

Optimally Guarding 2-Reflex Orthogonal Polyhedra by Reflex Edge Guards

Giovanni Viglietta*

October 27, 2017

Abstract

We study the problem of guarding an orthogonal polyhedron having reflex edges in just two directions (as opposed to three) by placing guards on reflex edges only.

We show that

$$\left\lfloor \frac{r - g}{2} \right\rfloor + 1$$

reflex edge guards are sufficient, where r is the number of reflex edges in a given polyhedron and g is its genus. This bound is tight for $g = 0$. We thereby generalize a classic planar Art Gallery theorem of O'Rourke, which states that the same upper bound holds for vertex guards in an orthogonal polygon with r reflex vertices and g holes.

Then we give a similar upper bound in terms of m , the total number of edges in the polyhedron. We prove that

$$\left\lfloor \frac{m - 4}{8} \right\rfloor + g$$

reflex edge guards are sufficient, whereas the previous best known bound was $\lfloor 11m/72 + g/6 \rfloor - 1$ edge guards (not necessarily reflex).

We also discuss the setting in which guards are *open* (i.e., they are segments without the endpoints), proving that the same results hold even in this more challenging case.

Finally, we show how to compute guard locations in $O(n \log n)$ time.

1 Introduction

Background

In discrete geometry, the *Art Gallery* problem asks to place a (preferably small) number of *guards* in a given enclosure, so that the guards collectively

*University of Ottawa, 800 King Edward Ave., Ottawa ON Canada, K1N 6N5.
gvigliet@uottawa.ca.

see the whole enclosure. Most classic results on the Art Gallery problem are surveyed in [5, 6, 9].

If the enclosure is a simple polygon with n vertices and the guards are points, then $\lfloor n/3 \rfloor$ guards are always sufficient, and there are polygons in which they are also necessary. If the polygon is *orthogonal* (i.e., its edges meet at right angles), then the optimal number of point guards reduces to $\lfloor n/4 \rfloor$. Furthermore, if the orthogonal polygon has h holes, it can be guarded by

$$\left\lfloor \frac{n + 2h}{4} \right\rfloor$$

point guards, as established by O'Rourke (see [5, Theorem 5.1]). Because $n = 2r - 4h + 4$, where r is the number of *reflex* vertices in the orthogonal polygon (as can be proved by a straightforward induction on h), we may express the same upper bound in terms of r , as

$$\left\lfloor \frac{r - h}{2} \right\rfloor + 1.$$

Even though O'Rourke does not mention this aspect, a careful analysis of his method (a decomposition of the polygon into L-shaped pieces, see [5, Sections 2.5, 2.6]) reveals that, if $r > 0$, all the guards can be chosen to lie on *reflex* vertices.

More recently, the attention has shifted to 3-dimensional enclosures, and especially to polyhedra. Point guards are much less effective in this setting: there exist polyhedra with n vertices where guards placed at every vertex do not see the whole interior, and where $\Omega(n^{3/2})$ non-vertex guards are required (refer to [5, Section 10.2.2]). This motivates the study of more powerful guards: an *edge guard* is a guard that has the extent of an entire edge; a point p is visible to an edge guard e if and only if there is a point of e that sees p .

In [3], Cano et al. proved that any polyhedron with m edges can be guarded by at most $\lfloor 27m/32 \rfloor$ *edge guards*. For *orthogonal* polyhedra (i.e., polyhedra whose faces meet at right angles), Urrutia discovered that $\lfloor m/6 \rfloor$ edge guards are sufficient, and conjectured that $\lfloor m/12 \rfloor + O(1)$ are optimal (see [9]). Urrutia's upper bound was later improved in [1] by Benbernou et al.: an orthogonal polyhedron with m edges and genus g can be guarded by $\lfloor 11m/72 + g/6 \rfloor - 1$ edge guards; if the polyhedron has r reflex edges, then $\lfloor 7r/12 \rfloor - g + 1$ edge guards are sufficient.

Another contribution of [1] was to show that the same upper bounds hold if the edge guards are deprived of their two endpoints: i.e., they hold for *open edge guards*. This weaker type of guard has been introduced by the author in [10] and has been studied also in [2] and [8]. For a more in-depth discussion on this topic and some motivations, refer to [10, Section 3.2].

Face guards have also been explored, first by Souvaine et al. in [7], and then by the author in [11].

Our contributions

In this paper we study the Art Gallery problem for *2-reflex* orthogonal polyhedra, which are orthogonal polyhedra having reflex edges oriented in just two directions, as opposed to three. These polyhedra have been introduced by the author and investigated in the context of face guards in [11]. In Section 2 we motivate the study of this special class of polyhedra, and we establish some terminology.

In Section 3 we prove that

$$\left\lfloor \frac{r-g}{2} \right\rfloor + 1$$

edge guards are sufficient to guard any 2-reflex orthogonal polyhedron with r reflex edges and genus g . This number of edge guards is optimal if $g = 0$. Our strategy places guards on reflex edges only, i.e., we use *reflex edge guards*. Observe that this is a generalization of O'Rourke's result for orthogonal polygons mentioned above.

In Section 4 we give a similar bound in terms of the total number of edges m : we show that

$$\left\lfloor \frac{m-4}{8} \right\rfloor + g$$

reflex edge guards are sufficient. For $g = 0$, this reduces the best known bound by roughly 18%. All the above results hold for both closed and open edge guards.

Finally, in Section 5 we show that guard locations achieving the above bounds can be computed in time $O(n \log n)$, where n is the size of a representation of the polyhedron.

A preliminary version of this paper has appeared in the author's Ph.D. thesis, [10].

2 2-reflex orthogonal polyhedra

This paper focuses on 2-reflex orthogonal polyhedra, i.e., orthogonal polyhedra whose reflex edges lie in at most two different directions. This is a case of intermediate complexity between the 1-reflex case (i.e., orthogonal prisms) and the full 3-reflex case (i.e., general orthogonal polyhedra).

Recall that simple orthogonal polygons with r reflex vertices can be guarded by $\lfloor r/2 \rfloor + 1$ guards. This obviously extends to simply connected orthogonal prisms with r reflex edges. Our main research question is whether the same bound in terms of r extends to the whole class of orthogonal polyhedra. We are still unable to fully answer this question, although we have evidence that this may be the case.

However, we can prove (as we will do in this paper) that the $\lfloor r/2 \rfloor + 1$ upper bound holds at least for 2-reflex orthogonal polyhedra. We regard

this as a very important sub-case, and a necessary step toward a proof for general orthogonal polyhedra. Indeed, we believe that repeatedly “cutting away” 2-reflex orthogonal subpolyhedra from a given orthogonal polyhedron eventually yields a “kernel” with enough structural properties to make it efficiently guardable.

