Edge Guards for Polyhedra in Three-Space

Javier Cano[*] Csaba D. Tóth[†] Jorge Urrutia[‡] Giovanni Viglietta[§]

Abstract

It is shown that every polyhedron in $\mathbb{R}^3$ can be guarded by at most $\frac{5}{6}$ of its edges. This result holds even if the boundary of the polyhedron is disconnected (i.e., if the polyhedron has "holes"), and regardless of the genus of each connected component of its boundary.

When a polyhedron is homeomorphic to a ball and all its faces are triangles, the bound can be improved slightly to $\frac{29}{36}$ of its edges.

Keywords: Art Gallery Problem; edge guard; polyhedron; edge cover; 1-skeleton

1 Introduction

A *polyhedron* P is a compact connected 3-manifold in $\mathbb{R}^3$ bounded by a piecewise linear 2-manifold. As such, P has finitely many vertices, edges, and faces. Faces need not be simply connected, as shown in Figure 1. Every edge of P is a (topologically closed) line segment between two vertices of P; each edge is shared by two faces of P as a common part of their boundary. Two points a and b are *visible* in a polyhedron P if the closed line segment ab is contained in P. For the edges of a polyhedron P, we adapt the notion of *weak visibility*: an edge e of P is visible to a point p if there is a point $q \in e$ such that p and q are visible in P. A set S of edges jointly *guard* P if every point $a \in P$ is visible to some edge in S. It is possible that a point $a \in P$ does not see any vertex of P [13, Section 10.2]; however, it can be proved that every point $a \in P$ sees at least six edges of P. It follows that every polyhedron with m edges can be guarded with at most $m - 5$ edges.

It was conjectured in [18] that every polyhedron of genus zero with m edges can be guarded with at most $\frac{m}{6}$ edge guards. This bound would be optimal apart from an additive constant: for every $k \in \mathbb{N}$, there is a polyhedron P_k in $\mathbb{R}^3$ with $6(k+1)$ edges that requires at least k edge guards (see Figure 1). The polyhedron P_k is the union of a flat tetrahedron T and k pairwise disjoint thin tetrahedra attached to one face of T such that none of their apexes can be seen from any of the edges of T. Since each thin tetrahedron has to be guarded by one of its edges, P requires k edge guards.

In the preliminary version of this paper [4], it was proved that every polyhedron with m edges (and arbitrary genus) can be guarded by at most $\frac{27}{32}m$ edges. Prior to the present paper, this was the only known nontrivial upper bound for the edge guard problem for general polyhedra. For every polyhedron P, it was shown how to choose a set of edges that jointly guard P as the

[*]Posgrado en Ciencia e Ingeniería de la Computación, Universidad Nacional Autónoma de México, Mexico City, Mexico. j_cano@uxmcc2.iimas.unam.mx

[†]Department of Mathematics, California State University Northridge, Los Angeles, CA; Department of Computer Science, Tufts University, Medford, MA, USA. csaba.toth@csun.edu

[‡]Instituto de Matemáticas, Universidad Nacional Autónoma de México, Mexico City, Mexico. urrutia@matem.unam.mx

[§]School of Information Science, JAIST, 1–1 Asahidai, Nomi City, Ishikawa 923–1211, Japan. johnny@jaist.ac.jp **(Corresponding author)**

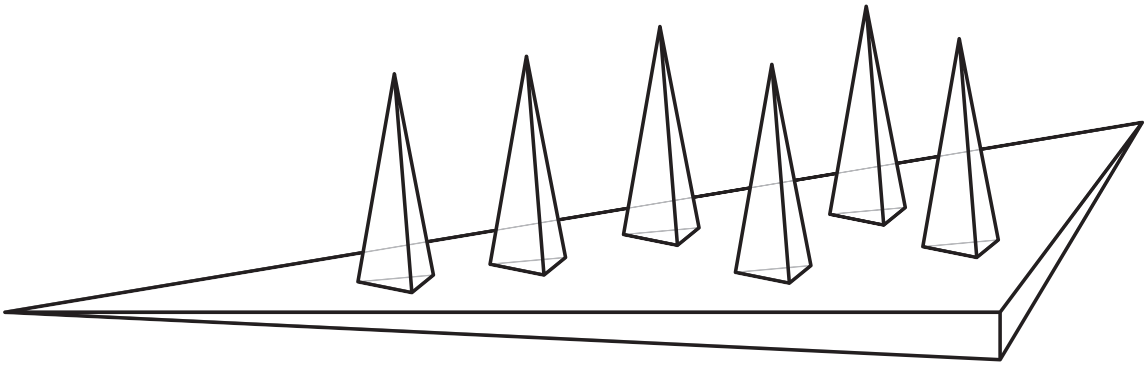

Figure 1: A polyhedron with m edges that requires $\frac{m}{6} - 1$ edge guards.

union of two sets: (i) a set of edges that cover all vertices of P, and (ii) at most 3/4 of the remaining edges. In the present paper, we will show how to refine this technique in order to improve the upper bound in [4] from $\frac{27}{32}m$ to $\frac{5}{6}m$ edges.

The *1-skeleton* of a polyhedron P is the graph defined by the vertices and edges of P. In Section 2, we observe that each connected component of the 1-skeleton of a polyhedron is a 2-edge-connected graph with minimum degree at least 3. An *edge cover* of a graph $G = (V, E)$ is a set of edges $E' \subseteq E$ such that every vertex in V is incident to an edge in E'. By placing guards at every edge in an edge cover of the 1-skeleton of P, we ensure that every point in P that sees a vertex is guarded. In Section 3, using classical matching theory, we give upper bounds on the size of a minimum edge cover in a 2-edge-connected graph with minimum degree at least 3.

In Section 4, we describe a partition of the edges of a polyhedron P into four classes, and prove some geometric properties of such a partition. Essentially, we show that every point in P sees either edges in certain classes or a vertex of P. These properties translate into four different strategies for guarding P. Finally, in Section 5, we prove that a combination of the four strategies and a careful choice of an edge cover of the 1-skeleton of P yield a set of at most $\frac{5}{6}m$ edges that jointly guard P.

