On Guarding Orthogonal Polygons with Sliding Cameras

Therese Biedl,[1*] Timothy M. Chan,[1*] Stephanie Lee,[1] Saeed Mehrabi,[1] Fabrizio Montecchiani,[2†] and Hamideh Vosoughpour[1]

[1] Cheriton School of Computer Science, University of Waterloo, Canada.
{biedl, tmchan}@cs.uwaterloo.ca, {s363lee, smehrabi, hvosough}@uwaterloo.ca
[2] Department of Engineering, University of Perugia, Perugia, Italy.
fabrizio.montecchiani@unipg.it

Abstract. A sliding camera inside an orthogonal polygon P is a point guard that travels back and forth along an orthogonal line segment γ in P. The sliding camera g can see a point p in P if the perpendicular from p onto γ is inside P. In this paper, we give the *first* constant-factor approximation algorithm for the problem of guarding P with the minimum number of sliding cameras. Next, we show that the sliding guards problem is linear-time solvable if the (suitably defined) dual graph of the polygon has bounded treewidth. On the other hand, we show that the problem is NP-hard on orthogonal polygons with holes even if only horizontal cameras are allowed. Finally, we study art gallery theorems for sliding cameras, thus, give upper and lower bounds in terms of the number of sliding cameras needed relative to the number of vertices n.

1 Introduction

Let P be a (not necessarily orthogonal) polygon with n vertices. The art gallery problem, posed by Victor Klee in 1973 [25], asks for the minimum number 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 [7] 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. [15] 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 [16, 23, 25] 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 [22], even on simple orthogonal polygons [28] or monotone polygons [21]. A number of results concerning approximation algorithms are also known [13, 20, 21].

* Research of the authors is supported by NSERC.

† Research of the author supported in part by the MIUR project AMANDA, prot. 2012C4E3KT_001. Research was done while the author was visiting the University of Waterloo.

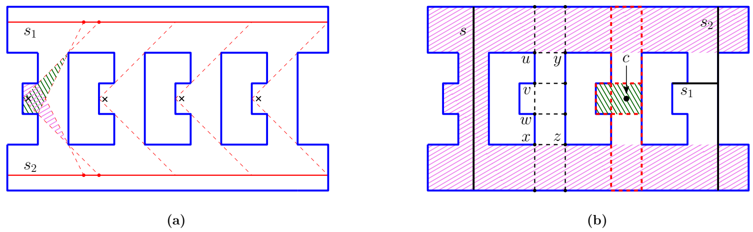

Fig. 1: (a) An orthogonal polygon P that can be guarded with two orthogonal mobile guards, but requires $\Theta(n)$ sliding cameras to be guarded since no two crosses can be seen by one sliding camera. (b) Sliding camera s sees the rising-shaded subpolygon of P. We also show parts of the pixelation induced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

Mobile Guards and Sliding Cameras. A *mobile guard* is a point guard that travels along a line segment γ inside P. Guard γ can see a point p in P if and only if there exists a point $g \in \gamma$ such that the line segment pg lies entirely inside P. If the line segment γ must be orthogonal, then we call it an *orthogonal mobile guard*. Moreover, if the line segment pg is required to be perpendicular to γ, then we call γ a *sliding camera*. Note that an orthogonal mobile guard travelling along γ may see a larger area of P than a sliding camera travelling along γ, see also Figure 1(a). The notion of mobile guards was introduced by Avis and Toussaint [2]. O'Rourke [26] proved that $\lfloor n/4 \rfloor$ (not necessarily orthogonal) 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].

In this paper we study the *Minimum Sliding Cameras (MSC)* problem, i.e., we want to guard an orthogonal polygon P with the minimum number of sliding cameras. We also consider the variant *Minimum Horizontal Sliding Cameras (MHSC)* where only horizontal cameras are allowed. These problems were introduced by Katz and Morgenstern [18], who proved that MHSC can be solved in polynomial time in the special case where the polygon is simple (has no holes). It was shown later that MSC is NP-hard in polygons with holes [12, 24]; NP-hardness in simple polygons is open. Durocher et al. [11] claimed a (3.5)-approximation algorithm for MSC 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 [18] gave a 2-approximation algorithm, which was later improved by De Berg et al. [4] to a linear-time exact algorithm.