Regardless of this theoretical aspect, 2-reflex orthogonal polyhedra are of interest in themselves, as they can already express a rich variety of shapes.

Structure and terminology

For the purpose of this paper, an orthogonal polyhedron is defined as a connected 3-manifold (with boundary) in \mathbb{R}^3 that can be obtained as the union of finitely many *boxes*, i.e., cuboids whose edges are parallel to the coordinate axes. Without loss of generality, we stipulate that every 2-reflex orthogonal polyhedron encountered in this paper has only horizontal reflex edges, and no vertical ones.

The intersection between a 2-reflex orthogonal polyhedron and a horizontal plane not containing any of its vertices is a collection of pairwise disjoint rectangles. As the plane is moved upwards or downwards, the intersection does not change, until a vertex is reached. Thus, a natural way to partition a 2-reflex orthogonal polyhedron is into maximal cuboids, any two of which are either disjoint or touch each other at their top or bottom faces. Each cuboid in such a partition is called a *brick*, and any non-empty intersection between two bricks is called their *contact rectangle*. Figure 1 shows all the possible ways in which two bricks can touch each other.

By “cutting” a 2-reflex orthogonal polyhedron along the contact rectangle of two adjacent bricks, all the reflex edges bordering the rectangle turn into convex edges, and several different things may happen to the edge set. A convex edge may split in two distinct edges, new edges may be created, and multiple edges may merge together.

Figure 1 also indicates the number of edges gained or lost after such a split, for each different configuration. m and m' are the number of edges in the polyhedron before and after the split, respectively. Similarly, r and r' are the number of reflex edges before and after the split, respectively.

In Figure 2, a type-(t) contact rectangle between two bricks is illustrated, before and after the cut. The polyhedron in the second picture has $m = 23$ edges, of which $r = 2$ are reflex. The polyhedra in the third picture have $m' = 24$ edges in total, of which $r' = 0$ are reflex (cf. Figure 1(t)).

Remark 2.1. *When cutting a polyhedron of positive genus along a contact rectangle, we may fail to disconnect it, but just lower its genus. The resulting polyhedron is degenerate, in that its boundary is self-intersecting. We will occasionally encounter such degeneracies in intermediate steps of inductive proofs, and we will tolerate them as long as their presence does not invalidate our reasoning.*




Figure 1: Possible contact rectangles of two adjacent bricks, viewed from above. The light-shaded bricks lie on top of the darker ones; thick lines denote reflex edges, and the dashed ones are those covered by the top brick (i.e., the ones invisible from above).

Referring again to Figure 1, we call each type-(a) or type-(f) contact rectangle a *collar*, because its boundary is made of four reflex edges “winding” around a smaller brick. Singling out collars to treat them as separate cases will often be needed in our proofs. The (technical) reason is that collars minimize the ratio

$$\frac{m - m' + 12}{r - r'}.$$

This ratio is 3 for collars, whereas it is at least 4 for any other contact type.

We also want to single out contact types (d) and (i), because they produce just one reflex edge each, and this will turn out to be the hardest case to handle. These two contact types will be called *primitive*, and a 2-reflex or-




Figure 2: Type-(t) contact rectangle between two bricks, before and after a cut

thogonal polyhedron having only primitive contact rectangles will be called a *stack*.

Observe that each brick of a stack may have up to two bricks attached to each of its horizontal faces. A stack where each brick has either zero or two other bricks attached to its top face is called a *castle* (Figure 3 shows an example). The bottom brick of a castle is called its *base brick*. It is easy to see that a castle has an even number of reflex edges, because they all come in parallel pairs.




Figure 3: Castle with dashed lines marking contact rectangles between bricks

If a castle is turned upside down and its base brick is attached to another castle's base brick via a primitive contact rectangle, the resulting shape is a stack called *double castle* (see Figure 4). It follows that a double castle has an odd number of reflex edges.

Castles and double castles will play a fundamental role in the next sec-




Figure 4: Double castle

tion.

3 Bounds in terms of r

In this section we are going to generalize O'Rourke's theorem on orthogonal polygons to 2-reflex orthogonal polyhedra, showing that the exact same upper bound holds. It turns out that the central part of O'Rourke's main argument can be considerably simplified, and then rephrased and generalized in terms of polyhedra (Lemmas 2.13–2.15 of [5] are condensed in our Lemma 3.4), while the other parts of the proof require more sophisticated constructions and some novel ideas.

Eventually, everything boils down to guarding castles and double castles. So we will resolve these first, and then we will use them as building blocks to prove our full theorem.

We say that an orthogonal polyhedron is *monotone* if it is a prism and if its intersection with any vertical line is either empty or a single line segment. In other terms, a monotone orthogonal polyhedron can be constructed by extruding a monotone orthogonal polygon (refer to [5]).

Lemma 3.1. *Any monotone orthogonal polyhedron with $r > 0$ reflex edges is guardable by at most*

$$\left\lfloor \frac{r}{2} \right\rfloor + 1 \tag{1}$$

(open or closed) reflex edge guards.

Proof. Without loss of generality, suppose that the reflex edges are parallel to the y axis (hence, all the points of a reflex edge have the same x coordinate). Sort all reflex edges by increasing x coordinate, breaking ties arbitrarily, and let $(e_i)_{1 \leq i \leq r}$ be the sorted sequence. Now assign a guard to each edge whose index is an odd number, plus a guard to e_r if r is even. Thus, the bound (1) is matched.

To show that the polyhedron is guarded, let x_i be the x coordinate of e_i , for $1 \leq i \leq r$, and let x_0 and x_{r+1} be the minimum and the maximum x coordinate of a vertex of the polyhedron, respectively. Then, it is straightforward to see that a guard lying on e_i guards at least the points whose x coordinate lies between x_{i-1} and x_{i+1} , as Figure 5 suggests. \square




Figure 5: Guarding a monotone orthogonal polyhedron. Thick reflex edges represent guards.

Of course, among the monotone orthogonal polyhedra there are all 1-reflex castles. In the next lemma we show that we can actually use one less guard for the remaining castles.