The gap between the $\frac{m}{6} - 1$ lower bound of Figure 1 and the $\frac{5}{6}m$ upper bound proved in this paper is still substantial. Further reducing this gap is left as an open problem.

Related work. The problem of guarding a polyhedron by choosing a subset of its edges is an instance of the well-known Art Gallery Problem [13, 18]. Most of the previous research in this field focused on polygons in the plane. For example, it is well known that every simple polygon with n vertices can be guarded by at most $\frac{n}{3}$ point guards [5], and that every simple orthogonal polygon with n vertices can be guarded by at most $\frac{n}{4}$ point guards [9]. Both bounds are tight. Shermer proved that every simple polygon with n vertices can be guarded by at most $\frac{3}{10}n + 1$ of its edges [17]. On the other hand, Toussaint constructed polygons where $\frac{n}{4}$ edges guards are necessary; aside from a few small n, this lower bound is widely believed to be tight [12, 14, 16].

Everett and Rivera-Campo [6] showed that every triangulated polyhedral terrain in $\mathbb{R}^3$ with n vertices can be guarded by $\frac{n}{3}$ edges, as this many edges can cover all faces of a plane triangulation with n vertices. They also proved that the faces of every plane graph with n vertices can be guarded by $\frac{2}{5}n$ edges. See also [3] for other variants of the Art Gallery Problem for polyhedral terrains.

For *orthogonal* polyhedra with m edges in $\mathbb{R}^3$, it was conjectured that $\frac{m}{12}$ edge guards are always sufficient: for every $k \in \mathbb{N}$, there are *orthogonal* polyhedra P_k in $\mathbb{R}^3$ with $12(k+1)$ edges that require at least k edge guards [18]. Benbernou et al. [1] showed that for every orthogonal polyhedron of genus g, $\frac{11}{72}m - \frac{g}{6} - 1$ edge guards are sufficient. Also, denoting by r the total number of *reflex* edges in an orthogonal polyhedron, $\frac{7}{12}r - g + 1$ edge guards are sufficient. More recently, Viglietta [19] improved these upper bounds to $\frac{m-4}{8} + g$ and $\frac{r-g}{2} + 1$ edge guards, respectively, for *2-reflex* orthogonal polyhedra; i.e., orthogonal polyhedra having reflex edges in only two directions (the latter bound is optimal for $g = 0$). All the aforementioned results about orthogonal polyhedra also hold for *open* edge guards: an open edge e of P is visible from a point p if there is a point q in the relative interior of e such that p and q are visible in P.

2 Polyhedra and 1-skeletons

A *polyhedron* P is defined as a compact and connected 3-manifold in $\mathbb{R}^3$ whose boundary ∂P is a piecewise linear 2-manifold. Since P is a compact set in $\mathbb{R}^3$, it is contained in some ball, and it contains its own boundary: $\partial P \subset P$. The piecewise linear structure of ∂P naturally yields a subdivision into finitely many polygonal *faces*. Namely, a face is any maximal planar subset of ∂P whose relative interior is connected. A face is bounded by line segments called *edges*, whose endpoints are called *vertices* of P. Edges and vertices are defined in such a way that no vertex lies in the relative interior of an edge. Each edge bounds exactly two (non-coplanar) faces. If the internal dihedral angle between these two faces is less than π, the edge is *convex*; otherwise, the internal dihedral angle must be greater than π, and the edge is *reflex*.[1] Each vertex of a polyhedron is shared by three or more faces, and is incident to the same number of edges. If all edges incident to a vertex are convex, the vertex is said to be *convex* as well. A vertex is a *saddle* if it has both convex and reflex incident edges.

The boundary ∂P of a polyhedron need not be connected. One connected component of ∂P, the *outer boundary*, includes all vertices of the convex hull of P (plus possibly other vertices). In addition, ∂P has one connected component for each *hole* of P. For example, $P = \{(x, y, z) \in \mathbb{R}^3 \mid 1 \le \max\{|x|, |y|, |z|\} \le 2\}$ is a polyhedron consisting of a cube C_1 with a central cubic hole C_2. Its boundary ∂P has two connected components, coinciding with ∂C_1 (the outer boundary) and ∂C_2, respectively.

Each connected component of ∂P is a compact orientable 2-manifold, which is either homeomorphic to a sphere or has one or more *handles* [8, Theorem 9.3.11]. The number of handles of a connected component of ∂P is called the *genus* of that component. For example, if the boundary of a polyhedron has the topology of a sphere, it has genus zero. On the other hand, $P = \{(x, y, z) \in \mathbb{R}^3 \mid 1 \le \max\{|x|, |y|\} \le 2 \wedge |z| \le 1\}$ is a polyhedron whose boundary has the topology of a torus, and hence it has genus one.

The *1-skeleton* of a polyhedron P is the graph defined by the vertices and edges of P. Note that even if a polyhedron has a connected boundary of genus zero, its 1-skeleton is not necessarily connected. An example is given in Figure 1, where the polyhedron depicted has neither handles nor holes, but its 1-skeleton has $\frac{m}{6}$ connected components; one for each tetrahedron.

Also, observe that the top face of the large tetrahedron in Figure 1 is a triangle with several triangular holes. In general, if the boundary of a polyhedron is connected but its 1-skeleton is not, there must be a face with some holes in it. Note that some of these holes may touch each other: this would happen for instance if the bases of two of the thin tetrahedra in Figure 1 shared a vertex.

A celebrated theorem by Steinitz characterizes the 1-skeletons of *convex* polyhedra as the 3-vertex-connected planar graphs with at least four vertices [7]. This characterization does not extend to non-convex polyhedra. In fact, there exist polyhedra whose 1-skeleton is planar and

[1] The definition of face prevents two faces sharing an edge from being coplanar, and therefore the dihedral angle between them cannot be π.

connected but not 2-vertex-connected: such would be, for instance, the polyhedron in Figure 1 if there were only a single thin tetrahedron, and one of its base vertices lay on the perimeter of the top face of the large tetrahedron.

However, it is easy to see that every connected component of the 1-skeleton of a polyhedron is 2-edge-connected. Indeed, a bridge in the 1-skeleton would correspond to an edge of the polyhedron that bounds only one face.

