Computational Geometry 44 (2011) 1–8

Contents lists available at ScienceDirect

Computational Geometry: Theory and Applications

www.elsevier.com/locate/comgeo

A note on the perimeter of fat objects

Prosenjit Bose [a,*,1], Otfried Cheong [b,2], Vida Dujmović [a,1]

[a] School of Computer Science, Carleton University, Ottawa, Canada
[b] Department of Computer Science, Korea Advanced Institute of Science and Technology, Daejeon, South Korea

ARTICLE INFO

Article history:
Received 15 January 2009
Accepted 30 June 2010
Available online 22 July 2010
Communicated by De

Keywords:
Realistic inputs
Fat objects
Visibility
Perimeter
Fractals

ABSTRACT

In this Note, we show that the size of the perimeter of (α,β)-covered objects is a linear function of the diameter. Specifically, for an (α,β)-covered object O, $\text{per}(O) \leqslant c\frac{\text{diam}(O)}{\alpha\beta\sin^2\alpha}$, for a positive constant c. One easy consequence of the result is that every point on the boundary of such an object sees a constant fraction of the boundary. Locally γ-fat objects are a generalization of (α,β)-covered objects. We show that no such relationship between perimeter and diameter can hold for locally γ-fat objects.

1. Introduction

Often, the worst case lower bound for a geometric algorithm occurs when the input consists of 'long' and 'skinny' objects. However, such artificial configurations do not occur in many 'realistic' inputs. This motivates the study of objects considered to be more likely to occur in real-life applications. A number of realistic models have been introduced and studied in the literature (see [5] for a survey). We consider four such classes of objects, namely (α,β)-covered objects, locally γ-fat objects, ε-area-good objects, and ε-boundary-good objects. (α,β)-covered objects and locally γ-fat objects are classes of *fat objects*, that is, objects that cannot be arbitrarily long and skinny. ε-area-good and ε-boundary-good objects are 'realistic' in a different sense which is expressed in terms of visibility. We briefly describe the four models.

(α,β)-*covered* objects were introduced by Efrat [6] as a generalization of convex fat objects to a class of non-convex objects. Roughly speaking, an object O is (α,β)-covered if for every point p on the boundary, ∂O, of O there is a large fat triangle contained in O that has p as a vertex. *Locally* γ-*fat* objects were introduced by de Berg [4] as a generalization of (α,β)-covered objects. An object O is locally γ-fat if for any disk D with center p in O and not completely containing O, the connected component of $D \cap O$ containing p has area at least γ times the area of D.

An ε-good object O, introduced by Valtr [10], has the property that every point $p \in O$ can see a constant fraction of the area of O. We require this only of the points $p \in \partial O$, and call such objects ε-*area-good* (this is a strictly larger class of objects, see Fig. 7). Kirkpatrick [8] introduced a similar class of ε-*boundary-good* objects O, where every point $p \in \partial O$ can see a constant fraction of the length of the boundary of O.

Table 1 shows the relations between these four classes of objects. The way to interpret the table is as follows. A YES entry in the table means that the class of objects quantified by a constant c in the row of Table 1 belongs to the class

* Corresponding author.
E-mail addresses: jit@cg.scs.carleton.ca (P. Bose), otfried@kaist.edu (O. Cheong), vida@cg.scs.carleton.ca (V. Dujmović).

1 Research supported in part by NSERC.
2 Research supported by the Korea Science and Engineering Foundation Grant R01-2008-000-11607-0 funded by the Korean government.

0925-7721/$ – see front matter
doi:10.1016/j.comgeo.2010.06.002

Table 1
Pairwise relationships between four classes of objects.

Object class	(α', β')-covered	Locally γ'-fat	ε'-area-good	ε'-boundary-good
(α, β)-covered	YES	YES [4]	YES (Observation 1)	YES (Corollary 1)
locally γ-fat	NO [4] (Fig. 6)	YES	NO (Fig. 6)	NO (Fig. 6)
ε-area-good	NO	NO	YES	NO (Observation 2)
ε-boundary-good	NO	NO	NO (Observation 2)	YES

Table 2
Properties of the four classes of objects.

