An Approximation Algorithm for the Art Gallery Problem

Édouard Bonnet * Tillmann Miltzow †

Abstract

Given a simple polygon $\mathcal{P}$ on n vertices, two points x, y in $\mathcal{P}$ are said to be *visible* to each other if the line segment between x and y is contained in $\mathcal{P}$. The *point-guard art gallery problem* asks for a minimum set S such that every point in $\mathcal{P}$ is visible from a point in S. Assuming integer coordinates and a special general position assumption, we present the first $O(\log \mathrm{OPT})$-approximation algorithm for the point guard art gallery problem. This algorithm combines ideas of Efrat and Har-Peled [7] and Deshpande et al. [3, 4]. In addition, we point out a mistake in the latter.

1 Introduction

In 1973, Victor Klee posed to Chvátal the art gallery problem as follows. Given a simple polygon $\mathcal{P}$ on n vertices, two points x, y in $\mathcal{P}$ are said to be visible to each other if the line segment between x and y is contained in $\mathcal{P}$. The point-guard art gallery problem asks for a minimum set S such that every point in P is visible from a point in S.

A huge amount of research is committed to the studies of combinatorial and algorithmic aspects of the art gallery problem, see the following surveys [10, 18–20]. Most of this research, however is not focused directly on the art gallery problem but on variants, based on different definitions of visibility, restricted classes of polygons, different shapes and positions of guards and so on. The arguably most natural definition of visibility is the one we defined above. One of the first combinatorial results is the elegant proof of Fisk that $\lfloor n/3 \rfloor$ guards are always sufficient and sometimes necessary for a polygon with n vertices [9].

On the algorithmic side, very few variants are solvable in polynomial time [5, 17], but most results are on approximating the minimum number of guards [3, 4, 7, 11, 14, 15]. Many of the approximation algorithms are based on the fact that the range space defined by the visibility regions has bounded VC-dimension [12, 13, 21] for simple polygons. This makes it easy to use the algorithmic ideas of Clarkson [1, 2].

On the lower bound side the paper of Eidenbenz et al. showed for most relevant variants NP-hardness and inapproximability [8]. In particular, they show for the main variants that there is no c-approximation algorithm for simple polygons, for some constant c. For polygons with holes, they can even show that there is no $o(\log n)$-approximation algorithm.

Very surprisingly, there is only one exact algorithm for the point guard art gallery problem running in $n^{O(k)}$ time attributed to Micha Sharir [7]. (Here, k is the size of the optimal solution.) And only, we recently the authors could give an almost matching lower bound by ruling out $n^{o(k/\log k)}$, assuming ETH [6].

Regarding approximation algorithms for the point guard variant, the results are very sparse. For monotone polygons and rectilinear polygons approximation algorithms are known [16]. For general polygons, Deshpande et al. gave a randomized pseudopolynomial time $O(\log n)$ approximation algorithm [3, 4]. However, we show that their algorithm is not correct. Efrat and Har-Peled gave a randomized polynomial time algorithm, by restricting guards to a very fine grid. They show that their grid solution S_{grid} is at most by a factor of $\log(OPT_{\text{grid}})$ away from the optimal grid solution OPT_{grid}. However, they could not prove that their OPT_{grid} is indeed an approximation of an optimal guard placement OPT. Developing the ideas of Deshpande et al. in combination of the algorithm of Efrat and Har-Peled we attain the first randomized polynomial time approximation algorithm for general simple polygons.

To keep the proof simple, we introduce a *special general position assumption.* We say a line is an *extension* of a polygon if it is the supporting line of two vertices of the polygon $v, w \in V(\mathcal{P})$. We say a polygon satisfies the special general position assumption, if no three points lie on a line and no three extensions meet in a point $p \in \mathcal{P} \setminus V(\mathcal{P})$.