Observation 1. *Each connected component of the 1-skeleton of a polyhedron is a 2-edge-connected graph with minimum degree at least 3.*

3 Edge covers

An *edge cover* of a graph $G = (V, E)$ is a set of edges $E' \subseteq E$ such that every vertex $v \in V$ is incident to some edge in E'. It is well known that a minimum edge cover is the union of a maximum matching $M \subseteq E$ and one extra edge for each vertex not covered by M. Hence the size of a minimum edge cover is $|V| - |M|$; this number is denoted by $\rho(G)$. The goal of this section is to study $\rho(G)$ when G is a 2-edge-connected graph with minimum degree at least 3. In Section 5, we will use these results to construct a minimum edge cover of the 1-skeleton of a polyhedron.

Nishizeki and Baybars [11] proved that a maximum matching in a connected planar graph with $n \geq 10$ vertices and minimum degree at least 3 has at least $\frac{n+2}{3}$ edges. So, every such graph has an edge cover of size at most $\frac{2n-2}{3}$, which can be computed in $O(n)$ time [15]. However, we seek upper bounds in terms of the number of *edges*, rather than the number of vertices.

We will now review some terminology and the Edmonds-Gallai Structure Theorem for maximum matchings [10, 20]. Let $G = (V, E)$ be a simple graph. A matching $M \subseteq E$ is *perfect* if it covers all vertices of G; it is *near-perfect* if it covers all but one vertex of G. If every subgraph obtained by deleting one vertex from G has a perfect matching, then G is said to be *factor-critical*.

According to the Edmonds-Gallai Structure Theorem, if $M \subseteq E$ is a maximum matching of G, then there is a vertex set $U \subseteq V$ (a Berge-Tutte witness set) with the following properties:

- M restricts to a perfect matching on every even component of $G[V \setminus U]$;
- every odd component of $G[V \setminus U]$ is factor-critical, and M restricts to a near-perfect matching on it;
- M matches all vertices of U to vertices in distinct odd components of $G[V \setminus U]$.

A minimum edge cover of G can be obtained by augmenting the maximum matching M with one extra edge for each odd component of $G[V \setminus U]$ that is not fully covered by M.

We are now in the position to prove the following lemma.

Lemma 2. *Let G be a 2-edge-connected graph with m edges and minimum degree at least 3. Then $\rho(G) \leq \lfloor \frac{m+1}{3} \rfloor$, and this bound is the best possible. Moreover, if $\rho(G) > \frac{m}{3}$, then G is factor-critical.*

Proof. We begin by observing that the bound $\rho(G) \leq \lfloor \frac{m+1}{3} \rfloor$ is the best possible. If $m \equiv 0 \pmod 3$, our lower-bound construction is the complete bipartite graph $K_{3,\frac{m}{3}}$. If $m \geq 9$, such a graph satisfies the hypotheses of the lemma and has a minimum edge cover of size $\frac{m}{3}$. For $m \equiv 1 \pmod 3$, we use the same construction, with the addition of a single edge between two vertices of maximum degree. For $m \equiv 2 \pmod 3$, the lower-bound construction is illustrated in Figure 2. For $m = 5$, we have the 1-skeleton of a pyramid with square base. The base of

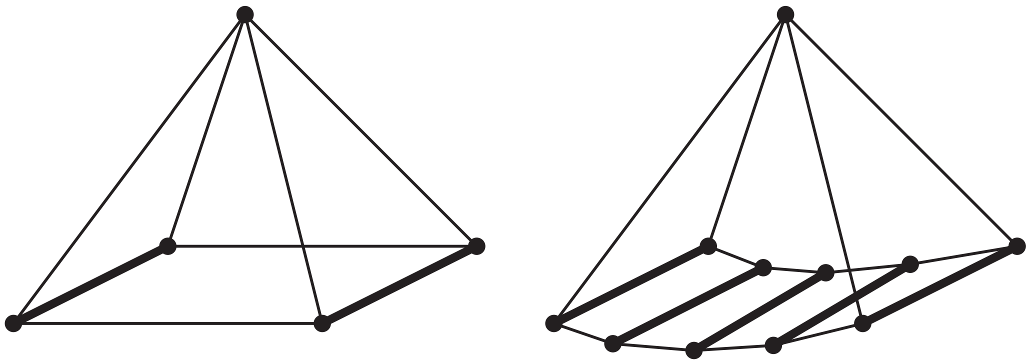

Figure 2: Lower-bound constructions for Lemma 2 when $m \equiv 2 \pmod 3$. Maximum matchings are highlighted.

the pyramid can be extended to a ladder for larger values of m. In these graphs, the size of a minimum edge cover is $\frac{m+1}{3}$.

Let $G = (V, E)$ be a 2-edge-connected graph with $n = |V|$ vertices, $m = |E|$ edges, and minimum degree at least 3. We will now prove that $\rho(G) \leq \lfloor \frac{m+1}{3} \rfloor$. Note that $3n \leq 2m$, since every vertex is incident to three or more edges, and each edge is incident to two vertices.

Let $M \subseteq E$ be a maximum matching of G. The Edmonds-Gallai Structure Theorem yields a Berge-Tutte witness set $U \subseteq V$.

If $U = \emptyset$, then $G[V \setminus U]$ has a unique connected component, which coincides with G itself. Assume that n is even; then M is a perfect matching on G with $\frac{n}{2}$ edges. Since M is also a minimum edge cover, we have $\rho(G) = \frac{n}{2} \leq \frac{m}{3}$.

On the other hand, if $U = \emptyset$ and n is odd, then G is factor-critical, and M is a near-perfect matching on G with $\frac{n-1}{2}$ edges. Hence, $\rho(G) = n - |M| = \frac{n+1}{2}$. Note that the inequality $3n \leq 2m$ cannot hold with equality, because the left-hand side is odd and the right-hand side is even. Thus, $3n \leq 2m - 1$. It follows that $\rho(G) \leq \lfloor \frac{m+1}{3} \rfloor$.

