Robustly Guarding Polygons

Rathish Das
University of Houston, TX, USA

Omrit Filtser
The Open University of Israel, Israel

Matthew J. Katz
Ben-Gurion University of the Negev, Beer-Sheva, Israel

Joseph S.B. Mitchell
Stony Brook University, NY, USA

Abstract

We propose precise notions of what it means to guard a domain "robustly", under a variety of models. While approximation algorithms for minimizing the number of (precise) point guards in a polygon is a notoriously challenging area of investigation, we show that imposing various degrees of robustness on the notion of visibility coverage leads to a more tractable (and realistic) problem for which we can provide approximation algorithms with constant factor guarantees.

2012 ACM Subject Classification Theory of computation → Computational geometry

Keywords and phrases geometric optimization, approximation algorithms, guarding

Digital Object Identifier 10.4230/LIPIcs.SoCG.2024.47

Related Version *Full Version*: `https://arxiv.org/abs/2403.11861`

Funding *Matthew J. Katz*: Partially supported by Grant 495/23 from the Israel Science Foundation and Grant 2019715/CCF-20-08551 from the US-Israel Binational Science Foundation/US National Science Foundation.
Joseph S.B. Mitchell: Partially supported by the National Science Foundation (CCF-2007275).

1 Introduction

A fundamental set cover problem that arises in geometric domains is the classic "art gallery" or "guarding" problem: Given a geometric domain (e.g., polygon P), place a set of points ("guards") within P, such that every point of P is seen by at least one of the guards. This problem has many variants and has been studied extensively from many perspectives, including combinatorics, complexity, approximation algorithms, and algorithm engineering for solving real instances to provable optimality or near-optimality.

Approximation algorithms for guarding have been extensively pursued for decades (see related work), where the various variants differ from one another in (i) the underlying domain, e.g., simple polygon vs. polygon with holes, (ii) the portion of the domain that must be guarded, e.g., only its boundary, the entire domain, or a discrete set of points in it (iii) the type of guards, e.g., static, mobile, or with various restrictions on their coverage area, (iv) the restrictions on the location of the guards, e.g., only at vertices (vertex guards) or anywhere (point-guards), and (v) the underlying notion of vision, e.g., line of sight or rectangle visibility. Despite all this work, even an $O(\log \mathrm{OPT})$-approximation algorithm for point-guarding a simple n-gon, without (weak) additional assumptions, is unknown (see below).

In this paper, we present and discuss a new and natural notion of vision called *robust vision*. Under the standard notion of vision, two points in a polygon P see each other if and only if the line segment between them is contained in P. However, in the context of guarding, where e.g. the guards are not necessarily stationary entities or perhaps their location is

40th International Symposium on Computational Geometry (SoCG 2024).
Editors: Wolfgang Mulzer and Jeff M. Phillips; Article No. 47; pp. 47:1–47:17
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

imprecise, it makes sense to require that a guard g that is responsible for guarding a point p can see p from any point in some vicinity of its specified location. In this case we say that g robustly guards p. It is also possible that the location of an entity to be guarded is imprecise or alternatively the entity may move in the vicinity of its specified location, and we would like to ensure that the guard in charge of this entity does not lose sight of it, i.e., we would like it to guard the entity robustly. (Here, we mostly focus on the former case of robust vision.) Robust guarding is a generalization of standard guarding in that when the requirement of robustness tends to zero, robust guarding reduces to standard guarding.

Robust guarding is especially important in light of recent results showing that there are polygons that can be guarded with 2 guards, but only if both guards are very precisely placed at points with irrational coordinates [34] (see also [1]). In our formulation of robust guarding, we explicitly model the fact that a guard may be imprecisely placed at locations within a polygon, and may in fact move around within some neighborhood; we insist, then, for a point p to be "seen" that it must be seen no matter where the guard may be within a disk that subtends at least some minimum angle when viewed from p. The breakthrough result [2] showing that the guarding problem is complete within the existential theory of the reals is strong evidence of the algebraic difficulty of computing exact optimal sets of guards, even in polygons in the plane.

Our techniques and results. We summarize our contributions and methods:

(1) We introduce notions of "robust vision" within a polygonal domain P and analyze the optimal guarding problem from this new perspective. In particular, for an appropriately small parameter $\alpha > 0$, we say that a guard at point g *α-robustly guards* a point $p \in P$ if p sees (under ordinary visibility) all points within a Euclidean disk of radius $\alpha\|g - p\|$ centered at g. In the figure below, g α-robustly guards p, but not p' or p''.

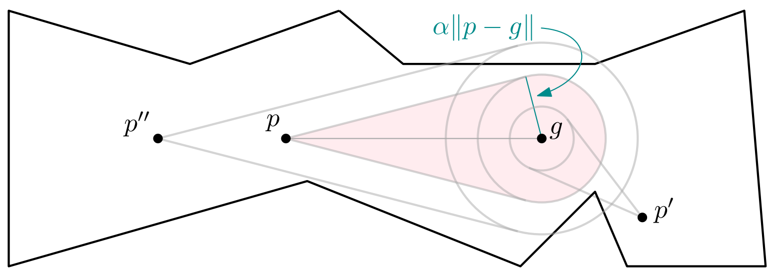

Note that as α approaches 0, the degree of robustness decreases, and at the limit we get standard guarding, where g guards p if and only if $\overline{gp} \subset P$. We characterize the α-robust visibility region, $\mathrm{Vis}_\alpha(g)$, of all points α-robustly visible from g (Section 2.1), as well as the region $\mathrm{Vis}_\alpha^{\mathrm{inv}}(p)$ of all points g from which p is α-robustly visible (Section 2.2). In particular, we prove that $\mathrm{Vis}_\alpha(g)$, which in general is not a polygon, is star-shaped and $O(\alpha)$-fat. Moreover, we show that both regions can be computed efficiently.

(2) We show that, as with ordinary guarding, the problem of computing a minimum cardinality set of guards in P that α-robustly see all of P is APX-hard, making it unlikely that there exists a PTAS or an exact polynomial-time algorithm for the problem.

(3) We present an $O(1)$-approximation algorithm for robustly guarding a general polygonal domain P. (The approximation factor depends on the robustness parameter α, and the algorithm is a bicriteria, allowing a slight relaxation of α.) This is to be contrasted with the situation for ordinary guarding, for which even finding a logarithmic-factor approximation algorithm for placing guards at points within a simple polygon requires some additional (weak) assumptions.

Specifically, Theorem 24 states that, given a polygon P with n vertices, one can compute in poly(n) time the cardinality of, and an implicit representation of, a set of $O(\alpha^{-6})|\mathrm{OPT}_\alpha|$ points that $\alpha/8$-robustly guard P, where OPT_α is a minimum-cardinality set of guards that α-robustly guard P. In additional time $O(|\mathrm{OPT}_\alpha|)$ one can output an explicit set of such points. We present this result by first presenting a result of a similar flavor for robustly guarding a discrete set S of points within P (Theorem 14). Critical to our main result (Theorem 24) is Theorem 21, which shows that one can compute a discrete set Q of points (candidate guards) that is guaranteed to contain a subset of $O(\alpha^{-4})|OPT_\alpha|$ points that $\alpha/4$-robustly guard P. This result is quite delicate and requires some technical geometric analysis, utilizing a medial axis decomposition and carefully placed grid points in portions of P. This is in contrast with the classic guarding problem in which the existence of a polynomial-size discrete candidate set that suffices for good approximation has been elusive.

(4) In the full version of the paper, we extend/generalize our definition of robust vision to include the possibility that p need not see *all* of the neighborhood of g, but only a *fraction* of that neighborhood, and we require that g sees a fraction of a neighborhood of p as well. Within this model we are able to obtain improved factors, with some tradeoffs. We defer many proofs and technical arguments to the full version of the paper, while attempting to convey intuition, and detailed figures, in this version.