Theorem 1 *There is an $O(\log |OPT|)$ approximation algorithm for* Point Guard Art Gallery *that runs in randomized polynomial time in the size of the input, given the following assumption:*

integer vertex representation: *Vertices are given by integers, represented in binary.*

*Institute for Computer Science and Control, Hungarian Academy of Sciences (MTA SZTAKI), `edouard.bonnet@lamsade.dauphine.fr`

†Institute for Computer Science and Control, Hungarian Academy of Sciences (MTA SZTAKI), `t.miltzow@gmail.com`

special general position assumption: *No three extensions meet in a point $p \in \mathcal{P} \setminus V(\mathcal{P})$.*

In this four page abstract we focus on the key ideas and henceforth omit the proofs of most lemmas and some technical details.

2 Preliminaries

Polygons and visibility. For any two distinct points v and w in the plane, we denote by $\mathrm{seg}(v, w)$ the segment whose two endpoints are v and w, by $\mathrm{ray}(v, w)$ the ray starting at v and passing through w, by $\ell(v, w)$ the supporting line passing through v and w. We also denote by $\mathrm{disk}(v, r)$ the disk centered in point v and whose radius is r, and by $\mathrm{dist}(a, b)$ the distance between object a and object b.

A polygon is *simple* if it is not self-crossing and has no holes. For any point x in a polygon $\mathcal{P}$, $V(x)$ denotes the *visibility region* of x within $\mathcal{P}$, that is the set of all the points $y \in \mathcal{P}$ such that segment $\mathrm{seg}(x, y)$ is entirely contained in $\mathcal{P}$.

3 Approximation

Given a polygon $\mathcal{P}$, we will always assume that all its vertices are given by *positive* integers in binary. We denote by M the largest integer, by D the diameter of the polygon and define $L = 10D$. Note that this implies $L \geq 5M$ and $\log L$ is linear in the input size. We set the width to $w = L^{-20}$ and define the grid Γ as $(w \cdot \mathbb{Z}^2) \cap \mathcal{P}$. Note that all vertices of $\mathcal{P}$ have integer coordinates and thus are included in Γ. We denote by OPT an optimal solution to the point guard art gallery problem and by k its size. We denote by OPT_{grid} an optimal solution to the grid guard art gallery problem, this is, guards are restricted to lie on the grid Γ, and denote by $k_{\text{grid}} = |OPT_{grid}|$.

The idea is to show that the algorithm of Efrat and Har-Peled gives an approximation algorithm under the integer vertex representation assumption, the grid Γ as described above and the special general position assumption.

Theorem 2 (Efrat, Har-Peled [7]) *Given a simple polygon $\mathcal{P}$ with n vertices, one can spread a grid Γ inside $\mathcal{P}$, and compute an $O(\log k_{\text{grid}})$-approximation to the smallest subset of Γ that sees $\mathcal{P}$. The expected running time of the algorithm is $O(n\, k_{\text{grid}}^2 \log k_{grid} \log(nk_{\text{grid}})\ \log^2 \Delta)$, where Δ is the ratio between the diameter of the polygon and the grid size.*

Note that the grid size equals $w = L^{-20}$, thus $\Delta \leq L^{21}$ and consequently $\log \Delta \leq 21 \log L$, which is linear in the size of the input. It remains to show the following lemma given the assumptions and notation mentioned above.

Lemma 3 *$\exists c \in \mathbb{N}$ such that $k_{\text{grid}} \leq c \cdot k$.*

The way we use the integer coordinate assumption is to infer distance lower bounds between various objects of interest.

Lemma 4 (Distances) *Let v and w be vertices of $\mathcal{P}$, ℓ_1 and ℓ_2 supporting lines of two vertices, and p and q intersections of supporting lines. Then the following holds:*