Assume now that $U \neq \emptyset$. We will prove that in this case $\rho(G) \leq \frac{m}{3}$ holds. Denote the connected components of $G[V \setminus U]$ by $G_i = (V_i, E_i)$, for $i = 1, 2, \ldots, \ell$. Hence, we have

$$|V| = |U| + \sum_{i=1}^{\ell} |V_i|.$$

For every $i = 1, \ldots, \ell$, if $|V_i|$ is even, then M restricts to a perfect matching on G_i with $|V_i|/2$ edges. If $|V_i|$ is odd, then M restricts to a near-perfect matching on G_i with $(|V_i| - 1)/2$ edges. Additionally, M contains one edge for each vertex of U. Thus,

$$|M| = |U| + \sum_{i=1}^{\ell} \left\lfloor \frac{|V_i|}{2} \right\rfloor.$$

Since $\rho(G) = |V| - |M|$, we have

$$\rho(G) = \sum_{i=1}^{\ell} \left\lceil \frac{|V_i|}{2} \right\rceil. \tag{1}$$

Let $\overline{E}_i \subseteq E$ denote the set of all edges incident to vertices in V_i; that is, all edges in E_i plus the edges between U and V_i. The edge sets $\overline{E}_i$, $i = 1, \ldots, \ell$, are pairwise disjoint, and so $\sum_{i=1}^{\ell} |\overline{E}_i| \leq m$. By (1), in order to prove that $\rho(G) \leq \frac{m}{3}$, it suffices to prove that the following holds for $i = 1, \ldots, \ell$:

$$\left\lceil \frac{|V_i|}{2} \right\rceil \leq \frac{|\overline{E}_i|}{3}. \tag{2}$$

Let x_i be the sum of degrees of the vertices in V_i. Since the minimum degree is at least 3, we have $x_i \geq 3|V_i|$. Also, at least two edges in $\overline{E}_i$ are incident to some vertices in U, because G is 2-edge-connected. Hence $x_i \leq 2|\overline{E}_i| - 2$, and therefore

$$3|V_i| \leq 2|\overline{E}_i| - 2. \tag{3}$$

If $|V_i|$ is even, we immediately obtain $|V_i|/2 < |\overline{E}_i|/3$, which implies (2). If $|V_i|$ is odd, then the two sides of (3) have opposite parity, and hence we can improve the inequality to $3|V_i| \leq 2|\overline{E}_i| - 3$, which in turn is equivalent to (2).

From the discussion above, it follows that $\rho(G) > \frac{m}{3}$ may hold only if $U = \emptyset$ and n is odd, which implies that G is factor-critical. □

Corollary 3. *Let G be a 2-edge-connected graph with m edges and minimum degree at least 3 such that $\rho(G) > \frac{m}{3}$. Then for every edge e of G,*

(i) there exists a minimum edge cover of G containing e, and

(ii) there exists a minimum edge cover of G not containing e.

Proof. Let $e = \{u, v\}$. By Lemma 2, G is factor-critical, and thus it has a (maximum) matching M that covers every vertex except u. Of course, M does not contain e. We can extend M to a minimum edge cover of G by adding any of the (at least three) edges incident to u: choosing e yields (i), and any other choice yields (ii). □

4 Edge classes and their properties

Let P be a polyhedron with m edges. (Recall from Section 2 that the boundary of P may be disconnected, and each connected component may have arbitrary genus.) Let $G = (V, E)$ denote the 1-skeleton of P. We may assume, by rotating P if necessary, that no edge in E is parallel to any coordinate plane. Similarly, we may assume that no face of P is parallel to any coordinate axis, and that all vertices of P have distinct x- (respectively, y- and z-) coordinates.

In order to help intuition, we visualize the xy-plane as "horizontal" and the z-axis as "vertical"; thus, any plane parallel to the z-axis is said to be vertical. Accordingly, if a point a has a larger y-coordinate (respectively, z-coordinate) than a point b, we say that a lies to the *right* of b (respectively, *above* b), and b lies to the *left* of a (respectively, *below* a). Recall that every edge of P is incident to exactly two faces of P.

We distinguish between four types of edges in E as follows. For every edge $e \in E$, let H_e denote the (unique) vertical plane containing e. The plane H_e divides $\mathbb{R}^3$ into two half-spaces, lying to the left and to the right of H_e, respectively. We say that e is a *left edge* (respectively, a *right edge*) if both faces incident to e lie to the left (respectively, to the right) of H_e, locally. If the two faces incident to e locally lie on opposite sides of H_e and the interior or P lies above (respectively, below) both faces, then e is a *lower edge* (respectively, an *upper edge*). Clearly, these four classes form a partition of E. See Figure 3 for examples.

We will now prove two lemmas that immediately yield four different strategies for guarding a polyhedron.

Proposition 4. *Every point on the boundary of a polyhedron sees a non-left edge.*