Object class	Good union-complexity	Good guardability	Good perimeter-length
(α, β)-covered	YES [6]	YES [1,2], also [8] and Corollary 1	YES (Theorem 1)
locally γ-fat	YES [4]	NO (Fig. 6)	NO (Theorem 2)
ε-area-good	NO	YES [1,10]	?
ε-boundary-good	NO	YES [1,8]	?

of objects quantified by c' in the column of Table 1, where c' is some function of c. The entry NO means that no such constant exists. Trivially, all the entries on the diagonal of the table are YES. The four NO entries that are not justified by any references are implied by the fact that thin triangles are ε-area-good and ε-boundary-good objects with $\varepsilon = 1$, yet they are neither (α, β)-covered, nor locally γ-fat for any α, β, γ.

Various properties have been proven for realistic objects under various models, in particular good *union complexity*, good *guardability*, and good *perimeter-length*.

The complexity of the union of n general (constant-complexity) objects can be $\Theta(n^2)$, but this can only be achieved with objects that are long and skinny (such as n triangles or line segments). The union complexity of a set of n pseudo-disks is $\mathcal{O}(n)$ [7], and we consider a class of objects to have *good union complexity* if the union of n objects from that class has near-linear size.

Fat objects behave more like disks as opposed to line segments, and indeed the union complexity of *n fat triangles*—that is, triangles having no angle smaller than some constant α—is $\mathcal{O}(n \log\log n)$ [9]. Efrat generalized this result by showing that the union complexity of n (α, β)-covered objects (of constant description complexity) is at most $\mathcal{O}(\lambda_s(n) \log^2 n \log\log n)$ [6]. This result was later improved and generalized to locally γ-fat objects by de Berg [4]. On the other hand, since every convex object is ε-area-good and ε-boundary-good, these classes do not have good union complexity. This is summarized in the second column of Table 2.

A convex object can be guarded with a single guard, and we say that a class of objects has *good guardability* if any object in the class can be guarded by a constant number of guards. Valtr [10] showed that ε-good objects have good guardability. His proof, together with the main theorem of Addario-Berry et al. [1], implies the same result for ε-area-good objects. Kirkpatrick [8] showed that ε-boundary-good objects have good guardability. Locally γ-fat objects are known not to have good guardability, while (α, β)-covered objects have good guardability, see the third column of Table 2.

The perimeter of a convex object is at most π times its diameter, while for general objects, the perimeter length cannot be bounded by a function of the diameter. We consider a class of objects to have *good perimeter-length* if the length of the perimeter of an object is at most a constant times its diameter.

Our main result is to show that (α, β)-covered objects have good perimeter length. As a corollary, we obtain that each point on the boundary of an (α, β)-covered object O sees a constant fraction of the length of the boundary of O, and so (α, β)-covered objects are ε-boundary-good for some ε that depends only on α and β. On the other hand, we show that a family of curves that converges to the Koch snowflake [11] defines a family of objects that is locally γ-fat for $\gamma = 3\sqrt{3}/(128\pi)$, has diameter one, but contains objects of arbitrarily large perimeter length. We leave open the question of whether ε-area-good and ε-boundary-good objects have good perimeter-length, see the last column of Table 2.

