Guarding Orthogonal Art Galleries with Sliding k-Transmitters: Hardness and Approximation

Therese Biedl[a], Timothy M. Chan[a], Stephanie Lee[a], Saeed Mehrabi[a,*], Fabrizio Montecchiani[b], Hamideh Vosoughpour[a], Ziting Yu[a]

[a] *Cheriton School of Computer Science, University of Waterloo, Waterloo, Canada.*
[b] *Department of Engineering, University of Perugia, Perugia, Italy.*

Abstract

A sliding k-transmitter inside an orthogonal polygon P, for a fixed $k \geq 0$, is a point guard that travels along an axis-parallel line segment s in P. The sliding k-transmitter can see a point $p \in P$ if the perpendicular from p onto s intersects the boundary of P in at most k points. In the Minimum Sliding k-Transmitters (ST_k) problem, the objective is to guard P with the minimum number of sliding k-transmitters.

In this paper, we give a constant-factor approximation algorithm for the ST_k problem on P for any fixed $k \geq 0$. Moreover, we show that the ST_0 problem is NP-hard on orthogonal polygons with holes even if only horizontal sliding 0-transmitters are allowed. For $k > 0$, the problem is NP-hard even in the extremely restricted case where P is simple and monotone. Finally, we study art gallery theorems; i.e., we give upper and lower bounds on the number of sliding transmitters required to guard P relative to the number of vertices of P.

Keywords: Art gallery problem, Sliding k-transmitters, ϵ-nets

1. Introduction

Let P be a (not necessarily orthogonal) polygon with n vertices. The art gallery problem, posed by Victor Klee in 1973 [28], asks for the minimum number

☆A preliminary version of this paper appeared in proceedings of the 11th International Conference and Workshop on Algorithms and Computation (WALCOM 2017) [5] and the 28th Canadian Conference on Computational Geometry (CCCG 2016) [8]. Research of TB and TC is supported in part by the Natural Sciences and Engineering Research Council of Canada (NSERC). Research of FM is supported in part by the project: "Algoritmi e sistemi di analisi visuale di reti complesse e di grandi dimensioni" - Ricerca di Base 2017, Dipartimento di Ingegneria dell'Universit degli Studi di Perugia". Research was done while FM was visiting the University of Waterloo.

*Corresponding author

Email addresses: `biedl@cs.uwaterloo.ca` (Therese Biedl), `tmchan@cs.uwaterloo.ca` (Timothy M. Chan), `s363lee@uwaterloo.ca` (Stephanie Lee), `smehrabi@uwaterloo.ca` (Saeed Mehrabi), `fabrizio.montecchiani@unipg.it` (Fabrizio Montecchiani), `hvosough@uwaterloo.ca` (Hamideh Vosoughpour), `z44yu@uwaterloo.ca` (Ziting Yu)

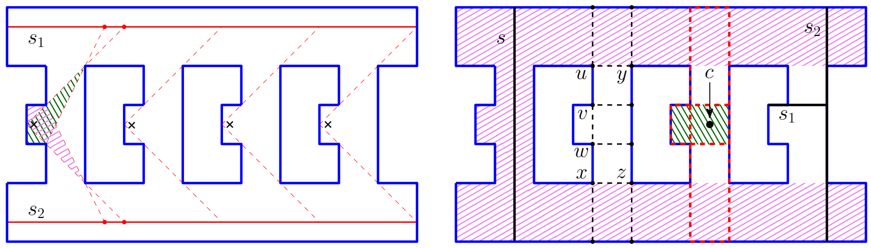

(a) **(b)**

Figure 1: (a) An orthogonal polygon P that can be guarded with two horizontal mobile guards, but requires $\Theta(n)$ sliding cameras to be guarded since no two crosses can be seen by one sliding camera. (b) Sliding camera s sees the rising-shaded subpolygon of P. We also show parts of the pixelation induced by rays from reflex vertices $\{u, v, w, x, y, z\}$, and the cross c whose supporting horizontal slice is downward shaded. Segments s_1 and s_2 are guard-segments.

of point guards required to guard P, where a point guard g sees a point $p \in P$ if the line segment connecting g to p lies inside P. Chvátal [11] was the first to answer the question by giving the tight bound $\lfloor n/3 \rfloor$ on the number of point guards that are needed to guard a simple polygon with n vertices. For polygons with holes, Hoffmann et al. [19] proved that $\lfloor (n+h)/3 \rfloor$ point guards are always sufficient and occasionally necessary, where h is the number of holes. For *orthogonal polygons*, it was proved multiple times [20, 26, 28] that $\lfloor n/4 \rfloor$ point guards are always sufficient and sometimes necessary to guard the interior of a simple orthogonal polygon with n vertices. Finding the minimum number of guards is NP-hard on simple polygons [25], even on simple orthogonal polygons [32] or monotone polygons [24]. A number of results concerning approximation algorithms are also known [17, 23, 24]. See the surveys by O'Rourke [31] or Urrutia [35] for a history of the art gallery problem.

Mobile guards. A *mobile guard* is a point guard that travels along a line segment s inside P. This guard can see a point p in P if and only if there exists a point $g \in s$ such that the line segment pg lies entirely inside P. The notion of mobile guards was introduced by Avis and Toussaint [3]. O'Rourke [30] proved that $\lfloor n/4 \rfloor$ mobile guards are sufficient for guarding arbitrary polygons with n vertices. For orthogonal polygons with n vertices, $\lfloor (3n+4)/16 \rfloor$ mobile guards are always sufficient and sometimes necessary [1].

Sliding cameras. The sliding camera model of visibility, which can be viewed as an *orthogonal mobile guard* visibility, was introduced by Katz and Morgenstern in 2011 [22]. In this model, a point guard travels along an axis-parallel line segment s inside P and it can see a point $p \in P$ if the perpendicular from p onto s lies entirely inside P. Note that an orthogonal mobile guard traveling along s may see a larger area of P than a sliding camera traveling along s, see also Figure 1(a). Katz and Morgenstern [22] proved that the problem of guarding a simple orthogonal polygon with the minimum number of sliding cameras can be solved in polynomial time, if only horizontal cameras are allowed.

Durocher et al. [15] showed that the problem (when both horizontal and vertical cameras are allowed) is NP-hard on polygons with holes (see also [16]); the NP-hardness on simple polygons remains open. Moreover, Durocher et al. [14] claimed a (3.5)-approximation algorithm for the problem on simple orthogonal polygons, but this was later discovered by the authors to be incorrect (private communication). For the special case of monotone orthogonal polygons, Katz and Morgenstern [22] gave a 2-approximation algorithm, which was later improved by de Berg et al. [13] to a linear-time exact algorithm. See Mehrabi's thesis [27] for a more detailed set of known results on sliding cameras.

k-transmitters. Motivated by covering a region with wireless transmitters, Aichholzer et al. [2] introduced variants where guards can see through a limited number of walls. To model this, they defined a *k-transmitter* as a point p in a polygon P that is considered to see all points q in P for which the line segment $\overline{pq}$ intersects the boundary of P at most k times. Notice that only cases of even k are interesting. Finding the minimum number of k-transmitters to guard a simple polygon is NP-hard [10], regardless of whether the k-transmitters are points or polygon edges. Numerous bounds are known on the number of k-transmitters that are necessary and sufficient, depending on the type of k-transmitters (point or edge) and the type of polygon [2, 4, 10].

Our model and results. In this paper, we combine the concept of a k-transmitter with that of a sliding camera and introduce *sliding k-transmitters*. More precisely, given any fixed integer $k \geq 0$, a sliding k-transmitter is a point guard g that travels along an axis-parallel (i.e., horizontal or vertical) line segment s inside an orthogonal polygon P. (We will sometimes omit "sliding" for the rest of this paper, as we study no other type of guards.) The sliding k-transmitter g can see a point $p \in P$ if and only if the perpendicular from p onto s intersects the boundary of P at most k times (see Section 2 for a more precise definition). For instance, the two sliding 2-transmitters s_1 and s_2 shown in Figure 1(a) guard P entirely, while viewing s_1 as a sliding 4-transmitter would suffice to guard P entirely. We allow sliding k-transmitters to travel along the edges of the polygon.[1] A sliding camera is the same as a sliding 0-transmitter. We define the *Minimum Sliding k-Transmitters (ST_k)* problem on an orthogonal polygon P as the problem of guarding P with the minimum number of sliding k-transmitters. When only horizontal sliding k-transmitters are allowed, we refer to this problem as *Minimum Horizontal Sliding k-Transmitters (HST_k)* problem.

In this paper, we give approximation algorithms and hardness results for both ST_k and HST_k problems for any fixed $k \geq 0$. Let P be any orthogonal polygon with n vertices. Our results are as follows.

1. We give a constant-factor approximation algorithm for the ST_k and HST_k problems on P for any fixed $k \geq 0$. Our algorithm works by constructing a

[1] With some minor modifications, the results in this paper also hold if guards must be strictly inside P except at their end.

small ε-net for the hitting set problem that naturally arises from the ST_k and HST_k problems. This gives then an $O(1)$-approximation algorithm for these problems. As opposed to previous attempts at such approximation algorithms [14], our algorithm works even if P has holes.

2. We prove that the HST_0 problem is NP-hard when P is allowed to have holes. The same proof also works for the ST_0 problem, and it is different and perhaps simpler than the previous NP-hardness proof for ST_0 [15]. Moreover, for all $k > 0$, we show that the ST_k problem is NP-hard, even if P is simple and monotone. This is, to our knowledge, the first NP-hardness results for guarding polygons that are both monotone and orthogonal, in any variant of the guarding problem.

3. We consider art gallery theorems for sliding k-transmitters, i.e., theorems that bound the number of sliding k-transmitters relative to the number of vertices. We present the following results for P: (i) $\lfloor (3n+4)/16 \rfloor$ sliding 0-transmitters are always sufficient and sometimes necessary to guard P entirely, (ii) For any $k > 0$, $\lfloor n/6 \rfloor$ sliding k-transmitters are sometimes necessary to guard P entirely, (iii) For any $k \geq 0$, $\lfloor n/4 \rfloor$ only-horizontal sliding k-transmitters are always sufficient and sometimes necessary to guard P. We also consider the natural restriction where sliding k-transmitters are not allowed to intersect each other, and show here that then $\lfloor (n+1)/5 \rfloor$ sliding 0-transmitters are always sufficient to guard P.

Organization. This paper is organized as follows. We first give some definitions and preliminary results in Section 2. We then give our constant-factor approximation algorithm in Section 3. Section 4 contains our hardness results, and art-gallery-theorem results are in Section 5. Finally, we conclude the paper with a discussion of open problems in Section 6.

2. Preliminaries

Throughout the paper, P is an orthogonal polygon with n vertices; we denote the boundary of P by δP. An orthogonal line segment s inside P is called *maximal* if both endpoints of s lie on the boundary of P. Given a fixed integer $k \geq 0$, a sliding k-transmitter g can see a point $p \in P$ if and only if the perpendicular from p onto s intersects the boundary of P at most k times. By "intersecting" the boundary of P in this definition, we require that the line of sight enters the exterior of P immediately after such an intersection. More precisely, we consider P to be closed and define the number of intersections of a line segment gp with δP to be twice the number of connected components of $gp \cap ext(P)$, where $ext(P)$ denotes the exterior of P. Our definition of sliding k-transmitters allows any horizontal or vertical line segment inside P to be used as such, but we show below that we can restrict the attention to a finite subset of line segments inside P and hence discretize the ST_k problem on P.