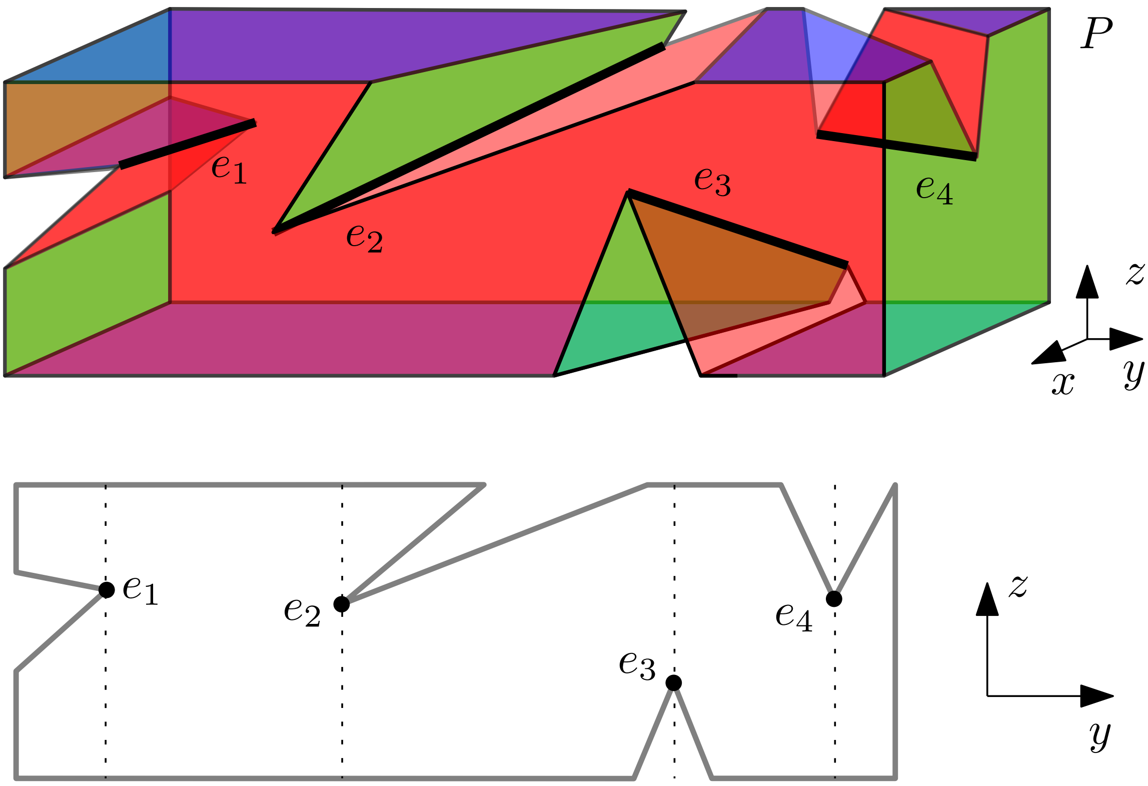

Figure 3: Top: A left edge e_1, a right edge e_2, a lower edge e_3, and an upper edge e_4 in a polyhedron P. Bottom: The cross-section of the polyhedron P with a plane parallel to the yz-plane, which is stabbed by edges $e_1, \dots, e_4$. Dotted lines indicate the vertical planes $H_{e_1}, \dots, H_{e_4}$.

Proof. Let a be a point on the boundary of a polyhedron P, and let F be a face of P containing a. The orthogonal projection of F onto the horizontal plane xy is a polygon F', possibly with holes (some of which may touch each other or the external boundary of F', cf. Section 2). An edge of F' is *left-facing* (respectively, *right-facing*) if the interior of F' locally lies on its left (respectively, right).

Let a' be the projection of a onto the xy-plane; obviously, $a' \in F'$. Shoot a ray from a' directed leftward, and let p be the first point where the ray hits the boundary of F' (possibly, $p = a'$). Observe that p lies on a right-facing edge e' of F' (if p is a vertex of F', it has at least one incident right-facing edge e'). The edge e' is the vertical projection of an edge e of P. Since F locally lies on the right of e, it cannot be a left edge of P. Now, since a' sees e' in F', it follows that a sees e in P. $\square$

Lemma 5. *In a polyhedron, every point sees a non-left edge.*

Proof. Suppose, to the contrary, that there is a polyhedron P with a point $a \in P$ that sees only left edges. Consider the cross-section of P with the plane H_a containing a and parallel to the yz-plane (refer to Figure 4). Due to Proposition 4, a cannot be a vertex of P, or else it would see a non-left edge. Furthermore, if H_a contains a vertex of P, we may slightly rotate P around a, so that all edges preserve their respective classes, and H_a no longer contains vertices of P.

The intersection $H_a \cap P$ may have several connected components; let P_a denote the component that contains a. Note that P_a is a 2-dimensional polygon, possibly with holes. Each vertex of P_a corresponds to a unique edge of P: since H_a contains no vertices of P, each vertex

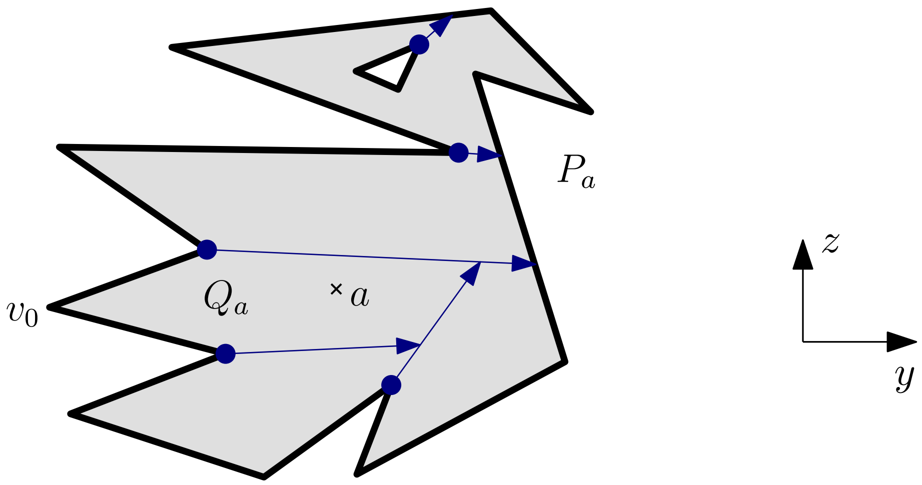

Figure 4: The polygon P_a is the cross-section of the polyhedron P with the plane H_a containing a and parallel to the yz-plane. The vertices in V_a^* are marked with large dots. P_a is decomposed into sub-polygons by rays emitted by the vertices in V_a^*. The sub-polygon Q_a contains a. Since Q_a is convex, a sees the leftmost vertex v_0 of Q_a.

of P_a is the unique intersection point between a single edge of P and the plane H_a. Because each edge of P is incident to exactly two faces, it follows that each vertex of P_a is incident to exactly two edges of P_a, and therefore its internal angle is well defined. Let V_a^* denote the set of reflex vertices of P_a that correspond to left edges of P.

Decompose the polygon P_a as follows. Consider the vertices in V_a^* in an arbitrary order. From each vertex $v \in V_a^*$ successively shoot a ray along its angle bisector, and draw a line segment ℓ_v along the ray from v to the point where the ray first hits the boundary of P_a or a previously drawn segment. If a ray hits a vertex of P_a, perturb the ray slightly so that it does not hit any vertex. Since $v \in V_a^*$, the two edges of P_a incident to v lie to the left of v, and so the line segment ℓ_v lies strictly to the right of v (in fact, v is the left endpoint of ℓ_v).