2. Preliminaries

In this paper, by an *object* we mean a closed region, O, of the plane, $\mathbb{E}^2$, such that O is connected and $\mathbb{E}^2 \setminus O$ is connected. For an object O, ∂O denotes the boundary of O. We say that an object O_1 is *contained* in object O_2, if $O_1 \subseteq O_2$. The *diameter* of O, denoted by $\mathsf{diam}(O)$, is the maximum Euclidean distance between any two points in O. If J is a collection of Jordan curves in the plane, the *perimeter* of J, denoted by $\mathsf{per}(J)$, is the sum of the lengths of all the curves in J. The perimeter, $\mathsf{per}(O)$, of an object O, is $\mathsf{per}(O) := \mathsf{per}(\partial O)$. If O is a simple polygon[3] then we refer to the line-segments and their intersection points in ∂O as edges and vertices of ∂O (and O) respectively. For a simple polygon P and a vertex $v \in \partial P$, the *angle at* v is the angle in the interior of P, determined by the two edges of ∂P that are incident to v. For an edge with endpoints v and w, vw denotes both the edge of ∂P, and the line segment $\overline{vw}$ in the plane. For two points, s and t, in the plane, $\overline{st}$ denotes the line segment with s and t as its endpoints.

[3] All the polygons considered in this paper are simple.

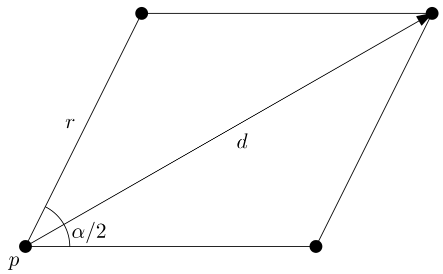

Fig. 1. Illustration of d-directional r-long tile at p.

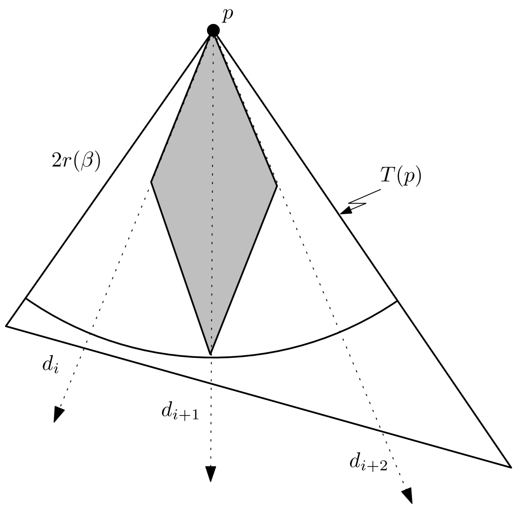

Fig. 2. Illustration for the proof of Lemma 1.

An object O is (α, β)-*covered* if for each point $p \in \partial O$ there exists a triangle $T(p)$, called a *good triangle* of p, such that:

1. p is a vertex of $T(p)$, $T(p) \subseteq O$, and
2. each angle of T is at least α, and the length of each edge of T is at least $\beta \cdot \mathsf{diam}(O)$.

Notice that this immediately implies that $\alpha \leqslant \pi/3$. For a point p in the plane, a *ray* at p is a half-line with endpoint at p that is oriented away from p. The *direction* of a ray R is the counter-clockwise angle between the positive x-axis (rooted at p) and ray R.

3. Perimeter of (α, β)-covered objects

Consider a parallelogram having as one of its vertices a point p in the plane and having angle $\frac{\alpha}{2}$ at p, $0 < \alpha \leqslant \pi/3$. Let d be the direction of the ray at p that bisects the angle at p. We call such a parallelogram a *d-directional tile at p*. If the four edges of a d-directional tile at p have the same length r—that is, it is a rhombus—then we call it a *d-directional r-long tile at p* (see Fig. 1). Let Γ denote the set of directions, $\Gamma = \{d_i := i\frac{\alpha}{5} \mid i \in \mathbb{Z}, 0 \leqslant i < \frac{10\pi}{\alpha}\}$. Let $\mathsf{r}(\alpha, \beta) := \frac{1}{2}\beta \sin\alpha \cdot \mathsf{diam}(O)$.