Related Work. Eidenbenz, Stamm, and Widmayer [23] have shown that optimally guarding a simple polygon P is not only NP-hard (a classical result [33, 35, 36]) but is APX-hard (there is no PTAS unless P=NP); if P has holes, they show that there is no $o(\log n)$ approximation algorithm unless $P = NP$. A recent breakthrough of Abrahamsen, Adamaszek and Miltzow [2] has shown that point guarding in general polygons is $\exists\mathbb{R}$-complete, making it unlikely the problem is in NP. Further, [1, 34] have shown that optimal solutions to even very small problems requiring 2-3 guards in simple polygons may require precise placement of guards at irrational points.[1] These results imply the necessity of algebraic methods to compute exact solutions. Efrat and Har-Peled [21] present a randomized $O(\log \mathrm{OPT}_{\mathrm{grid}})$-approximation algorithm where the placement of guards is restricted to a fine grid. However, they do not prove that their approximation ratio holds when compared to general point guard placement of optimal solution. Building on [21] and on Deshpande et al. [18], Bonnet and Miltzow [13] gave a randomized $O(\log \mathrm{OPT})$-approximation algorithm for point guards within a simple polygon P under mild assumptions: vertices have integer coordinates, no three vertices are collinear, and no three extensions meet in a point within P that is not a vertex, where an extension is a line passing through two vertices of P. The problem has also been examined from the perspective of smoothed analysis [19, 24] and parameterized complexity [8, 4, 14, 3]. For guards that must be placed at discrete locations on the boundary of a simple polygon P, King and Kirkpatrick [29, 31] obtained an $O(\log \log OPT)$-approximation, by building ε-nets of size $O((1/\varepsilon) \log \log(1/\varepsilon))$ for the associated hitting set instances, and applying [17]. If the disks bounding the visibility polygons at these discrete locations are shallow (i.e., every point in the domain is covered by $O(1)$ disks), then a local search based PTAS exists [7].

For simple polygons with special structures, such as monotone polygons, terrains, and weakly-visible polygons [32, 11, 26, 25, 10, 12, 9], there are improved approximation ratios (constant, or even PTAS), but these algorithms utilize the very special structures of these

[1] It is an interesting open problem to determine whether or not an optimal set of robust guards might require irrational coordinates for guards, for input polygons P with integer coordinates.

classes of polygons. Aloupis et al [5] considered guarding in "(α, β)-covered" polygons (which are intuitively "fat" polygons, see [20]), and showed that the boundary of such polygons can be guarded by $O(1)$ guards. Polygons in which no point sees a particularly small area also have special analysis and algorithms [28, 30, 38, 39].

Some prior work has addressed guarding from a robustness perspective, though the perspectives are significantly different from ours. Efrat, Har-Peled, and Mitchell [22] consider a definition of robust guarding in which a point is robustly guarded if it is seen by at least 2 guards from significantly different angles. Hengeveld and Miltzow [27] consider a notion of "robust" vision by examining the impact that certain changes to "visibility" has on the optimal number of guards needed to cover a domain P. They introduce the notion of *vision-stability*: A polygon P has vision-stability δ if the optimal number of "enhanced guards" (who can see an additional angle δ around corners) is equal to the optimal number of "diminished guards" (whose visibility region is decreased by an angle δ at each shadow-casting corner).

Practical methods employing heuristics, algorithm engineering and combinatorial optimization have been successful for computing exact or approximately optimal solutions in many instances [6, 15, 27, 37], though some of the exact methods can potentially fail or run forever on contrived instances, e.g., those requiring irrational guards. ([27] is able to detect such contrived instances and to solve exactly instances that are vision-stable.)

2 Robust guarding

In this section we introduce and discuss a new and natural notion of vision called *robust vision*. Let P be a polygonal domain (a polygon, possibly with holes) in the plane having in total n vertices. Under the standard notion of vision, a point $p \in P$ sees another point $q \in P$ if the segment $\overline{pq}$ is contained in P. In the context of guarding, where the guards are not necessarily stationary or their locations are imprecise, one would like to ensure that a guard does not lose eye contact with any of the points it is responsible to see, when it moves in some vicinity of its specified location. The following definition attempts to capture these common situations. Here, we denote by $D(p, r)$ the disk of radius r centered at a point p.

▶ **Definition 1** (Robust Guarding)**.** *Given a polygon P and parameter $0 < \alpha \leq 1$, we say that a point $g \in P$* ***α-robustly guards*** *a point $p \in P$ if p sees $D(g, \alpha \cdot \|p - g\|)$.*

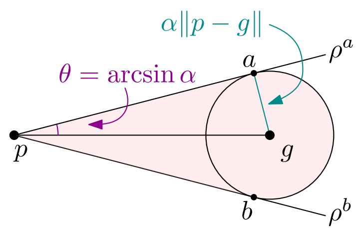

Figure 1 A point g that α-robustly guards another point p. The pink "ice cream cone" is contained in the polygon P.

Let g, p be two points in a polygon P, such that g α-robustly guards p. By definition, p sees the disk $D(g, \alpha\|p - g\|)$. Consider the two rays ρ^a, ρ^b from p tangent to $D(g, \alpha\|p - g\|)$. Let a (resp. b) be the tangency point of ρ^a (resp. ρ^b) and $D(g, \alpha\|p - g\|)$. Notice that the segments $\overline{ap}, \overline{bp}$ with the disk $D(g, \alpha\|p - g\|)$ create an "ice cream cone". Formally, we define the **ice cream cone** from p to g as the union of the triangle $\triangle apb$ and the disk $D(g, \alpha\|p - g\|)$. Since p must see the entire ice cream cone, we have the following observation.

▶ **Observation 2.** *A point $g \in P$ α-robustly guards a point $p \in P$ if and only if the ice cream cone from p to g is contained in P.*

Consider the triangle $\triangle agp$. We have $\|a - g\| = \alpha\|p - g\|$ and thus $\angle apg = \arcsin \alpha = \theta$. Symmetrically, $\angle bpg = \theta$, and we get $\angle apb = 2\theta$. We obtain the following observation.

▶ **Observation 3.** *The angle $\angle apb$ of the ice cream cone from p to g is $2\theta = 2 \arcsin \alpha$.*

Note that as α approaches 0, the degree of robustness decreases, and at the limit we get standard guarding, where g guards p if and only if $\overline{gp} \subset P$. On the other hand, if P has a vertex v with internal angle smaller than 2θ, then v cannot be α-robustly guarded by any other point $g \neq v$ in P (v can guard itself). Nevertheless, if all internal angles in P are at least 2θ, then there is always a finite set of guards G that robustly guard P (this will become clear in Section 3). We therefore assume that $\alpha \leq \sin \frac{\phi}{2}$, where ϕ is the smallest internal angle of P, and thus $2\theta = 2 \arcsin \alpha \leq \phi$, so P can be guarded by a finite number of α-robust guards. Of course, if ϕ is very "small", then, as mentioned above, the degree of robustness decreases, in the sense that the vicinity in which a guard may move while guarding a point p becomes more limited.

Also note that P does not have to be fat in order to be guarded robustly; however, the number of robust guards required may depend on geometric features of P (see Figure 2).

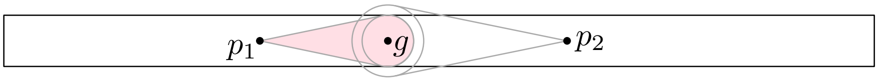

Figure 2 In a thin rectangle P, the point g α-robustly guards p_1, but not p_2. Here, the number of robust guards required depends on the aspect ratio of the rectangle P.

The geometric observation below will be very useful in the following sections. A proof is given in the full version of the paper, see Figure 3 for an illustration.

▶ **Observation 4.** *Consider a (convex) cone K defined by two rays ρ_0, ρ_1 from a point p, such that the small angle between them is $\theta = \arcsin \alpha$. Let q be a point in K, such that the smaller angle between $\overline{pq}$ and one of the rays ρ_0, ρ_1 is $c \cdot \theta$ (for some $c < 1$). Then the disk $D(q, c\alpha\|p - q\|)$ is contained in K.*

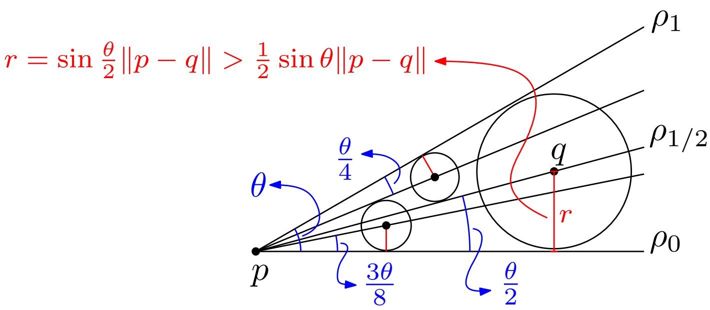

Figure 3 The cone K with angle θ. For q on $\rho_{1/2}$, the disk $D(q, \frac{1}{2}\alpha\|p - q\|)$ is contained in K.