Lemma 3.2. *Any castle with $2r$ reflex edges that is not a prism is guardable by at most r (open or closed) reflex edge guards.*

Proof. We prove our claim by well-founded induction on r . So, suppose the claim is true for all castles that are not prisms and have fewer than $2r$ reflex edges, and let a non-prism castle \mathcal{C} be given, having exactly $2r$ reflex edges. Note that \mathcal{C} cannot be a cuboid, and hence r is strictly positive and two castles \mathcal{C}_1 and \mathcal{C}_2 must lie on top of \mathcal{C} . Let e_1 (resp. e_2) be the reflex edge bordering the contact rectangle between the base brick of \mathcal{C} and the base brick of \mathcal{C}_1 (resp. \mathcal{C}_2). Let $2r_1$ and $2r_2$ be the numbers of reflex edges of \mathcal{C}_1 and \mathcal{C}_2 , respectively. It follows that

$$r = r_1 + r_2 + 1.$$

We have three cases.

- If neither \mathcal{C}_1 nor \mathcal{C}_2 is a prism, they both satisfy the inductive hypothesis and can be guarded by r_1 and r_2 reflex edge guards, respectively. Now we just assign a guard to e_1 in order to guard the base brick of \mathcal{C} , and our upper bound of r guards is matched.
- If \mathcal{C}_1 is a prism and \mathcal{C}_2 is not (the symmetric case is analogous), then the induction hypothesis applies to \mathcal{C}_2 , and we place r_2 reflex edge guards accordingly. Now, because \mathcal{C}_1 is a prism, its reflex edges are all parallel, and two sub-cases arise.
 - If the reflex edges of \mathcal{C}_1 are parallel to e_1 (or \mathcal{C}_1 has no reflex edges), then \mathcal{C}_1 and the base brick of \mathcal{C} together form a monotone orthogonal polyhedron with $2r_1 + 1$ reflex edges in total. By Lemma 3.1, such a polyhedron can be guarded by $r_1 + 1$ guards. Along with the previously assigned r_2 guards, this yields r guards, as desired.
 - If the reflex edges of \mathcal{C}_1 are not parallel to e_1 , then they are orthogonal to e_1 . As a consequence, all of \mathcal{C}_1 is visible to e_1 , and so is the base brick of \mathcal{C} , as Figure 6 illustrates. It follows that $r_2 + 1$ guards are sufficient in this case, which is less than r .




Figure 6: Edge e_1 guards all of \mathcal{C}_1 plus the base brick.

- If both \mathcal{C}_1 and \mathcal{C}_2 are prisms, at least one of them (say, \mathcal{C}_1) must have reflex edges that are orthogonal to e_1 , otherwise \mathcal{C} would be a prism, too. Therefore $r_1 \geq 1$, and \mathcal{C}_1 is guarded by assigning a guard to e_1 . Again, two sub-cases arise.
 - If the reflex edges of \mathcal{C}_2 are parallel to e_2 , then \mathcal{C}_2 and the base brick of \mathcal{C} form a monotone orthogonal polyhedron with $2r_2 + 1$

reflex edges, which is guardable by $r_2 + 1$ reflex edge guards, by Lemma 3.1. Overall, we assigned $1 + r_2 + 1 \leq r_1 + r_2 + 1 = r$ guards, as required.

- If the reflex edges of \mathcal{C}_2 are orthogonal to e_2 , then e_2 guards \mathcal{C}_2 , along with the base brick of \mathcal{C} . We assigned only two guards, and $2 \leq r_1 + 1 \leq r$, concluding the proof.

□

The two previous lemmas enable us to prove that our upper bound holds at least for double castles.

Lemma 3.3. *Any double castle with r reflex edges is guardable by at most*

$$\left\lfloor \frac{r}{2} \right\rfloor + 1 \tag{2}$$

(open or closed) reflex edge guards.

Proof. Let \mathcal{D} be a double castle with r reflex edges, made of two castles \mathcal{C}_1 and \mathcal{C}_2 having $2r_1$ and $2r_2$ reflex edges, respectively. Let e be the reflex edge lying on the contact rectangle between the two castles. Because $r = 2r_1 + 2r_2 + 1$, (2) can be rewritten as

$$r_1 + r_2 + 1. \tag{3}$$

We distinguish three cases.

- If neither \mathcal{C}_1 nor \mathcal{C}_2 is a prism, by Lemma 3.2 they can be guarded by at most r_1 and r_2 reflex edge guards, respectively. Hence, (3) is satisfied.
- If \mathcal{C}_1 is a prism and \mathcal{C}_2 is not (the symmetric case is analogous), by Lemma 3.2 \mathcal{C}_2 can be guarded by at most r_2 reflex edge guards. We have two sub-cases.
 - If $r_1 = 0$, \mathcal{C}_1 is a cuboid and can be guarded by e . In total, $r_2 + 1$ guards are assigned, which matches (3).
 - If $r_1 > 0$, \mathcal{C}_1 can be guarded by $r_1 + 1$ guards, by Lemma 3.1. Together with the previous r_2 guards, these match the upper bound (3).
- If both \mathcal{C}_1 and \mathcal{C}_2 are prisms, we have three sub-cases.
 - If the reflex edges of \mathcal{C}_1 and \mathcal{C}_2 are parallel to e , then \mathcal{D} is a monotone orthogonal polyhedron with $2r_1 + 2r_2 + 1$ reflex edges, and according to Lemma 3.1 it is guardable by $r_1 + r_2 + 1$ reflex edge guards, which agrees with (3). This holds also if $r_1 = 0$ or $r_2 = 0$.

- If the reflex edges of \mathcal{C}_1 are parallel to e and the reflex edges of \mathcal{C}_2 are orthogonal to e (the symmetric case is analogous), then we may assume that $r_2 \geq 1$. We assign one guard to e in order to guard \mathcal{C}_2 , and $r_1 + 1$ guards to reflex edges in \mathcal{C}_1 , in accordance with Lemma 5. In total, we assigned $r_1 + 1 + 1 \leq r_1 + r_2 + 1$ guards, thus satisfying (3).
- Finally, if the reflex edges of both \mathcal{C}_1 and \mathcal{C}_2 are orthogonal to e , then a guard assigned to e sees all of \mathcal{D} , which obviously satisfies (3).

□

Now that we know how to guard double castles, we can move on to more general shapes, such as stacks. This is the last step before our main theorem, and it generalizes [5, Lemmas 2.13–2.15].

Recall that a stack has only primitive contact rectangles. If two neighboring bricks \mathcal{B} and \mathcal{B}' share a type-(d) (resp. type-(i)) contact rectangle (refer to Figure 1), and \mathcal{B} lies below (resp. above) \mathcal{B}' , then we say that \mathcal{B} is a *parent* of \mathcal{B}' , and \mathcal{B}' is a *child* of \mathcal{B} . It follows that a brick in a stack can have at most two children above and two children below, and shares exactly one reflex edge with each child (see Figure 7). Moreover, if a brick has one parent above (resp. below), then it has no other neighboring bricks above (resp. below). Finally, if a brick \mathcal{B} has exactly one child \mathcal{B}' on one side (regardless of the number of neighboring bricks on the other side), then the reflex edge shared by \mathcal{B} and \mathcal{B}' is said to be *isolated*. It follows that the number of reflex edges in a stack and the number of isolated reflex edges have the same parity.

