Available online at www.sciencedirect.com
SCIENCE DIRECT®
Theoretical Computer Science 354 (2006) 379–390
Theoretical Computer Science
www.elsevier.com/locate/tcs

On boundaries of highly visible spaces and applications☆

John H. Reif[a], Zheng Sun[b,*]

[a]Department of Computer Science, Duke University, Box 90129, Durham, NC 27708-0129, USA
[b]Google Inc., 1600 Amphitheatre Parkway, Mountain View, CA 94043, USA

Abstract

The purpose of this paper is to investigate the properties of a certain class of highly visible spaces. For a given geometric space $\mathscr{C}$ containing *obstacles* specified by disjoint subsets of $\mathscr{C}$, the *free space* $\mathscr{F}$ is defined to be the portion of $\mathscr{C}$ not occupied by these obstacles. The space is said to be *highly visible* if at each point in $\mathscr{F}$ a viewer can see at least an ε fraction of the entire $\mathscr{F}$. This assumption has been used for robotic motion planning in the analysis of random sampling of points in the robot's configuration space, as well as the upper bound of the minimum number of guards needed for art gallery problems. However, there is no prior result on the implication of this assumption to the geometry of the space under study. For the two-dimensional case, with the additional assumptions that $\mathscr{C}$ is bounded within a rectangle of constant aspect ratio and that the volume ratio between $\mathscr{F}$ and $\mathscr{C}$ is a constant, we use the proof technique of "charging" each obstacle boundary segment by a certain portion of $\mathscr{C}$ to show that the total length of all obstacle boundaries in $\mathscr{C}$ is $O(\sqrt{n\mu(\mathscr{F})/\varepsilon})$, if $\mathscr{C}$ contains polygonal obstacles with a total of n boundary edges; or $O(\sqrt{n}\mu(\mathscr{F})/\varepsilon)$, if $\mathscr{C}$ contains n convex obstacles that are piecewise smooth. In both cases, $\mu(\mathscr{F})$ is the volume of $\mathscr{F}$. For the polygonal case, this bound is tight as we can construct a space whose boundary size is $\Theta(\sqrt{n\mu(\mathscr{F})/\varepsilon})$. These results can be partially extended to three dimensions. We show that these results can be applied to the analysis of certain probabilistic roadmap planners, as well as a variant of the art gallery problem. We also propose a number of conjectures on the properties of these highly visible spaces.

Keywords: Motion planning; Art gallery problem; Computational geometry; Approximation algorithms; Robotics

1. Introduction

Computational geometry is now a mature field with a multiplicity of well-defined foundational problems associated with, for many cases, efficient algorithms as well as well-established applications over a broad range of areas including computer vision, robotic motion planning and rendering. However, as compared to some other fields, the field of computational geometry has not yet explored as much the methodology of looking at reasonable sub-cases of inputs that appear in practice for practical problems. For example, in matrix computation, there is a well-established set of specialized matrices, such as sparse matrices, structured matrices, and banded matrices, for which there are especially efficient algorithms.

☆ This work was supported in part by NSF ITR Grant EIA-0086015, DARPA/AFSOR Contract F30602-01-2-0561, NSF EIA-0218376, and NSF EIA-0218359.
* Corresponding author.
E-mail addresses: reif@cs.duke.edu (J.H. Reif), sunzheng@gmail.com (Z. Sun).

0304-3975/$ - see front matter
doi:10.1016/j.tcs.2005.11.027

One assumption that has been used in a number of previous works in computational geometry is the ε-visibility assumption. Let $\mathcal{C}$ be a geometric space containing a specified set of *obstacles* defined to be compact subsets of $\mathcal{C}$. We define the *obstacle space* $\mathcal{B}$ to be the union of all obstacles, and define the *free space* $\mathcal{F}$ to be $\mathcal{C}\backslash\mathcal{B}$, the portion of $\mathcal{C}$ not occupied by the obstacles. A point v in $\mathcal{F}$ is said to be ε-visible if a viewer at v can see at least an ε fraction of $\mathcal{F}$. The space $\mathcal{C}$ is said to be ε-visible if all points in $\mathcal{F}$ are ε-visible. Please note that some of the prior authors called this assumption instead ε-goodness.

1.1. Probabilistic roadmap planners

The ε-visibility assumption, in particular, has been used in the analysis of the probabilistic roadmap (PRM) methods for planning obstacle-avoiding motions for rigid or articulated robots. The robotic motion planning problems are known to be intractable (e.g. [21,24,8,25,22]). Except for a few special cases with small degrees of freedom (dof), the complete and general motion planning algorithms found are very costly. Therefore, an increasing attention has been paid during the last decade to efficient randomized heuristics that can arrive at, with a good probability, an approximated solution for any instance of a motion planning problem.

Instead of directly searching the whole free space $\mathcal{F}$ of the robot's configuration space $\mathcal{C}$ for a feasible path, a classic PRM planner [19,12,20,14] proceeds in two phases. It first randomly picks in $\mathcal{F}$ a set of points, called *milestones*. With these milestones, it constructs a *roadmap* by connecting each pair of milestones between which a collision-free path can be computed using a simple local planner. For any given query with initial and goal configurations s and t, the planner first finds two milestones s' and t' such that a simple collision-free path can be found connecting s (t, respectively) with s' (t', respectively) and then searches the roadmap for a path connecting s' and t'. The PRM planners have proved to be very effective in practice, capable of solving robotic motion planning problems with many dof. They also find applications in other areas such as virtual prototyping, computer animation, computational biology, etc.

The performance of a PRM planner depends on two key features of the roadmaps it constructs, *visibility* and *connectivity*. Firstly, for any given (initial or goal) configuration v, there should exist in the roadmap a milestone v' such that a local planner can find a path connecting v and v'. Since in practice most PRM planners use local planners that connect configurations by straight line segments, this implies that the milestones collectively need to see the entire (or at least a significant portion of) free space.

Secondly, the roadmap should capture the connectivity of the free space it represents. Any two milestones in the same connected component of the free space should also be connected via the roadmap, or otherwise the planner would give "false negative" answers to some queries. In recent years several PRM planners have been proposed (e.g. [9,1,2,17,18,10]) to improve the connectivity of the roadmap constructed by boosting the sampling density in the narrow passages of $\mathcal{F}$.

The earlier PRM planners pick milestones with a uniform distribution in the free space. The success of these planners motivated Kavraki et al. [13] to establish a theoretical foundation for the effectiveness of this sampling method. They showed that, for an ε-visible configuration space, $O((1/\varepsilon)\log 1/\varepsilon)$ milestones uniformly sampled in the free space are needed to *adequately* cover the free space with a *high* probability. ⟨*Be specific about high probability and adequate*.⟩ Hsu et al. [9] used the ε-visibility assumption along with other assumptions to prove the connectivity of the roadmap generated by uniform sampling.

Recently some new PRMs [1,2] that (randomly) pick milestones close to boundaries of obstacles have been proposed. These planners have shown to be more efficient than the earlier PRMs based on uniform sampling in the free space by better capturing narrow passages in the configuration space; that is, the roadmaps they construct have better connectivity. However, there has been no prior theoretical result on the visibility of the roadmaps constructed using the sampled boundary points.