Lemma 1. *Let O be an (α, β)-covered object. For every point $p \in \partial O$, there exists a direction $d \in \Gamma$ such that the d-directional $\mathsf{r}(\alpha, \beta)$-long tile at p is contained in $\{O \setminus \partial O\} \cup \{p\}$.*

Proof. Let $T(p)$ be a good triangle at p. Since each angle of $T(p)$ is at least α, there are three consecutive directions d_i, d_{i+1}, d_{i+2} in Γ such that for each of the three directions, the ray at p with that direction intersects the interior of $T(p)$. Consider the intersection, I, of the disk of radius $2\mathsf{r}(\alpha, \beta)$ centered at p and the region of the plane bounded by two rays at p with directions d_i and d_{i+2}, as illustrated in Fig. 2. The height of $T(p)$ is greater than $2\mathsf{r}(\alpha, \beta)$. Thus $I \subset \{T(p) \setminus \partial T(p)\} \cup \{p\}$. This completes the proof, since the d_{i+1}-directional $\mathsf{r}(\alpha, \beta)$-long tile at p is contained in I. $\square$

For a fixed direction $d \in \Gamma$, let O_d denote the set of points in ∂O, such that for each point $p \in O_d$, the d-directional $\mathsf{r}(\alpha, \beta)$-long tile at p is contained in $\{O \setminus \partial O\} \cup \{p\}$. Since O is (α, β)-covered object, Lemma 1 implies that $\bigcup_{d \in \Gamma} O_d = \partial O$.

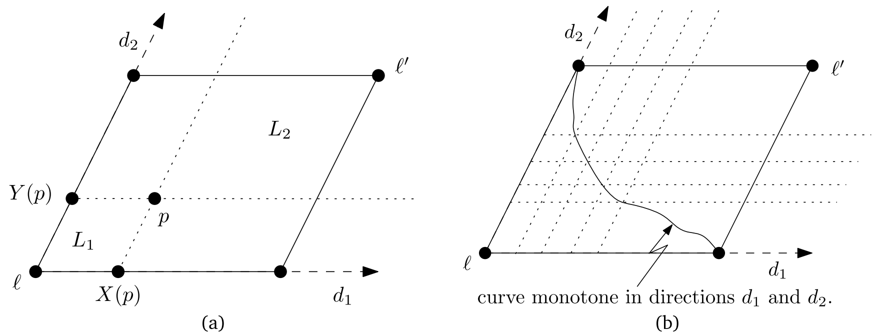

Fig. 3. Illustration for the proof of Lemma 2.

Lemma 2. *Let O be an (α, β)-covered object. For any $d \in \Gamma$ and any point ℓ in the plane, let L be a d-directional $\mathrm{r}(\alpha, \beta)$-long tile at ℓ. Let ℓ' be the vertex of ∂L that is not adjacent to ℓ. Then the following two statements hold.*

1. *If $O_d \cap L \neq \emptyset$, then $\ell' \in O$ and for every point $p \in O_d \cap L$, ℓ' is in a good triangle of p.*
2. $\mathrm{per}(O_d \cap L) \leqslant 2\mathrm{r}(\alpha, \beta)$.

Proof. Both statements are trivial if $O_d \cap L = \emptyset$. Assume now that $O_d \cap L \neq \emptyset$ and consider a point $p \in O_d \cap L$. Since there are at most two intersection points of the boundaries of translates of two congruent convex polygons, the d-directional $\mathrm{r}(\alpha, \beta)$-long tile at p, denoted S_p contains ℓ'. By definition of O_d, $S_p \subset O$ and thus $\ell' \in O$. Furthermore, S_p is contained in a good triangle of p, and thus ℓ' is in a good triangle of p. This completes the proof of the first claim. We now prove the second claim.