Our Results. In this paper, we give hardness results and algorithms for both MSC and MHSC. Specifically, we give two (conceptually very different) algorithms. The first works by constructing a small ε-net for the hitting set problem that naturally arises from MSC. This gives then an $O(1)$-approximation algorithm for the MSC problem on orthogonal polygons. Note that no constant-factor approximation algorithm was known previously, and, as opposed to previous attempts at such approximation-algorithms [11], our algorithm works even on orthogonal polygons with holes. The second algorithm uses a tree-decomposition approach. We show that if the dual graph of the so-called pixelation of the polygon has bounded treewidth, then MSC can be solved in polynomial time. In particular this holds in so-called thin polygons that have no holes.

Both the above approaches also work (and become even simpler) for MHSC where only horizontal cameras are allowed. We also establish NP-hardness of MHSC for polygons with holes. The same proof also works for MSC and is different, and perhaps simpler, than the previous NP-hardness proof for MSC [12].

Finally, we consider art gallery theorems for sliding cameras, i.e., theorems that bound the number of guards relative to the number of vertices. We present the following results for an orthogonal polygon P with n vertices: (i) $\lfloor (3n+4)/16 \rfloor$ sliding cameras are always sufficient and sometimes necessary to guard P entirely, (ii) $\lfloor n/4 \rfloor$ only-horizontal sliding cameras are always sufficient and sometimes necessary to guard P, and (iii) if sliding cameras are not allowed to intersect each other, then $\lfloor (n+1)/5 \rfloor$ cameras are always sufficient to guard P.

Due to space constraints, some proofs will be given in the full version of this paper.

2 Preliminaries

Throughout the paper, P denotes an orthogonal polygon with n vertices. The *horizontal* (respectively *vertical*) *segmentation* of P consists of extending a horizontal (vertical) ray inward from any reflex vertex of P until it hits another vertex or edge. The rectangles in the resulting partition of P are called the horizontal (vertical) *slices* of P. Each slice can be represented by the horizontal (vertical) line segment that halves the slice; we call these the *slice-segments* and denote them by Σ.

The *pixelation* of P is obtained by doing both the horizontal and the vertical segmentation of P. The resulting rectangles are called *pixels*. The pixelation may well have $\Theta(n^2)$ pixels. Notice that the pixels are in 1-1-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 σ_H and σ_V of these two slices cross. We say that σ_H and σ_V *support* c. Denote the set of crosses by X.

A *sliding camera* γ is a horizontal or vertical line segment that is inside P. (We will frequently omit "sliding", as we study no other type of camera.) The region visible from γ is the set of all points p such that the perpendicular from p to γ is inside P. Note that doing a *parallel shift* (i.e., translating a horizontal camera vertically or a vertical camera horizontally) does not change its visibility region for as long as we stay inside P. We may hence assume that any camera runs along pixel-edges. We may also restrict our attention to cameras that are maximal line segments within P (all others would see a subset). Let Γ be the set of *guard-segments* which are maximal horizontal and vertical line segments within P that run along pixel edges. See Figure 1(b) for an illustration.

The following lemma (whose proof is given in the full version of the paper) is a straightforward re-formulation of what guarding means, but casts the problem into a discrete framework that will be crucial later.

Lemma 1. *A set S of k sliding cameras guards polygon P if and only if there exists a set of k guard-segments $S' \subseteq \Gamma$ such that for every cross $c \in X$, at least one of the slice-segments that support c is intersected by some $\gamma \in S'$.*

We say that a guard-segment γ *hits* a cross c if and only if γ intersects one of the slice-segments supporting c. Lemma 1 can then be re-stated as that S' hits all crosses. In fact, the algorithms we design later will allow further restrictions: we can specify exactly which crosses should be hit and which cameras may be used as guards. So assume we are given some $X' \subseteq X$ and some $\Gamma' \subseteq \Gamma$. The (X', Γ')-*sliding cameras problem* consists of finding a minimum subset of cameras in Γ' that hit all crosses in X'. Note that with a suitable choice of Γ' this encompasses both MSC and MHSC.