This prompts us to propose a new notion of visibility, which we call *ε'-boundary-visibility*. For a given $\varepsilon' > 0$, a configuration space is said to be ε'-boundary-visible if every point in the free space can view at least ε' fraction of the entire obstacle boundary. Using the same proof technique as [13], [1] one can show that, in an ε'-boundary-visible configuration space, $O((1/\varepsilon')\log 1/\varepsilon')$ milestones uniformly sampled on obstacle boundaries are needed to adequately cover the free space with a high probability.

[1] The difference is that, in our proof, every point v in the free space sees at least ε' fraction of obstacle boundaries, and therefore the probability that k points uniformly sampled on obstacle boundaries cannot see v is $(1-\varepsilon')^k$.

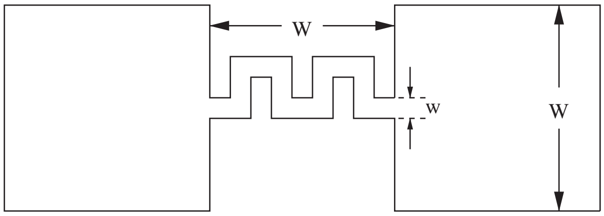

Fig. 1. $\varepsilon' = \sqrt{\varepsilon}$.

With the existence of narrow passages, it appears that $\varepsilon' > \varepsilon$ for most of the cases. Fig. 1 shows a space with two relatively large chambers connected by a narrow corridor. Let W (w, respectively) be the width of the chambers (corridor, respectively). This space is ε-visible when $\varepsilon = \Theta((w/W)^2)$, and also ε'-boundary-visible when $\varepsilon' = \Theta(w/W) = \Theta(\sqrt{\varepsilon})$. However, one can also construct an ε-visible and ε'-boundary-visible space where ε' is much smaller than ε, as shown in Section 3.

This raises the following question: under the ε-visibility assumption, what is the lower bound of ε' with respect to ε?

1.2. Art gallery problems

The ε-visibility assumption has also been used in bounding the number of guards needed for art gallery problems [11,27,28,15]. Potentially, this assumption might also allow for much more efficient algorithms in this case. The assumption appears to be reasonable in large number of practical cases as long as the considered area is within a closed area (such as a room).

The original art gallery problem was first proposed by V. Klee, who described the problem as the following: how many guards are necessary, and how many guards are sufficient, to guard the paintings and works of art in an art gallery with n walls? Later, Chvátal [3] showed that $\lfloor n/3 \rfloor$ guards are always sufficient and occasionally necessary to guard a simple polygon with n edges. Since then, there have been numerous variants of the art gallery problem proposed, including vertex guard problem, edge guard problem, fortress and prison yard problems, etc. (See [26] for a comprehensive review of various art gallery problems.)

Although for the worst case the number of guards needed is $\Theta(n)$ for polygonal galleries with n edges, intuitively, one would expect that galleries that are ε-visible should require much fewer guards. By translating the result of Kavraki et al. [13] into the context of art gallery problems, a uniformly random placement of $O((1/\varepsilon)\log 1/\varepsilon)$ guards is very likely to guard an adequate portion of the gallery. Kavraki et al. [13] also conjectured that in d dimensions any ε-visible polygonal gallery with h holes can be guarded by at most $f_d(h, 1/\varepsilon)$ guards, for some polynomial function f_d. Following some ideas of an earlier work by Kalai and Matoušek [11], Valtr [27] confirmed the $2\mathcal{D}$ version of the conjecture by showing that $f_2(h, 1/\varepsilon) = (2 + o(1))(1/\varepsilon)\log(1/\varepsilon)\log(h + 2)$. However, Valtr [28] disapproved the $3\mathcal{D}$ version of the conjecture by constructing for any integer k a $\frac{5}{9}$-visible art gallery that cannot be guarded by k guards. Kirkpatrick [15] later showed that $64(1/\varepsilon)\log\log 1/\varepsilon$ vertex guards are needed to guard all vertices of a simply connected polygon P that has the property that each vertex of P can see at least ε fraction of the other vertices of P. He also gave a similar result for boundary guards.

It has been proved that, for various art galleries problems, finding the *minimum* number of guards is difficult. Lee and Lin [16] proved that the minimum vertex guard problem for polygons is NP-hard. Schuchardt and Hecker [23] further showed that even for orthogonal polygons, whose edges are parallel to either the x-axis or the y-axis, the minimum vertex and point guard problems are NP-hard.

There have not been many approximation algorithms for art gallery problems. Ghosh [7] presented an $O(n^5 \log n)$ algorithm that can compute a vertex guard set whose size is at most $O(\log n)$ times the minimum number of vertex guards needed. However, Eidenbenz et al. [5,6] showed that vertex guard, edge guard, and point guard problems for polygons with holes cannot be approximated by any polynomial algorithm within a ratio of $(1 - \varepsilon)/12 \ln n$ for any $\varepsilon > 0$, unless $NP \subseteq TIME(n^{O(\log\log n)})$. Later Eidenbenz [4] further proved that, for polygons without holes, vertex guard, edge guard, and point guard problems are *APX*-hard, which means that there exists an $\varepsilon_0 > 0$ such that no polynomial time approximation algorithm can guarantee to solve any of these problems with an approximation ratio of $1 + \varepsilon_0$, unless $NP = P$.

With the assumption of ε-visibility, however, one can use a simple and efficient randomized approximation algorithm based on the result of Kavraki et al. [13] for the original art gallery problem. Moreover, this approximation algorithm does not require the assumption that the space is polygonal.

1.3. Our result

Intuitively, for an ε-visible space, the total size of all obstacle boundaries cannot be arbitrarily large; an excessive size of obstacle boundaries would inevitably cause a point in $\mathcal{F}$ to lose ε-visibility by blocking a significant portion of its view. Our main result of this paper is an upper bound of the boundary size of ε-visible spaces in two and (in some special cases) three dimensions. The upper bound of the boundary size not only is a fundamental property for the geometric spaces of this type, but also may have implications to other applications that use this assumption.

We show that, for an ε-visible $2\mathcal{D}$ space, the total length of all obstacle boundaries is $O(\sqrt{n\mu(\mathcal{F})/\varepsilon})$, if the space contains polygonal obstacles with a total of n boundary edges; or $O(\sqrt{n\mu(\mathcal{F})}/\varepsilon)$, if the space contains n convex obstacles that are piecewise smooth. In both cases, $\mu(\mathcal{F})$ is the area of $\mathcal{F}$. For the case of polygonal obstacles, this bound is tight as one can construct an ε-visible space containing obstacle boundaries with a total length of $\Theta(\sqrt{n\mu(\mathcal{F})/\varepsilon})$.

Our result can be used to bound the number of guards needed for the following variant of the original art gallery problem: given a space with a specified set of obstacles, how to put points on boundaries of obstacles so that these points see the entire (or a significant portion of) space. We call this problem *boundary art gallery problem*. Using Lin's proof technique [16] it is easy to show that this problem is also NP-hard. This problem can find applications in practical situations where the physical constraints would only allow points to be placed on obstacle boundaries. For example, one might need to install lights on the walls to enlighten a closed space consisting of rooms and corridors.

If this result can be extended to higher dimensions, we can also apply it to bounding the number of milestones needed to adequately cover the free space for PRM planners [1,2] that place milestones "pseudo-uniformly" on the boundary of the free space using various techniques.

2. Bounding boundary size for $2\mathcal{D}$ and $3\mathcal{D}$ ϵ-visible spaces

In this section we prove an upper bound of the boundary size of $2\mathcal{D}$ ε-visible spaces. We also show that this result can be partially extended to $3\mathcal{D}$ ε-visible spaces.

