A linear time 8/3-approximation for r-star guards in simple orthogonal art galleries

Tamás Róbert Mezei
mezei.tamas.robert@renyi.mta.hu

Alfréd Rényi Institute of Mathematics, Hungarian Academy of Sciences, Reáltanoda u. 13–15, 1053 Budapest, Hungary

March 31, 2018

Abstract

We study the problem of covering simple orthogonal art galleries with rectangular stars. The problem has been shown to be polynomial [9], but to our knowledge, the exponent of the running time is still in the double digits. A linear-time 3-approximation algorithm using a partitioning into staircase shaped regions has been discovered by [8]. This is a follow-up paper to our recent theoretical result [5] linking point guards to horizontal mobile guards and vertical mobile guards (vision is restricted to rectangular vision). The result of this paper is that the algorithm implicitly described by our theoretical result can in fact be run in linear time. The novelty of the approach is the sparse representation of the pixelation graph of simple orthogonal polygons and the heavy reliance on so-called horizontal and vertical R-trees. After translating the problem into graph theory, geometrical insight is barely needed to verify the correctness of the algorithm.

1 Introduction

Art gallery problems in general ask the minimum number of guards with given power (for example, static or mobile) and type of vision (line of sight, rectangular vision, etc.) required to control the gallery. A point guard is a point in the interior of a polygon, and it covers any point in the closed polygon (the gallery) to which the guard can be joined by a line segment contained by the closed polygon (line of sight vision). The art gallery theorem due to Chvátal states that given an n-vertex simple polygon, $\lfloor \frac{n}{3} \rfloor$ point guards are sufficient and sometimes necessary to cover the closed polygon. In 1980 the sharp bound for the special case of n-vertex simple orthogonal polygons was determined to be $\lfloor \frac{n}{4} \rfloor$ by Kahn, Klawe, and Kleitman.

Mobile guards were introduced by Avis and Toussaint. A mobile guard patrols a line segment inside the gallery and sees every point in the gallery which can be seen from at least one point on its patrol. O'Rourke proved that to cover an n-vertex simple polygon, $\lfloor \frac{n}{4} \rfloor$ mobile guards are sufficient and sometimes necessary. To cover simple orthogonal polygons, Aggarwal proved that the extremal bound on the number of mobile guards is $\lfloor \frac{3n+4}{16} \rfloor$.

Throughout the paper, we restrict the meaning of orthogonal polygons to axis-parallel orthogonal polygons. Győri and O'Rourke independently proved that there is a stronger combinatorial theorem behind the simple orthogonal art gallery theorem: any n-vertex simple orthogonal polygon can be partitioned into at most $\lfloor \frac{n}{4} \rfloor$ simple orthogonal polygons of at most 6 vertices. Recently, Győri and Mezei proved that Aggarwal's theorem can be stated in a stronger form as well: any n-vertex simple orthogonal polygon can be partitioned into at most $\lfloor \frac{3n+4}{16} \rfloor$ simple orthogonal polygons of at most 8 vertices.

Two points in an orthogonal polygon (using axis-parallel sides) have rectangular or r-vision of each other, if there is an axis-parallel rectangle containing both points such that the rectangle is contained in the closed polygon. A region covered by a point guard inside the gallery is called an r-star. The mentioned partitioning theorems have an important feature in common: the bound on the number of guards required does not change if vision is restricted to r-vision. We mention that Katz and Morgenstern [6] introduced the notion of sliding cameras, which are maximal horizontal or vertical line segments in the gallery equipped with r-vision.

For simple orthogonal polygons an optimal point guard system using r-vision can be found in polynomial time, a result due to Worman and Keil [9]. The degree of the polynomial bounding the running time was originally 17, which may be too high for practical applications. Subsequently, Lingas, Wasylewicz, and Żyliński [8] gave a linear time 3-approximation algorithm for the problem. Their proof uses partition into staircase shaped regions, where the problem can be solved exactly in linear time.