1. *$dist(v, w) > 0 \Rightarrow dist(v, w) \geq 1$.*
2. *$dist(v, \ell_1) > 0 \Rightarrow dist(v, \ell_1) \geq 1/L$.*
3. *$dist(p, \ell_1) > 0 \Rightarrow dist(p, \ell_1) \geq 1/L^5$.*
4. *$dist(p, q) > 0 \Rightarrow dist(p, q) \geq 1/L^4$.*
5. *Let $\ell_1 \neq \ell_2$ be parallel. Then $dist(\ell_1, \ell_2) \geq 1/L$.*
6. *Let $a \in \mathcal{P}$ be a point and ℓ_1 and ℓ_2 be some lines with $dist(\ell_i, a) < d$, for $i = 1, 2$. Then ℓ_1 and ℓ_2 intersect in a point p with $dist(a, p) \leq dL^2$.*

Proof. The idea of the proof is very simple. We look up the formula for each claimed distance. This formula is in most cases a fraction. By the assumption the nominator is at least one. The variables in the denominator can be safely upper bounded by L. The claim follows immediately. ☐

Grid points. See Figure 1 for the following description. Each point x of the optimal solution is in some grid cell $\mathrm{grid}(x)$. For the sake of brevity, we assume that the grid cell does not contain any point of the boundary of the polygon, as in Figure 1 b) and c).

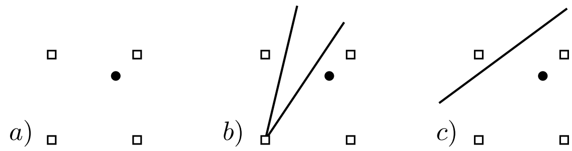

Figure 1: The way that the polygon boundary might interact with the grid cell.

Local visibility containment property. We say a point x in the grid cell formed by g_1, g_2, g_3, g_4 has the *local visibility containment property* (LVCP) if $V(x) \subseteq V(g_1) \cup V(g_2) \cup V(g_3) \cup V(g_4)$.

Cones. Given a point x and two points r_1, r_2 in the plane, we define the cone of x with respect to r_1 and r_2 as the unique cone $C(x)$ with apex x that is bounded by $\mathrm{ray}(x, r_1)$ and $\mathrm{ray}(x, r_2)$ and forms an angle smaller than π.

Opposite reflex vertices and bad regions. Given a polygon $\mathcal{P}$ and two reflex vertices r_1 and r_2, consider the supporting line $\ell = \ell(r_1, r_2)$ restricted to $\mathcal{P}$. The supporting line defines two halfplanes ℓ^+ and ℓ^-. We say r_1 is *opposite* to r_2 if $\text{disk}(r_1, \varepsilon) \cap \partial\mathcal{P} \subset \ell^+$ and $\text{disk}(r_2, \varepsilon) \cap \partial P \subset \ell^-$ for some $\varepsilon > 0$.

Given two opposite reflex vertices r_1 and r_2, we define their *bad regions* as the stripe around $\ell(r_1, r_2)$ with width $8wL$, see Figure 2.

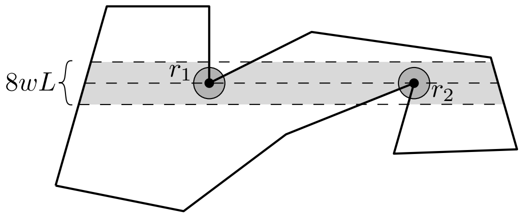

Figure 2: Illustration of opposite reflex vertices and bad regions.

Lemma 5 (Loc. Visib. Containment Prop.) *Let $x \in \mathcal{P}$ be outside any bad region then x has the local visibility property.*

Proof. Here, we give only the idea. Let $x \in \mathcal{P}$ be some point and g_1, g_2, g_3, g_4 be the grid points surrounding x. Further let r_1 and r_2 be any two reflex vertices visible from x. the same fashion. We can restrict ourselves to show $C(x) \subseteq C(g_1) \cup C(g_2) \cup C(g_3) \cup C(g_4)$. Further, we only have to show this for the region behind $\ell(r_1, r_2)$. In case that there is some $g_i \in C(x)$, this is trivially true. For the other case, see Figure 3. In Figure 3, we see that the dark red