2.1. Preliminaries

Suppose $\mathcal{C}$ is the $2\mathcal{D}$ space bounded inside a rectangle $\mathcal{R}$. We let $\partial\mathcal{B}$ denote the boundaries of all obstacles. For each point $v \in \mathcal{F}$, we define the *visibility set* $\mathcal{V}_v$ of v to be $\{v' \mid \text{line segment } \overline{vv'} \subset \mathcal{F}\}$. That is, $\mathcal{V}_v$ is the set of all free space points that can be seen from v.

We assume that the *aspect ratio* λ of $\mathcal{R}$, defined to be the ratio between the lengths of the shorter and longer sides of $\mathcal{R}$, is no less than λ_0, where λ_0 is a constant between 0 and 1. We also assume that the *free space ratio* $\rho = \mu(\mathcal{F})/\mu(\mathcal{C})$ is no less than ρ_0 for some constant $\rho_0 > 0$.

The boundary size $|\partial\mathcal{B}|$ can be arbitrarily large if the aspect ratio λ is not lower-bounded. As shown in Fig. 2a, a rectangular space with length $1/\sqrt{t}$ and width $\sqrt{t}$, for some $0 < t < 1$, is divided into $n = 1/\varepsilon$ chambers. This space is ε-visible, as for each point $v \in \mathcal{F}$, the visibility set of v (which is the chamber containing v) is exactly ε fraction of $\mathcal{F}$. For this space, we have $\lambda = t$, $\rho = 1$, $\mu(\mathcal{F}) = 1$, and $|\partial\mathcal{B}| = \Theta(1/\varepsilon\sqrt{t})$, and therefore the boundary size $|\partial\mathcal{B}|$ cannot be bounded with respect to $\mu(\mathcal{F})$.

Fig. 2b shows that, even if the aspect ratio is 1, the boundary size $|\partial\mathcal{B}|$ still cannot be bounded with respect to $\mu(\mathcal{F})$ unless the free space ratio ρ is also lower-bounded. In this example, a rectangular space exactly the same as the one in Fig. 2a is padded with a rectangle obstacle of length $1/\sqrt{t}$ and width $(1-t)/\sqrt{t}$ (so that the entire space is a $1/\sqrt{t} \times 1/\sqrt{t}$ square). This space is ε-visible, with $\lambda = 1$, $\rho = t$, $\mu(\mathcal{F}) = 1$, and $|\partial\mathcal{B}| = \Theta(1/\varepsilon\sqrt{t})$. Therefore, the boundary size $|\partial\mathcal{B}|$ can be arbitrarily large if we let $t \to 0$.

A segment of the boundary (which we call *sub-boundary*) of an obstacle is said to be *smooth* if the curvature is continuous along the curve defining the boundary. The boundary of an obstacle is said to be *piecewise smooth* if it

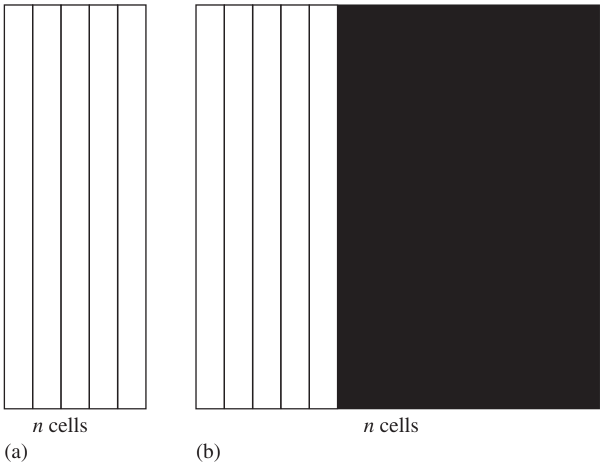

Fig. 2. Bad cases where the boundary size can be arbitrarily large: (a) Small λ and constant ρ; (b) small ρ and constant λ.

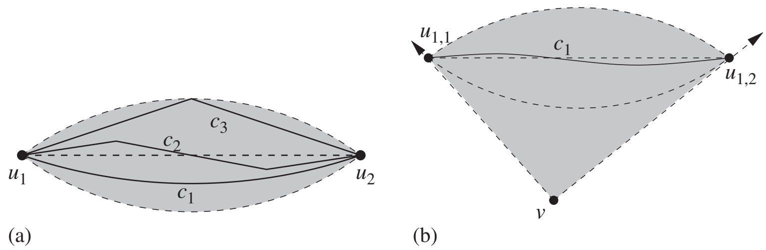

Fig. 3. Lines and curves are not drawn proportionally: (a) Various ε-flat sub-boundaries bounded between two arcs; (b) blocked visibility near ε-flat sub-boundary.

consists of a finite number of smooth sub-boundaries. In this section we assume that the boundaries of all obstacles inside $\mathcal{R}$ are piecewise smooth.

For a smooth sub-boundary c, the *turning angle*, denoted by $\mathcal{A}(c)$, is defined to be the integral of the curvature along c. For a piecewise sub-boundary c, the turning angle is defined to be the sum of the turning angles of all smooth sub-boundaries of c, plus the sum of the instantaneous angular changes at the joint points. Observe that the turning angle of the boundary of an obstacle is 2π if the obstacle is convex, or greater than 2π if it is non-convex. In some sense, the turning angle of the boundary of an obstacle reflects the geometric complexity of the obstacle.

For each sub-boundary c, we use $|c|$ to denote the length of c, and use $c[u_1, u_2]$ to denote the part of c between two points u_1 and u_2 on c. For any point $v \notin c$, we let u_1 and u_2 be the two points on c such that c is lying between the two rays $\overrightarrow{vu_1}$ and $\overrightarrow{vu_2}$. We call u_1 and u_2 *bounding points* of c by v. We define the *viewing angle of c from v* to be $\angle u_1 v u_2$.

For each obstacle, we decompose its boundary into minimum number of ε-flat *sub-boundaries*. A sub-boundary c is said to be ε-flat if $\mathcal{A}(c) \leqslant \pi - \theta_\varepsilon$, where $\theta_\varepsilon = \lambda_0\rho_0/16(1+\lambda_0^2) \cdot \varepsilon$. Let u_1 and u_2 be the two endpoints of c.

Property 1. *Any* ε-flat *sub-boundary c is bounded between two minor arcs each with chord* $\overline{u_1u_2}$ *and angle* $2\theta_\varepsilon$ (*as shown in Fig.* 3*a*).

Property 2. *The width of an* ε-flat *sub-boundary c, defined by* $|\overline{u_1u_2}|$, *is no less than* $|c| \cdot \cos\theta_\varepsilon/2$.

Property 3. *The height of an* ε-flat *sub-boundary c, defined by the maximum distance between any point* $u \in c$ *and line segment* $\overline{u_1u_2}$, *is no more than* $|c|/2 \cdot \sin\theta_\varepsilon/2$.

Since ε-flat sub-boundaries are "relatively" flat, any point $v \in \mathcal{F}$ "sandwiched" between two ε-flat sub-boundaries will have a limited visibility, as we show in the following lemma:

Lemma 1. *If $v \in \mathcal{F}$ is a point between two ε-*flat *sub-boundaries c_1 and c_2 and the total viewing angle of c_1 and c_2 from v is more than $2\pi - 6\theta_\varepsilon$, then v is not ε-*visible.