Lemma 3.4. *Any stack with $r > 0$ reflex edges and genus g is guardable by at most*

$$\left\lfloor \frac{r - g}{2} \right\rfloor + 1 \tag{4}$$

(open or closed) reflex edge guards.

Proof. Our proof is by well-founded induction on r . Following O'Rourke (see [5, Section 2.5]), we say that a contact rectangle R in a simply connected stack \mathcal{S} yields an *odd cut* if R partitions \mathcal{S} into two stacks \mathcal{S}_1 and \mathcal{S}_2 , one of which has an odd number of reflex edges. The presence of odd cuts in a simply connected stack is very desirable, in that it permits to successfully apply the inductive hypothesis. Indeed, let $2r_1 + 1$ and r_2 be the number of reflex edges of \mathcal{S}_1 and \mathcal{S}_2 (the symmetric case is analogous). Then $r = 2r_1 + r_2 + 2$, because the cut resolves exactly one reflex edge. Recall that \mathcal{S} is simply connected, and so are \mathcal{S}_1 and \mathcal{S}_2 . Since \mathcal{S}_1 is non-convex, we can apply the inductive hypothesis on it, guarding it with at most

$$\left\lfloor \frac{2r_1 + 1}{2} \right\rfloor + 1 = r_1 + 1$$




Figure 7: Different configurations for the upper children of a brick in a stack

reflex edge guards. Similarly, if \mathcal{S}_2 is non-convex, we can guard it with at most

$$\left\lfloor \frac{r_2}{2} \right\rfloor + 1 \quad (5)$$

reflex edge guards. On the other hand, if \mathcal{S}_2 is a cuboid, we can guard it by assigning a guard to the reflex edge lying on R . In this case, $r_2 = 0$, and hence (5) still holds. As a result, we have guarded all of \mathcal{S} with at most

$$r_1 + 1 + \left\lfloor \frac{r_2}{2} \right\rfloor + 1 = \left\lfloor \frac{2r_1 + r_2 + 2}{2} \right\rfloor + 1 = \left\lfloor \frac{r}{2} \right\rfloor + 1$$

reflex edge guards. We stress that it makes sense to talk about odd cuts only for stacks of genus 0.

Now, let \mathcal{S} be a stack with exactly $r > 0$ reflex edges and genus g , and assume that the lemma's statement holds for all non-convex stacks with fewer than r reflex edges. There are four cases to consider.

- Let $g > 0$. Note that each cut along a contact rectangle either disconnects \mathcal{S} or lowers its genus by 1, and each cut resolves exactly one contact rectangle without creating or modifying other contact rectangles. Because \mathcal{S} is partitioned by contact rectangles into cuboids, whose genus is zero, it follows that there exists at least one contact rectangle R such that cutting through R yields a (degenerate, see Remark 2.1) stack \mathcal{S}' with $r' = r - 1$ reflex edges and genus $g' = g - 1$. Observe that \mathcal{S}' is non-convex, because it is made of at least two bricks,

and so we can apply the inductive hypothesis on it and guard it with at most

$$\left\lfloor \frac{r' - g'}{2} \right\rfloor + 1 = \left\lfloor \frac{r - g}{2} \right\rfloor + 1$$

reflex edge guards.

- If $g = 0$ and $r > 0$ is even, then any contact rectangle yields an odd cut. Indeed, cutting \mathcal{S} along a contact rectangle resolves exactly one reflex edge and produces two stacks (because \mathcal{S} is simply connected). The amounts of reflex edges in these two stacks must have opposite parity, because their sum must be odd. Hence, one of them is odd.
- Let $g = 0$, and let \mathcal{S} have a brick \mathcal{B} with exactly one neighboring brick above and exactly one below (Figure 8 sketches one possible configuration for \mathcal{B}). We show that one of the two contact rectangles bordering \mathcal{B} yields an odd cut. Let R_1 be the upper contact rectangle and R_2 be the lower one. If R_1 does not yield an odd cut, the stack above R_1 has an even number of reflex edges. But then, the stack above R_2 , which additionally includes \mathcal{B} and the reflex edge belonging to R_1 , has an odd number of reflex edges. It follows that either R_1 or R_2 yields an odd cut.

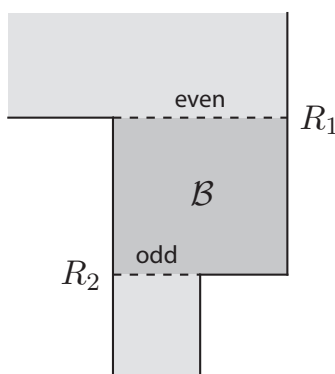


Figure 8: Sketch of a brick with one neighbor above and one below

- If none of the above are satisfied, we show that \mathcal{S} must be a double castle, and hence our claim holds by Lemma 3.3. Let then r be odd, so that \mathcal{S} has at least one isolated reflex edge corresponding to a contact rectangle R . Additionally, let no brick of \mathcal{S} have exactly one neighbor above and one neighbor below. We show that \mathcal{S}' , i.e., the stack above R , is a castle. (By a symmetric argument, the stack below R will be an upside-down castle, and \mathcal{S} will then be a double castle.) Although this is easy to see (build the castle from bottom to top, without adding bricks that have only one neighbor below and one above), here follows a formal proof.

Let $d(\mathcal{B})$ be the minimum number of bricks one has to traverse to reach brick \mathcal{B} from R (while always staying inside \mathcal{S}'), and let \mathcal{S}'_h be the set of bricks \mathcal{B} in \mathcal{S}' such that $d(\mathcal{B}) \leq h$. We prove by induction on h that \mathcal{S}'_h is a castle whose base brick contains R .

The claim is true for $h = 0$, because the brick just above R (call it $\tilde{\mathcal{B}}$) is the only one in contact with R , and a brick is indeed a castle.

Observe that no brick in \mathcal{S}' can be attached to the bottom face of $\tilde{\mathcal{B}}$, because the reflex edge of \mathcal{S} corresponding to R is isolated. Hence, $\tilde{\mathcal{B}}$ must have either zero or two neighbors in \mathcal{S}' , and both of them are above. It follows that \mathcal{S}'_1 is a castle, as well.