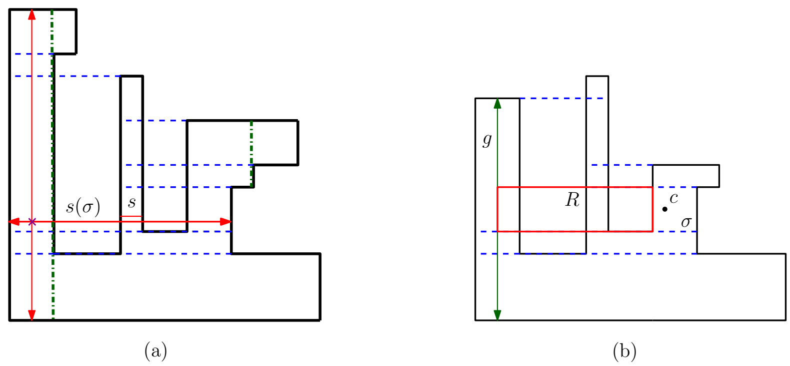

Figure 2: (a) Horizontal partition-segments (blue dashed), one horizontal slice-segment (red solid) and two vertical guard-segments (green dot-dashed) for $k = 2$. A cross is the intersection point of a horizontal and a vertical slice segment. (b) An illustration in support of the proof of Lemma 1

Slices and slice segments. Define horizontal *partition segments* as follows. Start with a horizontal edge e. Extend e leftwards until it intersects the boundary of P (in the above sense) for the kth time. Likewise extend e rightwards. If there are not enough such intersections, then stop at the last one. See Figure 2(a). The horizontal partition segments split the interior of the polygon into rectangles. Any maximal rectangle inside the polygon that does not intersect a horizontal partition segment (except at its boundaries) is called a *horizontal slice*. Since any edge gives rise to one partition segment, and any partition segment intersects the interior of $O(k+1)$ vertical edges, we have $O((k+1)n)$ horizontal slices. The following lemma argues that this partitioning is "correct" in the sense that any maximal transmitter either guards all or nothing of the interior of a slice.

Lemma 1. *Let σ be a horizontal slice and let c be a point in its interior. If a maximal vertical sliding k-transmitter g sees c, then it sees all points of σ.*

Proof. Assume without loss of generality that g lies to the left of σ. Let R be the rectangle bounded from above by the line containing the top edge of σ, bounded from below by the line containing the bottom edge of σ, bounded from the right by the left edge of σ, and bounded from the left by the vertical line through g. See Figure 2(b). We consider R to be open on the top and bottom, and closed on the left and right. Since g can see c, there is some maximal horizontal segment inside R that intersects the boundary of P at most k times. Note that R cannot contain a vertex of P: if it did, the horizontal edge incident to the rightmost such vertex would generate a partition segment that crosses σ, contradicting that σ is a horizontal slice. Therefore any maximal horizontal segment inside R intersects the boundary of P at most k times, which implies that that g sees every point in σ. □

A slice is a 2D region, but it will be convenient to represent it using a 1D segment that possibly extends even beyond the slice. Fix a horizontal slice σ.

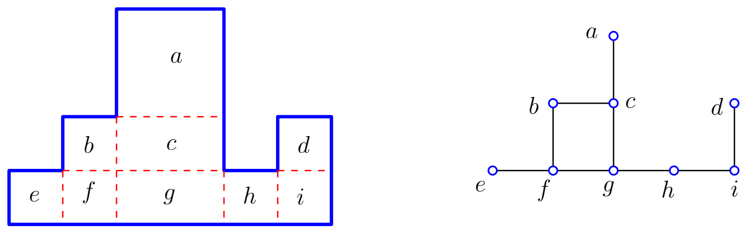

Figure 3: A polygon P with its pixelation (left) and the weak dual graph D of P (right) for $k = 0$.

Let s be the maximal horizontal segment inside σ that is equidistant from the top and bottom. Extend s (much like we did for partition segments) to the left until it intersects the boundary of P for the kth time, and also extend s to the right until it intersects the boundary of P for the kth time. We call the resulting segment the *slice segment* of σ and denote it by $s(\sigma)$. See Figure 2 for an example. Observe that the slice segment of a horizontal slice captures the slice "being guarded" in the sense that any maximal vertical line segment that intersects the slice segment would guard the slice entirely. We define *vertical slices* of P and *vertical slice segments* in an analogous fashion. There are $O((k+1)n)$ vertical slice segments since there are $O((k+1)n)$ vertical slices.

Observation 1. A slice σ is guarded entirely if and only if there is a guard intersecting the slice-segment $s(\sigma)$.

Pixelation, the dual graph, and crosses. The *pixelation* of P is obtained by taking the union of horizontal and vertical segmentations of P. The resulting rectangles inside P are called *pixels*. There are $O(n^2)$ pixels since the boundary of each is defined by one of the $O(n)$ edges. There may well be $\Omega(n^2)$ pixels. Notice that the pixels are in one-to-one correspondence with pairs of slices that cross. We can hence identify each pixel with a *cross* c, which is the point where the two slice segments $s(\sigma_H)$ and $s(\sigma_V)$ (corresponding to horizontal slice σ_H and vertical slice σ_V, respectively) cross. We say that $s(\sigma_H)$ and $s(\sigma_V)$ *support* c. See also Figure 1. Recall that $s(\sigma_H)$ and $s(\sigma_V)$ were defined starting with a segment strictly inside the slice; in consequence c is strictly inside the pixel. Denote the set of crosses by X.

Define the *dual graph* D of P to be the weak dual graph of the pixelation of P for $k = 0$; i.e., D has a vertex for every pixel and two pixels are adjacent in D if and only if they have a common side. See Figure 3 for an example. Some previous papers [34, 7] studied guarding problems in the special case where this dual graph D is a tree. Generalizing this as in [7], we call a polygon P *thin* if all the four corners of every pixel lies on δP. Any polygon for which D is a tree is also thin, but thin polygons are more general in that they include some polygons with holes.

Guard segments. Let g be a sliding k-transmitter in P. Notice that doing a *parallel shift*; i.e., translating g vertically (if it is horizontal) or horizontally (if it is vertical) while staying inside P, does not change its visibility region as long as we stay inside P. We may hence assume that every sliding k-transmitter runs along pixel edges. We may also restrict our attention to sliding k-transmitters that are maximal line segments within P (because all others would see a subset). More formally, let s be a horizontal edge of P. Define a sliding k-transmitter s' to be the segment obtained by extending s until we intersect the boundary of P for the kth time. (As usual, overlaps with horizontal edges do not count as intersection.) The resulting line segments are called *horizontal-guard segments*. We denote the set of all such horizontal guard segments by Γ_H. Define *vertical-guard segments* Γ_V similarly, and set $\Gamma = \Gamma_H \cup \Gamma_V$ to be the set of guard segments. Note that there are at most n guard segments, one per edge.

We say that a guard segment γ *hits* a cross c if and only if γ intersects one of the slice segments supporting c. We now show that in order to solve the ST_k problem on P it is sufficient to restrict our attention to hitting crosses using guard segments.

Lemma 2. *There exists a set of m sliding k-transmitters that guards P if and only if there exists a set $S' \subseteq \Gamma$ of m guard segments such that every cross c is hit by some guard segment $\gamma \in S'$.*

Proof. ($\Rightarrow$) Suppose that we have a set of m sliding k-transmitters that guards P entirely. Fix one sliding k-transmitter s. Translate s in parallel (i.e., move it horizontally if s is vertical, move it vertically if s is horizontal) until we reach δP. Thus, s is now overlapping an edge of P. Extend s so that it is maximal while still within P. Both operations can only increase the region seen. The resulting segment s' is a guard segment. After doing this to all sliding k-transmitters, we now have a set of guard segments S' that sees all of P. Now, consider any cross $c \in X$. Since c is a point in P, it is guarded by some guard segment $\gamma \in S'$. Since c is in the interior of its pixel, therefore there exists a point $g \in \gamma$ such that the line segment $\overline{gc}$ is normal to γ and intersects δP in at most k points. But, $\overline{gc}$ is part of the slice segment that supports c and is perpendicular to γ. So g is the intersection point between that slice segment and guard segment γ.

($\Leftarrow$) Fix an arbitrary point q inside P, say it belongs to horizontal slice σ_H and vertical slice σ_V (choosing the slice arbitrarily if q lies on the boundary between two). Let c be the cross corresponding to pixel $\sigma_H \cap \sigma_V$. There exists a guard-segment $g \in S'$ that hits c; say g is vertical. Then, by Lemma 1 segment g sees all of σ_H and therefore also q. $\square$

Lemma 2 means that we have discretized the problem: it suffices to find a minimum-cardinality set $S' \subseteq \Gamma$ that hits all crosses in X. In fact, the algorithms we design later will allow further generalization: we can specify exactly which crosses should be hit and which sliding k-transmitters may be used as guards. To make this formal, suppose that we are given some $\Gamma' \subseteq \Gamma$ and some $X' \subseteq X$. Then, the (Γ', X')-*ST_k problem* consists of finding a minimum-cardinality subset of sliding k-transmitters in Γ' that hit all crosses in X', or to report that no such set exists. Notice that with a suitable choice of Γ' this encompasses both ST_k and HST_k.

3. Approximation Algorithms via ε-nets

In this section, we give a polynomial-time $O(1)$-approximation algorithm for both HST_k and ST_k problems for any $k \geq 0$. We first give some definitions.

Hitting sets. A *set system* is a pair $\mathcal{R} = (\mathcal{U}, \mathcal{S})$, where $\mathcal{U}$ is a universe set of objects and $\mathcal{S}$ is a collection of subsets of $\mathcal{U}$. A *hitting set* for the set system $(\mathcal{U}, \mathcal{S})$ is a subset of $\mathcal{U}$ that intersects every set in $\mathcal{S}$. For the (Γ', X')-ST_k problem, we construct a set system as follows. Let $\mathcal{U} = \Gamma'$ be all potential sliding k-transmitters. For each cross $c \in X'$ that needs to be hit, define S_c to be all the sliding k-transmitters in $\mathcal{U}$ that hit c, and let $\mathcal{S}$ be the collection of these sets. From the definitions, finding a minimum hitting set for this set system is the same as solving the (Γ', X')-ST_k problem.

An ε-*net* for a set system $\mathcal{R} = (\mathcal{U}, \mathcal{S})$ is a subset N of $\mathcal{U}$ such that every set S in $\mathcal{S}$ with size at least $\varepsilon \cdot |\mathcal{U}|$ has a non-empty intersection with N. Brönnimann and Goodrich [9] showed that ε-nets can be used to derive approximation algorithms as follows. Define a *net finder* to be a (poly-time) algorithm that, for a given set system $\mathcal{R} = (\mathcal{U}, \mathcal{S})$ and any given $r > 0$, computes an $(1/r)$-net of $\mathcal{R}$ whose size is at most $s(r)$ for some function s. Also, a *verifier* is a poly-time algorithm that, given a subset $H \subset \mathcal{U}$, states (correctly) that H is a hitting set, or returns a non-empty set $R \in \mathcal{S}$ such H does not hit S.

Lemma 3. *[9] Let $\mathcal{R}$ be a set system that admits both a poly-time net finder and a poly-time verifier. Then there is a poly-time algorithm that computes a hitting set of size at most $s(4 \cdot OPT)$, where OPT stands for the size of an optimal hitting set, and $s(r)$ is the size of the $(1/r)$-net.*

Clearly, the hitting set problems corresponding to the HST_k and ST_k problems both have a polynomial-time verifier. Therefore, the lemma gives an $O(1)$-approximation algorithm for as long as we can find an ε-net whose size is $O(1/\varepsilon)$. In the following, we first give a small ε-net for the hitting set corresponding to the HST_k problem and will then show how to re-use this net in order to get an ε-net for the ST_k problem.