Proof. For each $i = 1, 2$, let $u_{i,1}$ and $u_{i,2}$ be the two endpoints of c_i. $\mathcal{V}_v$ is the union of the following three regions:

Region I: the region bounded by sub-boundary c_1, $\overline{vu_{1,1}}$ and $\overline{vu_{1,2}}$;
Region II: the region bounded by sub-boundary c_2, $\overline{vu_{2,1}}$ and $\overline{vu_{2,2}}$;
Region III: the region not inside either $\angle u_{1,1}vu_{1,2}$ or $\angle u_{2,1}vu_{2,2}$.

First, we claim that $\angle u_{1,1}vu_{1,2} \leqslant \pi + \theta_\varepsilon$, or otherwise one can find a point v' on c_1 such that $\angle u_{1,1}v'u_{1,2} < \pi - \theta_\varepsilon$, a contradiction to the assumption that c_1 is ε-flat. Similarly, we have $\angle u_{2,1}vu_{2,2} \leqslant \pi + \theta_\varepsilon$.

Since the total viewing angle of v blocked by c_1 and c_2 is more than $2\pi - 6\theta_\varepsilon$, we have $\angle u_{1,1}vu_{1,2} > \pi - 7\theta_\varepsilon$ and $\angle u_{2,1}vu_{2,2} > \pi - 7\theta_\varepsilon$. Since c_1 is ε-flat, the volume of Region I is bounded by the union of $\Delta u_{i,1}vu_{i,2}$ and the arc with chord $|c_1|$ and angle $2\theta_\varepsilon$, as shown in Fig. 3b. Since $|c_1| \cdot \cos(\theta_\varepsilon/2) \leqslant |\overline{u_{1,1}u_{1,2}}| \leqslant L_\mathcal{R} \leqslant \sqrt{(\lambda_0^2+1)/\lambda_0\rho_0\mu(\mathcal{F})}$, where $L_\mathcal{R}$ is the length of the diagonal of $\mathcal{R}$, the volume of Region I is bounded by $\mathrm{O}(\varepsilon\mu(\mathcal{F}))$. Similarly, the volume of Region II is also bounded by $\mathrm{O}(\varepsilon\mu(\mathcal{F}))$.

Region III is the union of two (possibly merged) cones with a total angle of $6\theta_\varepsilon$, and therefore the volume of Region III is also $\mathrm{O}(\varepsilon\mu(\mathcal{F}))$. Hence, the region visible from v has a total volume of $\mathrm{O}(\varepsilon\mu(\mathcal{F}))$.

We chose the constant factor $\lambda_0\rho_0/16(1+\lambda_0^2)$ of θ_ε in such a way that the total volume of Regions I, II, and III (i.e. the volume of $\mathcal{V}_v$) is bounded by $\varepsilon\mu(\mathcal{F})$. Therefore, v is not ε-visible. □

In the rest of this section we will prove the following theorem:

Theorem 1. *If the boundaries of all obstacles can be divided into n ε-*flat *sub-boundaries, the total length of all obstacle boundaries is bounded by* $\mathrm{O}(\sqrt{n\mu(\mathcal{F})/\varepsilon})$.

However, to prove Theorem 1 we need two lemmas, which we will prove in the next subsection. In Section 2.3 we will show the proof of this theorem as well as its corollaries.

2.2. *Forbidden neighborhoods of ε-flat sub-boundaries*

For each ε-flat sub-boundary c with endpoints u_1 and u_2, we divide it into 15 equal-length segments, and let u'_1 and u'_2 be the two endpoints of the middle segment. The *ε-neighborhood* of c, denoted by $\mathcal{N}_\varepsilon(c)$, is defined to be the union of points from each of which the viewing angle of $c[u'_1, u'_2]$ is greater than $\pi - \theta_\varepsilon$, as shown in Fig. 4a. It is easy to see that, for any $v \in \mathcal{N}_\varepsilon(c)$, the distance between v and line segment $\overline{u'_1u'_2}$ is no more than $|c[u'_1, u'_2]|/2 \cdot \tan\theta_\varepsilon = |c|/30 \cdot \tan\theta_\varepsilon$. The distance between v and line segment $\overline{u_1u_2}$ is no more than the sum of the distance between u and $\overline{u'_1u'_2}$ and the maximum distance between $\overline{u'_1u'_2}$ and $\overline{u_1u_2}$, which is $|c|/30 \cdot \tan\theta_\varepsilon + |c|/2 \cdot \sin\theta_\varepsilon/2$.

These neighborhoods are "forbidden" in the sense that they do not overlap with each other if the corresponding sub-boundaries are roughly the same length, as we will show in Lemma 2. By "charging" a certain portion of $\mathcal{C}$ to each ε-flat sub-boundary, we show that the total length of all ε-flat sub-boundaries, that is, the length of $\partial\mathcal{B}$, can be upper-bounded.

Lemma 2. *The ε-neighborhoods of two sub-boundaries c_1 and c_2 do not overlap if $|c_1|/2 \leqslant |c_2| \leqslant 2|c_1|$.*

Proof. Suppose for the sake of contradiction $v \in \mathcal{C}$ is a point inside $\mathcal{N}_\varepsilon(c_1) \cap \mathcal{N}_\varepsilon(c_2)$, where the length ratio between c_1 and c_2 is between $\frac{1}{2}$ and 2. For each $i = 1, 2$, we let $u_{i,1}$ and $u_{i,2}$ be the two endpoints of c_i, and let $u'_{i,1}$ and $u'_{i,2}$ be the endpoints of the portion of c_i incident to the ε-neighborhood of c_i. Let v_i be the projection of v on line segment $\overline{u_{i,1}u_{i,2}}$, and let v'_i be the intersection of c_i and the straight line that passes both v_i and v.

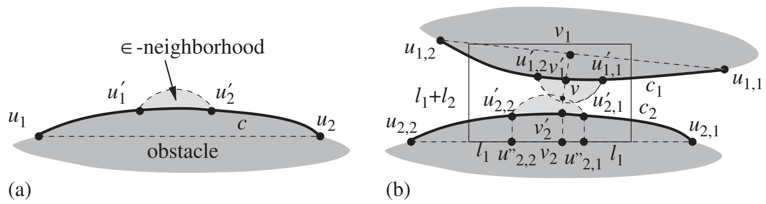

Fig. 4. Lines and curves are not drawn proportionally: (a) ε-neighborhood; (b) ε-neighborhoods are non-overlapping for sub-boundaries with similar lengths.

The intuition here is as the following: since c_1 and c_2 are "relatively" flat, non-intersecting, and about the same length, for $\mathcal{N}_\varepsilon(c_1)$ and $\mathcal{N}_\varepsilon(c_2)$ to overlap, $\overline{u_{1,1}u_{1,2}}$ and $\overline{u_{2,1}u_{2,2}}$ have to be "almost" parallel and also close to each other. That way, we can find in the free space between c_1 and c_2 a point that can only see less than $\varepsilon\mu(\mathcal{F})$ of the free space as its visibility is mostly "blocked" by c_1 and c_2, leading to a contradiction to the assumption that $\mathcal{C}$ is ε-visible.

There are a number of cases corresponding to different geometric arrangements of the points, line segments and curves (sub-boundaries). In the following we assume that $\overline{u_{1,1}u_{1,2}}$ and $\overline{u_{2,1}u_{2,2}}$ do not intersect, v lies between $\overline{u_{1,1}u_{1,2}}$ and $\overline{u_{2,1}u_{2,2}}$, and v_1' (v_2', respectively) lies between v and v_1 (v_2, respectively), as shown in Fig. 4b. The other cases can be analyzed in an analogous manner.