2.1 The robust visibility region

Let P be a polygon in the plane, possibly with holes, having in total n vertices. Denote by $\text{Vis}(p)$ the standard visibility polygon of p, i.e., $\text{Vis}(p)$ is the set of points $q \in P$ such that q is visible from p. Let $\text{Vis}_\alpha(p)$ denote the **α-robust visibility region** of p, i.e. the set of points in P that are α-robustly guarded by p. Interestingly, unlike $\mathrm{Vis}(p)$, the robust visibility region is rarely a polygon. Nonetheless, in this subsection we present an efficient algorithm to compute it. First, we characterize $\mathrm{Vis}_\alpha(g)$ and reveal some of its interesting geometric properties. We begin by showing that $\mathrm{Vis}_\alpha(g)$ is fat, under the following standard definition of fatness.

Fatness. For a disk $D(p, r)$ of radius r centered at point $p \in P$, let $C(p, r)$ denote the connected component of $D(p, r) \cap P$ in which p lies. We say that a polygon P is **γ-fat** if for any point $p \in P$ and radius r such that $D(p, r)$ does not fully contain P, the area of $C(p, r)$ is at least $\gamma \cdot \pi r^2$, i.e., at least γ times the area of $D(p, r)$.

We use Claim 5 below to show that $\mathrm{Vis}_\alpha(g)$ is star-shaped and $O(\alpha)$-fat. In order not to interrupt the flow of reading with rather technical details, the proof is provided in the full version of the paper.

▷ Claim 5. Let $\mathcal{K}$ be a set of (convex) γ-fat kites (quadrilaterals with reflection symmetry across a diagonal), all having a common point p. Then the union $U = \bigcup_{K \in \mathcal{K}} K$ is $\gamma/4$-fat.

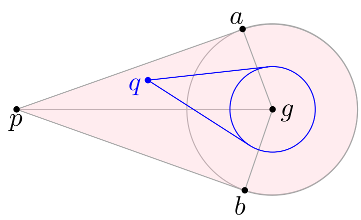

Figure 4 The ice cream cone C from p to g is shaded in pink. For any point $q \in C$, the ice cream cone from q to g is contained in C. Therefore, $\mathrm{Vis}_\alpha(g)$ contains C.

▶ **Lemma 6.** *For any $0 < \alpha < 1$ and $g \in P$, $\mathrm{Vis}_\alpha(g)$ is star-shaped and $O(\alpha)$-fat.*

Proof. Let p be a point in $\mathrm{Vis}_\alpha(g)$, and denote by C the ice cream cone from p to g. Notice that for any point $q \in C$ we have $\|q - g\| \le \|p - g\|$, and thus the ice cream cone from q to g is contained in C (see Figure 4). Since C is contained in P, q sees a disk of radius $\alpha\|q - g\|$ centered at g, and thus $q \in \mathrm{Vis}_\alpha(g)$. We conclude that $C \subseteq \mathrm{Vis}_\alpha(g)$, and that $\mathrm{Vis}_\alpha(g)$ is star-shaped.

Let a, b be the tangency points defining the cone C. We have $\angle pag = \angle pbg = \frac{\pi}{2}$, $\angle apb = 2\theta$, and $\angle agb = \pi - 2\theta$, and thus the kite $pagb$ is $O(\alpha)$-fat. Since $C \subseteq \mathrm{Vis}_\alpha(g)$, the union of kites over all $p \in \mathrm{Vis}_\alpha(g)$ is exactly $\mathrm{Vis}_\alpha(g)$. By applying Claim 5 on the set of kites (all have a common vertex g), we conclude that $\mathrm{Vis}_\alpha(g)$ is $O(\alpha)$-fat. ◀

Given a point $g \in P$, denote by D_g the disk with maximum radius centered at g and contained in P. Denote by R_g the radius of D_g. If g α-robustly guards p, then $D(g, \alpha\|g - p\|)$ is contained in P, and thus $\alpha\|g - p\| \le R_g$. We obtain the following observation.

▶ **Observation 7.** *If g α-robustly guards p, then $\|g - p\| \le R_g/\alpha$.*

By Observation 7, we know that $\mathrm{Vis}_\alpha(g)$ is contained in $D(g, R_g/\alpha)$. Moreover, by definition, $\mathrm{Vis}_\alpha(g) \subseteq \mathrm{Vis}(g)$. A point $p \in \mathrm{Vis}(g)$ belongs to $\mathrm{Vis}_\alpha(g)$ if and only if the ice cream cone from p to g is contained in $\mathrm{Vis}_\alpha(g)$. Hence, to compute $\mathrm{Vis}_\alpha(g)$ exactly, we need to take into account each of the reflex vertices that may "block" a potential ice cream cone.

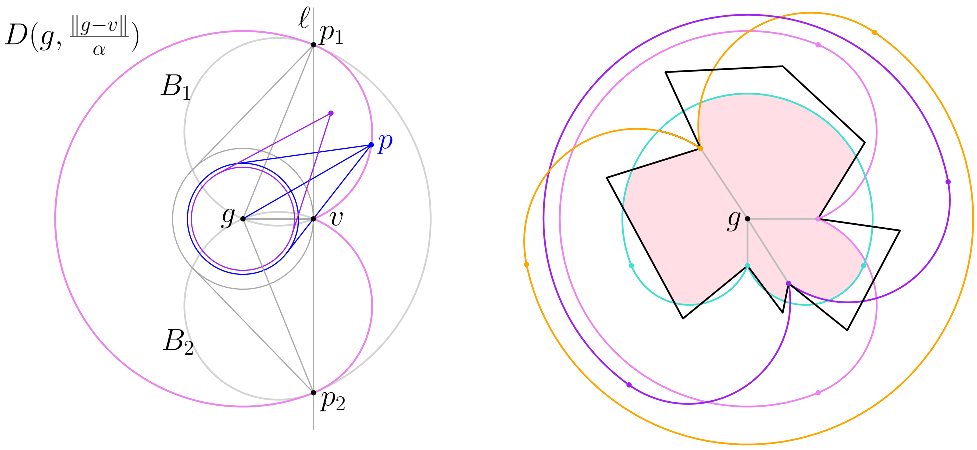

Figure 5 Computing $\mathrm{Vis}_\alpha(g)$ as the intersection of heart shapes for every reflex vertex of $\mathrm{Vis}(g)$. Left: the construction of a single heart shape (in violet). Right: $\mathrm{Vis}_\alpha(g)$ is the area shaded in pink.

We thus compute for each reflex vertex v the locus of all points p such that the ice cream cone from p to g is touching v. The ice cream cone may touch v either at a point on its circular arc or on one of its edges, so we get three relevant circles for each vertex v (see Figure 5, left). We then compute $\mathrm{Vis}_\alpha(g)$ as the intersection of those heart-shaped regions (one for each reflex vertex), $D(g, R_g/\alpha)$, and the boundary of P (see Figure 5, right). A full proof of Lemma 8 is given in the full version of the paper.

▶ **Lemma 8.** *Computing $Vis_\alpha(g)$ can be done in polynomial time.*

2.2 The robust inverse visibility region

The definition of α-robust guarding is not bidirectional; it is possible that g α-robustly guards p, but g is not α-robustly guarded by p. (In the full version of the paper, we discuss a more general notion of visibility that is bidirectional.) We therefore define the **robust inverse visibility region** of a point as follows. For a point $p \in P$, denote by $\mathrm{Vis}_\alpha^{\mathrm{inv}}(p)$ the set of points $g \in P$ such that p is α-robustly guarded by g. Although $\mathrm{Vis}_\alpha(g)$ is fat (as shown in Lemma 6), $\mathrm{Vis}_\alpha^{\mathrm{inv}}(p)$ is not necessarily fat; in fact, the robust inverse visibility region may be a single line segment (see Figure 6).

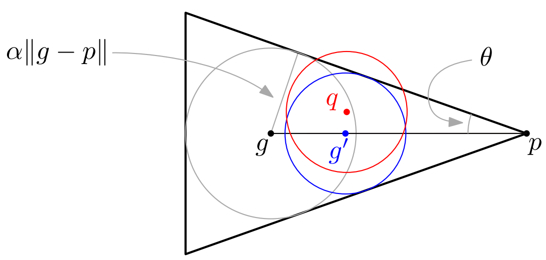

Figure 6 An isosceles triangle polygon with angle exactly 2θ at the point p. For any point g' on the segment $\overline{gp}$, the (blue) disk $D(g', \alpha\|p-g'\|)$ is tangent to the legs of the triangle, and thus g' α-robustly guards p. However, any point not on the segment does not α-robustly guard p. Thus $\mathrm{Vis}_\alpha^{\mathrm{inv}}(p) = \overline{gp}$.