3 Approximation Algorithms via ε-nets

In this section, we give approximation algorithms for MSC and MHSC that are based on phrasing the problem as a hitting set problem and then using ε-nets. We do this first for MHSC, and then later re-use those ε-nets for MSC.

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 (X', Γ')-sliding camera problem, we construct a set system as follows. Let $\mathcal{U} = \Gamma'$ be all potential sliding cameras. For each cross $c \in X'$ that needs to be hit, define S_c to be all the cameras in $\mathcal{U}$ that hit c, and let $\mathcal{S}$ be the collection of these sets. From the definitions, finding a hitting set for this set system is the same as solving the (X', Γ')-sliding-camera 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 [6] 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 2. *[6] 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.*

Thus, the lemma gives an $O(1)$-approximation algorithm for as long as we can find an ε-net whose size is $O(1/\varepsilon)$. (Clearly the hitting set problems defined by MHSC and MSC both have a polynomial-time verifier.)

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

Lemma 3. *A set S of horizontal guard-segments hits all crosses in a set U' if and only if S intersects all the vertical slice-segments that support crosses in U'.*

Proof. If camera 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 MHSC, 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 and was shown to be NP-complete [17]. We hence have:

Corollary 1. *MHSC reduces to the Orthogonal Segment Covering problem.*

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 for MHSC.

Lemma 4. *The Orthogonal Segment Covering problem has a $(1/r)$-finder with size-function $s(r) \in O(r)$.*

Proof. Here, we sketch the proof; the full proof appears in the full version of the paper. Observe that a horizontal line segment $[x, x'] \times y$ intersects a vertical line segment $a \times [b, b']$ if and only if the point (x, y, x') lies in the range $(-\infty, a] \times [b, b'] \times [a, \infty)$. The union of these ranges forms a geometric object that (as one can argue) has complexity $O(n)$. Clarkson and Varadarajan [8] showed that ε-nets of small size can be found for hitting set problems in such a geometric object, using random sampling. □

Combining the above results gives:

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

An ε-net for the MSC Problem. Using the ε-net for MHSC, we can easily find one for MSC and hence have an approximation algorithm for this as well.

Theorem 2. *There exists a poly-time $O(1)$-approximation algorithm for the MSC problem.*

Proof. Fix a polygon P and consider the (X', Γ')-sliding camera 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 of MHSC for P, X' and the horizontal cameras in Γ'. Let T_V be a $1/2r$-net for the hitting set of MVSC (i.e., when we want to guard the polygon using only vertical sliding guards) for P, X' and the vertical cameras in Γ'. Set $T := T_H \cup T_V$. We claim that T is a $1/r$-net for $\mathcal{R}$.

So assume 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 Polygons with Bounded-Treewidth Pixelation

Recall that the *pixelation* of a polygon is obtained by cutting the polygon horizontally and vertically at all reflex vertices. The *dual graph D of the pixelation* is obtained by interpreting the pixelation as a planar graph and taking its weak dual (i.e., dual graph but omit the outer face). Thus, D has a vertex for every pixel of P, and two pixels are adjacent in D if and only if they share a side. We now show how to solve MSC and MHSC under the assumption that D has small treewidth. (Our approach was inspired by a similar result for a different guarding problem [5], but the construction here is simpler.)

2-dominating Set in an Auxiliary Graph. By Lemma 1, the (X', Γ')-sliding camera problem is equivalent to finding a set of guard-segments that hits at least one supporting slice-segment of each cross. This naturally gives rise to an auxiliary graph H as follows: The vertices of H are $X' \cup \Sigma \cup \Gamma'$. For any $c \in X'$, add an edge from c to each of its two supporting slice-segments. For any guard-segment γ, add an edge to any slice-segment that it intersects. From Lemma 1, and since there are no edges from X' to Γ', one immediately sees the following:

Lemma 5. *The minimum guard set for the (X', Γ')-sliding-cameras problem corresponds to a subset $S \subseteq \Gamma'$ of vertices in H such that all vertices in X' are within distance 2 from S.*

The above lemma means that the sliding-camera problem reduces to a graph-theoretic problem that is quite similar to the 2-dominating set (the problem of finding a set S such that all other vertices have distance at most 2 from S); the only change is that we restrict which vertices may be used for S and which vertices must be within distance 2 from S. 2-dominating set is an NP-hard problem in general, but is easily shown to be polynomial in graphs that have bounded treewidth, which we define next.