Since line segments $\overline{u_{1,1}u_{1,2}}$ and $\overline{u_{2,1}u_{2,2}}$ do not intersect, either both $\overline{v_1u_{2,1}}$ and $\overline{v_1u_{2,2}}$ lie between $\overline{u_{1,1}u_{1,2}}$ and $\overline{u_{2,1}u_{2,2}}$, or both $\overline{v_2u_{1,1}}$ and $\overline{v_2u_{1,2}}$ lie between $\overline{u_{1,1}u_{1,2}}$ and $\overline{u_{2,1}u_{2,2}}$. Without loss of generality we assume that it is the former case. Let $l_1 = |\overline{vv_1}|$ and $l_2 = |\overline{vv_2}|$. Let $u''_{2,1}$ ($u''_{2,2}$, respectively) be the projection of $u'_{2,1}$ ($u'_{2,2}$, respectively) on $\overline{u_{2,1}u_{2,2}}$. Observe that v_1' lies inside the small rectangle of width $|\overline{u''_{2,1}u''_{2,2}}| + 2l_1$ and height $l_1 + l_2$ (the solid rectangle in Fig. 4b). Since $|\overline{u_{2,2}u''_{2,2}}| = |\overline{u_{2,2}u_{2,1}}| - |\overline{u''_{2,2}u_{2,1}}| > |\overline{u_{2,2}u_{2,1}}| - |c[u'_{2,2}, u_{2,1}]|$, we have

$$\tan \angle v_1'u_{2,1}u_{2,2} \leqslant \frac{l_1 + l_2}{|\overline{u_{2,2}u_{2,1}}| - |c[u'_{2,2}, u_{2,1}]| - l_1}$$
$$\leqslant \frac{\left(\frac{1}{30}\cdot\tan\theta_\varepsilon + \frac{1}{2}\cdot\sin\frac{\theta_\varepsilon}{2}\right)\cdot(|c_1| + |c_2|)}{|c_2|\cdot\cos\frac{\theta_\varepsilon}{2} - \frac{8|c_2|}{15} - \left(\frac{1}{30}\cdot\tan\theta_\varepsilon + \frac{1}{2}\cdot\sin\frac{\theta_\varepsilon}{2}\right)\cdot|c_1|}.$$

Applying $|c_1| \leqslant 2|c_1|$ and $\theta_\varepsilon < \frac{1}{12}$, we now have

$$\tan \angle v_1'u_{2,1}u_{2,2} \leqslant \frac{\theta_\varepsilon\cdot\left(\frac{1}{30\cos\theta_\varepsilon} + \frac{1}{4}\right)\cdot 3|c_2|}{\left(\cos\frac{\theta_\varepsilon}{2} - \frac{8}{15} - \left(\frac{1}{15}\cdot\tan\theta_\varepsilon + \sin\frac{\theta_\varepsilon}{2}\right)\right)\cdot|c_2|} \leqslant \frac{5\theta_\varepsilon}{2} \leqslant \frac{5}{2}\tan\theta_\varepsilon \leqslant \tan\frac{5\theta_\varepsilon}{2}.$$

It follows that $\angle v_1'u_{2,1}u_{2,2} \leqslant 5\theta_\varepsilon/2$. Similarly, we can show that $\angle v_1'u_{2,2}u_{2,1} \leqslant 5\theta_\varepsilon/2$, and therefore $\angle u_{2,1}v_1'u_{2,2} \geqslant \pi - 5\theta_\varepsilon$. Since v_1' is on c_1, $\angle u_{1,1}v_1'u_{1,2} \geqslant \pi - \theta_\varepsilon$. Therefore, the viewing angle from v_1' not blocked by c_1 and c_2 is no more than $2\pi - (\pi - \theta_\varepsilon) - (\pi - 5\theta_\varepsilon) = 6\theta_\varepsilon$. According to Lemma 1 v_1' is not ε-visible. Therefore, we can find a point $v_1^* \in \mathcal{F}$ close to v_1' which is also not ε-visible, a contradiction to the assumption that $\mathcal{C}$ is ε-visible. □

Next we give a lower bound of the volume of the ε-neighborhood of any ε-flat sub-boundary with the following lemma:

Lemma 3. *For any* ε-flat *sub-boundary* c, *the volume of* $\mathcal{N}_\varepsilon(c)$ *is* $\Omega(\theta_\varepsilon \cdot |c|^2)$.

Proof. We will show that, the ε-neighborhood of c has a volume no less than $\mu_0 = (\theta_\varepsilon|c[u_1', u_2']|^2)/18\kappa_1$, for some constant $\kappa_1 > 1$. (We will explain later how this constant κ_1 is chosen.)

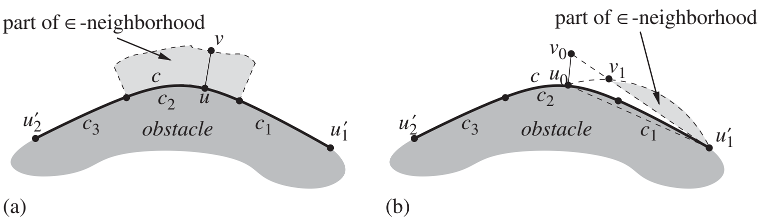

Fig. 5. In the figures we only show the portion of sub-boundary c between u_1' and u_2': (a) ε-flat sub-boundary: case I; (b) ε-flat sub-boundary: case II.

We divide $c[u_1', u_2']$ into three equal-length segments, c_1, c_2, and c_3. For any point u on $c[u_1', u_2']$, we say that $v \in \mathcal{F}$ is the *lookout point* of u if line segment $\overline{vu}$ is normal to $c[u_1', u_2']$ and the viewing angle of $c[u_1', u_2']$ from v is $\pi - \theta_\varepsilon$. We call the length of $\overline{uv}$ the *lookout distance* of $c[u_1', u_2']$ at u.

We first consider Case I, where for each point $u \in c_2$ the length of the lookout distance of c at u is at least $l = \theta_\varepsilon |c[u_1', u_2']|/3\kappa_1$, as shown in Fig. 5a. In this case, the volume of the ε-neighborhood of c outside c_2 is at least $|c_2| \cdot l - l^2 \cdot \theta_\varepsilon/2 = |c[u_1', u_2']|^2 \cdot \theta_\varepsilon/9\kappa_1 \cdot (1 - (\theta_\varepsilon^2/2\kappa_1)) \geqslant |c[u_1', u_2']|^2 \cdot \theta_\varepsilon/18\kappa_1 = \mu_0$, and therefore the volume of the ε-neighborhood of c is no less than μ_0.

Now we consider Case II, where there exists a point $u_0 \in c_2$ such that the lookout distance at u_0 is less than l, as shown in Fig. 5b. Let v_0 be the lookout point of u_0. Since $\mathcal{A}(c[u_1', u_2']) \leqslant \mathcal{A}(c) \leqslant \theta_\varepsilon$, v_0 will see at least one of the two endpoints of $c[u_1', u_2']$, or otherwise the viewing angle of v_0 is less than $\pi - \theta_\varepsilon$. Without loss of generality we let u_1' be an endpoint of $c[u_1', u_2']$ that is visible from v_0. $c[u_0, u_1']$, the part of c between u_0 and u_1', lies below line segments $\overline{v_0 u_1'}$. Since $u_0 \in c_2$, we have $|c[u_0, u_1']| \geqslant |c_1| = |c[u_1', u_2']|/3$.