The segments ℓ_v collectively decompose P_a into sub-polygons. Denote by $Q_a \subseteq P_a$ a sub-polygon containing the point a (note that a may lie on the boundary of more than one sub-polygon). We claim that Q_a is a convex polygon. Indeed, if this were not the case, a would see a reflex vertex v_i of Q_a. Both edges of Q_a incident to v_i must lie on ∂P, because none of the previously drawn segments ℓ_v can start or end at a reflex vertex of one of the sub-polygons. In particular, v_i does not correspond to a left edge of P, which contradicts our initial assumption that a only sees left edges. Hence, Q_a is convex.

Let v_0 be the leftmost vertex of Q_a. Both edges of Q_a incident to v_0 are on the right side of v_0, or else there would be another vertex of Q_a to the left of v_0. Moreover, both edges incident to v_0 lie on ∂P. Indeed, if neither of them lay on ∂P, they would be sub-segments of ℓ_w and $\ell_{w'}$, for some $w, w' \in V_a^*$. In this case, v_0 would be the left endpoint of, say, ℓ_w, implying that $v_0 = w$. But then, $\ell_{w'}$ would be incident to $v_0 = w$, a vertex of P_a other than w', which is impossible by construction. Similarly, if only one of the edges incident to v_0 lay on ∂P, then v_0 would correspond to a non-left edge of P. The other edge incident to v_0 would then be a sub-segment of ℓ_w with $w \neq v_0$ (there is no such segment as ℓ_{v_0} because $v_0 \notin V_a^*$). Again, we would have a segment ℓ_w incident to a vertex other than w, which is impossible by construction.

Since both edges incident to v_0 lie on ∂P and to the right of v_0, it follows that v_0 corresponds to a right edge of P. Because Q_a is convex, a sees v_0, which contradicts the assumption that a only sees left edges. $\square$

Of course, by a symmetric argument, we also obtain the following.

Corollary 6. *In a polyhedron, every point sees a non-right edge.* □

Lemma 7. *In a polyhedron, every point sees a vertex or a non-lower reflex edge.*

Proof. Let P be a polyhedron and let $a \in P$ be a point that sees no vertices of P. We will prove that a sees a reflex edge of P that is not a lower edge.

Shoot a ray from a directed vertically downward, and let b be the first point where the ray hits the boundary ∂P. Let F be a face of P containing b. Fix a triangulation of F, let T be a triangle containing b, and let v be any vertex of T. For all $0 \leq \lambda \leq 1$, we denote by p_λ the point $\lambda v + (1 - \lambda)b$. Observe that the points p_λ span the line segment $bv \subset T$, with $p_0 = b$ and $p_1 = v$.

Define $X = \{\lambda \in [0, 1] \mid a \text{ does not see } p_\lambda\}$. Note that $0 \notin X$, because a sees $b = p_0$, and $1 \in X$, because $v = p_1$ is a vertex of P, and a does not see any vertex. Since X is bounded and not empty, the number $x = \inf X$ is a well-defined strictly positive real.

Because ∂P is piecewise linear, there is a left (respectively, right) neighborhood of x corresponding to points p_λ that are visible (respectively, not visible) to a. This implies that the segment ap_x intersects ∂P without "crossing" it; i.e., without exiting P. One such point of intersection is therefore a point of strict non-convexity of ∂P, which must lie on a reflex edge e (and is visible to a). Also, since $p_{x-\varepsilon}$ is visible to a for all $\varepsilon \in [0, x]$, it follows that there are points of P that lie directly below e, and thus e is not a lower edge of P. □

Again, by a symmetric argument, we obtain the following.

Corollary 8. *In a polyhedron, every point sees a vertex or a non-upper reflex edge.* □

We give one last lemma, which will be useful in the next section. Recall that a saddle vertex of a polyhedron is a vertex with both convex and reflex incident edges.

Lemma 9. *If a connected component of the 1-skeleton of a polyhedron contains no saddle vertices, then it contains both a left edge and a right edge.*

Proof. Let P be a polyhedron, and let $G = (V, E)$ be a connected component of its 1-skeleton such that no vertex in V is a saddle vertex of P. Then either all edges in E are convex or all of them are reflex. Indeed, if there existed $e, e' \in E$ with e convex and e' reflex, then there would be a saddle vertex in G along a path connecting e and e' (such a path exists, since G is connected).

Suppose that all edges in E are convex. Let $v \in V$ be the vertex of G with smallest x-coordinate, and let L be the plane through v that is parallel to the yz-plane. Let L_ε denote the plane obtained by translating L by the vector $(\varepsilon, 0, 0)$. Since all edges of P incident to v are in E, they are all convex. It follows that for a sufficiently small $\varepsilon > 0$, the intersection between L_ε and the faces of P incident to v is a (convex) polygon Q. Now it is easy to see that the leftmost (respectively, rightmost) vertex of Q lies on an edge of G that is a left (respectively, right) edge of P.

If all edges in E are reflex, the proof is identical, except that the leftmost vertex of Q lies on a right edge of P, and vice versa. □

5 Obtaining a set of guards

The results in Section 4 lead to different ways of choosing edge guards in a polyhedron P. For instance, Lemma 5 implies that the set of all non-left edges guards P. An alternative strategy is to pick an edge cover E' of the 1-skeleton of P, and then pick all the non-lower reflex edges that are not in E'. Indeed, every point that sees a vertex of P also sees (an endpoint of) an edge in E'; the other points must see a non-reflex lower edge, by Lemma 7.

In the following, we will show how to combine such strategies in order to obtain our final upper bounds on the number of edge guards.

Lemma 10. *Every polyhedron with m edges whose 1-skeleton has an edge cover E' of size m' can be guarded with at most*

$$\frac{3m + m' - \ell - c}{4}$$

edge guards, where ℓ is the number of left edges in E', and c is the number of convex edges that are not in E'.