Consider two lines that intersect at a point $p \in O_d \cap L$, one with slope $d_1 := d - \frac{\alpha}{5}$ and the other with slope $d_2 := d + \frac{\alpha}{5}$. The two lines divide L into four regions. Denote by L_1 the region that contains ℓ and L_2 the region that contains ℓ', as illustrated in Fig. 3(a). (Note that when p lies on the boundary of L, any of the four regions may have zero area.) Since $L_2 \subseteq S_p \subset \{O \setminus \partial O\} \cup \{p\}$, it follows that for all points $q \in L_2 \setminus p$, $q \notin O_d$ (and in fact, $q \notin \partial O$). Similarly, for any point $q \in L_1 \setminus p$, p is contained in the d-directional $\mathrm{r}(\alpha, \beta)$-long tile at q, and thus $q \notin O_d$. We refer to this observation as the *empty region property of* p. Each point in $O_d \cap L$ has the empty region property. Notice that this implies that $O_d \cap L$ is a collection of Jordan curves such that there is at most one intersection between a line in direction d_1 with $O_d \cap L$ and similarly there is at most one intersection between a line in direction d_2 with $O_d \cap L$.

For a point $p \in L$, let $\mathrm{X}(p)$ be the intersection point of the line of slope d_2 that contains p, and the ray at ℓ with direction d_1. Similarly, $\mathrm{Y}(p)$ is the intersection point of the line with slope d_1 that contains p and the ray at ℓ with direction d_2, as illustrated in Fig. 3(a). We now show that the sum of the lengths of these Jordan curves in $O_d \cap L$ is at most $2\mathrm{r}(\alpha, \beta)$.

Consider any set of $n \geqslant 2$ distinct points $\{v_1, \ldots, v_n\}$ on the curves $O_d \cap L$, sorted by increasing order of their d_1-coordinate, where v_1 is the point in $O_d \cap L$ with smallest d_1-coordinate and v_n is the one with largest d_1-coordinate. By the above, we know that they are also sorted in decreasing d_2-coordinate. Let $\overline{C}_n$ be the polygonal chain with vertex set $\{v_1, \ldots, v_n\}$ and edge set $\overline{v_i v_{i+1}}$, $1 \leqslant i < n$, as illustrated in Fig. 3(b). In what follows, we prove that for all integers $n \geqslant 2$ the length of $\overline{C}_n$ is at most $2\mathrm{r}(\alpha, \beta)$. As n goes to infinity, $\overline{C}_n$ tends to a Jordan curve C that contains $O_d \cap L$, thus providing an upper bound of $2\mathrm{r}(\alpha, \beta)$ on the length of the curves in $O_d \cap L$ as well.

By the triangle inequality, for each $1 \leqslant i < n$,

$$\mathrm{per}(\overline{v_i v_{i+1}}) \leqslant \mathrm{per}\big(\overline{\mathrm{X}(v_i)\mathrm{X}(v_{i+1})}\big) + \mathrm{per}\big(\overline{\mathrm{Y}(v_i)\mathrm{Y}(v_{i+1})}\big).$$

Since each point in $O_d \cap L$ has the empty region property, each point on $\overline{C}_n$ has a unique d_1- and d_2-coordinate, that is, $\forall_{s,t \in \overline{C}_n : s \neq t} \mathrm{X}(s) \neq \mathrm{X}(t)$ and $\forall_{s,t \in \overline{C}_n : s \neq t} \mathrm{Y}(s) \neq \mathrm{Y}(t)$. Therefore, since the sides of L have length $\mathrm{r}(\alpha, \beta)$,

$$\mathrm{per}(\overline{C_n}) \leqslant \sum_{i=1}^{n-1} \big(\mathrm{per}\big(\overline{\mathrm{X}(v_i)\mathrm{X}(v_{i+1})}\big) + \mathrm{per}\big(\overline{\mathrm{Y}(v_i)\mathrm{Y}(v_{i+1})}\big)\big) \leqslant 2\mathrm{r}(\alpha, \beta),$$

which completes the proof. □

Lemma 3. *Let O be an (α, β)-covered object. For each $d \in \Gamma$, $\mathrm{per}(O_d) \leqslant c \frac{\mathrm{diam}(O)}{\beta \sin^2 \alpha}$, for some positive constant c.*

Proof. By Lemma 2, the portion of O_d that is inside a d-directional $\mathrm{r}(\alpha, \beta)$-long tile, has perimeter at most $2\mathrm{r}(\alpha, \beta)$. Thus to bound the perimeter of O_d, it is enough to bound the number of d-directional $\mathrm{r}(\alpha, \beta)$-long tiles that cover O_d.