Treewidth. A *tree decomposition* $\mathcal{T} = (I, \mathcal{X})$ of a graph $G = (V, E)$ consists of a tree I and an assignment $\mathcal{X} : V(I) \to 2^V$ of *bags* of vertices of G to nodes of I such that the following holds: (a) for every vertex $v \in V$, the set of bags containing v forms a non-empty connected subtree of I, (b) for every edge $e \in E$, at least one bag contains both ends of e. The *width* of a tree decomposition is the maximum bag-size minus 1, and the *treewidth* $tw(G)$ of a graph G is the smallest possible width over all tree decompositions of G. In particular, a tree has treewidth 1. We prove the following lemma in the full version of the paper.

Lemma 6. *Let P be a polygon whose dual graph D of the pixelation has treewidth at most k. Then for any choice of $X' \subseteq X$ and $\Gamma' \subseteq \Gamma$, the auxiliary graph H has treewidth at most $7k + 6$.*

To apply this treewidth-result, we must show that the problem can be expressed as a suitable logic-formula. (See e.g. [10, Chapter 7.4] for more details.) In particular, the following formula will do: A set S of guard-segments that guards X' satisfies $S \subseteq \Gamma' \quad \wedge \quad \forall u \in X' (\exists \sigma \in \Sigma \quad \mathrm{adj}(u, \sigma) \wedge \exists \gamma \in S \quad \mathrm{adj}(\sigma, \gamma))$ (where adj is a logic-formula to encode that the two parameters are adjacent in H). Since H has bounded treewidth, we can find the smallest set S that satisfies this (or report that no such S exists if Γ' was too small) in linear time using Courcelle's theorem [9]. Putting everything together, we hence have:

Theorem 3. *If P is a polygon whose dual graph has bounded treewidth, then the (X', Γ')-sliding-cameras problem can be solved in linear time.*

We give one application of this result. A *thin* polygon is a polygon for which no pixel-corner is in the interior. MSC and MHSC are NP-hard even for thin polygons with holes (as we will see in the next subsection). However, for thin polygons without holes, the dual graph of the pixelation is clearly a tree, hence has bounded treewidth, and both MSC and MHSC can be solved in linear time.

Corollary 2. *If P is a thin polygon without holes, then MSC and MHSC can be solved in linear time.*

This result is not directly comparable to existing results [19, 4]: it is stronger than these since it works for MSC and does not require monotonicity, but it is weaker than these since it requires a thin polygon. A natural question is whether this result for bounded treewidth could be used to generate a PTAS, by splitting the polygon (hence the planar graph) suitably and applying the "shifting technique" (see [3] or [10, Chapter 7.7.3]). We have not been able to develop such a PTAS, principally because the cameras are not "local" in the sense that they can guard pixels that have arbitrarily large distance in D. Creating a PTAS (or proving APX-hardness) hence remains an open problem.

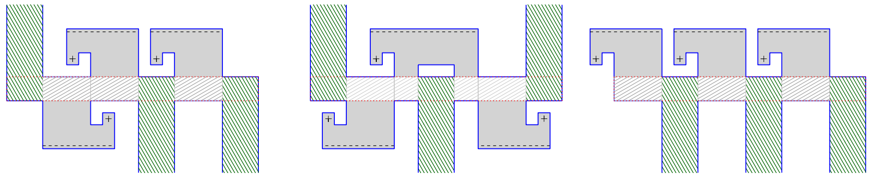

Fig. 2: The NP-hardness construction. Vertex-bars are red (dotted). Edge-strips are green (falling pattern). At each vertex-bar we attach an "elephant-gadget" (gray) that requires a sliding camera of its own (dashed) to guard the point inside the "trunk".

5 NP-hardness of MHSC

Recall that the MHSC problem is polynomial-time solvable on simple orthogonal polygons [18]. We show in this section that the problem becomes NP-hard on orthogonal polygons with holes. We note that the hardness proof of Durocher and Mehrabi [12] does not apply to the MHSC problem because they require both horizontal and vertical sliding cameras.