Let now $h \geq 1$, assume that \mathcal{S}'_h is a castle whose base brick $\tilde{\mathcal{B}}$ contains R , and let us show that the same holds for \mathcal{S}'_{h+1} . Any brick \mathcal{B} such that $d(\mathcal{B}) = h + 1$ must be attached to some brick \mathcal{B}' of \mathcal{S}'_h such that $d(\mathcal{B}') = h$ (see Figure 9). It is straightforward to see that any such \mathcal{B}' has one parent brick below and no bricks above, in \mathcal{S}'_h . Hence, \mathcal{B} can only be attached on top of \mathcal{B}' . Because \mathcal{B}' cannot have only one top and one bottom neighbor in \mathcal{S}' , it follows that \mathcal{B} cannot be a parent of \mathcal{B}' , but a child. Any other brick \mathcal{B}'' that is attached on top of \mathcal{B}' in \mathcal{S}' must also belong to \mathcal{S}'_{h+1} , by definition. As a consequence, one such brick \mathcal{B}'' must indeed be in \mathcal{S}'_{h+1} , otherwise \mathcal{B}' would once again have only one top and one bottom neighbor in \mathcal{S}' . No new brick is attached to $\tilde{\mathcal{B}}$, which remains the base of \mathcal{S}'_{h+1} . Thus the base brick of \mathcal{S}'_{h+1} contains R , proving our claim and concluding the inductive proof.




Figure 9: Sketch of a brick in a castle

Since \mathcal{S}' is connected and contains only a finite number of bricks, for a large-enough h we have $\mathcal{S}' = \mathcal{S}'_h$, implying that \mathcal{S}' is a castle.

□

In the next theorem, we give an upper bound on the number of guards for a 2-reflex orthogonal polyhedron in terms of r , g and b , where b is the number

of collars. The presence of b may look redundant (as it just contributes to lowering the bound), but we will actually have to carry this parameter along to the next section, in order to prove Theorem 4.2. The casual reader may ignore this term and read (6) simply as

$$\left\lfloor \frac{r-g}{2} \right\rfloor + 1.$$

Theorem 3.5. *Any 2-reflex orthogonal polyhedron with $r > 0$ reflex edges, b collars and genus g is guardable by at most*

$$\left\lfloor \frac{r-g}{2} \right\rfloor - b + 1 \tag{6}$$

(open or closed) reflex edge guards.

Proof. We proceed by induction on the number of non-primitive contact rectangles. The base case is given by non-convex stacks, for which (6) holds due to Lemma 3.4 and the fact that $b = 0$.

For the inductive step, let \mathcal{P} be a 2-reflex orthogonal polyhedron with $r > 0$ reflex edges, b collars, genus g , and a non-primitive contact rectangle R . We cut \mathcal{P} through R , thus resolving one non-primitive contact rectangle, and we distinguish two cases.

- If the cut does not disconnect \mathcal{P} , then it lowers its genus by 1. Let \mathcal{P}' be the resulting polyhedron (which is degenerate, see Remark 2.1), and let $g' = g - 1$ be its genus. By inductive hypothesis, \mathcal{P}' is guardable by

$$\left\lfloor \frac{r' - g'}{2} \right\rfloor - b' + 1 \tag{7}$$

guards, where r' and b' are, respectively, the number of reflex edges and collars of \mathcal{P}' . Two sub-cases arise.

- If R is a collar, then $r' = r - 4$ and $b' = b - 1$. Plugging these into (7), we obtain that $\mathcal{P} \setminus R$ is guardable by

$$\left\lfloor \frac{r-4-g+1}{2} \right\rfloor - b + 1 + 1 \leq \left\lfloor \frac{r-g}{2} \right\rfloor - b + 1$$

reflex edge guards.

- If R is not a collar, then $b' = b$. Because R is not primitive, $r' \leq r - 2$ (refer to Figure 1). Hence, \mathcal{P} is guardable by at most

$$\left\lfloor \frac{r-2-g+1}{2} \right\rfloor - b + 1 \leq \left\lfloor \frac{r-g}{2} \right\rfloor - b + 1$$

reflex edge guards.

- If the cut disconnects \mathcal{P} into \mathcal{P}_1 and \mathcal{P}_2 , then $g_1 + g_2 = g$, where g_1 (resp. g_2) is the genus of \mathcal{P}_1 (resp. \mathcal{P}_2). Let r_1 and b_1 (resp. r_2 and b_2) be, respectively, the number of reflex edges and collars of \mathcal{P}_1 (resp. \mathcal{P}_2). Two sub-cases arise.

- If R is a collar, then $r_1 + r_2 = r - 4$ and $b_1 + b_2 = b - 1$. If $r_1 > 0$, then we can apply the inductive hypothesis to \mathcal{P}_1 and guard it with at most

$$\left\lfloor \frac{r_1 - g_1}{2} \right\rfloor - b_1 + 1 \quad (8)$$

guards. Otherwise, \mathcal{P}_1 is a cuboid, and we can guard it by assigning a guard to any edge of R (they are all reflex). In this case, $r_1 = g_1 = b_1 = 0$, and (8) is still satisfied. We proceed similarly with \mathcal{P}_2 , and thus we have assigned a combined number of guards that is at most

$$\left\lfloor \frac{r_1 - g_1}{2} \right\rfloor + \left\lfloor \frac{r_2 - g_2}{2} \right\rfloor - b_1 - b_2 + 2 \leq \left\lfloor \frac{r - 4 - g}{2} \right\rfloor - b + 3 = \left\lfloor \frac{r - g}{2} \right\rfloor - b + 1.$$

This many guards are then sufficient to guard \mathcal{P} .

- If R is not a collar, then $b_1 + b_2 = b$. Because R is not primitive, $r_1 + r_2 \leq r - 2$ (refer to Figure 1). Once again, if $r_1 > 0$, then we can apply the inductive hypothesis to \mathcal{P}_1 and guard it with a number of guards that is bounded by (8). Otherwise, \mathcal{P}_1 is a cuboid, and we can guard it by assigning a guard to any reflex edge of R (there is at least one). In this case, $r_1 = g_1 = b_1 = 0$, and (8) is still satisfied. Again, we do the same with \mathcal{P}_2 , and thus we have assigned a combined number of guards that is at most

$$\left\lfloor \frac{r_1 - g_1}{2} \right\rfloor + \left\lfloor \frac{r_2 - g_2}{2} \right\rfloor - b_1 - b_2 + 2 \leq \left\lfloor \frac{r - 2 - g}{2} \right\rfloor - b + 2 = \left\lfloor \frac{r - g}{2} \right\rfloor - b + 1,$$

and \mathcal{P} is guarded.