Proof. Let P be a polyhedron, and let $G = (V, E)$ be its 1-skeleton, with $|E| = m$. Assume that G has an edge cover E', with $|E'| = m'$. Let c and r be, respectively, the number of convex and reflex edges in $E \setminus E'$. We denote by m_{left}, m_{right}, m_{lower}, m_{upper} the number of left, right, lower, and upper edges in E, respectively. Also, we denote by m'_{left} the number of left edges in E', and by c_{left} (respectively, r_{left}) the number of convex (respectively, reflex) left edges in $E \setminus E'$. Similarly, we define m'_{right}, m'_{lower}, and so on.

Let g be the size of a minimum set of edges that collectively guard P. By Lemma 5 and Corollary 6, we have the inequalities

$$g \le m - m_{\text{left}},$$
$$g \le m - m_{\text{right}}.$$

Furthermore, Lemma 7 and Corollary 8 yield

$$g \le m' + r - r_{\text{lower}},$$
$$g \le m' + r - r_{\text{upper}} = m - c - r_{\text{upper}},$$

where we used the identity $m = m' + c + r$. By adding up these four inequalities and observing that $r - r_{\text{lower}} - r_{\text{upper}} = r_{\text{left}} + r_{\text{right}}$, we obtain

$$\begin{aligned} 4g &\le 3m + m' - m_{\text{left}} - m_{\text{right}} + r - r_{\text{lower}} - r_{\text{upper}} - c \\ &= 3m + m' + (r_{\text{left}} - m_{\text{left}}) + (r_{\text{right}} - m_{\text{right}}) - c. \end{aligned}$$

Note that $r_{\text{right}} - m_{\text{right}} \le 0$, so we can drop this term. On the other hand, $r_{\text{left}} - m_{\text{left}} = -m'_{\text{left}} - c_{\text{left}} \le -m'_{\text{left}}$, which yields

$$g \le \frac{3m + m' - m'_{\text{left}} - c}{4}.$$

Now, renaming m'_{left} to ℓ gives the desired bound. □

Finally, we prove our main result.

Theorem 11. *Every polyhedron with m edges can be guarded with at most $\frac{5}{6}m$ edge guards.*

Proof. Let P be a polyhedron with m edges (whose boundary has any number of connected components, each of which has arbitrary genus), and let $G = (V, E)$ be its 1-skeleton. We will construct an edge cover E' of G of size m' in such a way that

$$m' - \ell - c \le \frac{m}{3}, \tag{4}$$

where ℓ and c are defined as in Lemma 10. Observe that plugging (4) into the upper bound given by Lemma 10 yields

$$\frac{3m + m' - \ell - c}{4} \leq \frac{3m + m/3}{4} = \frac{5}{6}m$$

edge guards, as desired.

We will construct E' on each connected component of G separately. For each component $G_i = (V_i, E_i)$ with $|E_i| = m_i$, we will choose an edge cover E'_i of size m'_i that satisfies an inequality analogous to (4):

$$m'_i - \ell_i - c_i \leq \frac{m_i}{3}, \tag{5}$$

where ℓ_i is the number of left edges in E'_i, and c_i is the number of convex edges in $E_i \setminus E'_i$. Then, adding up the instances of (5) corresponding to all connected components of G, we will obtain (4).

Let G_i be a connected component of G; recall that G_i is 2-edge-connected and has minimum degree at least 3 (cf. Observation 1). Lemma 2 states that $\rho(G_i) \leq \frac{m_i+1}{3}$. If G_i has an edge cover of size $m'_i \leq \frac{m_i}{3}$, then (5) is satisfied, and we are done. So, let us assume that

$$\frac{m_i}{3} < \rho(G_i) \leq \frac{m_i + 1}{3}.$$

Suppose that G_i has a left edge e; by Corollary 3(i), G_i has a minimum edge cover E'_i containing e. Hence, $m'_i \leq \frac{m_i+1}{3}$, $\ell_i \geq 1$, and (5) is satisfied. On the other hand, if G_i has no left edges, then it has a saddle vertex, due to Lemma 9. In particular, G_i has a convex edge e'. Then, by Corollary 3(ii) there is a minimum edge cover of G_i not containing e'. In this case, $m'_i \leq \frac{m_i+1}{3}$, $c_i \geq 1$, and again (5) is satisfied. □

We can also give a better upper bound for a special class of polyhedra.

Theorem 12. *Every polyhedron homeomorphic to a ball (i.e., with a connected boundary of genus zero) with triangular faces and m edges can be guarded with at most $\frac{29}{36}m$ edge guards.*

Proof. Let P be such a polyhedron. Its 1-skeleton G is isomorphic to a plane triangulation, and hence it is a 3-vertex-connected planar graph. Also, P satisfies Euler's formula $f + n = m + 2$, where f is the number of faces of P, and n is the number of its vertices. Since each face has three edges and each edge is shared by two faces, we have $3f = 2m$, which yields $3n = m + 6$.

Assume that $n \geq 10$. It is known [2, Theorem 5] that a 3-vertex-connected planar graph with $n \geq 10$ vertices has a matching of size at least $\frac{n+4}{3}$. This yields an edge cover for G of at most $n - \frac{n+4}{3} = \frac{2n-4}{3} = \frac{2}{9}m$ edges. By taking $m' \leq \frac{2}{9}m$ in Lemma 10 (and dropping the non-positive term $-\ell - c$), we obtain an upper bound of $\frac{29}{36}m$ edge guards for P, as desired.

Assume now that $n \leq 9$. Of course there are no polyhedra with fewer than four vertices. Also, if $n = 4$, then $m = 6$, $f = 4$, and thus P is a tetrahedron, which is convex. In this case one edge guard suffices, which is less than $\frac{29}{36}m$. In the following, we will assume that $5 \leq n \leq 9$.

By combining [2, Theorem 3] and [2, Lemma 4], we have that any 3-vertex-connected planar graph with $n \geq 4$ vertices has a matching of size at least $\min\{\frac{n-1}{2}, \frac{n+4}{3}\}$. If $n \leq 9$, then $\frac{n-1}{2} < \frac{n+4}{3}$, and therefore G has a matching of size at least $\frac{n-1}{2}$, which is perfect or near-perfect. Such a matching can be extended to an edge cover of at most $\frac{n+1}{2} = \frac{m+9}{6}$ edges. Lemma 10 with $m' \leq \frac{m+9}{6}$ gives an upper bound of

$$\frac{19m + 9 - 6\ell - 6c}{24} \tag{6}$$

edge guards. For $n = 5, 6, 7, 8$, we have $m = 9, 12, 15, 18$; in each of these cases, $\lfloor \frac{19m+9}{24} \rfloor$ is equal to $\lfloor \frac{29}{36} m \rfloor$, as desired.

