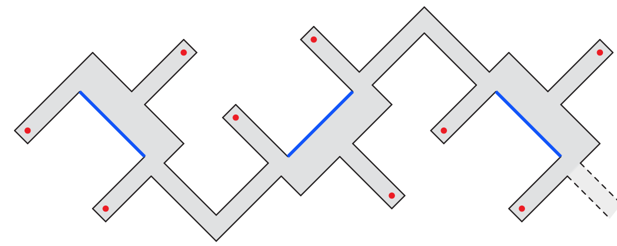

Figure 6: Matching ratio in Theorem 4. Notice that the same example also solves the corresponding problem for 2D polygons.

Note that the above analysis does not hold in the case of closed polyhedra, i.e., when the boundary does not obstruct visibility, since we can no longer argue that a single closed edge guard is locally dominated by 3 open edge guards.

4 Upper bound

We now establish an upper bound on the number of open edge guards required to guard an orthogonal polyhedron.

Theorem 5 *Every orthogonal polyhedron with e edges in total and r reflex edges is guardable by $\lfloor \frac{e+r}{12} \rfloor$ open edge guards.*

Proof. Let e_x and r_x be the number of X-parallel edges and reflex edges, respectively; e_y, e_z, r_y, r_z are similarly defined. Without loss of generality, assume X is the direction that minimizes the sum $e_x + r_x$, so that $e_x + r_x \leqslant \frac{e+r}{3}$. Of course, a guard on every X-parallel edge suffices to cover all of $\mathcal{P}$, but we can do much better with a selected subset of these edges. We argue below that selecting the three types of X-parallel edges circled in Figure 7 suffice (as do three other symmetric configurations). Let the number of X-edges of each of the eight types shown be $\alpha, \ldots, \delta'$ as labeled in Figure 7.

Hence we could place $\alpha+\beta'+\delta'$ guards, or $\gamma+\beta'+\delta'$ guards, or $\beta+\alpha'+\gamma'$ guards, or $\delta+\alpha'+\gamma'$ guards.

By choosing the minimum of these four sums, we place at most

$$(\alpha+\beta+\gamma+\delta+2\alpha'+2\beta'+2\gamma'+2\delta')/4$$

$$= \frac{e_x+r_x}{4} \leqslant \frac{e+r}{12}$$

guards.

Next we prove that our guard placement works.

We consider any point p in $\mathcal{P}$ and show that p is guarded by the edges selected in Figure 7. Let ω be the

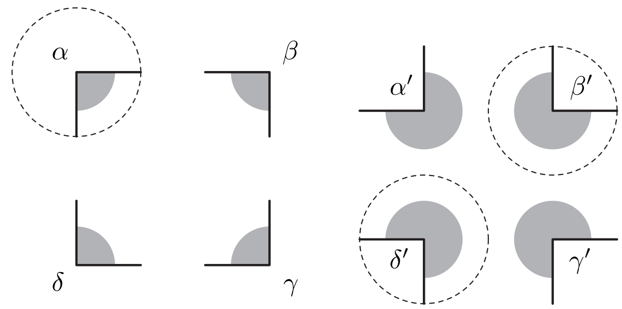

Figure 7: Possible configurations of X-edges. The X-axis is directed toward the reader. The circled configurations are those selected in the proof of Theorem 5.

X-orthogonal plane containing p and let Q be the intersection of the (open) polyhedron $\mathcal{P}$ with ω. To prove that p is guarded, we first shoot an axis-parallel ray from p. For our choice of guarding edges, the ray is directed upward. Let q be the intersection point of the ray and the boundary of Q that is nearest to p. Next, grow leftwards a rectangle whose right side is pq until it hits a vertex v of Q. If it hits several simultaneously, let v be the topmost. There are three possible configurations for v, shown in Figure 8, and each corresponds to a selected configuration in our placement of guards (Figure 7). If v lies in the interior of the guarding edge, then p is guarded. If v is an endpoint of such an edge, then we show that p is guarded by a sufficiently small neighborhood of v that belongs to the guarding edge. Every face of $\mathcal{P}$ that does not intersect ω has a positive distance from ω. Let d be the smallest such distance. Then, the points of the guarding edge at distance strictly less than d from v see p.

If a different triplet of guarding edges is chosen, the above construction is suitably rotated by a multiple of 90°. □

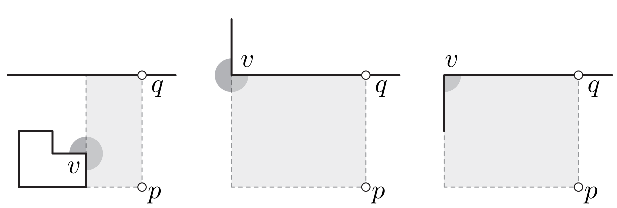

Figure 8: An illustration of the proof of Theorem 5.

Our placement of guards in the single slices resembles a construction given in [1], in a slightly different model.

By combining the results of Theorem 5 with those of Theorem 2, we obtain two corollaries.

Corollary 6 *Let e denote the number of edges of an orthogonal polyhedron and let g denote its genus. Then $\frac{11}{72}e - \frac{g}{6} - 1$ open edge guards are sufficient to guard the interior of the polyhedron.*

Corollary 7 *Let r denote the number of reflex edges of an orthogonal polyhedron and let g denote its genus. Then $\frac{7}{12}r - g + 1$ open edge guards are sufficient to guard the interior of the polyhedron.*

5 Conclusions

We have elucidated the relationship between the required number of closed edge guards and open edge guards. We also improved the current state of the art and obtained a better upper bound ($\frac{11}{72}e$ vs. the previously known $\frac{e}{6}$) on the number of edge guards that suffice for coverage.

We remark that, due to the observation following Theorem 2 , our methods do not improve on the $\frac{e}{6}$ upper bound when applied to orthogonal shapes with disconnected boundary. Indeed, in this case the $\frac{e+r}{12}$ given by Theorem 5 still holds, but the r to e ratio can be arbitrarily close to 1.

We conclude with a few possible future directions. The same construction used in Theorem 5 could be analyzed more closely to achieve a tighter upper bound. In contrast with the fact that the polyhedra with highest r to e ratio are responsible for the worst cases in our analysis, such polyhedra are nonetheless intuitively easy to guard by selecting a small fraction of their reflex edges. Isolating these cases and analyzing them separately may yield an improved overall bound.

We also conjecture that suitably placing guards on roughly half of the (open) reflex edges solves our Art Gallery Problem in any orthogonal polyhedron, which would imply that $\frac{1}{2}r + O(1)$ guards suffice (this many are needed in Figure 4 when $g = 0$).

Observe that our construction in Theorem 5 places guards in just one direction. It would be interesting to investigate this restriction of the Art Gallery Problem (i.e., with the additional constraint that edge guards are mutually parallel), perhaps showing that the lower bound given in Figure 1 can be improved in this more restrictive scenario.

On the other hand, refining our construction by placing guards in all three directions, according to some local properties of the boundary, is likely to yield better upper bounds.

References

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

[2] S. Devadoss and J. O'Rourke. *Discrete and Comptuational Geometry.* Princeton University Press, 2011.

[3] J. Urrutia. Art gallery and illumination problems. In J.-R. Sack and J. Urrutia, editors, *Handbook of Computational Geometry*, pages 973–1027. North-Holland, 2000.