An ε-net for the HST_k problem. We now show the existence of such a small ε-net for the HST_k problem. To this end, we need (yet another) reformulation that simplifies the problem.

Lemma 4. *A set S of horizontal-guard segments hits all crosses in a set M if and only if S intersects all the vertical slice segments that support crosses in M.*

Proof. If k-transmitter γ hits cross c, then it intersects either its horizontal supporting slice segment σ_H or its vertical supporting slice segment σ_V. But if γ intersects σ_H, then since both are horizontal and γ is maximal we have $\sigma_H \subseteq \gamma$, in case of which γ also contains point c and therefore intersects σ_V. So, either way γ intersects σ_V. □

For the HST_k problem, it hence suffices to represent every cross by its vertical slice segment and so reduce the problem to the following. Given a set of horizontal line segments $\mathcal{H}$ and a set of vertical line segments $\mathcal{V}$, find a minimum set $S \subseteq \mathcal{H}$ such that every line segment in $\mathcal{V}$ is intersected by S. This problem is also known as the *Orthogonal Segment Covering* problem. We hence have:

Corollary 1. The HST_k problem reduces to the Orthogonal Segment Covering problem.

The Orthogonal Segment Covering problem is known to be NP-complete [21]. The following lemma shows that the Orthogonal Segment Covering problem has a small ε-net; by the above this immediately implies a small ε-net for the hitting set problem corresponding to the HST_k problem.

Lemma 5. *The Orthogonal Segment Covering problem has a* $(1/r)$*-finder with size-function* $s(r) \in O(r)$*.*

PROOF. We employ a result of Clarkson and Varadarajan [12] that shows that ε-nets of small size can be found for hitting set problems in geometric objects, using random sampling. The size is as desired under the assumption that these geometric objects have small union complexity, i.e., the union of n of these geometric objects has complexity $O(n)$.

Observe that a horizontal line segment $\gamma = [x, x'] \times y$ intersects a vertical line segment $\sigma = a \times [b, b']$ if and only if the point $q = (x, y, x')$ lies in the range $Q = (-\infty, a] \times [b, b'] \times [a, \infty)$.

We call Q an extruded translated quadrant since it is a quadrant of the xz-plane that has been translated to (a, b, a) and then extruded in y-direction, i.e., repeated for all $y \in [b, b']$. Hence, we can map every horizontal segment into a point and every vertical segment into an extruded translated quadrant such that the segments intersect if and only if the point lies in such a quadrant.

Viewing the y-coordinate as *time*, each vertical segment hence is mapped to a translated quadrant that exists for some time, and the union complexity becomes the number of changes to a dynamic two-dimensional union of translated quadrants. In general there could be a quadratic number of such changes, but our translated quadrants are special: the corners of Q are on the line $\{(a, y, a) : y \in \mathbb{R}\}$ and hence all corners of all quadrants are on the plane $\{(x, y, z) : x = z\}$. As such, adding or deleting one such quadrant to the union causes a constant number of changes. So the union complexity of our translated quadrants is linear. By [12], there exists a $(1/r)$-net with size $O(r)$ for hitting these translated quadrants by points, hence hitting vertical segments by horizontal segments. □

Putting everything together, we have the following theorem.

Theorem 6. *There exists a polynomial-time* $O(1)$*-approximation algorithm for the Orthogonal Segment Covering problem and the* HST_k *problem.*

An ε-net for the ST_k problem. We now show that using the ε-net for the HST_k problem, we can easily find one for the ST_k and hence have an $O(1)$-approximation algorithm for this problem as well.

Theorem 7. *There exists a polynomial-time $O(1)$-approximation algorithm for the ST_k problem on any orthogonal polygon P.*

Proof. Fix a polygon P and consider the (Γ', X')-ST_k problem for P. It suffices to show that for any $r > 0$ there exists a $1/r$-net T of size $O(r)$ for the corresponding hitting set problem $\mathcal{R}$. Let T_H be a $1/2r$-net for the hitting set corresponding to the HST_k problem for P, X' and the horizontal sliding k-transmitters in Γ'. Let T_V be a $1/2r$-net for the hitting set of VST_k (i.e., when we want to guard the polygon using only-vertical sliding k-transmitters) for P, X' and the vertical k-transmitters in Γ'. Set $T := T_H \cup T_V$. We claim that T is a $1/r$-net for $\mathcal{R}$.

Suppose that some set S_c in the hitting set problem satisfies $|S_c| \geq |\mathcal{U}|/r$. Translating back, this means that some cross $c \in X'$ is hit by at least $|\Gamma'|/r$ guard segments. Assume w.l.o.g. that at least half of these hitting guard segments are horizontal. Then, the vertical slice segment σ_V that supports c intersects at least $|\Gamma'|/2r$ horizontal guard segments in Γ'. By definition of a $(1/2r)$-net, therefore there is a line segment $\gamma \in T_H$ that intersects σ_V. Therefore $\gamma \in T$ hits c as required. □

4. Hardness Results

In this section, we give two hardness results. First, we show that the HST_0 problem (i.e., guarding polygon P with the minimum number of only-horizontal sliding 0-transmitters) is NP-hard on orthogonal polygons with holes (Section 4.1). The same construction also proves NP-hardness of ST_0, though this result was known before (with a different proof). The construction could easily be adapted to $k > 0$, but rather than doing this, we give a different construction for $k > 0$ which results in NP-hardness of ST_k and HST_k even if the polygon is simple, connected and monotone (Section 4.2). This shows that $k > 0$ is substantially harder than $k = 0$ because ST_0 is linear-time solvable on monotone orthogonal polygons [13].

4.1. HST_0 Problem

Recall that the HST_0 problem is polynomial-time solvable on simple orthogonal polygons [22]. We show in this section that the HST_0 problem becomes NP-hard on orthogonal polygons with holes. Note that the hardness proof of Durocher and Mehrabi [16] does not apply to the HST_0 problem because they require both horizontal and vertical 0-transmitters.

The reader may recall that we showed that the HST_0 problem reduces to the Orthogonal Segment Covering problem, which is known to be NP-hard [21]. However, this does not prove NP-hardness of the HST_0 problem, because not every instance of Orthogonal Segment Covering can be expressed as an instance of the HST_0 problem. Instead, we give a different reduction from Minimum Vertex Cover on maximum degree-3 planar graphs. This problem (which is NP-hard [18]) consists of, given a planar graph $G = (V, E)$ with at most 3 incident edges at each vertex, find a minimum set $C \subseteq V$ such that for every edge at least one endpoint is in C.

Given a maximum degree-3 planar graph G, we first compute a *bar visibility representation* of G, that is, we assign to each vertex a horizontal line segment (called *bar*) and to each edge (v, w) a vertical *strip* of positive width that joins the corresponding bars and that is disjoint from all other bars and strips. It is well-known that every planar graph has such a representation (see e.g. Tamassia and Tollis [33]), and it can be found in linear time. By making strips sufficiently thin, we can ensure that no two strips of edges occupy the same x-range.

From this visibility representation, we construct an orthogonal polygon P such that G has a vertex cover of size k if and only if P can be guarded by $3N + k$ horizontal sliding 0-transmitters, where N is the number of the vertices of G. Polygon P will have a hole for every inner face of G (and perhaps some additional holes).

An initial version of the polygon consists of the visibility representation, with bars thickened into positive height. Clearly if we have a vertex cover of G, then we can place horizontal sliding 0-transmitters in the bars of all corresponding vertices and guard all edge-strips. This is, however, not sufficient to also guard all vertex-bars, which is why we add small gadgets that require additional sliding 0-transmitters, and these 0-transmitters guard none of the edge strips but all other parts of the vertex bars.

We add three gadgets, called *elephant-gadgets*, around every vertex-bar u in Γ. This might be done in one of the three different ways, up to symmetry, depending on how the edge-strips are connected to u; see Figure 4 for an illustration. Recall that we assume that no two edge-strips share an x-range, so that one of these constructions is always possible if the vertex has degree 3. For vertices of smaller degrees, we omit some of the edge-strips at the ends of the bars (but not the elephant-gadgets) appropriately.

Lemma 8. *The following statements are equivalent: (i) G has a vertex cover of size k, (ii) P can be guarded with $k + 3N$ horizontal sliding 0-transmitters, (iii) P can be guarded with $k + 3N$ sliding 0-transmitters.*

Proof. Given a vertex cover C of size k, we can place one horizontal "vertex 0-transmitter" inside the bar of each vertex of C, and one horizontal "spine 0-transmitter" inside each elephant-gadget. Then the vertex 0-transmitters cover all edge-strips, and the spine 0-transmitters cover all the elephants-gadgets and the parts of vertex-bars not in edge-strips. So we can guard P with $k + 3N$ horizontal sliding 0-transmitters.

Clearly, if we can guard P with $k + 3N$ horizontal sliding 0-transmitters, then it can also be guarded with $k + 3N$ arbitrary ones.

Finally, assume we can guard P with $k + 3N$ (not necessarily horizontal) sliding 0-transmitters. Define a vertex set C as follows: if any 0-transmitter lies

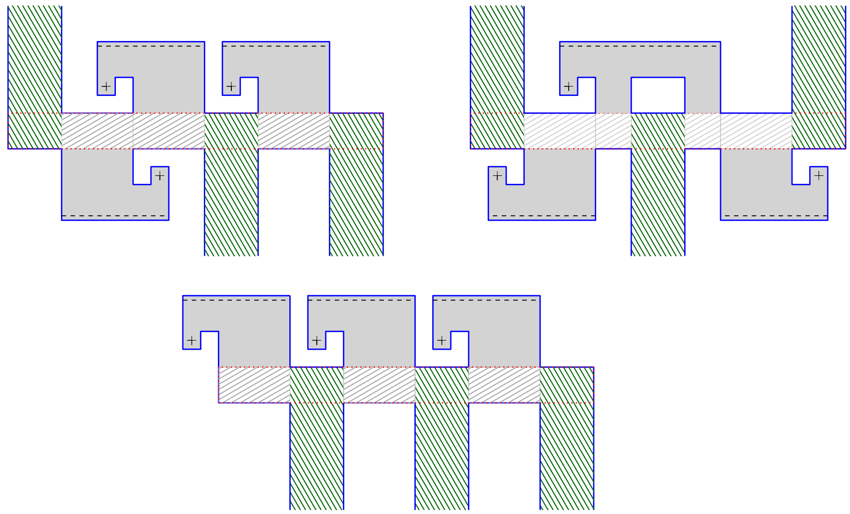

Figure 4: The NP-hardness construction. Vertex-bars are red (dotted). Edge-strips are green (falling pattern). Elephant-gadgets are gray; the dashed segment indicates the spine 0-transmitter and the cross indicates the trunk. Gray (rising pattern) regions indicates the part of the vertex-box guarded by the spine 0-transmitters.

fully inside the bar of vertex v, then add v to C. If any 0-transmitter intersects the edge-strip of (v, w) outside of vertex bars, then arbitrarily add one of v and w to C. We claim that $|C| \leq k$. This holds because we have $3N$ elephant-gadgets, and each elephant-gadget has a "trunk"-pixel that must be guarded, but can be guarded only by a 0-transmitter that is entirely within that elephant-gadgets. So at least $3N$ 0-transmitters are neither inside a vertex-bar nor intersect an edge-strip, and $|C| \leq k$. Finally we claim that C is a vertex cover. Indeed, the strip of any edge e contains a pixel that is not in a vertex-bar. The 0-transmitter that guards this pixel resides inside the edge-strip of e or inside the bars of one of its ends. Therefore one endpoint of e has been added to C. $\square$

Clearly, P can be constructed in polynomial time and so by Lemma 8 NP-hardness follows. Also, notice that P is thin (recall that a polygon is thin if all the four corners of every pixel lies on the boundary of the polygon). The NP-hardness of guarding problems in thin polygons (albeit with other models of guards and visibility) have been studied before [34, 7]. The NP-hardness holds for both HST_0 and ST_0 problems; note that the NP-hardness of the ST_0 problem is already known [16], but the constructed polygon is not thin. We summarize:

Theorem 9. *The HST_0 and ST_0 problems are both* NP*-hard on thin orthogonal polygons with holes.*

4.2. ST_k Problem for $k > 0$

In this section, we show that the ST_k problem is NP-hard, for any $k > 0$, even if the polygon is orthogonal and monotone (hence simple). We first prove this for $k = 2$ and then extend it to any $k > 2$.