Finally, if $n = 9$, we know that G has a matching M of all vertices except one, say v. Let e be an edge incident to v; by rotating P if necessary, we can assume that e is a left edge of P. Since the set $E' = M \cup \{e\}$ is an edge cover of G containing a left edge, we can plug $\ell \geq 1$ in (6), obtaining an upper bound of $\lfloor \frac{19m+3}{24} \rfloor$ edge guards. For $m = 21$, this is equal to $16 = \lfloor \frac{29}{36} m \rfloor$. $\square$

Acknowledgments. The authors would like to thank Joseph O'Rourke for comments and suggestions that improved the readability of this paper. This research was partially supported by the NSERC grant RGPIN 35586 and PAPIIT IN105221 Programa de Apoyo a la Investigación e Innovación Tecnológica, UNAM, Mexico.

References

[1] Nadia M. Benbernou, Erik D. Demaine, Martin L. Demaine, Anastasia Kurdia, Joseph O'Rourke, Godfried T. Toussaint, Jorge Urrutia, and Giovanni Viglietta. Edge-guarding orthogonal polyhedra. In *Proceedings of the 23rd Canadian Conference on Computational Geometry (CCCG)*, pages 461–466, Toronto, ON, 2011. URL: `http://www.cccg.ca/proceedings/2011/papers/paper50.pdf`.

[2] Therese Biedl, Erik D. Demaine, Christian A. Duncan, Rudolf Fleischer, and Stephen G. Kobourov. Tight bounds on maximal and maximum matchings. *Discrete Mathematics*, 285(1–3):7–15, 2004. `doi:10.1016/j.disc.2004.05.003`.

[3] Prosenjit Bose, David G. Kirkpatrick, and Zaiqing Li. Worst-case-optimal algorithms for guarding planar graphs and polyhedral surfaces. *Computational Geometry*, 26(3):209–219, 2003. `doi:10.1016/S0925-7721(03)00027-0`.

[4] Javier Cano, Csaba D. Tóth, and Jorge Urrutia. Edge guards for polyhedra in three-space. In *Proceedings of the 24th Canadian Conference on Computational Geometry (CCCG)*, pages 155–159, Charlottetown, PEI, 2012. URL: `http://2012.cccg.ca/papers/paper27.pdf`.

[5] Václav Chvátal. A combinatorial theorem in plane geometry. *Journal of Combinatorial Theory, Series B*, 18(1):39–41, 1975. `doi:https://doi.org/10.1016/0095-8956(75)90061-1`.

[6] Hazel Everett and Eduardo Rivera-Campo. Edge guarding polyhedral terrains. *Computational Geometry*, 7(3):201–203, 1997. `doi:10.1016/0925-7721(95)00051-8`.

[7] Branko Grünbaum. *Convex Polytopes*, volume 221 of *Graduate Texts in Mathematics*. Springer-Verlag, 2nd edition, 2003. `doi:10.1007/978-1-4613-0019-9`.

[8] Morris W. Hirsch. *Differential Topology*, volume 33 of *Graduate Texts in Mathematics*. Springer-Verlag, 1976. `doi:10.1007/978-1-4684-9449-5`.

[9] Jeff Kahn, Maria Klawe, and Daniel Kleitman. Traditional galleries require fewer watchmen. *SIAM Journal on Algebraic Discrete Methods*, 4(2):194–206, 1983. `doi:10.1137/0604020`.

[10] László Lovász and Michael D. Plummer. *Matching Theory*, volume 121 of *Annals of Discrete Mathematics*. North-Holland, 1986.

[11] Takao Nishizeki and Ilker Baybars. Lower bounds on the cardinality of the maximum matchings of planar graphs. *Discrete Mathematics*, 28(3):255–267, 1979. `doi:10.1016/0012-365X(79)90133-X`.

[12] Joseph O'Rourke. Galleries need fewer mobile guards: A variation on Chvátal's theorem. *Geometriae Dedicata*, 14(3):273–283, 1983. `doi:10.1007/BF00146907`.

[13] Joseph O'Rourke. *Art Gallery Theorems and Algorithms*. Oxford University Press, 1987.

[14] Joseph O'Rourke. Visibility. In Jacob E. Goodman, Joseph O'Rourke, and Csaba D. Tóth, editors, *Handbook of Discrete and Computational Geometry*, chapter 47. CRC Press, Boca Raton, FL, 3rd edition, 2017.

[15] Ignaz Rutter and Alexander Wolff. Computing large matchings fast. *ACM Transactions on Algorithms*, 7(1):1:1–1:21, 2010. `doi:10.1145/1868237.1868238`.

[16] Thomas C. Shermer. Recent results in art galleries. *Proceedings of the IEEE*, 80(9):1384–1399, 1992. `doi:10.1109/5.163407`.

[17] Thomas C. Shermer. A tight bound on the combinatorial edge guarding problem. In David Avis and Prosenjit Bose, editors, *Snapshots of Computational and Discrete Geometry*, volume 3, pages 191–223. School of Computer Science, McGill University, Montreal, 1994.

[18] Jorge Urrutia. Art gallery and illumination problems. In Jörg-Rüdiger Sack and Jorge Urrutia, editors, *Handbook of Computational Geometry*, pages 973–1027. North Holland / Elsevier, 2000. `doi:10.1016/b978-044482537-7/50023-1`.

[19] Giovanni Viglietta. Optimally guarding 2-reflex orthogonal polyhedra by reflex edge guards. *Computational Geometry*, 86, 2020. `doi:10.1016/j.comgeo.2019.101589`.

[20] Douglas B. West. A short proof of the Berge-Tutte formula and the Gallai-Edmonds structure theorem. *European Journal of Combinatorics*, 32(5):674–676, 2011. `doi:10.1016/j.ejc.2011.01.009`.