The reader may recall that we showed that MHSC reduces to the Orthogonal Segment Covering problem, which is known to be NP-hard [17]. However, this does not prove NP-hardness of MHSC, because not every instance of Orthogonal Segment Covering can be expressed as MHSC. Instead we give a different reduction from Minimum Vertex Cover on max-deg-3 planar graphs. This problem (which is NP-hard [14]) 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 max-deg-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 strips. It is well-known that every planar graph has such a representation (see e.g. Tamassia and Tollis [29]), 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 can construct in polynomial time an orthogonal polygon P such that the following holds (see Figure 2; the proof of the following lemma will appear in the full version of the paper):

Lemma 7. *The following are equivalent: (i) G has a vertex cover of size k; (ii) P can be guarded with $k + 3N$ horizontal sliding cameras; (iii) P can be guarded with $k + 3N$ sliding cameras.*

The constructed polygon is thin. NP-hardness of guarding problems in thin polygons (albeit with other models of guards and visibility) have been studied before [30, 5]. The NP-hardness holds for both MSC and MSHC; NP-hardness of MSC was known before [12], but the constructed polygon was not thin. We summarize:

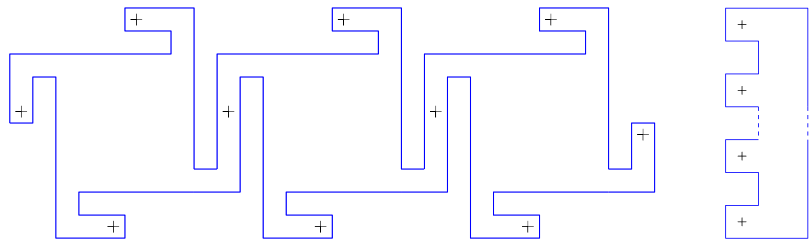

Fig. 3: (Left) A polygon that requires $(3n+4)/16$ cameras. (Right) A polygon that requires $n/4$ horizontal cameras.

Theorem 4. *MSC and MHSC is* NP*-hard on thin orthogonal polygons with holes.*

6 Art Gallery Theorems

We now consider the art gallery theorems for the MSC and MHSC problems; that is, we give tight bounds, depending on n, on the number of sliding cameras needed to guard an orthogonal polygon P with n vertices.

Recall that 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 3) actually works for sliding cameras, since no two of the $(3n+4)/16$ pixels marked with a cross can be guarded by one camera. The upper bound, indeed, also works for sliding guards. We very briefly review the approach taken in [1]. The idea is to guard first a small portion of P using one or two mobile guards, cutting a 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 indeed a sliding camera would have achieved the same as the mobile guard used. So we have the following result.

Theorem 5 (Based on [1]). *Given a simple orthogonal polygon P with n vertices, $\lfloor(3n+4)/16\rfloor$ sliding cameras are always sufficient and sometimes necessary to guard P entirely.*

For the MHSC problem, Figure 3 shows a polygon that requires $\lfloor n/4\rfloor$ horizontal sliding cameras. We show in the full version of the paper that this is tight.

Theorem 6. *Given an orthogonal polygon P with n vertices, $\lfloor n/4\rfloor$ horizontal sliding cameras are always sufficient and sometimes necessary to guard P entirely.*

Non-Crossing Sliding Cameras. There are cases in the upper bound approach of Aggarwal (Theorem 5 and [1]) in which the trajectories of mobile guards intersect. We show using a different approach that $\lfloor(n+1)/5\rfloor$ non-crossing sliding cameras are always sufficient to guard a simple orthogonal polygon P with n vertices that is in *general position* in the sense that no two vertical edges of P have the same x-coordinate.

Recall that earlier we considered the dual graph of the pixelation of a polygon P. In this section, we again consider a dual graph, but this time of one of the segmentations of P. Thus, consider (say) the vertical segmentation obtained after extending vertical rays from all reflex vertices. 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 [27]. In the following, we show that $\lfloor 2/5 \cdot |T| + 3/5 \rfloor$ non-crossing sliding cameras 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 partitiion $T - T(u)$ recursively. Let T_0 be the tree remaining in the base case (when no 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 $2/5 \cdot |T(u)|$ cameras.