We remark that also the rectangle R is guarded in every case, because it is contained in \mathcal{P}' , \mathcal{P}_1 , and \mathcal{P}_2 , as defined above. \square

For $g = 0$, the upper bound given in Theorem 3.5 is tight, as Figure 10 implies. By contrast, for arbitrary genus, the situation is still unclear, and even the corresponding 2-dimensional problem is open (i.e., optimally guarding orthogonal polygons with holes, see [5, 9]).

We believe that the bound given in Theorem 3.5 holds also for general (i.e., 3-reflex) orthogonal polyhedra.




Figure 10: 2-reflex orthogonal polyhedron requiring $\lfloor r/2 \rfloor + 1$ reflex edge guards.

Conjecture 3.6. *Any orthogonal polyhedron with $r > 0$ reflex edges and genus g is guardable by at most*

$$\left\lfloor \frac{r - g}{2} \right\rfloor + 1$$

(open or closed) reflex edge guards.

4 Bounds in terms of m

In this section we bound the number of reflex edge guards required to guard a 2-reflex orthogonal polyhedron in terms of m , as opposed to r . Rather than providing a radically new construction, we bound m with respect to r , and then we apply Theorem 3.5.

Note that a naive application of this method would not improve on the state of the art. Indeed, the sharpest possible inequality between m and r (involving also the genus g) is

$$m \geq 3r - 12g + 12, \tag{9}$$

which yields an upper bound of

$$\frac{m}{6} + O(g)$$

guards when applied to Theorem 3.5. This result was already obtained by Urrutia in [9] and improved in [1].

To get around this, we will refine (9) by introducing the number of collars as an additional parameter.

Lemma 4.1. *In every 2-reflex orthogonal polyhedron, the number of edges m , the number of reflex edges r , the number of collars b , and the genus g satisfy the inequality*

$$m \geq 4r - 12g - 4b + 12.$$

Proof. We will prove that, for any collection of k (pairwise internally disjoint) 2-reflex orthogonal polyhedra,

$$m \geq 4r - 12g - 4b + 12k \quad (10)$$

holds. Here, m (resp. r , g , b) is the sum of the edges (resp. reflex edges, genera, collars) of the k polyhedra. Then, by plugging $k = 1$, we will obtain our claim.

Our proof proceeds by induction on r . For $r = 0$ we have a collection of k cuboids, each of which has 12 edges, so $m = 12k$, $g = b = 0$, and therefore (10) holds as desired.

If $r > 0$, there is at least one (horizontal) reflex edge, which is a side of the contact rectangle R of two adjacent bricks, both belonging to the same polyhedron \mathcal{P} of the collection. We can resolve this reflex edge (and up to three others) by separating the two bricks with a horizontal cut through R . As a consequence, either \mathcal{P} gets partitioned in two polyhedra (in which case the new number of polyhedra is $k' = k + 1$), or the genus of \mathcal{P} decreases by 1 (in which case the new total genus is $g' = g - 1$). Either way,

$$k' - g' = k - g + 1. \quad (11)$$

By inductive hypothesis,

$$m' \geq 4r' - 12g' - 4b' + 12k', \quad (12)$$

where m' (resp. r' , b') is the new number of edges (resp. reflex edges, collars), after the cut.

We have two cases. If R is a collar, then $b' = b - 1$, $m' = m$, and $r' = r - 4$ (see Figure 1). Plugging these into (12) and combining the result with (11) immediately yields (10), as claimed.

Otherwise (i.e., R is not a collar), by inspection of Figure 1, it is clear that

$$m - m' \geq 4(r - r') - 12 \quad (13)$$

(recall that types (a) and (f) must be ignored, because they correspond to a collar). By combining (11), (12), (13) and plugging $b' = b$, we obtain again (10), concluding the proof. \square

Theorem 4.2. *Any non-convex 2-reflex orthogonal polyhedron with m edges and genus g is guardable by at most*

$$\left\lfloor \frac{m - 4}{8} \right\rfloor + g \quad (14)$$

(open or closed) reflex edge guards.

Proof. Let $r > 0$ be the number of reflex edges in the polyhedron. By Lemma 4.1,

$$r \leq \frac{m}{4} + 3g + b - 3, \quad (15)$$

where b is the number of collars. Applying (15) to Theorem 3.5, we obtain that the number of guards is bounded by

$$\left\lfloor \frac{r-g}{2} \right\rfloor - b + 1 \leq \left\lfloor \frac{m}{8} + \frac{3}{2}g + \frac{b}{2} - \frac{3}{2} - \frac{g}{2} \right\rfloor - b + 1 = \left\lfloor \frac{m-4}{8} - \frac{b}{2} \right\rfloor + g \leq \left\lfloor \frac{m-4}{8} \right\rfloor + g.$$

□

We remark that our Theorem 4.2 is an improvement upon the previous state of the art, in that it lowers the upper bound provided by Urrutia in [9] by roughly 25% and the one in [1] by roughly 18% (for $g = 0$). Also, it shows how guards can be chosen to lie on reflex edges, as opposed to arbitrary edges.

We are unable to raise the lower bound of $\lfloor m/12 \rfloor + 1$ guards given in [9], even if just reflex edge guards are to be employed. On the other hand, our upper bound in terms of m seems far from optimal: indeed, we could treat separately other types of contact rectangles between bricks, rather than just the collars, and this could be enough to improve Lemma 4.1 and lower the bound in (14). Therefore, we formulate a stronger version of Urrutia's Conjecture.

Conjecture 4.3. *Any non-convex orthogonal polyhedron with m edges is guardable by at most*

$$\frac{m}{12} + O(1)$$

(open or closed) reflex edge guards.

We already know that this statement holds for simply connected prisms, and Lemma 3.4 implies that it holds more generally for simply connected stacks. Specifically, we have the following.

Theorem 4.4. *Any non-convex stack with m edges and genus g is guardable by at most*

$$\left\lfloor \frac{m}{12} + \frac{g}{2} \right\rfloor \quad (16)$$

(open or closed) reflex edge guards.

Proof. A straightforward induction on the number of reflex edges r , based on Figure 1, reveals that

$$m = 6r - 12g + 12$$

(recall that stacks only have type-(d) and type-(i) contact rectangles). Solving for r and substituting in (4) immediately yields (16). □

5 Time complexity

To conclude the chapter, we show how to efficiently compute guard positions matching the upper bounds given in Theorems 3.5 and 4.2. Notice that both bounds refer to the very same construction, described in Section 3. In the present section, we will merely translate such a construction into an algorithm that runs in $O(n \log n)$ time.