4.2.1. Sliding 2-Transmitters

Again, we use a reduction from Minimum Vertex Cover in a planar graph G. This time we can allow arbitrary degrees in the graph, but we require that G is 2-connected; the problem is NP-hard under these restrictions (see e.g. [6]).

Given a planar 2-connected graph G with n vertices and m edges, we first compute its bar visibility representation similar to Section 4.1. We may move vertex-bars up and down slightly as needed so that all vertex-bars have distinct y-coordinates. As before we can make edge-strips thin enough such that no two of them have overlapping x-range. Since the graph is 2-connected, the construction in [33] guarantees that all vertices except the bottommost one have a neighbor below, and all vertices except the topmost one have a neighbor above.

Gadgets. We start by thickening each vertex-bar into a box, and place three copies of this box above each other with the same x-range. These three boxes are connected to each other by *channels*, which are thin vertical corridors (thin enough so that their x-range is strictly within that of the vertex-box, and does not intersect an edge-strip). We place these two channels at opposite ends of the vertex-boxes, resulting in a Z-shape or an S-shape (the choice between the two is arbitrary for now, but will be determined later). We call the result a *vertex-gadget*; see Figure 5. By making the height of boxes small enough, we may assume that no two vertex-gadgets have overlapping y-range.

For each edge e, the *edge-gadget* of e is a small axis-aligned box placed strictly within the strip representing e in such a way that its y-range intersects no y-range of another (vertex- or edge-) gadget. See Figure 5. Notice that from any edge-gadget there are vertical lines-of-sight to the vertex-gadgets of the endpoints of the edge.

The reduction. Let P' be the polygon obtained by replacing all vertex-bars and edge-strips with these gadgets. P' is y-monotone (i.e., any horizontal line intersects it in one interval), but not connected (for now we allow the polygon to be disconnected, but we will discuss the modifications to make it connected later). Since no y-ranges overlap, one can easily verify that vertical 2-transmitters are never required.

Observation 2. Any vertical sliding 2-transmitter in P' can be replaced by a horizontal sliding 2-transmitter that guards at least as much.

See also Figure 5. We call the three boxes of a vertex-gadget the *top, middle* and *bottom* box, and also use *outer boxes* to mean the top and bottom box.

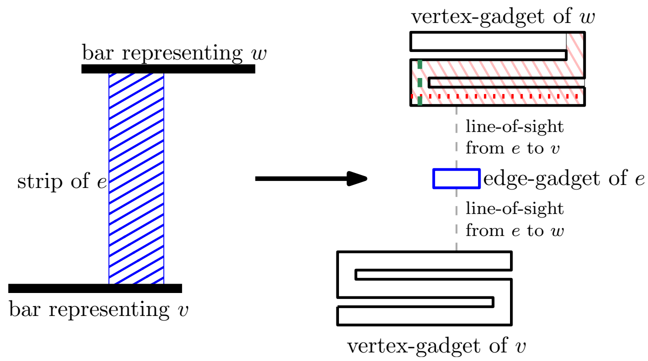

Figure 5: Vertex- and edge-gadgets. The pink (falling pattern) region is guarded by the red (dotted) horizontal 2-transmitter. Note that it includes everything that the green (dashed) vertical 2-transmitter can see.

Lemma 10. *For any set S of horizontal sliding 2-transmitters that guard P' entirely, there exists a set S' of horizontal sliding 2-transmitters that guard P' entirely such that $|S'| \leq |S|$ and no sliding 2-transmitter of S' is inside an edge-gadget.*

Proof. Let $s \in S$ be a sliding 2-transmitter that lies in an edge-gadget B corresponding to edge $e = (v, w)$. After possible renaming, assume that the vertex-gadget corresponding to v is below e and w is above e.

Assume first that one of v, w (say v) has a horizontal sliding 2-transmitter s' in the outer box facing e. After possibly extending s' we may assume that it spans the entire outer box of v. Since the x-range of B is within the x-range of v, s' sees everything that s saw and that was below s. So we can replace s by a sliding 2-transmitter in the outer box of w facing e, and this can only increase the guarded region.

Now, assume that neither v nor w has a horizontal sliding 2-transmitter in the outer box facing e. Consider a point p in the top box of v that is just outside the x-range of B, but still within the x-range of w. The only horizontal sliding 2-transmitters that could guard p are in the bottom box of w or in the middle box of v. By assumption we therefore have a sliding 2-transmitter in the middle box of v. Likewise w must have a sliding 2-transmitter in the middle box of w. We can thus move the sliding 2-transmitter in B to the bottom box of w without decreasing the guarded region. □

Lemma 11. *Let S be a set of horizontal sliding 2-transmitters that guard P' entirely and that do not lie inside edge-gadgets. Then, for any vertex v, there must be at least one sliding 2-transmitter inside the vertex-gadget of v. If there is exactly one such sliding 2-transmitter, then it must be in the middle box of v.*

Proof. Pick a point p in the middle box of v that is not in the x-range of the channels. Let s be a horizontal sliding 2-transmitter that guards p. Then s must be in one of the three boxes of v.

Assume now that exactly one sliding 2-transmitter is inside the vertex-gadget of v, and it is not in the middle box. Say the sliding 2-transmitter is in the bottom box. If v has any neighbor w above, then let p be a point in the top box of v and in the same x-range as the edge-gadget of (v, w). To guard p, we need either a sliding 2-transmitter in the edge-gadget (which was excluded) or in the top or middle box of v (which was also excluded). So v cannot have any neighbor above. By construction that means that v is the topmost of all vertices. To guard the top box of v, we then must have a sliding 2-transmitter in the top or middle box of v; again contradiction. □

Lemma 12. *The following statements are equivalent: (i) G has a vertex cover of size k, (ii) P' can be guarded by $n+k$ sliding 2-transmitters, and (iii) P' can be guarded by $n+k$ horizontal sliding 2-transmitters.*

Proof. Given a vertex cover C of G, we place horizontal transmitters as follows: if $v \in C$, then place a maximal horizontal sliding 2-transmitter in both outer boxes of v, else place a maximal horizontal sliding 2-transmitter in the middle box of v. Clearly, we have $n + |C|$ sliding 2-transmitters and every vertex-gadget is guarded. For every edge e, one endpoint v is in C, and hence both bottom and top box of v contain sliding 2-transmitters. The one in the outer box of v that faces e then guards the edge-gadget of e.

Vice versa, assume that set S of sliding 2-transmitters guards P'. By the above results, we may assume that they are all horizontal and none are in an edge-gadget. Define C to be all those vertices whose vertex-gadgets are intersected by at least two sliding 2-transmitters. Since every vertex-gadget intersects at least one sliding 2-transmitter we have $|C| \leq |S| - n$. For every edge (v, w), the edge-gadget must be guarded by a sliding 2-transmitter that is in an outer box of v or w, say v. Then v must contain at least two sliding 2-transmitters by Lemma 11, so $v \in C$. Hence, C is a vertex cover. □

Connecting the polygon. We now explain how to make the polygon connected while staying monotone. Let $g_1, \dots, g_{m+n}$ be the gadgets in P', sorted in bottom-to-top order (since y-ranges are disjoint, this is well-defined). The idea is to connect each g_i to g_{i+1} using a *connector-gadget*. This is an S-shaped or Z-shaped gadget much like a vertex-gadget, except that the top and bottom box both add a zig-zag near the end. Also, one of the channels has flexible height, so that the connector-gadget can have arbitrary height. We attach the ends of the connector-gadget C to corners of g_i and g_{i+1}. Figure 6 shows how to do this if the x-range of C is disjoint (except at the ends) from the ones of g_i and g_{i+1}, and the inset in Figure 7 shows how to do this if C shares x-range with them (in case of which we push the zig-zag to the very end to avoid overlap.)

However, we cannot connect consecutive gadgets if the connector-gadget would cross a line-of-sight. To avoid doing this, we will subdivide edges.

Observation 3 (Folklore). If G^s results from graph G by subdividing one edge twice, then G has a vertex cover of size k if and only if G^s has a vertex cover of size $k+1$.

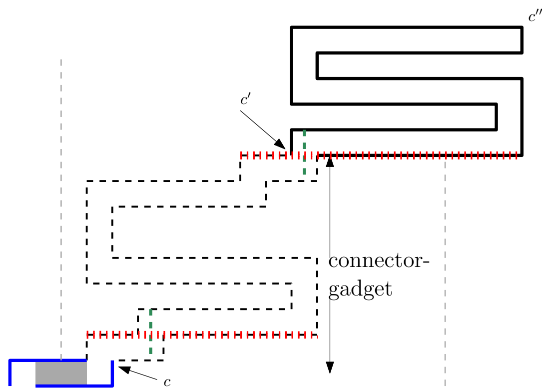

Figure 6: Connecting an edge-gadget to a vertex-gadget if there is no line of sight between them. We again show how some vertical transmitters can be replaced by horizontal transmitters.

We proceed as follows. First "parse" the bottommost gadget g_1: use an S-shape for it and fix as current corner its top right corner. Assume now we have parsed gadget g_i already, and fixed one top corner c of it as current corner. Let g_{i+1} be the next gadget above g_i. Considering its two bottom corners, we choose the corner c' so that $\overline{cc'}$ crosses as few lines-of-sight as possible. If line segment $\overline{cc'}$ crosses no line-of-sight, then attach a connector-gadget between c and c', using as shape (i.e., S or Z) the one that has c and c' at its ends. Let c'' be the diagonally opposite corner from c' in gadget g_{i+1}. If g_{i+1} is a vertex-gadget, then use as shape (i.e., S or Z) for g_{i+1} the one that has c' and c'' at its ends. This finishes (in this case) the parsing of gadget g_{i+1}, and we continue to connect to the next gadget with current corner c''.

Now, assume that $\overline{cc'}$ crosses some lines-of-sight, say $l_1, \ldots, l_\ell$ in order from c to c'. For all j, line-of-sight l_j represents an edge e_j. Add two new vertex-gadgets and two new edge-gadgets that we place along l_j, in the y-range between g_i and g_{i+1}. We make their height small enough and move them up and down suitably (while staying between g_i and g_{i+1}), so that all their y-ranges are disjoint and the ones of l_j are below the ones of l_{j+1} for all j. We order the new gadgets such that along line-of-sight l_j we alternate between vertex-gadgets and edge-gadgets, so e.g. if l_j had a vertex-gadget at the bottom-end and an edge-gadget at the top-end, then we insert, from bottom-to-top, an edge-gadget, a vertex-gadget, an edge-gadget and a vertex-gadget. Effectively this corresponds to having subdivided edge e_j twice.