To guard one such $T(u)$, we consider it to consist of the following components (see also Figure 4): (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. One can show that X can have at most one vertex of degree 4 since it corresponds to an orthogonal polygon (we formally prove this in the full version). Hence, 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 $y \leq x/2 + 1$. The following is shown in the full version:

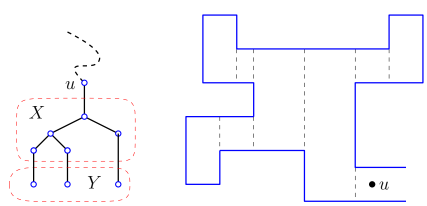

Fig. 4: An example of tree $T(u)$.

Lemma 8. *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 $X \cup \{u\}$.*

Since every leaf in Y can be covered using a single sliding camera, the subpolygon corresponding to $T(u)$ can hence be guarded with $y + 1$ sliding cameras. In fact, one can show that y sliding cameras suffice if y is close to the maximum, i.e., $y = x/2+1$ or $y = (x+1)/2$ (we formally prove this in the full version of the paper). Elementary calculations show that with this we use at most $2/5 \cdot |T(u)|$

cameras for $T(u)$ if $x \geq 6$. For $x < 6$, the only cases where the number of cameras is too large is $(x, y) = (1, 1)$ or $(4, 2)$ which can be dealt with by analyzing their structure directly. Finally tree T_0 may have an empty X, it then can always be guarded with $1 \leq 2/5 \cdot |T_0| + 3/5$ cameras. We hence have:

Theorem 7. *Given a simple orthogonal polygon P in general position with n vertices, $\lfloor (n+1)/5 \rfloor$ sliding cameras are always sufficient to guard P such that no two sliding cameras intersect each other.*

7 Conclusion

In this paper, we studied the problem of guarding an orthogonal polygon with the minimum number of sliding cameras. We gave the first constant-factor approximation algorithm for this problem, which works even if the polygon has holes. We also showed how to solve the problem optimally if the polygon is thin and has no holes, and we gave art-gallery-type results bounding the number of sliding cameras that are always sufficient and sometimes required. The most interesting remaining question is whether guarding an orthogonal polygon with sliding guards is polynomial if the polygon has no holes. Also, the factor in our $O(1)$-approximation algorithm (which we did not compute since it is hidden in the machinery of [6, 8]) is likely large. Can it be improved? Even better, could we find a PTAS or is the problem APX-hard?

References

1. Aggarwal, A.: The Art Gallery Theorem: its variations, applications and algorithmic aspects. Ph.D. thesis, Johns Hopkins University (1984), A summary can be found in [27], Chapter 3.
2. Avis, D., Toussaint, G.T.: An optimal algorithm for determining the visibility of a polygon from an edge. IEEE Trans. Computers 30(12), 910–914 (1981)
3. Baker, B.: Approximation algorithms for NP-complete problems on planar graphs. J. ACM 41(1), 153–180 (1994)
4. de Berg, M., Durocher, S., Mehrabi, S.: Guarding monotone art galleries with sliding cameras in linear time. In: Combinatorial Optimization and Applications (COCOA 2014). LNCS, vol. 8881, pp. 113–125 (2014)
5. Biedl, T., Mehrabi, S.: On r-Guarding Thin Orthogonal Polygons. In: 27th International Symposium on Algorithms and Computation (ISAAC 2016). LIPIcs, vol. 64, pp. 17:1–17:13 (2016)
6. Brönnimann, H., Goodrich, M.T.: Almost optimal set covers in finite VC-dimension. Discrete & Computational Geometry 14(4), 463–479 (1995)
7. Chvátal, V.: A combinatorial theorem in plane geometry. Journal of Combinatorial Theory, Series B 18, 39–41 (1975)
8. Clarkson, K.L., Varadarajan, K.R.: Improved approximation algorithms for geometric set cover. Discrete & Computational Geometry 37(1), 43–58 (2007)
9. Courcelle, B.: The monadic second-order logic of graphs. I. Recognizable sets of finite graphs. Information and Computation 85(1), 12–75 (1990)