Of course, to reason about running time, we have to specify how polyhedra are represented as data structures. We will assume that a polyhedron is stored as the array of its vertices, together with the array of its faces. Each face is a sequence of indices into the vertex array. The outer boundary of each face will be given in counterclockwise order with respect to a normal vector pointing outward, while its holes will be given in clockwise order.

We assume that numbers expressing vertex coordinates, as well as indices into the vertex array, have constant size. Note that, in an orthogonal polyhedron, each vertex is shared by at most six faces. So, if an orthogonal polyhedron has n vertices, the size of its representation is $\Theta(n)$. Since the running time of our algorithm is given in big-oh notation, we do not have to make any distinction between the number of vertices and the size of the representation of a given orthogonal polyhedron.

O'Rourke's algorithm for *simple* orthogonal polygons, detailed in [5, Section 2.6], also runs in $O(n \log n)$ time. As observed by Urrutia in [9], it could be optimized to run in $O(n)$ time, if Chazelle's linear-time triangulation algorithm were used (see [4]). Unfortunately, Urrutia's speedup is only applicable to orthogonal polygons without holes.

In principle, we could rephrase O'Rourke's algorithm in the language of 2-reflex orthogonal polyhedra and obtain a new $O(n \log n)$ algorithm. However, four issues arise that require additional care:

1. O'Rourke's algorithm works on simply connected polygons, while our algorithm should be applied to polyhedra of any genus.
2. O'Rourke's algorithm may assign guards to convex vertices, whereas we insist on having guards only on reflex edges.
3. O'Rourke's method to find horizontal cuts in polygons does not trivially extend to polyhedra.
4. O'Rourke's algorithm relies on guarding double histograms, whereas we need guard double castles.

We will now describe our modified algorithm, showing that each of the above issues has a relatively simple solution. Our algorithm takes as input a 2-reflex orthogonal polyhedron \mathcal{P} with n vertices, and outputs the set of reflex edges to which guards are assigned.

Preprocessing: $O(n)$

First of all, we do some preprocessing on \mathcal{P} to construct *adjacency tables* of faces, edges and vertices, which allow us to efficiently navigate the polyhedron’s surface. We mark each edge as reflex or convex and, if needed, we turn \mathcal{P} by 90° so that it contains no vertical reflex edges.

Finding contact lines: $O(n \log n)$

We first compute a structure that is very similar to the *horizontal visibility map* (also known as *trapezoidization*, see [5]) of each vertical face of \mathcal{P} . This is a well-studied 2-dimensional problem that consists in partitioning a polygon into trapezoids by drawing horizontal lines at vertices. In our case, faces are orthogonal polygons (perhaps with holes), trapezoids are actually rectangles, and cut lines are drawn at reflex vertices only. For reasons that will be explained later, these cut lines are called *contact lines*.

Note that some vertical faces of \mathcal{P} may have “degenerate” vertices lying on horizontal edges. For instance, in the middle picture of Figure 2 there is a polyhedron made of two bricks, the top of which has a right-facing rectangular face F . However, five vertices of the polyhedron lie on the perimeter of F : the one lying in the middle of the bottom edge of F is a degenerate vertex for F . In our algorithm, a degenerate vertex is treated as a pair of coincident vertices, one reflex and one convex: it is as if a degenerate vertex were actually a (degenerate) vertical edge of length 0.

Let F_i be a vertical face of \mathcal{P} with n_i vertices. We sort all the vertical edges of F_i by the z coordinate of their upper vertex (in $O(n_i \log n_i)$ time), and we “scan” F_i from top to bottom with a sweep line. We maintain a horizontally sorted list of all the vertical edges of F_i pierced by the sweep line, in which insertion and deletion take $O(\log n_i)$ time. This list also includes the degenerate edges defined above. Every time our sweep line hits a vertex v that is reflex in F_i or belongs to a reflex edge of \mathcal{P} , we draw a contact line from v to the next edge or the previous edge in the sorted list (according as the vertical edge containing v is facing right or left), we add a *dummy vertex* there (unless there is already a vertex), and we proceed with our sweep.

This process takes $O(n_i \log n_i)$ time and, letting $\tilde{n} = \sum_i n_i$, finding the contact lines on every vertical face of \mathcal{P} takes $O(n \log n)$ in total, because

$$\sum_i n_i \log n_i \leq \sum_i n_i \log \tilde{n} = \tilde{n} \log \tilde{n} = \Theta(n \log n).$$

Every time we find a new contact line, we also update the face, edge, and vertex data we computed in the preprocessing step. That is, as soon as a new contact line is found, the corresponding face gets a new edge and is perhaps partitioned in two coplanar faces (this step takes constant time).

Moreover, whenever we create a new dummy vertex w for F_i , we add it to the other face sharing it (say, F_j). If we still have to process F_j , we treat w as a reflex vertex and draw a contact line at w in F_j when we process it. Otherwise, F_j is now a rectangle, and we just draw an additional contact line in it at w , in constant time: we call this procedure *extension* of the contact line. However, if the extension generates yet another dummy vertex, there is no need to perform an “extension of the extension” to the next face, as we will see in a moment.

Finding bricks and contact rectangles: $O(n)$

Notice that the contact lines constructed in the previous step are precisely the edges of the contact rectangles of \mathcal{P} . Indeed, we did not draw contact lines only at the reflex vertices of each face, but also at the dummy vertices created while processing other faces and at convex vertices that lie on reflex edges of \mathcal{P} .

Observe that a dummy vertex may be created only on a contact rectangle of type (c), (d), (h), or (i) (refer to Figure 1). Indeed, in all other cases a vertex of a contact rectangle is also a vertex of the polyhedron. Moreover, only the contact rectangles of type (d) and type (i) have an edge with no endpoint on a vertex of \mathcal{P} : these are the only cases where the extension procedure is really needed, and performing it once per contact rectangle is enough to generate all its edges.

It is easy now to identify all the contact rectangles and all the bricks, navigating the boundary of \mathcal{P} using the adjacency tables that we precomputed and then updated every time a cut line was drawn. While we identify the contact rectangles, we also build a *brick graph* G , having a node for each brick and an edge connecting each pair of bricks sharing a contact rectangle.

Observe that issue (3) above is now solved.

Resolving non-primitive contact rectangles: $O(n)$

The non-primitive contact rectangles are those that are surrounded by more than one reflex edge of \mathcal{P} and, as such, are easy to find. As proven in Theorem 3.5, it is safe to cut \mathcal{P} at a non-primitive contact rectangle and place guards in the resulting polyhedra (or polyhedron of lower genus).

Instead of actually cutting \mathcal{P} and updating all the data structures, we merely delete the edges of G corresponding to all non-primitive contact rectangles.