All these gadgets can be connected with line segments that do not cross a line-of-sight. See Figure 7. We can hence connect all these gadgets as explained above. The only difference is that the next current corner c'' must be chosen to be the end of the line segment connecting to the next gadget. Normally, c'' will again be diagonally opposite from the previous corner c', but there is one exception per set of gadgets added for subdivisions. With that we have

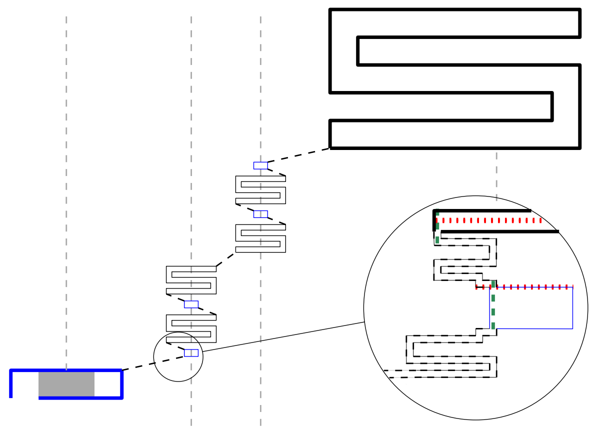

Figure 7: Connecting an edge-gadget to a vertex-gadget if there are lines of sight between them.

connected to g_{i+1}, and we repeat from there (after choosing its shape and the current corner as before).

Reduction revisited. Let N_s be the total number of times that we crossed some line-of-sight while connecting two vertically consecutive gadgets. Thus we have placed $2N_s$ new vertex-gadgets, $2N_s$ new edge-gadgets, and the total number of connector-gadgets is $N_c = n{+}m{+}4N_s{-}1$. Let G' be the graph with $n' = n{+}2N_s$ vertices obtained from G by subdividing edges, where we subdivide an edge 2ℓ times if its two lines-of-sight were crossed ℓ times in total when connecting gadgets. We know that G has a vertex cover C of size k if and only if G' has a vertex cover C' of size $k + N_s$. From C', we can easily construct a set of $n' + N_c + |C'| = n + 3N_s + N_c + k$ 2-transmitters that guard the entire polygon by placing guards in the middle box of all vertex-gadgets and connector-gadgets, except that for any vertex in C' we use two guards in the outer boxes.

Now assume we have a set S of $n{+}3N_s{+}N_c{+}k$ 2-transmitters that guard the entire polygon. With the addition of connector-gadgets, Observation 2 (vertical 2-transmitters can be replaced by horizontal ones) is not as obvious anymore, but still holds as long as sliding 2-transmitters may run along polygon-edges. See Figure 6. With this, Lemma 10, Lemma 11, and the equivalent of Lemma 11 for connector-gadgets, also hold. Using this, we can as before extract from S a vertex cover C' of G' size $|S| - n' - N_c$, and from it, a vertex cover C of G of size $|C'| - N_s = |S| - n - N_c - 3N_s = k$.

It remains to argue that the reduction has polynomial size, i.e., to bound N_s. We started with $n + m \in O(n)$ vertex-gadgets and edge-gadgets and $2m \in O(n)$ lines-of-sight. Connecting two gadgets hence contributes $O(n)$ to N_s, and therefore $N_s \in O(n^2)$ is polynomial. With that, the reduction is complete for $k = 2$. Note that the constructed polygon is connected and y-monotone (and therefore simple).

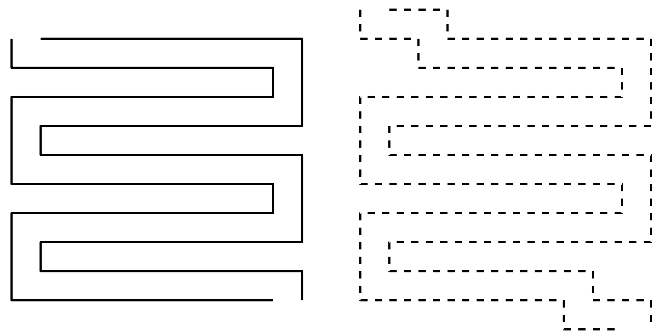

Figure 8: Vertex- and connector-gadget for $k = 4$.

4.2.2. Sliding k-Transmitters for $k > 2$

We now generalize to sliding k-transmitters for any fixed $k > 2$. The reduction is exactly the same as before, with the exception of the definition of vertex-gadgets and connector-gadgets.

The vertex-gadget now consists of $k + 1$ copies of the thickened bar in the visibility representation (earlier we had $3 = 2 + 1$ copies). They are connected with k channels at alternate ends, resulting in a zig-zag line. The connector-gadget is a vertex-gadget with additional small zig-zags in the top and bottom box (possibly pushed towards the end). See Figure 8. We can verify that again vertical sliding k-transmitters are never better than horizontal ones. Define for a vertex-gadget the *middle* box to be the $(k/2+1)$st box (recall that k is even), and the *outer boxes* to be the top and bottom box as before. With that, the proofs of Lemmas 10 and 11 carry almost verbatim, and the reduction holds again. We hence conclude:

Theorem 13. *For any $k > 0$, guarding a polygon with the minimum set of sliding k-transmitters is* NP*-complete, even if (i) the polygon is a simple y-monotone orthogonal polygon, and (ii) only horizontal sliding k-transmitters are allowed.*

The constructed polygon is thin. In fact, every gadget is a *thickened path* obtained by sliding a unit square along an orthogonal path. With suitable rescaling, in fact the entire polygon can be made into a thickened path, with one exception: whenever we subdivide edges, we must (at one edge-gadget) attach both connecting gadgets on the same (left or right) side, hence have a "leg" sticking out (this could perhaps be called a *thickened caterpillar*). We suspect that the construction could be modified to become a thickened path, but have not been able to work out the details yet.

5. Art Gallery Theorems

We now consider art gallery theorems for the ST_k and HST_k problems; that is, we give bounds (depending on n) on the number of k-transmitters required to guard an orthogonal polygon P with n vertices.

5.1. Horizontal guards

We first study the HST_k problem where only horizontal guards are allowed. Figure 9 shows a polygon that requires $n/4$ horizontal sliding transmitters, simply because no horizontal sliding transmitter can see two of the points marked $+$. Notice that this result holds regardless of the value of k, because in this polygon (as well as in any x-monotone polygon) a horizontal sliding k-transmitter can see no more than a horizontal sliding 0-transmitter at the same location. We next show that this bound is tight.

Theorem 14. *Given an orthogonal polygon P with n vertices and any $k \geq 0$, $\lfloor n/4 \rfloor$ horizontal sliding k-transmitters are always sufficient and sometimes necessary to guard P entirely.*

Proof. Necessity was argued above. Sufficiency follows directly from a result by O'Rourke [29] that any orthogonal polygon can be cut into $\lfloor n/4 \rfloor$ rectangles and L-shapes, each of which can be trivially guarded by a single horizontal sliding 0-transmitter. □

5.2. Arbitrary sliding transmitters

Aggarwal showed a tight bound $\lfloor (3n + 4)/16 \rfloor$ for the number of mobile guards necessary and sufficient to guard P [1]. Closer inspection reveals that the lower bound construction (see Figure 9) actually works for sliding 0-transmitters, since no two of the $(3n+4)/16$ pixels marked with a cross can be guarded by one sliding 0-transmitter. The upper bound also works for sliding 0-transmitters; in fact, the proof of the upper bound for sliding 0-transmitters becomes almost identical to mobile guards and so we only very briefly review the approach taken in [1]. The idea is to first guard a small portion of P using one or two mobile guards, cutting the guarded region out of P, and then guarding the rest of P by an induction hypothesis. There are numerous cases, but in all of them one can establish that a sliding 0-transmitter would have achieved the same as the mobile guard used. So, we have the following result.

Theorem 15 (Based on [1]). *Given a simple orthogonal polygon P with n vertices, $\lfloor (3n+4)/16 \rfloor$ sliding 0-transmitters are always sufficient and sometimes necessary to guard P entirely.*

For sliding k-transmitters for $k > 0$, the lower-bound example does not hold, but the staircase from Figure 9 can be used to show that $n/6$ sliding k-transmitters are sometimes required, because no two points marked $\times$ can be seen by one k-transmitter.

Corollary 2. Given any integer $k > 0$, $\lfloor n/6 \rfloor$ sliding k-transmitters are sometimes necessary to guard a simple orthogonal polygon P with n vertices.

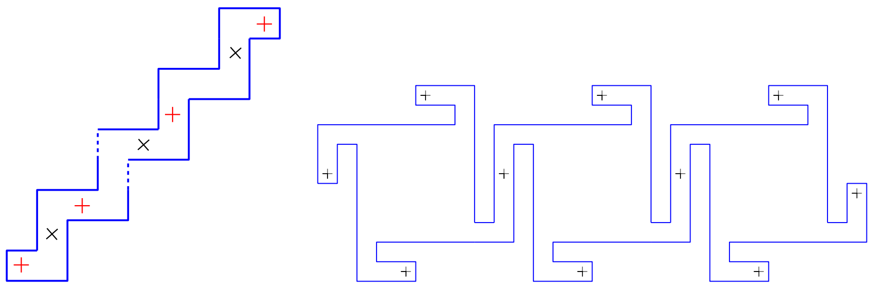

Figure 9: (Left) A polygon that requires $n/4$ horizontal sliding transmitters. (Right) A polygon that requires $(3n+4)/16$ sliding 0-transmitters.

5.3. Non-crossing sliding 0-transmitters

A natural variant of the art gallery problem might be to consider the problem under the restriction that no two sliding 0-transmitters can intersect each other. More precisely, given a simple orthogonal polygon P, we are interested in a result analogous to Theorem 15 such that the trajectories of no two sliding 0-transmitters would intersect each other; that is, the sliding 0-transmitters are pairwise *non-crossing*. Unfortunately, there are cases in the upper bound approach of Aggarwal (Theorem 15 and [1]) in which the trajectories of mobile guards intersect. Here, we show using a different approach that $\lfloor (n+1)/5 \rfloor$ non-crossing sliding 0-transmitters are always sufficient to guard a simple orthogonal polygon P with n vertices. To this end, we consider the dual graph of P and show how we can achieve the bound by guarding the subpolygons of P corresponding to a special subtree structure of the dual in a bottom-up fashion (see below for more details). While our bound is weaker than the one by Aggarwal, our proof gives an alternate (and arguably simpler) approach to prove such a bound.

We first show this result under the following *general position* assumption: no two vertical edges can be connected by a vertical line segment that lies inside the polygon. We will discuss how to remove this assumption at the end of this section.