Since curve $c[u_0, u_1']$ is also ε-flat, we have $|\overline{u_0 u_1'}| \geqslant |c[u_0, u_1']| \cdot \cos\theta_\varepsilon/2 > |c[u_1', u_2']|/6$. We use $\overline{u_0 u_1'}$ as the chord to draw a minor arc of angle $2\theta_\varepsilon$ outside $\overline{u_0 u_1'}$. The radius of this arc is $r_0 = |\overline{u_0 u_1'}|/2\sin\theta_\varepsilon \geqslant |c[u_1', u_2']|/12\theta_\varepsilon$. Let v_1 be the point where arc $\widehat{u_0 u_1'}$ intersects $\overline{v_0 u_1'}$. We claim that any point v' inside the closed region bounded by arc $\widehat{u_0 u_1'}$ and chord $\overline{u_1' v_1}$ belongs to the ε-neighborhood of c. First of all, v' is outside $c[u_0, u_1']$, as $c[u_0, u_1']$ lies below $\overline{v_0 u_1'}$. Secondly, the viewing angle of $c[u_1', u_2']$ from v' should be no less than the viewing angle of $c[u_0, u_1']$ from v', which is at least $\pi - \theta_\varepsilon$.

Now we consider the volume of the region bounded by $\widehat{u_0 u_1'}$ and $\overline{u_1' v_1}$. This is actually an arc $\widehat{u_1' v_1}$ with angle $\theta_0 = 2\theta_\varepsilon - 2\angle u_0 u_1' v_0$ and radius r_0. Since $\angle u_0 u_1' v_0 < |\overline{u_0 v_0}|/|\overline{u_0 u_1'}| < l/|c[u_1', u_2']|/6 = 2\theta_\varepsilon/\kappa_1$. As long as we choose κ_1 large enough, we can have $\angle u_0 u_1' v_0 < \theta_\varepsilon/2$ and therefore $\theta_0 > \theta_\varepsilon$. The volume of arc $\widehat{u_1' v_1}$, therefore, is $r_0^2/2(\theta_0 - \sin\theta_0) \geqslant r_0^2 \cdot \theta_0^3/14 \geqslant |c[u_1', u_2']|^2 \theta_\varepsilon/14 \cdot 12^2$. Once again, if we choose κ_1 large enough, we can have $\mu(\widehat{u_1' v_1}) \geqslant \theta_\varepsilon |c[u_1', u_2']|^2/18\kappa_1 = \mu_0$, and therefore the volume of the ε-neighborhood of c is greater than μ_0.

Since $|c[u_1', u_2']| = |c|/15$, we have $\mu(\mathcal{N}_\varepsilon(c)) = \Omega(\theta_\varepsilon \cdot |c|^2)$. □

2.3. *Putting it together*

With the lemmas established in the last subsection, we are ready to prove Theorem 1:

Proof of Theorem 1. Let $L_{\max}$ be the maximum length of all ε-flat sub-boundaries inside $\mathcal{R}$. We divide all ε-flat sub-boundaries into subsets $S_1, S_2, \ldots, S_k$. For each i, S_i contains the boundaries edges whose lengths are between $L_{\max}/2^i$ and $L_{\max}/2^{i-1}$.

We let $c_{i,1}, c_{i,2}, \ldots, c_{i,n_i}$ be the n_i sub-boundaries in S_i. By Lemma 2, $\mathcal{N}_\varepsilon(c_{i,j}) \cap \mathcal{N}_\varepsilon(c_{i,j'}) = \emptyset$, for any j and j', $1 \leqslant j, j' \leqslant n_i$. By Lemma 3, there exists a constant $K > 0$ such that $\mu(\mathcal{N}_\varepsilon(c_{i,j})) \geqslant K \cdot \theta_\varepsilon \cdot |c_{i,j}|^2$ for all i and j. Therefore, we have

$$\frac{\mu(\mathcal{F})}{\rho_0} \geqslant \mu(\mathcal{C})$$

$$\geqslant \mu\left(\bigcup_{j=1}^{n_i} \mathcal{N}_\varepsilon(c_{i,j})\right)$$
$$= \sum_{j=1}^{n_i} \mu(\mathcal{N}_\varepsilon(c_{i,j}))$$
$$= \sum_{j=1}^{n_i} K \cdot \theta_\varepsilon \cdot |c_{i,j}|^2$$
$$\geqslant n_i \cdot K \cdot \theta_\varepsilon \cdot \frac{L_{\max}^2}{4^i}.$$

Hence we have $n_i \leqslant 4^i \cdot \mu(\mathcal{F})/K \cdot \theta_\varepsilon \cdot L_{\max}^2 \cdot \rho_0$. Let $K' = \mu(\mathcal{F})/K \cdot \theta_\varepsilon \cdot L_{\max}^2 \cdot \rho_0$. Now we are to give an upper bound of $|\partial\mathcal{B}|$, which is defined to be $\sum_{i=1}^{k}\sum_{j=1}^{n_i}|c_{i,j}|$, the sum of all ε-flat sub-boundaries. Since $|c_{i,j}| \leqslant L_{\max} 2^{i-1}$, we have $|\partial\mathcal{B}| \leqslant L_{\max} \cdot \sum_{i=1}^{k} n_i \cdot 2^{-i+1}$. Observe that $\sum_{i=1}^{k} n_i = n$, $\sum_{i=1}^{k} n_i \cdot 2^{-i+1}$ is maximized when $n_i = K' \cdot 4^i$ for $i < \log_4 3n/K'$ and $n_i = 0$ for $i \geqslant \log_4 3n/K'$. Therefore, we have

$$\sum_{i=1}^{k} n_i \cdot 2^{-i+1} \leqslant \sum_{i=1}^{\log_4(3n/K')-1} K' \cdot 4^i \cdot 2^{-i+1}$$
$$= 2K' \sum_{i=1}^{\log_4(3n/K')-1} 2^i$$
$$< 2K' \cdot 2^{\log_4(3n/K')}$$
$$= \sqrt{12n \cdot K'}$$
$$= \sqrt{\frac{12n \cdot \mu(\mathcal{F})}{K \cdot \theta_\varepsilon \cdot L_{\max}^2 \cdot \rho_0}}.$$

Therefore, $|\partial\mathcal{B}|$ is no more than $\sqrt{12n \cdot \mu(\mathcal{F})/K \cdot \theta_\varepsilon \cdot \rho_0}$. Recall that K and ρ_0 are constants and that $\theta_\varepsilon = \Theta(\varepsilon)$, we have $|\partial\mathcal{B}| = \mathrm{O}(\sqrt{n\mu(\mathcal{F})/\varepsilon})$. □

If all the obstacles inside $\mathcal{C}$ are polygons, each boundary edge is an ε-flat sub-boundary, and therefore we have the following corollary:

Corollary 1. *If $\mathcal{C}$ contains polygonal obstacles with a total of n edges*, $|\partial\mathcal{B}|$ *is* $\mathrm{O}(\sqrt{n\mu(\mathcal{F})/\varepsilon})$.

If all obstacles inside $\mathcal{C}$ are convex, the boundary of each obstacle can be decomposed into $2\pi/\theta_\varepsilon$ ε-flat sub-boundaries, and therefore we have:

Corollary 2. *If $\mathcal{C}$ contains n convex obstacles that are piecewise smooth*, $|\partial\mathcal{B}|$ *is* $\mathrm{O}(1/\varepsilon\sqrt{n\mu(\mathcal{F})})$.