Ensuring simple connectedness: $O(n)$

By this point, \mathcal{P} has been partitioned into several, possibly not simply connected, stacks. As proven in Theorem 3.4, it is safe to further cut the stacks until they all become simply connected. To do so, we again process only

G , turning it into a forest. Such a task is accomplished by a depth-first traversal, starting at each connected component and deleting edges leading to already visited nodes. Recall that bricks in stacks have at most four neighbors, hence the time complexity of this traversal is indeed linear.

Observe that this step also solves issue (1) above.

Adjusting brick parity: $O(n)$

For reasons that will be clear shortly, we insist on having only stacks with an even number of bricks. O'Rourke achieves this by adding an extra “chip” to the polygon, in case he wants to change the parity of its reflex vertices. Then he applies his algorithm to the new polygon, and later removes the chip. If the chip happens to host a guard, then that guard is reassigned to the nearest convex vertex, after the chip is removed. Observe how this choice causes issue (2) above.

In order to avoid placing guards on convex edges, we proceed as follows. We compute the size of each connected component of G by a simple traversal. Then, in each component with an odd number of nodes, we find one leaf (recall that G is a forest) and delete the edge attached to it (if the component is already an isolated node, we leave it as it is). Finally, we collect each isolated node, remove it from G , find its corresponding brick \mathcal{B} in \mathcal{P} , find a contact rectangle bordering \mathcal{B} (one must exist), find one reflex edge e on it and assign it a guard. Referring to Figure 7, it is obvious that e guards all of \mathcal{B} .

The correctness of this step follows from the remarks contained in the proof of Lemma 3.4, that every contact rectangle in a stack with an odd number of bricks yields an odd cut, and that it is always safe to make odd cuts.

Identifying odd cuts: $O(n)$

We are left with stacks having an even number of bricks, and we want to further partition them with odd cuts. In order to identify odd cuts, we pick each connected component of G and we do a depth-first traversal, rooted anywhere. During the traversal, we compute the parity of the size of the subtree dangling from each edge of G we traverse. The parity is even if and only if that edge of G corresponds to an odd cut.

Now, it is easy to observe that cutting a stack having an even number of bricks at an odd cut yields two stacks that again have an even number of bricks. Additionally, this operation does not change the parity of the cuts in the two resulting stacks. Hence, there is no need to re-identify the odd cuts after a cut is made. In contrast, stacks with an odd number of bricks do not have such a property, and this motivates our previous step.

It follows that we may remove all the edges of G corresponding to odd cuts, without worrying about side effects.

Guarding double castles: $O(n)$

At this point, only non-convex stacks without odd cuts are left. As a consequence of the observations in Lemma 3.4, these are all double castles, which we now have to guard in linear time. In contrast, O’Rourke’s algorithm was left at this point with double histograms (cf. issue (4) above).

Our algorithm is based on the proofs of Lemmas 3.1, 3.2, and 3.3, which naturally yield a procedure that cuts along certain contact rectangles and selects guards in monotone orthogonal polyhedra.

The only non-trivial aspect is that, occasionally in the procedure, we need know if some castle (or upside-down castle) is a prism, and what the orientation of its reflex edges is. To efficiently answer this question, we precompute this information for every “sub-castle” of each double castle that we have. We identify the two castles constituting each double castle (in linear time), then we do a depth-first traversal of the subgraphs of G corresponding to those castles, starting from their base bricks. When we reach an internal node v , we recursively check if the subtrees dangling from its two children correspond to prisms, and if their reflex edges are oriented in the same direction. Then, after inspecting also the brick corresponding to v , we know if the subtree dangling from it corresponds to a prism and, if so, the direction of its reflex edges. Leaves are trivial to handle, in that they always correspond to prisms with no reflex edges.

Summarizing, and recalling the upper bounds given in Theorems 3.5 and 4.2, we have the following.

Theorem 5.1. *Given a 2-reflex orthogonal polyhedron of genus g with m edges, of which $r > 0$ are reflex, a guarding set of at most*

$$\min \left\{ \left\lfloor \frac{r-g}{2} \right\rfloor + 1, \left\lfloor \frac{m-4}{8} \right\rfloor + g \right\}$$

(open or closed) reflex edge guards can be computed in $O(n \log n)$ time. \square

Observe that the only superlinear step is the vertical sweep that finds the contact lines in every vertical face of the polyhedron. Whether a more efficient algorithm exists remains an open problem.

References

- [1] N.M. Benbernou, E.D. Demaine, M.L. Demaine, A. Kurdia, J. O’Rourke, G.T. Toussaint, J. Urrutia, and G. Viglietta. Edge-

- guarding orthogonal polyhedra. In *Proceedings of the 23rd Canadian Conference on Computational Geometry (CCCG)*, pp. 461–466, 2011.
- [2] S. Cannon, D.L. Souvaine, and A. Winslow. Hidden mobile guards in simple polygons. In *Proceedings of the 24th Canadian Conference on Computational Geometry (CCCG)*, pp. 161–166, 2012.
 - [3] J. Cano, C. D. Tóth, and J. Urrutia. Edge guards for polyhedra in three-space. In *Proceedings of 24th Canadian Conference on Computational Geometry (CCCG)*, pp. 163–167, 2012.
 - [4] B. Chazelle. Triangulating a simple polygon in linear time. *Discrete and Computational Geometry*, vol. 6, pp. 485–524, 1991.
 - [5] J. O’Rourke. *Art Gallery Theorems and Algorithms*. Oxford University Press, New York, 1987.
 - [6] T. Shermer. Recent results in art galleries. In *Proceedings of the IEEE*, vol. 80, pp. 1384–1399, 1992.
 - [7] D. L. Souvaine, R. Veroy, and A. Winslow. Face guards for art galleries. In *Proceedings of the 14th Spanish Meeting on Computational Geometry*, pp. 39–42, 2011.
 - [8] C. D. Tóth, G. T. Toussaint, and A. Winslow. Open guard edges and edge guards in simple polygons. In *Proceedings of the 23rd Canadian Conference on Computational Geometry (CCCG)*, pp. 449–454, 2011.
 - [9] J. Urrutia. Art gallery and illumination problems. In J.-R. Sack and J. Urrutia, editors, *Handbook of Computational Geometry*, pp. 973–1027. North Holland, 1999.
 - [10] G. Viglietta. *Guarding and Searching Polyhedra*. Ph.D. thesis, University of Pisa, 2012.
 - [11] G. Viglietta. Face-guarding polyhedra. *Computational Geometry: Theory and Applications*, vol. 47, no. 8, pp. 833–846, 2014.