5.3.1. Splitting the tree of the segmentation

Recall from Section 2 the dual graph of the pixelation of a polygon P. In this section, we again consider a dual graph, but this time that of only one of the segmentations of P. Thus, consider the vertical segmentation obtained after extending vertical rays from all reflex vertices. By the general position assumption, no such ray hits the boundary at a reflex vertex. Interpret the segmentation as a planar graph, and let G be its weak dual graph obtained by defining a vertex for every vertical rectangle and connecting two rectangles if and only if they share (part of) a side. If P is simple, then this dual graph is a tree T; we know that $|T| = n/2 - 1$ since P is in general position [31]. Let $R(w)$ be the rectangle of the vertical segmentation that corresponds to a node w in T, and let `rangeY`(w) be the range of y-coordinates of $R(w)$. We say that two

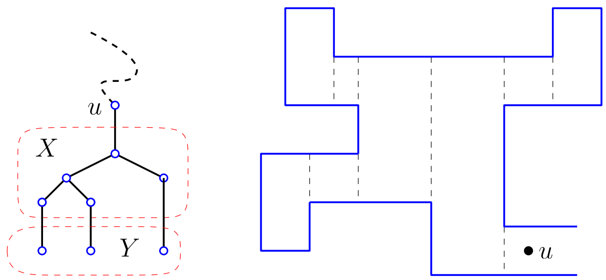

Figure 10: An example of tree $T(u)$.

nodes v, w *share the top* (respectively *share the bottom*) if the maximal (minimal) y-coordinates of $R(v)$ and $R(w)$ are the same. In what follows, for ease of notation we will sometimes say that "a 0-transmitter guards a node/subtree" when we mean to say that "a 0-transmitter guards the rectangle/subpolygon corresponding to the nodes/subtree". In the following, we show that $\lfloor 2/5 \cdot |T| + 3/5 \rfloor$ non-crossing sliding 0-transmitters are sufficient to guard P entirely, which therefore gives the desired bound.

We partition T into a set of disjoint subtrees as follows. Root T at a leaf. Let u be the lowest node in T that has degree two (i.e., u has only one child) and u is not the parent of a leaf. Let $T(u)$ be the subtree rooted at u, and partition $T - T(u)$ recursively. Let T_0 be the tree remaining in the base case (when no such u exists). T_0 may have just a single node; this will be treated separately. Any other subtree has the form $T(u)$ for some node u and at least 3 nodes, and we will argue now that we can guard it with at most $2/5 \cdot |T(u)|$ 0-transmitters.

5.3.2. Guarding one subtree: observations

To guard one such $T(u)$, we consider it to consist of the following components (see also Figure 10): (i) vertex u is the root of $T(u)$, (ii) let Y be all those leaves of $T(u)$ whose parent have only one child, and set $y = |Y|$, (iii) let $X = T(u) - \{u\} - Y$ and set $x = |X|$. Since u had only one child and Y consists of leaves, X forms a tree. By choice of u and Y, no interior node of X has degree 2. We next show that X can have at most one vertex of degree 4 (see Lemma 20). Before that, we first need a few observations.

Lemma 16. *Let v and w be two adjacent nodes of T. Then $\texttt{rangeY}(v) \subset \texttt{rangeY}(w)$ or $\texttt{rangeY}(w) \subset \texttt{rangeY}(v)$. Further, v and w share the top or the bottom.*

Proof. The vertical line segment $R(v) \cap R(w)$ must contain exactly one reflex vertex, for it must contain at least one (because it is a line of the vertical segmentation), and it cannot contain two (else P would not be in general position). Up to symmetry and renaming, we may assume that the reflex vertex is at the bottom end of $R(v) \cap R(w)$, with the horizontal edge incident to $R(v)$. See also Figure 10. Then the bottom side of $R(v)$ is above the bottom side of $R(w)$. The top side of $R(v)$ must be aligned with the top side of $R(w)$ (since there is no vertex here by general position), so $\texttt{rangeY}(v) \subset \texttt{rangeY}(w)$ and v and w share the top. □

Lemma 17. *If v is adjacent to w_1 and w_2 in such a way that $R(w_1)$ and $R(w_2)$ are attached on the same side (left or right) of $R(v)$, then $\texttt{rangeY}(w_1) \subset \texttt{rangeY}(v)$ and $\texttt{rangeY}(w_2) \subset \texttt{rangeY}(v)$.*

Proof. If instead we had (say) $\texttt{rangeY}(w_1) \supset \texttt{rangeY}(v)$, then $R(w_1)$ would occupy the entire region to the side of $R(v)$. This makes it impossible for $R(w_2)$ to attach on the same side without overlap. □

Lemma 18. *T is a maximum degree-4 tree, and for any node v of degree 4, exactly two rectangles of neighbors attach on each side (left or right) of $R(v)$.*

Proof. If T had a node v of degree 5 or more, or if v has degree 4 but on one side of $R(v)$ at most one rectangle attaches, then on some side (left or right) of $R(v)$ there must be at least three adjacent rectangles $R(w_1), R(w_2), R(w_3)$. Each of w_1, w_2, w_3 shares the top or the bottom with v by Lemma 16. This is impossible without overlap among $R(w_1), R(w_2), R(w_3)$. □

Lemma 19. *Let $w_0, w_1, \ldots, w_k$ be a path such that $\texttt{rangeY}(w_0) \supset \texttt{rangeY}(w_1)$ and none of $w_1, \ldots, w_{k-1}$ has degree 2. Then, $\texttt{rangeY}(w_0) \supset \texttt{rangeY}(w_1) \supset \cdots \supset \texttt{rangeY}(w_k)$.*

Proof. We proceed by induction; the base case ($\texttt{rangeY}(w_0) \supset \texttt{rangeY}(w_1)$) holds by assumption. Now assume that we know $\texttt{rangeY}(w_{i-1}) \supset \texttt{rangeY}(w_i)$ for some $i < k$. Since w_i has degree 3 or more, but one of the sides (left or right) of $R(w_i)$ is a subset of $R(w_{i-1})$, rectangle $R(w_{i+1})$ must attach at a side of $R(w_i)$ that contains two neighbors of w_i. By Lemma 17, we have $\texttt{rangeY}(w_i) \supset \texttt{rangeY}(w_{i+1})$. □

Lemma 20. *Let v_1, v_2 be two degree-4 nodes in T. Then, there exists a node of degree 2 on the (unique) path between v_1 and v_2 in T.*

Proof. Assume not, so all interior nodes of the path $v_1{=}w_0, w_1, \ldots, w_k{=}v_2$ between them have degree at least 3. Since v_1 has degree 4, both sides of $R(v_1)$ must be incident to exactly two other rectangles. By Lemma 17, hence $\texttt{rangeY}(w_0) \supset \texttt{rangeY}(w_1)$. By Lemma 19 then $\texttt{rangeY}(v_1) \supset \texttt{rangeY}(v_2)$. But, by the same reasoning applied to the reverse path also $\texttt{rangeY}(v_2) \supset \texttt{rangeY}(v_1)$. Contradiction. □

By Lemma 20, X can have at most one degree-4 vertex, which means that X forms a tree that is a rooted binary tree except that one node may have three children. Thus, X has at most $x/2 + 1$ leaves, and $y \leq x/2 + 1$.

Lemma 21. *Let c be the node of $X \cup \{u\}$ whose corresponding rectangle $R(c)$ has the maximum height. Let s be a maximal vertical line segment inside $R(c)$. Then s guards all rectangles corresponding to $X \cup \{u\}$.*

Proof. Let w be an arbitrary node in $X \cup \{u\}$ and consider the path $c{=}w_0, w_1, \ldots, w_k{=}w$ from c to w inside $X \cup \{u\}$. By choice of c and Lemma 16, we know $\texttt{rangeY}(c) \supset \texttt{rangeY}(w_1)$. By Lemma 19, we have $\texttt{rangeY}(c) \supset \texttt{rangeY}(w_1) \supset \cdots \supset \texttt{rangeY}(w_k)$ since $w_1, \ldots, w_{k-1}$ have degree at least 3 by choice of X. Therefore, s guards all of $R(w_1), \ldots, R(w_k)$ and in particular $R(w)$. $\square$

By Lemma 21 and the fact that Y can be guarded by at most $y = |Y|$ vertical 0-transmitters, we have the following.

Corollary 3. $T(u)$ can be guarded by at most $y{+}1$ vertical sliding 0-transmitters.

Guarding a subtree $T(u)$. Let $T(u)$ be a subtree obtained by partitioning T and split it into $T(u) = \{u\} \cup X \cup Y$ as before, where $x = |X|$ and $y = |Y|$. By Corollary 3, $y + 1$ 0-transmitters are sufficient to guard $T(u)$. We now describe some conditions under which y guards suffice; these will handle most of the cases.

Lemma 22. *If there exists a node $\ell \in Y$ such that $\texttt{rangeY}(\ell) \subset \texttt{rangeY}(p)$, where $p \in X$ is the parent of ℓ, then y vertical sliding 0-transmitters suffice to guard $T(u)$.*

Proof. Consider the 0-transmitter s inside the maximum-height rectangle $R(c)$ from Lemma 21. The proof of this lemma showed that $\texttt{rangeY}(c) \supset \texttt{rangeY}(p)$, and by assumption therefore $\texttt{rangeY}(c) \supset \texttt{rangeY}(\ell)$. Therefore, 0-transmitter s guards $R(\ell)$ as well. For each other node $\ell' \neq \ell \in Y$, add one more 0-transmitter inside $R(\ell')$ to guard it. This gives a total of $1 + (y - 1) = y$ non-crossing 0-transmitters as desired. $\square$

Lemma 23. *If $y = x/2 + 1$, then y non-crossing sliding 0-transmitters are sufficient to guard $T(u)$ entirely.*

Proof. We are done if Lemma 22 applies, so assume this is not the case. If y is the maximum possible value $x/2 + 1$, then X must contain a node of degree 4, and every leaf of X must have a node of Y attached to it. Let v be the node of degree 4. By Lemmas 17-19 we know that $\texttt{rangeY}(v) \supset \texttt{rangeY}(w)$ for all $w \in X \cup \{u\}$, and so v is the node c of Lemma 21, and the maximal vertical line segment s inside $R(v)$ guards $X \cup \{u\}$. We will now show that there exists one other 0-transmitter s' that guards at least two nodes in Y. This guard, together with s and with (at most) $y - 2$ other guards for all other nodes of Y satisfies the claim after retracting s a bit so that it does not cross s'. See also Figure 11.

By Lemma 18, we know that $R(v)$ has two neighbors on its right side, and one of them, say w_1, shares the top with v. Now either $w_1 = u$, or w_1 is a leaf of X, or w_1 has degree 3. In the last case, we proceed as in the induction step of Lemma 19 to see that $R(w_1)$ must have two neighbors on its right side, and one of them, say w_2, shares the top with w_1 and v. Continue this argument until we have a chain $w_1, \ldots, w_k$ of nodes that all share the top with v and that are to the right of $R(v)$; see Figure 11 for an illustration. Further, w_k is either u or a leaf

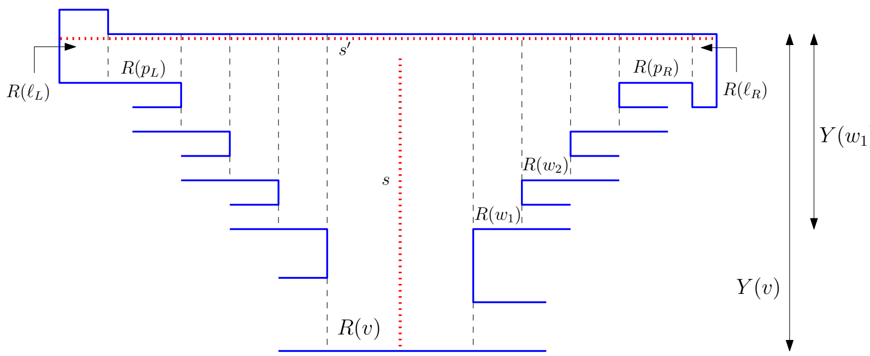

Figure 11: If v has degree 4, then 0-transmitter s is located inside $R(v)$, and two leaves in its subtree can be guarded by one 0-transmitter s'.