Nevertheless, we will show how to construct a star-shaped $O(\alpha)$-fat polygon F_p that contains $\mathrm{Vis}_\alpha^{\mathrm{inv}}(p)$ with the property that any $g \in F_p$ $\alpha/2$-robustly guards p. First we need the following claim, which is illustrated in Figure 7. The proof is in the full version of the paper.

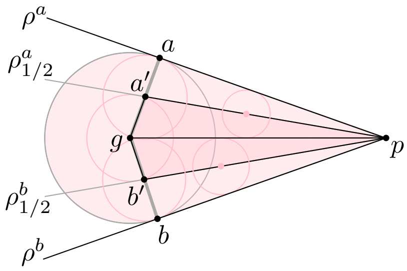

Figure 7 The kite $K = ga'pb'$ is $O(\alpha)$-fat, and any $q \in K$ $\alpha/2$-robustly guards p.

▷ Claim 9. Let g and p be points in P, such that g α-robustly guards p. Then there exists an $O(\alpha)$-fat kite K containing g and p, such that every point $q \in K$ $\alpha/2$-robustly guards p.

The lemma below now follows by taking the union of kites corresponding to every $g \in \mathrm{Vis}_\alpha^{\mathrm{inv}}(p)$. The complete proof is in the full version of the paper.

▶ **Lemma 10.** *Given a polygon P and a point $p \in P$, there exists a star-shaped $O(\alpha)$-fat polygon F_p that contains $Vis_\alpha^{inv}(p)$, and such that any $g \in F_p$ $\alpha/2$-robustly guards p. The size (radius of the smallest enclosing disk centered at p) of F_p is equal to that of $Vis_\alpha^{inv}(p)$.*

By definition, $\mathrm{Vis}_\alpha^{\mathrm{inv}}(p) \subseteq \mathrm{Vis}(p)$. Computing $\mathrm{Vis}_\alpha^{\mathrm{inv}}(p)$ can also be done in polynomial time, by computing a constant number of “constraints” per edge of $\mathrm{Vis}(p)$ (see Figure 8). A point $g \in P$ belongs to $\mathrm{Vis}_\alpha^{\mathrm{inv}}(p)$ if and only if the disk $D(g, \alpha\|p - g\|)$ is contained in $\mathrm{Vis}(p)$, or in other words, does not intersect any edge of $\mathrm{Vis}(p)$. For each edge $e = \{u, v\} \in P$, the locus of all points g such that $D(g, \alpha\|p - g\|)$ touches e, can be described by two disks (one per vertex) and a hyperbola (for the interior of the edge). A full proof of Lemma 11 is given in the full version of the paper.

▶ **Lemma 11.** *The region $Vis_\alpha^{inv}(p)$ can be computed in polynomial time.*

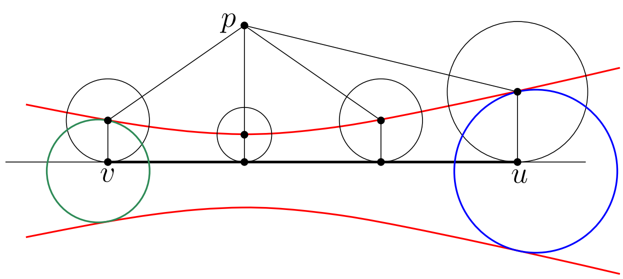

Figure 8 Three constraints defining the points g with $D(g, \alpha\|p - g\|)$ intersecting the edge $\{u, v\}$: the green disk containing v, the blue disk containing u, and the red hyperbola.

2.3 Hardness

The classic Art Gallery Problem is APX-hard by a reduction from the Hitting Lines problem [16]: Given a set $\mathcal{L}$ of lines, one is to find a minimum set of points that "hit" all the lines. The polygon constructed in the reduction is a "spike box" – a rectangle containing all the intersection points between lines in $\mathcal{L}$, and having a thin spike going out of it for each line. In order to guard the tip of a spike, one must place a guard in a small neighbourhood of the line segment (corresponding to a line in $\mathcal{L}$) generating the spike. Thus, hitting all of the lines is equivalent to guarding all spikes. In the full version of the paper we show a similar construction for the problem of α-robust guarding, and thereby obtain the following theorem.

▶ **Theorem 12.** *The α-robust guarding problem is APX-hard.*

2.4 Robustly guarding a discrete set of points

In this section we consider a discrete version of the robust guarding problem, where we are given a set S of m points in a polygon P, and the goal is to find a minimum set of α-robust guards for S. Besides being interesting in its own, the solution that we present will be used in the next section, where the goal is to α-robustly guard the entire polygon P.

Before we can present our algorithm, we need one more ingredient, regarding the fatness of our robust visibility polygons. The lemma below is a version of the well-known "fat-collection theorem" (see, e.g. [40]). Here, we define the **size** of a star-shaped object P with respect to a given center point o as the radius of its minimum enclosing ball centered at o. For completeness, we provide a proof in the full version of the paper.

▶ **Lemma 13.** *For any disk D of radius R, there exist a set C of $O(\alpha^{-2})$ points such that any α-fat star-shaped polygon that intersects D and has size at least R w.r.t. a given center point o, contains a point from C.*

The algorithm. Consider Algorithm 1, which gets as input the polygon P and the set S. In each iteration, the algorithm finds the point $g \in S$ with smallest $\mathrm{Vis}_\alpha^{\mathrm{inv}}(g)$, removes from S all the points s for which $\mathrm{Vis}_\alpha^{\mathrm{inv}}(s)$ intersects $\mathrm{Vis}_\alpha^{\mathrm{inv}}(g)$, and adds to the solution the corresponding set of hitting points from Lemma 13.

Algorithm 1 DiscreteRobustGuarding(P, S).

foreach $s \in S$ **do**
  Compute $\mathrm{Vis}_\alpha^{\mathrm{inv}}(s)$
  Compute $D(s)$, the minimum enclosing disk of $\mathrm{Vis}_\alpha^{\mathrm{inv}}(s)$ centered at s
$G \leftarrow \emptyset$
while $S \neq \emptyset$ **do**
  $g \leftarrow \mathrm{argmin}_{\mathrm{s} \in \mathrm{S}}\{\mathrm{size}(\mathrm{Vis}_\alpha^{\mathrm{inv}}(\mathrm{s}))\}$ ▷ `g is the point from S with smallest Vis`$_\alpha^{\mathrm{inv}}$
  $S(g) \leftarrow \{s \in S \mid \mathrm{Vis}_\alpha^{\mathrm{inv}}(s) \cap D(g) \neq \emptyset\}$
  $S \leftarrow S \setminus S(g)$
  Let $H(g)$ be the set of hitting points from Lemma 13 that correspond to $D(g)$, with fatness parameter $c \cdot \alpha$ (for a sufficiently small constant c).
  $G \leftarrow G \cup H(g)$
return G

▶ **Theorem 14.** *Given a polygon P with n vertices, and a set S of m points in P, one can compute in* $\mathrm{poly}(n, m)$ *time a set of $O(\alpha^{-2})|OPT_\alpha^S|$ points that $\alpha/2$-robustly guard S, where OPT_α^S is a minimum set of guards that α-robustly guard S.*

Proof. We show that the set G returned by Algorithm 1 satisfies the theorem.

Let $g_1, \ldots, g_k$ be the points from S that were found in line 6 of the algorithm, and consider the sequence of inverse visibility regions $\mathrm{Vis}_\alpha^{\mathrm{inv}}(g_1), \ldots, \mathrm{Vis}_\alpha^{\mathrm{inv}}(g_k)$. Any two of these regions are disjoint, because in line 8 we remove all points $s \in S$ for which $\mathrm{Vis}_\alpha^{\mathrm{inv}}(s) \cap \mathrm{Vis}_\alpha^{\mathrm{inv}}(g_i) \neq \emptyset$. Thus, the set $g_1, \ldots, g_k$ is a set of witnesses in P, i.e., no α-robust guard can guard both g_i and g_j for any $1 \le i, j \le k$. Therefore, in order to α-robustly guard S, one needs to put a guard in each $\mathrm{Vis}_\alpha^{\mathrm{inv}}(g_i)$, and hence $k \le |\mathrm{OPT}_\alpha^S|$.