In some sense, the upper bound stated in Corollary 1 is tight, as one can construct an ε-visible space (as shown in Fig. 6) inside a square consisting of $n = \frac{1}{\varepsilon}$ rectangular free space "cells," each with length $\sqrt{\mu(\mathcal{F})}$ and width $\varepsilon \cdot \sqrt{\mu(\mathcal{F})}$. The total length of obstacle boundaries is $\Theta(1/\varepsilon)\sqrt{\mu(\mathcal{F})} = \Theta(\sqrt{n\mu(\mathcal{F})/\varepsilon})$.

Nonetheless, we still conjecture that the best bound should be the following:

Conjecture 1. $|\partial\mathcal{B}|$ *is* $\mathrm{O}(1/\varepsilon\sqrt{\mu(\mathcal{F})})$.

2.4. Extension to three dimensions

In this subsection we show how to generalize our proof of Theorem 1 to $3\mathcal{D}$ spaces. For simplicity, we assume that the boundary (surface) of each obstacle is smooth, meaning that the curvature is continuous everywhere on the surface.

To replicate the proofs of Lemmas 1, 2, and 3 for the $3\mathcal{D}$ case, we first need to define the ε-flat surface patch, the $3\mathcal{D}$ counterpart of ε-flat sub-boundary. A surface patch s is said to be ε-flat if, for any point $u \in s$ and any plane p

that contains the line $l_{s,u}$, the curve $c = p \cap s$ is ε-flat. Here $l_{s,u}$ is the line that passes through u and is normal to s. Moreover, we also need the surface patch to be "relatively round." More specifically, we require that for each ε-flat surface patch s there exists a "center" v_s such that, $\max\{|\overline{v_s v}||v \in \partial s\}/\min\{|\overline{v_s v}||v \in \partial s\}$ is bounded by a constant. Here ∂s is the closed curve that defines the boundary of s. We call $R_{s,v_s} = \min\{|\overline{v_s v}||v \in \partial s\}$ the *minimum radius* of s at center v_s.

We define the ε-neighborhood $\mathcal{N}_\varepsilon(s)$ for an ε-flat surface patch similarly to the case of ε-flat sub-boundary. We choose a small "sub-patch" s' of s at the center of s so that the distance between v_s and every point on the boundary of s' is $k_1 \cdot R_{s,v_s}$, for some constant $k_1 < 1$. For any point v outside the obstacle that s is bounding, $v \in \mathcal{N}_\varepsilon(s)$ if and only if there exist two points $u_1, u_2 \in s'$ such that $\angle u_1 v u_2 > \pi - k_2\varepsilon$ for some constant $k_2 > 0$.

We use a sequence of planes each containing l_{v,s_v} to "sweep" through the volume of $\mathcal{N}_\varepsilon(s)$. Each such plane p contains a "slice" of $\mathcal{N}_\varepsilon(s)$ with an area of no less than $\Theta(\varepsilon \cdot R^2_{s,v_s})$, following the same argument of the proof of Lemma 5. Therefore, the total volume of $\mathcal{N}_\varepsilon(s)$ is $\Theta(\varepsilon \cdot R^3_{s,v_s}) = \Theta(\varepsilon \cdot \mu(s)^{3/2})$. Using a proof technique similar to the one used in the proof of 1, we can prove the following:

Theorem 2. *If $\mathcal{C}$ contains convex obstacles bounded by a total of n* ε-flat *surface patches*, $|\partial\mathcal{B}|$ *is* $O((n\mu(\mathcal{F})^2/\varepsilon^2)^{1/3})$.

3. Applications and open problems

As mentioned in Section 1.1, the number of milestones randomly sampled on obstacle boundaries needed to adequately cover $\mathcal{F}$ with a high probability can be bounded using the ε′-boundary-visibility assumption, and therefore, for all ε-visible spaces, we are interested in finding the lower bound of ε′ with respect to ε.

It is easy to see that for $2\mathcal{D}$ ε-visible spaces, we have the following lemma:

Lemma 4. $|\partial\mathcal{B}_v| = \Omega(\varepsilon\sqrt{\mu(\mathcal{F})})$ *for any* $v \in \mathcal{F}$.

Therefore, we can arrive at a lower bound of the fraction of all obstacle boundaries that each free space point can see for various cases by using Corollaries 1 and 2.

Corollary 3. *If $\mathcal{C}$ contains polygonal obstacles with a total of n edges*, $|\partial\mathcal{B}_v|$ *is* $\Omega(\sqrt{\varepsilon^3/n} \cdot |\partial\mathcal{B}|)$.

Corollary 4. *If $\mathcal{C}$ contains n convex obstacles that are piecewise smooth*, $|\partial\mathcal{B}_v|$ *is* $\Omega(\varepsilon^2/\sqrt{n} \cdot |\partial\mathcal{B}|)$.

In particular, if Conjecture 1 holds, we could have the following:

Conjecture 2. $|\partial\mathcal{B}_v|$ *is* $\Omega(\varepsilon^2 \cdot |\partial\mathcal{B}|)$ *for any* ε-visible $2\mathcal{D}$ *space*.

Conjecture 2 seems to be tight, as one can also construct a worst-case example where $|\partial\mathcal{B}_v| = O(\varepsilon^2 \cdot |\partial\mathcal{B}|)$. As shown in Fig. 6b, on the bottom of a square space with width d, $\Theta(\varepsilon^{-2})$ little "buckets" can be placed, each with height εd

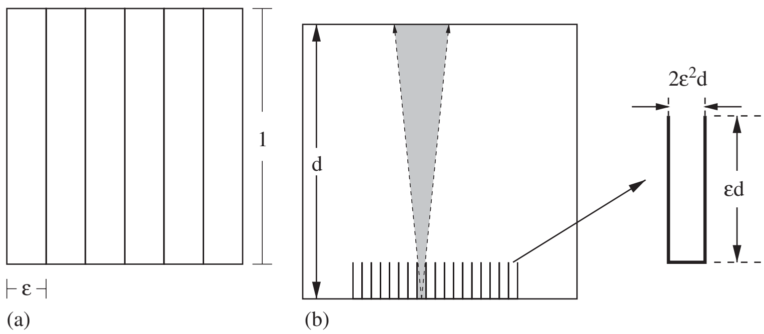

Fig. 6. Worst case examples: (a) Corollary 1 is tight; (b) conjecture 2 is tight.

and width $2\varepsilon^2 d$. This space is ε-visible, as any point in the bucket can at least see a portion of free space with volume εd^2. It is, however, ε'-boundary-visible only when $\varepsilon' = O(\varepsilon^2)$, as a point deep in the bucket can only see a boundary with length $\Theta(\varepsilon d)$, while the total boundary length is $\Theta(1/\varepsilon^2 \cdot (\varepsilon d)) = \Theta(d/\varepsilon)$.

Then, using the same proof technique as [13],[2] we can show that $O((1/\varepsilon^2) \log 1/\varepsilon)$ randomly sampled boundary points can view a significant portion of $\mathcal{F}$ with a high probability. These results can be applied to the boundary art gallery problem to provide an upper bound of the number of boundary guards needed to adequately guard the space.