In this paper we present the sketch of an algorithm for guarding simple orthogonal polygons with point guards using r-vision. An advantage of our approach is that it requires minimal geometrical insight, computation is mostly done on graphs, and there are only 3 subcases to verify.

Let P be a simple orthogonal polygon. Let m_H be the minimum size of a horizontal mobile guard system of P using r-vision. Define m_V analogously for vertical mobile guards. Lastly, let p be the minimum size point guard system of P using r-vision. Then

Theorem 1 (Győri and Mezei [5])**.**

$$\left\lfloor \frac{4(m_V + m_H - 1)}{3} \right\rfloor \geq p.$$

In [5] it has also been shown that both m_H and m_V (and the respective optimal guard system) can be computed in linear time (if holes are allowed, the problem is *NP*-hard [1]). The algorithm only relies on linear time triangulation of polygons by Chazelle [2] and an efficient least common ancestors algorithm in trees [3]. Using the trivial $p \geq m_V, m_H$ inequalities, it is clear that the left hand side is an $\frac{8}{3}$-approximation of p. The proof Theorem 1 is constructive. In the following sections we argue that the following holds.

Theorem 2. *There is a linear time algorithm which computes a covering set of point guards of P of cardinality at most $\lfloor \frac{4(m_H+m_V-1)}{3} \rfloor$.*

2 Preliminaries

In this 4 page extended abstract some technical details such as vision along degenerate rectangles will be neglected. Let S_H be the set of internally disjoint rectangles obtained by cutting horizontally at each reflex vertex of a simple orthogonal polygon P. Similarly, let S_V be defined analogously for vertical cuts of P. We may refer to the elements of these sets as **horizontal and vertical slices**, respectively. The horizontal R-tree T_H of P is equal to

$$T_H = \left(S_H, \left\{ \{h_1, h_2\} \subseteq S_H \ : \ h_1 \neq h_2,\ h_1 \cap h_2 \neq \emptyset \right\}\right),$$

i.e., T_H is the intersection graph of the horizontal slices of P. The graph T_H is indeed a tree, and we can think of it as a sort of dual of the planar graph determined by the union of P and its horizontal cuts. Similarly, T_V is the intersection graph of the vertical slices of P. The concept of R-trees were introduced by [4].

Let G be the intersection graph of S_H and S_V, i.e.,

$$G = \left(S_H \cup S_V, \left\{ \{h, v\} \ : \ h \in S_H,\ v \in S_V,\ \mathrm{int}(h) \cap \mathrm{int}(v) \neq \emptyset \right\}\right).$$

In other words, a horizontal and a vertical slice are joined by an edge iff their interiors intersect. We may also refer to G as the **pixelation graph** of P. This structure was already studied in [7]. Clearly, the **set of pixels** $\{\cap e = h \cap v \mid e = \{h, v\} \in E(G)\}$ is a cover of P. A cornerstone of the proof is the following lemma (without proof here).

Lemma 3. *G is a connected chordal bipartite graph (any cycle of length at least 6 has a chord).*

Definition 4 (r-vision of edges)**.** For any $e_1, e_2 \in E(G)$ we say that e_1 and e_2 have r-vision of each other iff $e_1 \cap e_2 \neq \emptyset$ or there exists a C_4 in G which contains both e_1 and e_2.

It is easy to see that two points $p_1 \in \mathrm{int}(\cap e_1)$ and $p_2 \in \mathrm{int}(\cap e_2)$ have r-vision of each other if and only if e_1 and e_2 have r-vision of each other in the above sense. Furthermore, every horizontal slice $h \in S_H$ can be mapped to a maximal horizontal mobile guard patroling the interior of h (by slightly modifying P, we may assume without loss of generality that every guard can be generated like this). The guard h covers exactly $\cup_{v \in N_G(h)} v$. Thus a covering set of horizontal mobile guards of P is a subset $M_H \subseteq S_H$, such that every $v \in S_V$ is covered in G by an element of M_H. Similarly, let $M_V \subseteq S_V$ be a covering set of vertical mobile guards of P. Without proof, we present any easy consequence of Lemma 3, and another almost trivial claim.