The set $H(g)$ is the set of $O(\alpha^{-2})$ hitting points obtained from Lemma 13. Since we chose g_i to be the point in S with minimum size (i.e., size of inverse visibility region), any $s \in S(g_i)$ has size larger than the radius of $D(g_i)$. By Lemma 10, for any $s \in S(g_i)$ there exists a star-shaped $O(\alpha)$-fat polygon F_s of size $\mathrm{size}(\mathrm{Vis}_\alpha^{\mathrm{inv}}(s))$ that contains $\mathrm{Vis}_\alpha^{\mathrm{inv}}(s)$, and any point in F_s, $\alpha/2$-robustly guards s. Thus by Lemma 13 each such F_s contains a point g_s from $H(g_i)$, and g_s $\alpha/2$-robustly guards s. We get that for any $1 \le i \le k$, the set $H(g_i)$ guards $S(g_i)$, and therefore G is a set of $O(\alpha^{-2})|\mathrm{OPT}_\alpha^S|$ points that $\alpha/2$-robustly guards S. (Note that we do not compute F_s, we use Lemma 10 only to show that the set of hitting points is sufficient.)

By Lemma 11, computing $\mathrm{Vis}_\alpha^{-1}(s)$ for every $s \in S$ can be done in $m \cdot \mathrm{poly}(n)$ time. Clearly, the while loop is executed for at most m rounds, each round runs in $\mathrm{poly}(n, m)$ time. Thus the total running time of Algorithm 1 is in $\mathrm{poly}(n, m)$. ◀

3 An $O(1)$-approximation for robustly guarding a polygon

Let P be a polygon with n vertices, and a parameter $\alpha \le \sin\frac{\phi}{2}$, where ϕ is the smallest internal angle of P. For technical reasons, we also assume that $\alpha \le 1/2$. Let OPT_α be a minimum set of points that α-robustly guard P. Our goal is to find a set of $O(\mathrm{poly}(\alpha^{-1})) \cdot |\mathrm{OPT}_\alpha|$ points that $c\alpha$-robustly guard P, for some constant $c \le 1$. Note that for a smaller radius we need less guards, i.e., for $\alpha' < \alpha$, $|\mathrm{OPT}_{\alpha'}| \le |\mathrm{OPT}_\alpha|$.

3.1 A medial axis based decomposition

Consider the medial axis of P (the set of all points in P having more than one closest point on ∂P – the boundary of P), and let M be the set of its vertices that do not lie on ∂P. The medial axis is a planar graph G, with some line-segment edges and some curved edges (subcurves of a parabola). Note that we also include in M vertices of degree 2 that represent, e.g., the intersection point of a line segment and a parabola that define the medial axis (see Figure 9). For each $v \in M$, let D_v be the *medial* disk centered at v. The disk D_v touches ∂P in at least two points, and we denote by $\mathcal{D}$ the set of all medial disks, i.e., $\mathcal{D} = \{D_v\}_{v \in M}$.

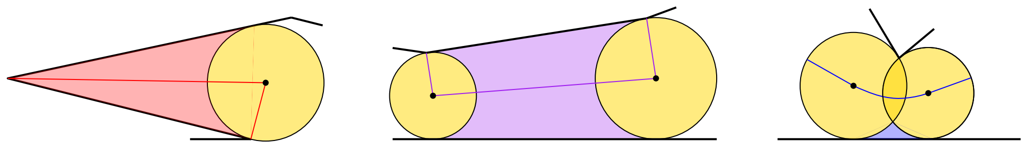

Figure 9 The three types of regions (cells) forming $P \setminus \mathcal{D}$; a red cell (left), a purple cell (middle), and a blue cell (right). The medial disks in $\mathcal{D}$ are shown in yellow.

Based on the structure of the medial axis, we decompose $P \setminus \mathcal{D}$ into 3 types of regions, Red, Purple, and Blue, as follows (see Figure 9).

- A Red region is a maximal connected region of $P \setminus \mathcal{D}$ which is bounded by two edges e_1, e_2 of P connected by a vertex w, and a single disk $D_v \in \mathcal{D}$, such that both e_1 and e_2 are tangent to D_v.
- A Purple region is a maximal connected region of $P \setminus \mathcal{D}$ which is bounded by two disjoint edges e_1, e_2 of P, and two disks $D_v, D_w \in \mathcal{D}$, such that each of e_1, e_2 is tangent to both D_v, D_w. Since there is no other feature of the polygon "between" D_v and D_w, the corresponding vertices $v, w \in M$ are connected by a line segment in the medial axis.
- A Blue region is any maximal connected region of $P \setminus \mathcal{D}$ which is neither red nor purple.

Adding medial disks in purple regions. In each purple region we add to M a set of vertices as follows. Let v, w be the two medial vertices that define a purple region Π, and assume that $R_v \geq R_w$ (see Figure 10). Consider the intersection $I = D(v, R_v/\alpha) \cap \Pi$, and notice that any point in I is α-robustly visible from v. If $D(v, R_v/\alpha)$ does not contain Π, then there are two intersection points, q_1, q_2, between $D(v, R_v/\alpha)$ and the edges defining Π. Let p_1 be the point on the segment $\overline{vw}$ such that the medial disk centered at p_1 touches the edges defining Π at the points q_1, q_2. We add p_1 to M, set $v = p_1$ and repeat the process, i.e. while $D(p_i, R_{p_i}/\alpha)$ does not contain the part of Π between $D(p_i, R_{p_i})$ and D_w, add to M the point p_{i+1} on the segment $\overline{vw}$, defined similarly but with respect to the disk $D(p_i, R_{p_i}/\alpha)$. Note that by adding the sequence $p_1, \ldots, p_k$ of additional medial vertices to M, we subdivide the purple region into $k+1$ smaller purple regions.

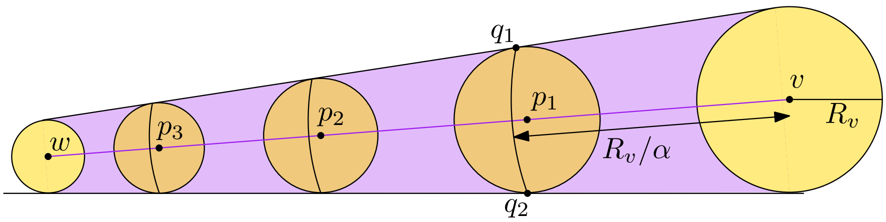

Figure 10 A purple region between D_w and D_v, and the added sequence of medial disks p_1, p_2, p_3.

Intuitively, those disks in the interior of a purple region Π were added in such a way that a single guard does not see too many of them robustly. More precisely, we have the following observation, which we prove in the full version of the paper.

▶ **Observation 15.** *For any $g \in P$, $Vis_\alpha(g)$ intersects at most four of the disks D_{p_i} that were added to* Π.

The observation below follows from the definitions of Purple and Red regions, and the observation that any blue region is bounded by two non-disjoint medial disks and one edge of P. A formal proof is given in the full version of the paper.

▶ **Observation 16.** *Any Red, Blue, or Purple region has at most two disks defining its boundary.*

Associating points with at most two medial disks. Given some point $p \in P$, we associate p with either one or two medial disks from $\mathcal{D}$ as follows. If p is contained in some disk from $\mathcal{D}$, we associate p with the largest disk from $\mathcal{D}$ that contains it. Otherwise, p belongs to either a Red, Blue, or Purple region, and we associate p with the disks from $\mathcal{D}$ defining that cell. (by Observation 16, there are at most two such disks).

▶ **Observation 17.** *Let p be a point in P.*

(i) *If $p \in D_v$ for some medial vertex v, then p is α-robustly visible from v.*

(ii) *If p is in a Red, Blue, or Purple region, then p is α-robustly visible from one of the centers of medial disks associated with p.*

Proof. The first statement is trivial. For the second statement, if p is in a Red region bounded by a medial disk D_v, then clearly p is visible from v (recall that we assume that convex angles in P are larger than 2θ). Else, if p is in a Blue region bounded by two non-disjoint medial disks D_v and D_u, then since $\alpha \leq \frac{1}{2}$ one of the following holds: (i) $\|p - v\| \leq R_v/\alpha$, and then p is α-robustly visible from v, or (ii) $\|p - u\| \leq R_u/\alpha$, and then p is α-robustly visible from u. Else, p is in a Purple region bounded by two medial disks D_v and D_u, such that $R_v \geq R_u$. By the construction of additional disks in purple regions, v sees the entire purple region α-robustly. ◀

▶ **Observation 18.** *Let g, p be two points in P such that g α-robustly guards p, and let D_v be the largest disk associated with g. Then $\alpha\|g - p\| \leq R_v$.*

Proof. For an edge $\{u, v\}$ of the medial axis with $R_u \leq R_v$, the radii of maximal disks with center on the edge is at most R_v. Points not on a medial edge clearly have smaller radius, and thus $R_g \leq R_v$ (recall that R_g is the radius of the largest disk centered at g and contained in P). By Observation 7, we get $\alpha\|p - g\| \leq R_g \leq R_v$. ◀

3.2 A set of candidate guards

For each $v \in M$ we place on D_v a set Q_v of $\Theta(\alpha^{-4})$ grid points, with edge length in $\Theta(\alpha^2 R_v)$. Denote $Q = M \cup \bigcup_{v \in M} Q_v$. For a point $g \in P$, let $Q(g) = \{Q_v \cup \{v\} \mid D_v \text{ is associated with } g\}$. As there are at most two disks associated with g, we have $|Q(g)| = O(\alpha^{-4})$. In this subsection we show that any α-robust guard g can be replaced by the set $Q(g)$ of $\alpha/4$-robust guards.