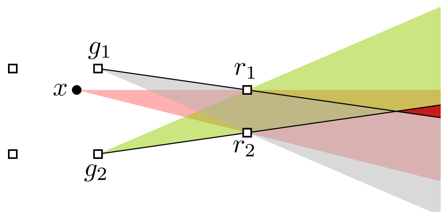

Figure 3: The cones $C(g_1)$ and $C(g_2)$ are not covering the same area as $C(x)$, behind $\ell(r_1, r_2)$. The reason is that ray(g_1, r_2) and ray(g_2, r_2) intersect.

area of $C(x)$ is not covered by by $C(g_1)$ and $C(g_2)$. It is not so difficult to see that it is sufficient to show that this situation will not appear. For this purpose it is sufficient to show that ray(g_1, r_1) and ray(g_2, r_2) do not cross. Roughly speaking, we compare the distance between the rays close to x and close to the reflex vertices. If the distance between the rays increases, we know that they will never meet. Very helpful for us is the distance of at least one between the reflex vertices and that x is outside of any bad region. $\square$

It is easy to believe that the local visibility containment property also holds inside the bad region. However this is not true. Here, we will briefly describe a counter-example. In particular, this example shows that the algorithm of Deshpande et al. is not correct as it is stated. However, we want to mention that their paper has ideas that motivated the algorithm presented in this preprint.

Deshpande et al. described an algorithm that worked in several steps [3, 4]. In the first step they generate a large number of points P that they store explicitly in memory. Thereafter, they find a solution for the point guard art gallery problem $S \subset P$. The crucial point is the claim that their point set P satisfies some variant of local visibility containment property. To be more precise, we say a set of points P has the *general local visibility containment property*, if for every point $x \in \mathcal{P}$ exists a finite collection $C \subset P \cap \text{disk}(x, 1)$, such that $V(x) \subseteq \bigcup_{p \in C} V(p)$. Our example shows that it is impossible to attain any finite point set that has this property.

Example See Figure 4, for the following description.

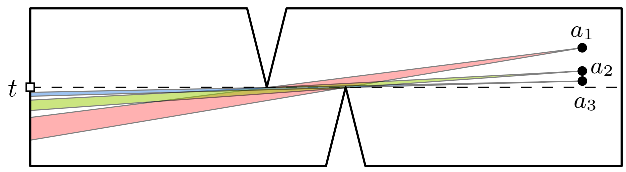

Figure 4: Illustration of the counter-example to the algorithm of Deshpande et al. [3, 4].

We have two opposite reflex vertices with supporting line ℓ. The points $(a_i)_{i \in \mathbb{N}}$ are chosen closer and closer to ℓ on the right side of the polygon. None of the a_i can see t, as this would require to be actually on ℓ. We choose the points $(a_i)_{i \in \mathbb{N}}$ in a way that their intervals will be all disjoint and arbitrarily close to t.

Consider now any finite set of points C in the "vicinity" of the $(a_i)_{i \in \mathbb{N}}$. We will show that there is some a_i, which sees some interval close to t, that is not seen by any point in C. Recall that no point sees the entire interval around t, but the visibility of the a_is come arbitrarily close to t. Thus, there is some a_i that sees something that is not visible by any point in C.

This shows that the general local visibility containment property cannot be attained already in this fairly straightforward polygon.

Despite the fact that it is not possible to achieve a general local visibility containment property for all points in $\mathcal{P}$, the exceptions are only for points in the bad regions. These cases we can handle in a different manner.

Lemma 6 *Let $x \in \mathcal{P}$ in three or more bad regions.*

Then there exists a reflex vertex r with $dist(x,r) \leq wL^3$. And r is one of the defining reflex vertices for all bad regions that x belongs to.

Proof. Let ℓ_1, ℓ_2, ℓ_3 be three different extensions. And we assume that x is in the bad region of all three of them. In case they all have on reflex vertex in common, it must be a defining reflex vertex and x cannot be too far away from r.