Claim 5. *If $G[M_H \cup M_V]$ is connected, then any edge $e_0 = \{h_0, v_0\} \in E(G)$ is r-visible from some edge of $G[M_H \cup M_V]$.*

Claim 6. *For any $h_1, h_2 \in S_H$ the following statements hold:*

- *$N_G(h_1)$ is the vertex set of a path in T_V, or in other words $N_G(h_1)$ induces a path in T_V.*
- *$N_G(h_1) \bigcap N_G(h_2)$ is either empty, contains exactly one slice, or induces a path in T_V.*
- *If G is 2-connected and h_1 is a neighbor of h_2 in the T_H, then $\left|N_G(h_1) \bigcap N_G(h_2)\right| \geq 2$.*

3 Sketch of the proof of Theorem 2

In [5] it was shown that M_H and M_V can be determined in linear time. G can also be determined in linear time by only storing the endpoints of $N_G(h)$ in T_V for a vertex $h \in S_H$ (we referred to this as the sparse representation). The proof of Theorem 1 is recursive and has three cases distinguished by the level of connectivity of $G[M_H \cup M_V]$: disconnected, connected but not 2-connected, and 2-connected.

Take an arbitrary slice $v_{\text{root}} \in M_V$, and make it the root of T_V (do this for T_H as well). Run the preprocessing of the least common ancestors (LCA) algorithm of [3] on T_V (and T_H, respectively). Observe that for any two $h_1, h_2 \in S_H$, we can compute $N_G(h_1) \cap N_G(h_2)$ by making 6 LCA queries in constant time. Indeed, $N_G(h_1) \cap N_G(h_2)$ is a path and its endpoints can be computed from the endpoints of the paths induced by $N_G(h_1)$ and $N_G(h_2)$ in T_V, which is stored in our representation of G.

Determining and storing $G[M_H \cup M_V]$ using sparse representation. Traverse T_V via a DFS started from v_{root}, and for each node $v \in S_V$ take note of the closest element of M_V which is on the search path between v_{root} and v. For each slice $h \in M_H$, determine the LCA of the endpoints of $N_G(h)$. If it is neither of the endpoints of $N_G(h)$, then the above labels allow us to determine the endpoints of $N_{G[M_H \cup M_V]}(h)$. If the LCA is one of the endpoints of $N_G(h)$, note this at the endpoint which is farther from v_{root}. By DFS traversing T_V one more time maintaining the subset of M_V contained in the search path, we can identify both ends of $N_{G[M_H \cup M_V]}(h)$.

Determining the R-forests on M_H and M_V. Join two slices $h_1, h_2 \in M_H$ by an edge if there exists a $v \in M_V$ such that $\{h_1, v\}, \{h_2, v\} \in E(G)$ and there does not exist $h_3 \in M_H$ which is between h_1 and h_2 in the path induced by $N_G(v)$ in T_H. We call the constructed graph the R-forest on M_H (since its components are trees). The definition for M_V goes analogously. It can be easily verified that if $G[M_H \cup M_V]$ is connected and we replace S_H, S_V, and G with M_H, M_V, and $G[M_H \cup M_V]$, Claim 6 still holds.

By the previously described structure, we can identify if $h_1, h_2 \in S_H$ have a common neighbor in M_V. For every $h \in M_H$ check if the closest element of $M_H \setminus \{h\}$ on the $h_{\text{root}} \to h$ path has a common M_V neighbor with h; if so, join them by an edge.

Let $v \in M_V$, and let $h_1, h_2 \in S_H$ be the endpoints of the path induced by $N_G(v)$ in T_H. Let the LCA of h_1 and h_2 be h_3, and suppose that $h_3 \neq h_1, h_2$. Let $h_4 \in M_H$ be the element closest to h_3 on the $h_3 \to h_1$ path, and let $h_5 \in M_H$ be the element closest to h_3 on the $h_3 \to h_2$ path. In the R-forest, h_4 and h_5 also need to be joined. Again, by DFS traversing T_H and maintaining the subset of M_H contained in the search path, we can identify h_4 and h_5.