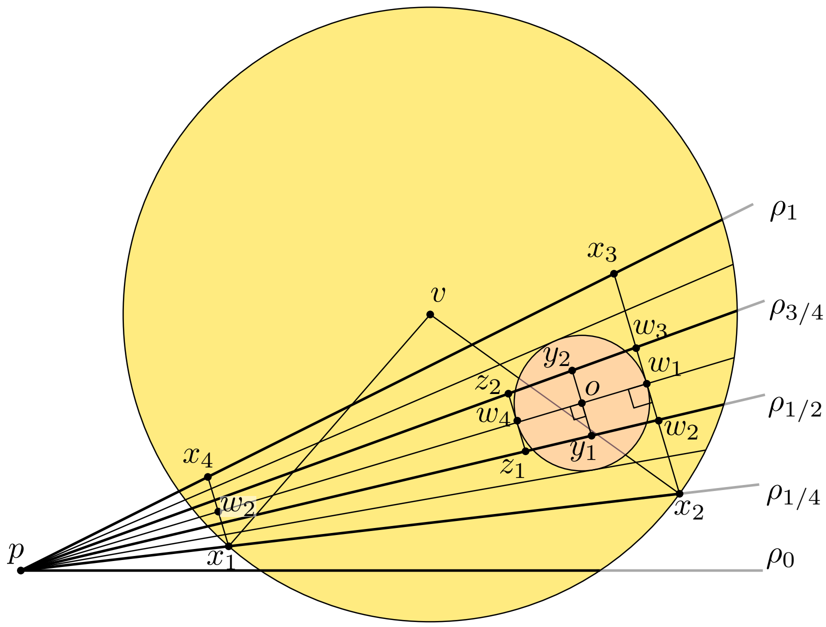

Figure 11 The construction for the proof of Lemma 19.

▶ **Lemma 19.** *Let K be a cone defined by two rays ρ_0, ρ_1 originated at p with small angle θ. If both ρ_0, ρ_1 intersect D_v, and p sees $K \cap D_v$, then there exists a grid point in Q_v that $\frac{\alpha}{4}$-robustly guards p.*

Proof. Assume for simplicity that ρ_0 lies on the x-axis, and ρ_1 lies above it. For $0 < \gamma < 1$, denote by ρ_γ the ray from p between ρ_0 and ρ_1 with angle $\gamma \cdot \theta$ from ρ_0. First, by Observation 4 for any point q that lies in K between $\rho_{1/4}$ and $\rho_{3/4}$, the disk $D(q, \frac{\alpha}{4}|p-q|)$ is contained in K. In addition, if v lies in K, then since both ρ_0, ρ_1 intersect D_v, we get that R_v is at least the distance between v and one of ρ_0, ρ_1. As in the proof of Observation 4, we get that $R_v \geq \sin\frac{\theta}{2}\|p-v\| > \frac{\sin\theta}{2}\|p-v\| = \frac{\alpha}{2}\|p-v\|$. Therefore, if v lies between $\rho_{1/4}$ and $\rho_{3/4}$, then clearly $D(v, \frac{\alpha}{4}\|p-v\|)$ is contained in $K \cap D_v$.

We thus assume w.l.o.g. that v lies above $\rho_{3/4}$ (the case when v lies below $\rho_{1/4}$ is symmetric). Our goal is to find a large enough square that lies in $K \cap D_v$ between $\rho_{1/4}$ and $\rho_{3/4}$, and such that for any point q in that square, $D(q, \frac{\alpha}{4}|p-q|)$ is contained in D_v.

In the following, we use some trigonometric identities, and the Maclaurin series expansions of some trigonometric functions (for $x \leq \frac{1}{2}$) to estimate distances. Specifically, $\tan(x) = \frac{\sin(x)}{\cos(x)}$, $\sin(2x) = 2\sin(x)\cos(x)$, $x \geq \sin(x) \geq x - \frac{x^3}{3!}$, $1 - \frac{x^2}{2} + \frac{x^4}{4!} \geq \cos(x) \geq 1 - \frac{x^2}{2}$, $2x \geq \arcsin(x) \geq x$, $2x \geq \tan(x) \geq x$.

Let x_1, x_2 be the two intersection points of $\rho_{1/4}$ and D_v (see Figure 11). We have

$$\|x_1 - x_2\| \geq 2\sin(\theta/4) \cdot R_v = \Theta(\theta) \cdot R_v$$

(we will get an equality when p is on ∂D_v, and ρ_0 is tangent to D_v). Consider the ray from x_2 perpendicular to $\rho_{5/8}$, and let w_1 be the intersection point with $\rho_{5/8}$, and x_3 the intersection point with ρ_1. Similarly, consider the ray from x_1 perpendicular to $\rho_{5/8}$, and let w_2 be the intersection point with $\rho_{5/8}$, and x_4 the intersection point with ρ_1. Since v is above $\rho_{1/2}$, the quadrilateral $x_1x_2x_3x_4$ is contained in D_v. Note that if p is on ∂P, we get $p = x_1 = x_4$, and $\triangle px_2x_3$ is a triangle contained in D_v.

We have

$$\|p - w_1\| = \cos(\frac{3\theta}{8})\|p - x_2\| \geq \cos(\frac{3\theta}{8})\|x_1 - x_2\| \geq (1 - \frac{9\theta^2}{64})\|x_1 - x_2\| = \Theta(\theta) \cdot R_v,$$

and

$$\|x_2 - w_1\| \geq \tan(\frac{3\theta}{8})\|p - w_1\| \geq \sin(\frac{3\theta}{8})\|p - x_2\| \geq \sin(\frac{3\theta}{8})\|x_1 - x_2\| = \Theta(\theta^2) \cdot R_v.$$

Let o be the point on $\overline{pw_1}$ such that $\|w_1 - o\| = \alpha/4\|o-p\|$, then the disk $D(o, \alpha/4\|o-p\|)$ is contained in $K \cap D_v$. Let y_1, y_2 be the points on $\rho_{1/2}, \rho_{3/4}$, respectively, such that the line through y_1, y_2 is the perpendicular to $\overline{pw_1}$ at o. Let z_1, z_2 be the points on $\rho_{1/2}, \rho_{3/4}$, respectively, such that the line through z_1, z_2 is the perpendicular to $\overline{pw_1}$ at the other intersection point, w_4, of $D(o, \alpha/4\|o-p\|)$ and $\overline{pw_1}$. For any point q in the quadrilateral z_1, y_1, y_2, z_2, the disk $D(q, \alpha/4\|p-q\|)$ is contained in D_v. We now claim that the quadrilateral $z_1y_1y_2z_2$ contains a disk of diameter in $\Theta(\|x_2 - w_1\|) = \Theta(\alpha^2) \cdot R_v$, and thus it must contain a grid point from Q_v. Intuitively, this is true because $\|z_1 - z_2\| = \Theta(\|x_2 - w_1\|)$ and $\|o - w_4\| = \alpha/4 \cdot \Theta(\|p - w_1\|)$. We provide detailed calculations for this claim in the full version of the paper, which finishes the proof. ◀

▶ **Lemma 20.** *For any $g \in P$, $Vis_\alpha(g) \subseteq \cup_{q \in Q(g)} Vis_{\alpha/4}(q)$.*

Proof. Let g, p be two points in P, such that p is α-robustly guarded by g. By Observation 17, if p is in one of the disks associated with g, or both p and g are in a Red, Blue, or Purple region, then p is α-robustly guarded by the centers of the disks associated with g. If this is not the case, then p, g must be in different cells of the arrangement $\mathcal{A}$, and therefore the segment $\overline{pg}$ must cross the boundary of some disk D_v associated with g.