10. Cygan, M., Fomin, F.V., Kowalik, L., Lokshtanov, D., Mark, D., Pilipczuk, M., Pilipczuk, M., Saurabh, S.: Parameterized Algorithms. Springer, Heidelberg, Germany (2015)
11. Durocher, S., Filtser, O., Fraser, R., Mehrabi, A.D., Mehrabi, S.: A (7/2)-approximation algorithm for guarding orthogonal art galleries with sliding cameras. In: Proceedings of Latin-American Symposium (LATIN 2014). LNCS, vol. 8392, pp. 294–305 (2014)
12. Durocher, S., Mehrabi, S.: Guarding orthogonal art galleries using sliding cameras: algorithmic and hardness results. In: Proceedings of Mathematical Foundations of Computer Science (MFCS 2013). LNCS, vol. 8087, pp. 314–324 (2013)
13. Eidenbenz, S., Stamm, C., Widmayer, P.: Inapproximability results for guarding polygons and terrains. Algorithmica 31(1), 79–113 (2001)
14. Garey, M.R., Johnson, D.S.: The rectilinear Steiner tree problem in NP-complete. SIAM Journal of Applied Mathematics 32, 826–834 (1977)
15. Hoffmann, F., Kaufmann, M., Kriegel, K.: The art gallery theorem for polygons with holes. In: Proceedings of Foundations of Computer Science (FOCS 1991). pp. 39–48 (1991)
16. Kahn, J., Klawe, M.M., Kleitman, D.J.: Traditional galleries require fewer watchmen. SIAM Journal on Algebraic Discrete Methods 4(2), 194–206 (1983)
17. Katz, M.J., Mitchell, J.S.B., Nir, Y.: Orthogonal segment stabbing. Comput. Geom. 30(2), 197–205 (2005)
18. Katz, M.J., Morgenstern, G.: Guarding orthogonal art galleries with sliding cameras. Inter. J. Comp. Geom. & App. 21(2), 241–250 (2011)
19. Katz, M.J., Roisman, G.S.: On guarding the vertices of rectilinear domains. Comput. Geom. 39(3), 219–228 (2008)
20. Kirkpatrick, D.G.: An $o(\lg\lg opt)$-approximation algorithm for multi-guarding galleries. Discrete & Computational Geometry 53(2), 327–343 (2015)
21. Krohn, E., Nilsson, B.J.: Approximate guarding of monotone and rectilinear polygons. Algorithmica 66(3), 564–594 (2013)
22. Lee, D.T., Lin, A.K.: Computational complexity of art gallery problems. IEEE Transactions on Information Theory 32(2), 276–282 (1986)
23. Lubiw, A.: Decomposing polygonal regions into convex quadrilaterals. In: Proceedings of the ACM Symposium on Computational Geometry (SoCG 1985). pp. 97–106 (1985)
24. Mehrabi, S.: Geometric optimization problems on orthogonal polygons: hardness results and approximation algorithms. Ph.D. thesis, University of Manitoba, Winnipeg, Canada (August 2015)
25. O'Rourke, J.: The complexity of computing minimum convex covers for polygons. In: 20th Allerton Conf. Commun. Control Comput. pp. 75–84 (1982)
26. O'Rourke, J.: Galleries need fewer mobile guards: a variation to Chvátal's theorem. Geometriae Dedicata 14, 273–283 (1983)
27. O'Rourke, J.: Art Gallery Theorems and Algorithms. The International Series of Monographs on Computer Science, Oxford University Press, New York, NY (1987)
28. Schuchardt, D., Hecker, H.: Two NP-hard art-gallery problems for ortho-polygons. Mathematical Logic Quarterly 41(2), 261–267 (1995)
29. Tamassia, R., Tollis, I.: A unified approach a visibility representation of planar graphs. Discrete & Computational Geometry 1, 321–341 (1986)
30. Tomás, A.P.: Guarding thin orthogonal polygons is hard. In: Fundamentals of Computation Theory (FCT 2013). LNCS, vol. 8070, pp. 305–316 (2013)