of X. In the same manner we can find a chain of nodes that share the top with v and are left of $R(v)$, and two more chains of rectangles that share the bottom with v and are left respectively right of $R(v)$. Only one of these four chains can end with u, so (up to symmetry) we may assume that neither chain at the top contains u. Combining the two chains that share the top with v, we obtain a path of nodes between two leaves p_L and p_R of X and including v, that all share the top with v. By $y = x/2+1$ both p_L and p_R are incident to nodes ℓ_L and ℓ_R in Y, and since Lemma 22 does not apply, we have $\texttt{rangeY}(\ell_L) \supset \texttt{rangeY}(p_L)$ and $\texttt{rangeY}(\ell_R) \supset \texttt{rangeY}(p_R)$. Therefore, a horizontal 0-transmitter s' along the top of $R(v)$ (and extended as far as possible) guards both $R(\ell_L)$ and $R(\ell_R)$ as desired (see Figure 11). □

Lemma 24. *If $y = (x+1)/2$ and $x > 1$, then y non-crossing sliding 0-transmitters are sufficient to guard $T(u)$ entirely.*

Proof. We are done if Lemma 22 applies, so assume this is not the case. Since $y = (x+1)/2$ and $x > 1$, all nodes of X have degree 3 and every leaf of X is incident to a node of Y. Let r be the unique child of u, hence the root of X. Let c and s be as in Lemma 21, i.e., $R(c)$ has maximum height among $X \cup \{u\}$, and s is a maximal vertical line segment inside $R(c)$. We distinguish cases:

1. There exists a leaf p of X such that some horizontal segment s' inside P intersects both $R(p)$ and $R(u)$: We know p is adjacent to some $\ell \in Y$. We also know $\texttt{rangeY}(\ell) \supset \texttt{rangeY}(p)$ since Lemma 22 does not apply. Hence extending the horizontal segment s' to the maximum inside P guards ℓ as well as u. For any other node $\ell' \neq \ell \in Y$, place a horizontal 0-transmitter inside the parent of ℓ' and extend it to the maximum. We claim that these y 0-transmitters suffice. All nodes of $Y \cup \{u\}$ are guarded by construction. Now consider any node $w \in X$, and observe that by degree 3 one side of $R(w)$ is incident to two neighbors of w, hence at least one child w' of w. Using Lemma 19, hence some leaf p in the subtree of w' satisfies $\texttt{rangeY}(p) \subset \texttt{rangeY}(w)$. The horizontal 0-transmitter placed inside p

therefore guards $R(w)$. Since all 0-transmitters are horizontal, they do not cross.

2. $c = u$: By choice of c, then $\texttt{rangeY}(u) \supset \texttt{rangeY}(r)$. By Lemma 19, then $\texttt{rangeY}(u) \supset \texttt{rangeY}(w)$ for all nodes $w \in X$, and the previous case applies for some (in fact, all) leaves of X.

3. $c = p$ for some leaf p of X: Let ℓ be the node in Y adjacent to p that exists by $y = (x+1)/2$. We know $\texttt{rangeY}(p) \subset \texttt{rangeY}(\ell)$ since Lemma 22 does not apply. Therefore a vertical 0-transmitter inside $R(\ell)$ would guard everything that s guards, and also guard $R(\ell)$. Using this 0-transmitter, plus one vertical 0-transmitter for each other node in Y, gives the required y 0-transmitters.

4. The remaining case is that c is an interior node of X, hence has degree 3. Up to symmetry, we may assume two of the neighbors share the top with c. By choice of c, $\texttt{rangeY}(c) \supset \texttt{rangeY}(w)$ for all neighbors w of c. As in the proof of Lemma 23, we can find a path between two nodes p_L and p_R that are either u or leaves of X such that all nodes on the path share the top with c. If one of p_L, p_R is u, then Case 1 applies. Otherwise, as in the proof of Lemma 23, we can guard the two nodes of Y attached to p_L and p_R with one horizontal 0-transmitter s', and this 0-transmitter, together with s (retracted a bit) and one 0-transmitter for all other nodes of Y, gives the desired set of y 0-transmitters.

□

5.3.3. Guarding one subtree: the cases

We can now show that any subtree $T(u)$ can be guarded with $2/5 \cdot |T(u)|$ 0-transmitters by distinguishing by x and y.

- Assume first that x is even and $y = x/2+1$. We can use y 0-transmitters to guards $x+y+1$ rectangles. Since $(x/2+1)/(x+x/2+2) \leq 2/5$ for $x \geq 2$, we achieve the desired ratio in this case.
- Assume next that x is even and $y \leq x/2$. We use $y+1$ 0-transmitters to guards $x+y+1$ rectangles. Since $y \leq x/2 \leq \frac{2}{3}x - 1$ if $x \geq 6$, we have $(y+1)/(x+y+1) \leq 2/5$ for $x \geq 6$. The cases $x = 4$ and $x = 2$ will be treated separately below.
- Assume next that $x > 1$ is odd and $y = (x+1)/2$. We can use y 0-transmitters to guards $x+y+1$ rectangles. Since $((x+1)/2)/(x+(x+1)/2+1) \leq 2/5$ for $x > 1$, we achieve the desired ratio in this case.
- Assume next that x is odd and $y \leq (x-1)/2$. We use $y+1$ 0-transmitters to guards $x+y+1$ rectangles. Since $y \leq (x-1)/2 \leq \frac{2}{3}x - 1$ for $x \geq 3$, we have $((x-1)/2+1)/(x+(x-1)/2+1) \leq 2/5$ for $x \geq 3$.

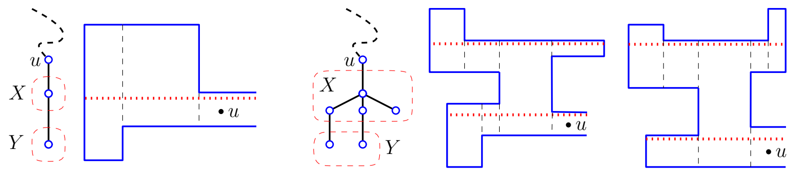

Figure 12: The tree and polygon(s) for (left) $x = y = 1$ and (right) $x = 4$ and $y = 2$.

- This leaves three cases: $x = 1$, $x = 2$ with $y \leq 1$, and $x = 4$ with $y \leq 2$. We first treat $x = 2$, which is the easiest. If $x = 2$, then the root of X has a single child, hence it has degree 2. By choice of u this means that $y = 0$. We use $y + 1 = 1$ 0-transmitter to guard 3 rectangles, which is a ratio of $1/3 < 2/5$.

- If $x = 1$, then $y \leq 1$. We cannot have $y = 0$, else the unique node in X would be a leaf and hence u would be a parent of a leaf, contradicting the choice of u. Hence $y = 1$. If Lemma 22 applies, then we can guard $T(u)$ with $1 < 2/5 \cdot |T(u)|$ 0-transmitters since $|T(u)| = 3$. Else we have $\texttt{rangeY}(\ell) \supset \texttt{rangeY}(r)$ (where ℓ and r are the unique nodes of Y and X) and a horizontal 0-transmitter through $R(r) \cap R(u)$ will guard everything. See also Figure 12.

- $x = 4$: If $y \leq 1$ then $y < \frac{2}{3}x - 1$, which as above gives a ratio of $< 2/5$. So assume $y = 2$. If Lemma 22 applies, then we can guard $T(u)$ with $2 < 2/5 \cdot |T(u)|$ 0-transmitters since $|T(u)| = 7$. Else, the polygon has a very restricted structure: We can have $x = 4$ only if the root r of X is a node of degree 4. At each of the four corners of $R(r)$ we have an incident rectangle with smaller y-range, and at two of them a rectangle with larger y-range (belonging to nodes in Y) attaches. See also Figure 12. Two horizontal 0-transmitters near the top and bottom of $R(r)$ will hence guard everything.

5.3.4. Results for general positions

Putting everything together, we have that every polygon can be guarded with $\lfloor (n + 1)/5 \rfloor$ sliding 0-transmitters if the polygon is in general position. In order to remove the general-position-assumption, we need a slightly stronger result that also restricts the shape and location of the 0-transmitters.

Lemma 25. *Given a simple orthogonal polygon P in general position with n vertices, $\lfloor (n+1)/5 \rfloor$ sliding 0-transmitters are always sufficient to guard P such that*

- *no two sliding 0-transmitters intersect each other, and*

- *any sliding 0-transmitter γ is strictly inside P in the sense that γ intersects the boundary of P only at its endpoints, if at all.*

Further, P can either be guarded with $\lfloor n/5 \rfloor$ such sliding 0-transmitters, or it can be guarded with $\lfloor (n+1)/5 \rfloor$ vertical sliding 0-transmitters.

PROOF. Recall that we split T into trees $T_0, T_1, \ldots, T_k$ for some $k \geq 0$, where T_0 contains the root of T. We guard T_i for $i > 0$ with at most $2|T_i|/5$ sliding 0-transmitters, and following the proofs, can see that all are strictly inside P. If $|T_0| \geq 3$, then likewise it is guarded with $2|T_0|/5$ sliding 0-transmitters. If $|T_0| = 2$, then it is guarded with $1 = \frac{2}{5}|T_0| + \frac{1}{5}$ sliding 0-transmitters. Otherwise ($|T_0| = 1$), we can guard it with $1 = \frac{2}{5}|T_0| + \frac{3}{5}$ sliding 0-transmitter, which we can choose to be strictly inside P. So, in total we use at most $\frac{2}{5}|T| + \frac{3}{5} = \frac{2}{5}(n/2-1) + \frac{3}{5} = \frac{n}{5} + \frac{1}{5}$ sliding 0-transmitter, and the upper bound holds by integrality.

Now, assume that we use exactly $(n+1)/5$ transmitters. This can happen only if we use exactly $\frac{2}{5}|T_i|$ transmitters for $i > 0$, and exactly $\frac{2}{5}|T_0| + \frac{3}{5}$ transmitters for T_0. This implies $|T_0| = 1$, and we use a vertical transmitter to guard T_0. It also implies that each T_i (for $i > 0$) has (x, y) in the set $\{(2,2), (6,3), (3,1)\}$ (in all other cases, some bound in the above case analysis was not tight). But $(2,2)$ is impossible as argued above. In the other two cases, we use $y+1$ sliding 0-transmitters, which are vertical by Corollary 3. $\square$

5.3.5. Removing the general position assumption

We now show that Lemma 25 holds even without the general position assumption. We could do this by applying the standard perturbation on the vertical edges having the same x-coordinate and then arguing that the resulting polygon can still be guarded using the same number of transmitters. However, such a perturbation creates new slices, so the dual tree changes, and it is non-trivial to argue how to transform a set of transmitters in the perturbed polygon to a set in the original polygon. For this reason, we prove the result using a different (and, in our opinion, cleaner) approach that does not move edges, but instead breaks the polygon into pieces, guards each separately, and combines guarding transmitter sets.

We say that a segment of the vertical segmentation of P is *critical* if both of its endpoints are reflex vertices of P; let M_P denote the set of all critical segments of the vertical segmentation of P. We prove the result (i.e., Lemma 25 with the general position assumption removed) by induction on $|M_P|$. The case $|M_P| = 0$ is trivial since then P is in general position, so assume $|M_P| > 0$. Take any critical segment s of P and let P_1 and P_2 be the subpolygons of P whose interior around s lies to the left and right of s, respectively. For $i = 1, 2$, since P_i has fewer critical segments than P, it can be guarded using $\lfloor (n_i+1)/5 \rfloor$ non-crossing sliding 0-transmitters by induction. If either of these polygons can actually be guarded with at most $n_i/5$ non-crossing sliding 0-transmitters, then we are done: combining the two sets of transmitters gives a set of at most $\lfloor (n+1)/5 \rfloor$ non-crossing sliding 0-transmitters since $n_1 + n_2 = n$.