Let ρ_0 be the ray from p in the direction of g, and assume for simplicity that ρ_0 lies on the x-axis. Denote by ρ^a (resp. ρ^b) the ray from p tangent to $D(g, \alpha\|p-g\|)$ above (resp. below) the x-axis. Denote by ρ^a_γ (resp. ρ^b_γ) the ray from p between ρ_0 and ρ^a (resp. ρ^b) with angle $\gamma \cdot \theta$ from ρ_0. Let a (resp. b) be the intersection point of ρ^a (resp. ρ^b) and $D(g, \alpha\|p-g\|)$.

Assume w.l.o.g. that v is above ρ_0, and let K be the cone defined by ρ_0 and ρ^a. Recall that D_v do not contain p, and $D_v \cap \overline{pg} \neq \emptyset$.

If $D_v \cap \overline{pa} \neq \emptyset$, then there exists a point $q_1 \in D_v \cap \overline{pa}$ and a point $q_2 \in D_v \cap \overline{pg}$ such that the triangle $\triangle pq_1q_2$ is contained in P, and since $q_1 \in \rho^a$ and $q_2 \in \rho_0$ we get that p sees $K \cap D_v$. Hence we can apply Lemma 19 on the rays ρ_0, ρ^a and get that there exists a grid point in Q_v that $\alpha/4$-robustly guards p.

Else, we are in the case when $D_v \cap \overline{pa} = \emptyset$. First, we claim that $g \in D_v$. Indeed, if $g \notin D_v$, then since $D_v \cap \overline{pg} \neq \emptyset$ it must be that v is between the vertical line trough g and the vertical line through p. Now, if v is above ρ^a, then clearly $D_v \cap \overline{pa} \neq \emptyset$. Else, if v is below ρ^a, then since by Observation 18 we have $R_v \geq \alpha\|p-g\|$, again we get that $D_v \cap \overline{pa} \neq \emptyset$.

Therefore, we are left with the following scenario: v is above ρ_0, $g \in D_v$, $a \notin D_v$, and $R_v \geq \alpha\|p-g\|$. Denote by w_1, w_2 the intersection points of D_v and $D(g, \alpha\|p-g\|)$, and consider the rays ρ_{w_1}, ρ_{w_2} from p to w_1, w_2, respectively. Since p sees both w_1, w_2, we get that p sees the cone between ρ_{w_1} and ρ_{w_2}. In the full version of the paper, we prove that $\angle w_1pw_2 \geq \theta$. This shows that we can apply Lemma 19 on this cone, and find a grid point in Q_v that $\alpha/4$-robustly sees p. ◀

By replacing each $g \in \mathrm{OPT}_\alpha$ with the set $Q(g)$, we get that the set $\bigcup_{g\in\mathrm{OPT}} Q(g)$ $\alpha/4$-robustly guards P. We obtain the following theorem.

▶ **Theorem 21.** *The set $Q = M \cup \bigcup_{v\in M} Q_v$ contains a set of $O(\alpha^{-4})|OPT_\alpha|$ points that $\alpha/4$-robustly guard P.*

In addition, we claim that the size of Q is linear in $n = |P|$ and $|OPT_\alpha|$.

▷ Claim 22. $|Q| = O(\alpha^{-4})(n + |OPT_\alpha|)$.

Proof. It is well known that the number of vertices that define the medial axis is $O(n)$. We only need to show that the number of vertices that we add in the purple regions is $O(|OPT_\alpha|)$. Indeed, by Observation 15, for any $g \in P$ $\mathrm{Vis}_\alpha(g)$ intersects at most four such consecutive disks, and thus the number of guards from OPT_α in a purple region with k additional disks is at least $\frac{k-8}{4}$ (4 from each side can be guarded by a point outside of the purple region). ◁

3.3 An $O(1)$-approximation greedy algorithm

Let Q be the set of candidate guards from Theorem 21, constructed with parameter $\alpha/8$ instead of α. Consider the arrangement $\mathcal{A}$ formed by the set of visibility regions $\{\mathrm{Vis}_{\alpha/8}(q) \mid q \in Q\}$. For each cell of this arrangement, we pick one sample point in the interior of the cell, and denote by S the set of these sample points.

▶ **Observation 23.** *If $Q' \subseteq Q$ $\frac{\alpha}{8}$-robustly guards S, then Q' $\frac{\alpha}{8}$-robustly guards all of P.*

Proof. Any guard that $\alpha/8$-robustly guards a point in the interior of a cell in $\mathcal{A}$ must $\alpha/8$-robustly guard the entire cell; otherwise, this cell would be subdivided. ◀

By Observation 23, it is enough to $\alpha/8$-robustly guard S from points in Q. We run Algorithm 1 on P and the set S, and by Theorem 14 we get a set G of $O(\alpha^{-2})|\mathrm{OPT}^S_\alpha| \leq O(\alpha^{-2})|\mathrm{OPT}_\alpha|$ points that $\alpha/2$-robustly guard S. However, in order to guard the entire polygon P, we need to guard S from points in Q only. So, we replace each point $g \in G$ by the set $Q(g)$ from Lemma 20. For any $g \in P$ we have $|Q(g)| = O(\alpha^{-4})$, so we obtain a set Q' of $O(\alpha^{-6})|OPT_\alpha|$ points from Q that $\alpha/8$-robustly guard S. By Observation 23, the set Q' $\alpha/8$-robustly guards P.

Computing the set Q, the arrangement $\mathcal{A}$, and the set S can be done in $\mathrm{poly}(n, |\mathrm{OPT}_\alpha|)$ time by Lemma 8 and Claim 22. By Theorem 14, the running time of Algorithm 1 is $\mathrm{poly}(n, |\mathrm{OPT}_\alpha|)$.

In fact, if we only want to return a constant factor approximation of $|\mathrm{OPT}_\alpha|$, we can do so in $\mathrm{poly}(n)$ time, as follows. By Observation 15, if the number of additional disks placed in a purple region Π is $k > 8$, then the number of guards placed in Π in an optimal solution is $\Omega(k)$. Moreover, except for $O(1)$ of these guards, none of them sees points outside Π. Therefore, we do not need to compute these guards explicitly in order to produce the rest of the guards, and we only record their number (which is a simple function of the dimensions of the purple region) in $O(1)$ time. Thus, we can cut the inner part (region inside a purple region excluding 4 disks at each of its ends; see Figure 12) of those purple regions from P, and obtain a set of disjoint subpolygons having in total $O(\alpha^{-4}) \cdot n$ candidate grid points. By applying the same algorithm separately on each subpolygon and then combining the solutions, we only loose a constant number of guards per purple region. Therefore we can output a constant factor approximation of $|\mathrm{OPT}_\alpha|$ in $\mathrm{poly}(n)$ time. To produce an explicit solution from this implicit representation, we only need to run the algorithm that computes the set of guards in each of the purple regions, in time linear in their number.

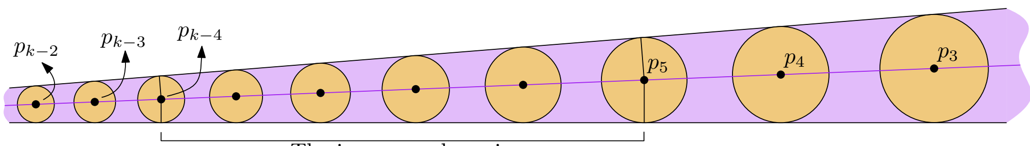

Figure 12 The inner part of a purple region with $k > 8$ added medial vertices.

▶ **Theorem 24.** *Given a polygon P with n vertices, one can compute in* $\mathrm{poly}(n)$ *time the cardinality of, and an implicit representation of, a set of $O(\alpha^{-6})|OPT_\alpha|$ points that $\alpha/8$-robustly guard P, where OPT_α is a minimum-cardinality set of guards that α-robustly guard P. In additional time $O(|OPT_\alpha|)$ we can output an explicit set of such points.*

References

1 Mikkel Abrahamsen, Anna Adamaszek, and Tillmann Miltzow. Irrational guards are sometimes needed. In *33rd International Symposium on Computational Geometry (SoCG 2017)*, volume 77, page 3. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2017.

2 Mikkel Abrahamsen, Anna Adamaszek, and Tillmann Miltzow. The art gallery problem is $\exists\mathbb{R}$-complete. *ACM Journal of the ACM (JACM)*, 69(1):1–70, 2021.