The remaining steps. Having completed the previous steps, we can identify connected and 2-connected components of $G[M_H \cup M_V]$. The method described in [5] can then be used to determine a subset of edges of each 2-connected component such that the chosen edges have r-vision of any edge induced by the neighborhood of the 2-connected component in G. The set of these induced neighborhoods cover every node in $S_H \cup S_H$, however, there are edges that join two slices that are in different neighborhoods. An extra guard has to be found for each component of $G[M_H \cup M_V]$, but the algorithm to find these is not discussed here. If $G[M_H \cup M_V]$ is connected, but not 2-connected, Claim 5 already implies that the union of the point guards constructed for the 2-connected components of $G[M_H \cup M_V]$ is indeed covering set of point guards for the whole polygon.

References

[1] Therese C. Biedl, Timothy M. Chan, Stephanie Lee, Saeed Mehrabi, Fabrizio Montecchiani, and Hamideh Vosoughpour. "On Guarding Orthogonal Polygons with Sliding Cameras". In: *WALCOM: Algorithms and Computation, 11th International Conference and Workshops, WALCOM 2017, Hsinchu, Taiwan, March 29-31, 2017, Proceedings.* Ed. by Sheung-Hung Poon, Md. Saidur Rahman, and Hsu-Chun Yen. Vol. 10167. Lecture Notes in Computer Science. Springer, 2017, pp. 54–65. ISBN: 978-3-319-53924-9. DOI: 10.1007/978-3-319-53925-6_5. arXiv: 1604.07099v1.

[2] Bernard Chazelle. "Triangulating a simple polygon in linear time". In: *Discrete Comput. Geom.* 6.5 (1991), pp. 485–524. ISSN: 0179-5376. DOI: 10.1007/BF02574703.

[3] Harold N. Gabow and Robert Endre Tarjan. "A linear-time algorithm for a special case of disjoint set union". In: *J. Comput. System Sci.* 30.2 (1985), pp. 209–221. ISSN: 0022-0000. DOI: 10.1016/0022-0000(85)90014-5.

[4] Ervin Győri, Frank Hoffmann, Klaus Kriegel, and Thomas Shermer. "Generalized guarding and partitioning for rectilinear polygons". In: *Comput. Geom.* 6.1 (1996), pp. 21–44. ISSN: 0925-7721. DOI: 10.1016/0925-7721(96)00014-4.

[5] Ervin Győri and Tamás Róbert Mezei. "Mobile vs. point guards". In: *Discrete Comput. Geom.* accepted in March (2018). arXiv: 1706.02619.

[6] Matthew J. Katz and Gila Morgenstern. "Guarding orthogonal art galleries with sliding cameras". In: *Internat. J. Comput. Geom. Appl.* 21.2 (2011), pp. 241–250. ISSN: 0218-1959. DOI: 10.1142/S0218195911003639.

[7] Adrian Kosowski and Paweł Małafiejski Michałand Żyliński. "Cooperative mobile guards in grids". In: *Comput. Geom.* 37.2 (2007), pp. 59–71. ISSN: 0925-7721.

[8] Andrzej Lingas, Agnieszka Wasylewicz, and Paweł Żyliński. "Linear-time 3-approximation algorithm for the *r*-star covering problem". In: *Internat. J. Comput. Geom. Appl.* 22.2 (2012), pp. 103–141. ISSN: 0218-1959. DOI: 10.1142/S021819591250001X.

[9] Chris Worman and J. Mark Keil. "Polygon decomposition and the orthogonal art gallery problem". In: *Internat. J. Comput. Geom. Appl.* 17.2 (2007), pp. 105–138. ISSN: 0218-1959. DOI: 10.1142/S0218195907002264.