So the only remaining case is when P_1 and P_2 both use exactly $(n_i+1)/5$ sliding 0-transmitters. In particular we then have (for $i = 1, 2$) $n_i = 5k_i + 4$

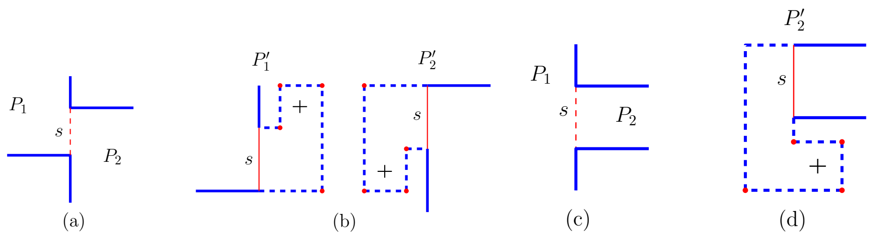

Figure 13: Removing critical segments. (a) and (b): Case 1. (c) and (d): Case 2.

for some integer $k_i \geq 0$, and P_i can be guarded with k_i vertical sliding 0-transmitters. Two cases arise:

Case 1. One of the horizontal edges at the ends of s belongs to P_1, the other belongs to P_2. Then, we modify P_1 and P_2 around s and create two new polygons P'_1 and P'_2; see Figure 13(a) and (b). For $i = 1, 2$, polygon P'_i has $n'_i = n_i + 4 = 5k_i + 8$ vertices. Since P'_i has fewer critical segments than P, it can be guarded with $\lfloor (n'_i + 1)/5 \rfloor = k_i + 1$ sliding 0-transmitters.

Consider the point marked + in P'_i. If it is guarded by a vertical sliding 0-transmitter γ, then we can replace γ by a horizontal sliding 0-transmitter that crosses s and guards at least as much. So we may assume that both guarding sets of P'_1 and P'_2 contain a horizontal sliding 0-transmitter that crosses s. Neither of them runs along the horizontal edge of P, since transmitters are strictly in the interior of polygons. Therefore we can obtain a guard set of P by combining the two guard sets of P'_1 and P'_2 and combining the two transmitters crossing s into one. This gives a set of $k_1 + k_2 + 1 = \lfloor n/5 \rfloor$ siding 0-transmitters since $n = n_1 + n_2 = 5(k_1 + k_2) + 8$. So the claim holds.

Case 2. Both horizontal edges at the ends of s belong to P_1, or both belong to P_2. Suppose w.l.o.g. that both horizontal edges at the ends of s belong to P_2; see Figure 13(a). Modify P_2 so as to obtain a new polygon P'_2 as shown in Figure 13(b). Note that P'_2 has $n'_2 = n_2 + 4 = 5k_2 + 8$ vertices and fewer critical segments than P, so it can be guarded using $\lfloor (n'_2 + 1)/5 \rfloor = k_2 + 1$ sliding 0-transmitters. Consider the point marked with + in P'_2. To guard this, there must exist a sliding 0-transmitter γ which cannot see anything of P_2 except points whose perpendicular onto s is in P_2.

Since P_1 was guarded by exactly $k_1 + 1 = (n_1 + 1)/5$ sliding 0-transmitters, we may assume that they are all vertical. In particular, segment s (in P_1) was guarded by some vertical sliding 0-transmitter γ' (or perhaps a set of such transmitters). Note that when re-combining P_1 and P_2, γ' can see all parts of P_2 that γ saw. As such, we can drop the transmitter γ from the combined set, and get a set of $k_1 + 1 + k_2 \leq \lfloor n/5 \rfloor$ sliding 0-transmitters. We can hence conclude:

Theorem 26. *Any simple orthogonal polygon P with n vertices can always be guarded using $\lfloor (n+1)/5 \rfloor$ sliding 0-transmitters such that no two of them intersect each other.*

6. Conclusion

In this paper, we studied the problem of guarding an orthogonal polygon with the minimum number of sliding k-transmitters. We gave a constant-factor approximation algorithm for this problem, which works even if the polygon has holes. Moreover, we showed two NP-hardness results for this problem. Furthermore, we gave art-gallery-theorem results bounding the number of sliding 0-transmitters that are always sufficient and sometimes required. The most interesting remaining question is whether guarding an orthogonal polygon with sliding 0-transmitters is polynomial-time solvable if the polygon has no holes.

References

[1] A. Aggarwal. *The Art Gallery Theorem: its variations, applications and algorithmic aspects.* PhD thesis, Johns Hopkins University, 1984. A summary can be found in [31], Chapter 3.

[2] O. Aichholzer, R. F. Monroy, D. Flores-Peñaloza, T. Hackl, J. Urrutia, and B. Vogtenhuber. Modem illumination of monotone polygons. *Comput. Geom.*, 68:101–118, 2018.

[3] D. Avis and G. T. Toussaint. An optimal algorithm for determining the visibility of a polygon from an edge. *IEEE Trans. Computers*, 30(12):910–914, 1981.

[4] B. Ballinger, N. Benbernou, P. Bose, M. Damian, E. D. Demaine, V. Dujmovic, R. Y. Flatland, F. Hurtado, J. Iacono, A. Lubiw, P. Morin, V. S. Adinolfi, D. L. Souvaine, and R. Uehara. Coverage with k-transmitters in the presence of obstacles. *J. Comb. Optim.*, 25(2):208–233, 2013.

[5] T. Biedl, T. M. Chan, S. Lee, S. Mehrabi, F. Montecchiani, and H. Vosoughpour. On guarding orthogonal polygons with sliding cameras. In S.-H. Poon et al., editor, *Proceedings of the 11th International Workshop on Algorithms and Computation (WALCOM 2017)*, volume 10167 of *Lecture Notes In Computer Science*, pages 1–12. Springer, 2017.

[6] T. Biedl, G. Kant, and M. Kaufmann. On triangulating planar graphs under the four-connectivity constraint. *Algorithmica*, 19(4):427–446, 1997.

[7] T. Biedl and S. Mehrabi. On r-Guarding Thin Orthogonal Polygons. In *27th International Symposium on Algorithms and Computation (ISAAC 2016)*, volume 64 of *LIPIcs*, pages 17:1–17:13, 2016.

[8] T. Biedl, S. Mehrabi, and Z. Yu. Sliding k-transmitters: Hardness and approximation. In *Proceedings of the 28th Canadian Conference on Computational Geometry (CCCG 2016), Vancouver, Canada*, 2016.

[9] H. Brönnimann and M. T. Goodrich. Almost optimal set covers in finite VC-dimension. *Discrete & Computational Geometry*, 14(4):463–479, 1995.

[10] S. Cannon, T. G. Fai, J. Iwerks, U. Leopold, and C. Schmidt. Combinatorics and complexity of guarding polygons with edge and point 2-transmitters. *Comput. Geom.*, 68:89–100, 2018.

[11] V. Chvátal. A combinatorial theorem in plane geometry. *Journal of Combinatorial Theory, Series B*, 18:39–41, 1975.

[12] K. L. Clarkson and K. R. Varadarajan. Improved approximation algorithms for geometric set cover. *Discrete & Computational Geometry*, 37(1):43–58, 2007.

[13] M. de Berg, S. Durocher, and S. Mehrabi. Guarding monotone art galleries with sliding cameras in linear time. *J. Discrete Algorithms*, 44:39–47, 2017.

[14] S. Durocher, O. Filtser, R. Fraser, A. D. Mehrabi, and S. Mehrabi. A (7/2)-approximation algorithm for guarding orthogonal art galleries with sliding cameras. In *Proceedings of Latin-American Symposium (LATIN 2014)*, volume 8392 of *LNCS*, pages 294–305, 2014.

[15] S. Durocher, O. Filtser, R. Fraser, A. D. Mehrabi, and S. Mehrabi. Guarding orthogonal art galleries with sliding cameras. *Comput. Geom.*, 65:12–26, 2017.

[16] S. Durocher and S. Mehrabi. Guarding orthogonal art galleries using sliding cameras: algorithmic and hardness results. In *Proceedings of Mathematical Foundations of Computer Science (MFCS 2013)*, volume 8087 of *LNCS*, pages 314–324, 2013.

[17] S. Eidenbenz, C. Stamm, and P. Widmayer. Inapproximability results for guarding polygons and terrains. *Algorithmica*, 31(1):79–113, 2001.

[18] M. R. Garey and D. S. Johnson. The rectilinear Steiner tree problem in NP-complete. *SIAM Journal of Applied Mathematics*, 32:826–834, 1977.

[19] F. Hoffmann, M. Kaufmann, and K. Kriegel. The art gallery theorem for polygons with holes. In *Proceedings of Foundations of Computer Science (FOCS 1991)*, pages 39–48, 1991.

[20] J. Kahn, M. M. Klawe, and D. J. Kleitman. Traditional galleries require fewer watchmen. *SIAM Journal on Algebraic Discrete Methods*, 4(2):194–206, 1983.

[21] M. J. Katz, J. S. B. Mitchell, and Y. Nir. Orthogonal segment stabbing. *Comput. Geom.*, 30(2):197–205, 2005.

[22] M. J. Katz and G. Morgenstern. Guarding orthogonal art galleries with sliding cameras. *Inter. J. Comp. Geom. & App.*, 21(2):241–250, 2011.

[23] D. G. Kirkpatrick. An $O(\lg\lg OPT)$-approximation algorithm for multi-guarding galleries. *Discrete & Computational Geometry*, 53(2):327–343, 2015.

[24] E. Krohn and B. J. Nilsson. Approximate guarding of monotone and rectilinear polygons. *Algorithmica*, 66(3):564–594, 2013.

[25] D. T. Lee and Arthur K. Lin. Computational complexity of art gallery problems. *IEEE Transactions on Information Theory*, 32(2):276–282, 1986.

[26] A. Lubiw. Decomposing polygonal regions into convex quadrilaterals. In *Proceedings of the ACM Symposium on Computational Geometry (SoCG 1985)*, pages 97–106, 1985.

[27] S. Mehrabi. *Geometric optimization problems on orthogonal polygons: hardness results and approximation algorithms.* PhD thesis, University of Manitoba, Winnipeg, Canada, August 2015.

[28] J. O'Rourke. The complexity of computing minimum convex covers for polygons. In *20th Allerton Conf. Commun. Control Comput.*, pages 75–84, 1982.

[29] J. O'Rourke. An alternate proof of the rectilinear art gallery theorem. *J. of Geometry*, 21:118–130, 1983.

[30] J. O'Rourke. Galleries need fewer mobile guards: a variation to Chvátal's theorem. *Geometriae Dedicata*, 14:273–283, 1983.

[31] J. O'Rourke. *Art Gallery Theorems and Algorithms.* The International Series of Monographs on Computer Science. Oxford University Press, New York, NY, 1987.

[32] D. Schuchardt and H. Hecker. Two NP-hard art-gallery problems for ortho-polygons. *Mathematical Logic Quarterly*, 41(2):261–267, 1995.

[33] R. Tamassia and I.G. Tollis. A unified approach to visibility representation of planar graphs. *Discrete & Computational Geometry*, 1:321–341, 1986.

[34] A. P. Tomás. Guarding thin orthogonal polygons is hard. In *Proceedings of Fundamentals of Computation Theory (FCT 2013)*, volume 8070 of *LNCS*, pages 305–316, 2013.

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