So assume that they do not have a reflex vertex in common. We want to show that the intersection of their corresponding bad regions is empty. These three lines form a triangle Δ with vertices of pairwise distance at least L^{-4} by Lemma 4 Item 4. Assume for the purpose of contradiction that x is in the bad region of all three lines. Then x has distance at most $4wL$ to all lines. By Lemma 4 Item 6 x has distance at most $4wL^3$ to the vertices of Δ. By the triangle inequality the vertices have pairwise distance at most $8wL^3 \ll L^{-4}$ – a contradiction. ☐

Proof. [Lemma 3] Let OPT be an optimal solution of size k. We construct a set of $6k$ guards $G \subset \Gamma$, which see the entire polygon. For each guard $x \in OPT$, we add the four grid points of x into G. Further, if x is in one or two bad regions, add the corresponding reflex vertex for each bad region into G. In case x is in more than two bad regions there is one reflex vertex that is defining all of them. Add it to G. For each $x \in OPT$ the local containment property holds, except for the bad regions it is in. These parts are seen by the reflex vertices we added. ☐

Acknowledgments

Both authors are supported by the ERC grant PARAMTIGHT: "Parameterized complexity and the search for tight complexity results", no. 280152. We thank Meirav Zehavi and Saeed Merhabi for useful discussions during a preliminary stage of the project. We also want to thank anonymous reviewers for their helpful feedback.

References

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

[2] K. L. Clarkson. Algorithms for polytope covering and approximation. In *WADS 1993*, pages 246–252, 1993.

[3] A. Deshpande. A pseudo-polynomial time $O(\log^2 n)$-approximation algorithm for art gallery problems. Master's thesis, Department of Mechanical Engineering, Department of Electrical Engineering and Computer Science, MIT, 2006.

[4] A. Deshpande, T. Kim, E. D. Demaine, and S. E. Sarma. A pseudopolynomial time $O(\log n)$-approximation algorithm for art gallery problems. In *WADS 2007*, pages 163–174, 2007.

[5] S. Durocher and S. Mehrabi. Guarding orthogonal art galleries using sliding cameras: algorithmic and hardness results. In *MFCS 2013*, pages 314–324. Springer, 2013.

[6] Édouard Bonnet and T. Miltzow. Parameterized hardness of art gallery problems. In *EuroCG 2016*, page to appear, 2016.

[7] A. Efrat and S. Har-Peled. Guarding galleries and terrains. *Inf. Process. Lett.*, 100(6):238–245, 2006.

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

[9] S. Fisk. A short proof of chvátal's watchman theorem. *J. Comb. Theory, Ser. B*, 24(3):374, 1978.

[10] S. K. Ghosh. *Visibility algorithms in the plane*. Cambridge University Press, 2007.

[11] S. K. Ghosh. Approximation algorithms for art gallery problems in polygons. *Discrete Applied Mathematics*, 158(6):718–722, 2010.

[12] A. Gilbers and R. Klein. A new upper bound for the vc-dimension of visibility regions. *Computational Geometry*, 47(1):61–74, 2014.

[13] G. Kalai and J. Matoušek. Guarding galleries where every point sees a large area. *Israel Journal of Mathematics*, 101(1):125–139, 1997.

[14] J. King. Fast vertex guarding for polygons with and without holes. *Comput. Geom.*, 46(3):219–231, 2013.

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

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

[17] R. Motwani, A. Raghunathan, and H. Saran. Covering orthogonal polygons with star polygons: The perfect graph approach. *J. Comput. Syst. Sci.*, 40(1):19–48, 1990.

[18] J. O'rourke. *Art gallery theorems and algorithms*, volume 57. Oxford University Press Oxford, 1987.

[19] T. C. Shermer. Recent results in art galleries [geometry]. *Proceedings of the IEEE*, 80(9):1384–1399, 1992.

[20] J. Urrutia et al. Art gallery and illumination problems. *Handbook of computational geometry*, 1(1):973–1027, 2000.

[21] P. Valtr. Guarding galleries where no point sees a small area. *Israel Journal of Mathematics*, 104(1):1–16, 1998.