It occurs to us that, although one can construct an example where there exists a free space point that can only see obstacle boundaries of size $\Theta(\varepsilon\sqrt{\mu(\mathcal{F})})$, the total volume of such points could be upper-bounded. In particular, we have the following conjecture:

Conjecture 3. *Every point in $\mathcal{F}$, except for a small subset of volume $O(\sqrt{\varepsilon}\mu(\mathcal{F}))$, can see obstacle boundaries of size $\Omega(\sqrt{\varepsilon\mu(\mathcal{F})})$.*

If we can prove both Conjectures 1 and 3, we can reduce the number of boundary points needed to adequately cover the space with high probability to $O((1/\varepsilon^{3/2}) \log 1/\varepsilon)$.

So far our results are limited to $2\mathcal{D}$ ε-visible spaces and some special cases of $3\mathcal{D}$ ε-visible spaces. If we can extend these results to higher dimensions, we will be able to provide a theoretical foundation for analyzing the effectiveness of the PRM planners [1,2] that (randomly) pick milestones close to boundaries of obstacles. These planners have shown to be more efficient than the earlier PRM planners based on uniform sampling in the free space by better capturing narrow passages in the configuration space; that is, the roadmaps they construct have better connectivity. However, there has been no prior theoretical result on the visibility of the roadmaps constructed using the sampled boundary points. With upper bound results analogous to the ones for $2\mathcal{D}$ and $3\mathcal{D}$ cases, we will be able to prove an upper bound of the number of milestones uniformly sampled on obstacle boundaries needed to adequately cover free space $\mathcal{F}$ with a high probability, an result similar to the one provided by Kavraki [13] for uniform sampling method.

4. Conclusion

In this paper we provided some preliminary results as well as several conjectures on the upper bound of the boundary size of ε-visible spaces in $2\mathcal{D}$ and $3\mathcal{D}$ spaces. These results can be used to bound the number of guards needed for the boundary art gallery problem. Potentially, they can also be applied to the analysis of a certain class of PRM planners that sample points close to obstacle boundaries.

References

[1] N.M. Amato, O.B. Bayazit, L.K. Dale, C. Jones, D. Vallejo, OBPRM: an obstacle-based PRM for 3d workspaces, in: Proc. Third Workshop on Algorithmic Foundations of Robotics, 1998, pp. 155–168.
[2] V. Boor, M.H. Overmars, A.F. Stappen, The Gaussian sampling strategy for probabilistic roadmap planners, in: Proc. 1999 IEEE Internat. Conf. on Robotics and Automation, 1999, pp. 1018–1023.
[3] V. Chvátal, A combinatorial theorem in plane geometry, J. Combin. Theory Ser. B 18 (1975) 39–41.
[4] S. Eidenbenz, Inapproximability results for guarding polygons without holes, in: Proc. Ninth Internat. Symp. on Algorithms and Computation, Lecture Notes in Computer Science, Vol. 1533, 1998, pp. 427–436.
[5] S. Eidenbenz, C. Stamm, P. Widmayer, Inapproximability of some art gallery problems, in: Proc. 10th Canadian Conf. on Computational Geometry, 1998, pp. 64–65.
[6] S. Eidenbenz, C. Stamm, P. Widmayer, Inapproximability results for guarding polygons and terrains, Algorithmica 31 (2001) 79–113.
[7] S.K. Ghosh, Approximation algorithms for art gallery problems, in: Proc. Canadian Information Processing Society Congress, 1987.
[8] J.E. Hopcroft, D.A. Joseph, S.H. Whitesides, Movement problems for 2-dimensional linkages, SIAM J. Comput. 13 (3) (1984) 610–629.
[9] D. Hsu, L.E. Kavraki, J.-C. Latombe, R. Motwari, S. Sorkin, On finding narrow passages with probabilistic roadmap planners, in: Proc. Third Workshop on Algorithmic Foundations of Robotics, 1998, pp. 141–153.
[10] D. Hsu, T. Jiang, J.H. Reif, Z. Sun, The bridge test for sampling narrow passages with probabilistic roadmap planners, in: Proc. 2003 IEEE Internat. Conf. on Robotics and Automation, 2003, pp. 4420–4426.
[11] G. Kalai, J. Matoušek, Guarding galleries where every point sees a large area, Israel J. Math. 101 (1997) 125–139.

[2] The difference is that, in our proof, every point v in the free space sees at least ε^2 fraction of obstacle boundaries, and therefore the probability that k points uniformly sampled on obstacle boundaries cannot see v is $(1-\varepsilon^2)^k$.

[12] L.E. Kavraki, J.-C. Latombe, Randomized preprocessing of configuration space for fast path planning, in: Proc. 1994 Internat. Conf. on Robotics and Automation, 1994, pp. 2138–2145.
[13] L.E. Kavraki, J.-C. Latombe, R. Motwani, P. Raghavan, Randomized query processing in robot motion planning, in: Proc. 27th Annu. ACM Symp. on Theory of Computing, 1995, pp. 353–362.
[14] L.E. Kavraki, P. Švestka, J.-C. Latombe, M.H. Overmars, Probabilistic roadmaps for path planning in high-dimensional configuration spaces, IEEE Trans. Robotics and Automation 12 (4) (1996) 566–580.
[15] D. Kirkpatrick, Guarding galleries with no nooks, in: Proc. 12th Canadian Conf. on Computational Geometry, 2000, pp. 43–46.
[16] D.T. Lee, A.K. Lin, Computational complexity of art gallery problems, IEEE Trans. Inform. Theory 32 (1986) 276–282.
[17] J.-M. Lien, S.L. Thomas, N.M. Amato, A general framework for sampling on the medial axis of the free space, in: Proc. 2003 IEEE Internat. Conf. on Robotics and Automation, 2003, pp. 4439–4444.
[18] M. Morales, S. Rodriguez, N.M. Amato, Improving the connectivity of PRM roadmaps, in: Proc. 2003 IEEE Internat. Conf. on Robotics and Automation, 2003, pp. 4427–4432.
[19] M. Overmars, A random approach to motion planning, Technical Report RUU-CS-92-32, Utrecht University, Utrecht, The Netherlands, 1992.
[20] M.H. Overmars, P. Švestka, A probabilistic learning approach to motion planning, in: Proc. First Workshop on Algorithmic Foundations of Robotics, 1994, pp. 19–37.
[21] J.H. Reif, Complexity of the mover's problem and generalizations, in: Proc. 20th IEEE Symp. on Foundations of Computer Science, 1979, pp. 421–427.
[22] J.H. Reif, Z. Sun, On frictional mechanical systems and their computational power, SIAM J. Comput. 32 (6) (2003) 1449–1474.
[23] D. Schuchardt, H. Hecker, Two NP-hard art-gallery problems for ortho-polygons, Math. Logic Quart. 41 (1995) 261–267.
[24] J.T. Schwartz, M. Sharir, On the piano mover's problem: II. General techniques for computing topological properties of real algebraic manifolds, Adv. Appl. Math. 4 (1983) 298–351.
[25] J. Sellen, Lower bounds for geometrical and physical problems, SIAM J. Comput. 25 (6) (1996) 1231–1253.
[26] J. Urrutia, Art gallery and illumination problems, in: J.-R. Sack, J. Urrutia (Eds.), Handbook of Computational Geometry, Elsevier Science Publishers B.V., North-Holland, Amsterdam, 2000, pp. 973–1026.
[27] P. Valtr, Guarding galleries where no point sees a small area, Israel J. Math. 104 (1998) 1–16.
[28] P. Valtr, On galleries with no bad points, Discrete Computat. Geom. 21 (3) (1999) 193–200.