3 Akanksha Agrawal, Kristine VK Knudsen, Daniel Lokshtanov, Saket Saurabh, and Meirav Zehavi. The parameterized complexity of guarding almost convex polygons. In *36th International Symposium on Computational Geometry, SoCG 2020*, pages LIPIcs–SoCG. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2020.

4 Akanksha Agrawal and Meirav Zehavi. Parameterized analysis of art gallery and terrain guarding. In *Computer Science–Theory and Applications: 15th International Computer Science Symposium in Russia, CSR 2020, Yekaterinburg, Russia, June 29–July 3, 2020, Proceedings 15*, pages 16–29. Springer, 2020.

5 Greg Aloupis, Prosenjit Bose, Vida Dujmovic, Chris Gray, Stefan Langerman, and Bettina Speckmann. Triangulating and guarding realistic polygons. *Comput. Geom.*, 47(2):296–306, 2014.

6 Yoav Amit, Joseph S. B. Mitchell, and Eli Packer. Locating guards for visibility coverage of polygons. *International Journal of Computational Geometry & Applications*, 20(05):601–630, 2010.

7 Rom Aschner, Matthew J. Katz, Gila Morgenstern, and Yelena Yuditsky. Approximation schemes for covering and packing. In Subir Kumar Ghosh and Takeshi Tokuyama, editors, *WALCOM: Algorithms and Computation, 7th International Workshop, WALCOM 2013, Kharagpur, India, February 14-16, 2013. Proceedings*, volume 7748 of *Lecture Notes in Computer Science*, pages 89–100, 2013.

8 Pradeesha Ashok and Meghana M Reddy. Efficient guarding of polygons and terrains. In *Frontiers in Algorithmics: 13th International Workshop, FAW 2019, Sanya, China, April 29–May 3, 2019, Proceedings 13*, pages 26–37. Springer, 2019.

9 Stav Ashur, Omrit Filtser, and Matthew J. Katz. A constant-factor approximation algorithm for vertex guarding a wv-polygon. *J. Comput. Geom.*, 12(1):128–144, 2021.

10 Stav Ashur, Omrit Filtser, Matthew J. Katz, and Rachel Saban. Terrain-like graphs: Ptass for guarding weakly-visible polygons and terrains. *Comput. Geom.*, 101:101832, 2022.

11 Boaz Ben-Moshe, Matthew J. Katz, and Joseph S. B. Mitchell. A constant-factor approximation algorithm for optimal 1.5d terrain guarding. *SIAM J. Comput.*, 36(6):1631–1647, 2007.

12 Pritam Bhattacharya, Subir Kumar Ghosh, and Bodhayan Roy. Approximability of guarding weak visibility polygons. *Discrete Applied Mathematics*, 228:109–129, 2017.

13 Édouard Bonnet and Tillmann Miltzow. An approximation algorithm for the art gallery problem. In *33rd Symposium on Computational Geometry (SoCG 2017)*, volume 77, page 20, 2017.

14 Édouard Bonnet and Tillmann Miltzow. Parameterized hardness of art gallery problems. *ACM Transactions on Algorithms (TALG)*, 16(4):1–23, 2020.

15 Dorit Borrmann, Pedro J de Rezende, Cid C de Souza, Sándor P Fekete, Stephan Friedrichs, Alexander Kröller, Andreas Nüchter, Christiane Schmidt, and Davi C Tozoni. Point guards and point clouds: Solving general art gallery problems. In *Proceedings of the 29th Symposium on Computational Geometry*, pages 347–348, 2013.

16 Björn Brodén, Mikael Hammar, and Bengt J Nilsson. Guarding lines and 2-link polygons is apx-hard. In *The Thirteenth Canadian Conference on Computational Geometry-CCCG'01, University of Waterloo, ON (13-15 August 2001)*, pages 45–48. Proc. 13th Canadian Conference on Computational Geometry, 2001.

17 Hervé Brönnimann and Michael T. Goodrich. Almost optimal set covers in finite vc-dimension. *Discret. Comput. Geom.*, 14(4):463–479, 1995.

18 Ajay Deshpande, Taejung Kim, Erik D. Demaine, and Sanjay E. Sarma. A pseudopolynomial time $O(\log n)$-approximation algorithm for art gallery problems. In *Algorithms and Data Structures, 10th International Workshop, WADS'07, Halifax, Canada, August 15-17, 2007, Proceedings*, pages 163–174, 2007.

19 Michael Gene Dobbins, Andreas Holmsen, and Tillmann Miltzow. Smoothed analysis of the art gallery problem. *arXiv preprint*, 2018. `arXiv:1811.01177`.

20 Alon Efrat. The complexity of the union of (alpha, beta)-covered objects. *SIAM J. Comput.*, 34(4):775–787, 2005.

21 Alon Efrat and Sariel Har-Peled. Guarding galleries and terrains. *Inf. Process. Lett.*, 100(6):238–245, 2006. `doi:10.1016/j.ipl.2006.05.014`.

22 Alon Efrat, Sariel Har-Peled, and Joseph S. B. Mitchell. Approximation algorithms for location problems in sensor networks. In *2nd International Conference on Broadband Networks (BROADNETS 2005), 3-7 October 2005, Boston, Massachusetts, USA*, pages 767–776, 2005. doi:10.1109/ICBN.2005.1589677.
23 Stephan Eidenbenz, Christoph Stamm, and Peter Widmayer. Inapproximability results for guarding polygons and terrains. *Algorithmica*, 31(1):79–113, 2001.
24 Jeff Erickson, Ivor Van Der Hoog, and Tillmann Miltzow. Smoothing the gap between np and er. *SIAM Journal on Computing*, pages FOCS20–102 – FOCS20–138, 2022.
25 Stephan Friedrichs, Michael Hemmer, James King, and Christiane Schmidt. The continuous 1.5D terrain guarding problem: Discretization, optimal solutions, and PTAS. *JoCG*, 7(1):256–284, 2016.
26 Matt Gibson, Gaurav Kanade, Erik Krohn, and Kasturi Varadarajan. Guarding terrains via local search. *Journal of Computational Geometry*, 5(1):168–178, 2014.
27 Simon B Hengeveld and Tillmann Miltzow. A practical algorithm with performance guarantees for the art gallery problem. In *37th International Symposium on Computational Geometry*, 2021.
28 Gil Kalai and Jiří Matoušek. Guarding galleries where every point sees a large area. *Israel Journal of Mathematics*, 101(1):125–139, 1997.
29 James King and David G. Kirkpatrick. Improved approximation for guarding simple galleries from the perimeter. *Discrete & Computational Geometry*, 46(2):252–269, 2011. doi:10.1007/s00454-011-9352-x.
30 David Kirkpatrick. Guarding galleries with no nooks. In *Proceedings of the 12th Canadian Conference on Computational Geometry*, pages 43–46. Citeseer, 2000.
31 David Kirkpatrick. An O(lg lg OPT)-approximation algorithm for multi-guarding galleries. *Discrete & Computational Geometry*, 53:327–343, 2015.
32 Erik A Krohn and Bengt J Nilsson. Approximate guarding of monotone and rectilinear polygons. *Algorithmica*, 66:564–594, 2013.
33 D Lee and Arthurk Lin. Computational complexity of art gallery problems. *IEEE Transactions on Information Theory*, 32(2):276–282, 1986.
34 Lucas Meijer and Tillmann Miltzow. Sometimes two irrational guards are needed. *arXiv preprint*, 2022. arXiv:2212.01211.
35 Joseph O'Rourke. Visibility. In *Handbook of discrete and computational geometry*, pages 875–896. Chapman and Hall/CRC, 2017.
36 Thomas C Shermer. Recent results in art galleries (geometry). *Proceedings of the IEEE*, 80(9):1384–1399, 1992.
37 Davi C Tozoni, Pedro J De Rezende, and Cid C De Souza. Algorithm 966: A practical iterative algorithm for the art gallery problem using integer linear programming. *ACM Transactions on Mathematical Software (TOMS)*, 43(2):1–27, 2016.
38 Pavel Valtr. Guarding galleries where no point sees a small area. *Israel Journal of Mathematics*, 104:1–16, 1998.
39 Pavel Valtr. On galleries with no bad points. *Discrete & Computational Geometry*, 21:193–200, 1999.
40 A. Frank van der Stappen and Mark H. Overmars. Motion planning amidst fat obstacles (extended abstract). In Kurt Mehlhorn, editor, *Proceedings of the Tenth Annual Symposium on Computational Geometry, Stony Brook, New York, USA, June 6-8, 1994*, pages 31–40. ACM, 1994.