Let D_1 and D_2 be two concentric disks with radii $\mathrm{diam}(O)$ and $\mathrm{diam}(O) + 2\mathrm{r}(\alpha, \beta)$, respectively. Since the radius of D_1 is $\mathrm{diam}(O)$, place D_1 such that $O \subseteq D_1$. Let S be a minimum cardinality set of d-directional $\mathrm{r}(\alpha, \beta)$-long tiles that covers

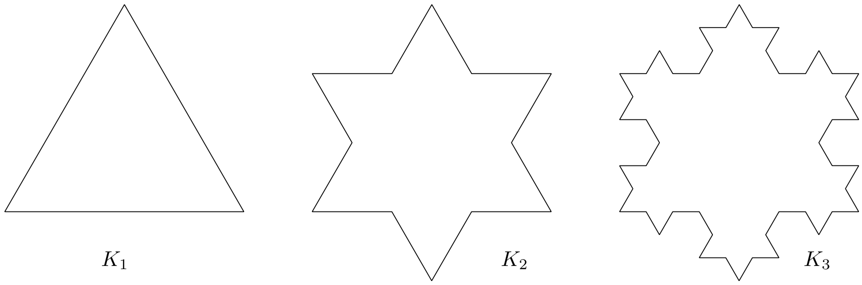

Fig. 4. The first three objects K_1, K_2, and K_3.

D_1 (and thus O). Note that any pair of distinct elements, S_i and S_j of S, must be non-overlapping, that is, the interiors of S_i and S_j do not have a point in common. Such a set S exists, since d-directional $\mathrm{r}(\alpha,\beta)$-long tiles can tile the plane. Thus $O \subset D_1 \subset \bigcup_{S_i \in S} S_i \subset D_2$. The latter inclusion follows from the fact that the diameter of a d-directional $\mathrm{r}(\alpha,\beta)$-long tile is at most $2\mathrm{r}(\alpha,\beta)$. The ratio of the area of D_2 and the area of a d-directional $\mathrm{r}(\alpha,\beta)$-long tile, gives the desired bound on the area $|S|$ of S. In particular, the smaller of the two angles in a d-directional $\mathrm{r}(\alpha,\beta)$-long tile is $\frac{2\alpha}{5}$, thus the area of each element in S is $\mathrm{r}(\alpha,\beta)^2 \sin\frac{2\alpha}{5}$. Since $O \subset \bigcup_{S_i \in S} S_i \subset D_2$, $|S| \leqslant \frac{\pi(2\mathrm{r}(\alpha,\beta)+\mathrm{diam}(O))^2}{\mathrm{r}(\alpha,\beta)^2 \sin(2\alpha/5)} = \frac{\pi(4+4\mathrm{diam}(O)/\mathrm{r}(\alpha,\beta)+\mathrm{diam}(O)^2/\mathrm{r}(\alpha,\beta)^2)}{\sin(2\alpha/5)}$.

Thus $\mathrm{per}(O_d) \leqslant 2\mathrm{r}(\alpha,\beta)|S| = \frac{2\pi(4\mathrm{r}(\alpha,\beta)+4\mathrm{diam}(O)+\mathrm{diam}(O)^2/\mathrm{r}(\alpha,\beta))}{\sin(2\alpha/5)}$, which is at most $c\frac{\mathrm{diam}(O)}{\beta \sin^2\alpha}$, for some positive constant c. $\square$

Lemmas 1 and 3 and the fact that $|\Gamma| = \lceil \frac{10\pi}{\alpha} \rceil$, imply the following theorem.

Theorem 1. *The perimeter of every (α,β)-covered object O is at most* $\mathrm{per}(O) \leqslant c\frac{\mathrm{diam}(O)}{\alpha\beta\sin^2\alpha}$, *for some positive constant c.*

This theorem in turn implies the following corollary. (A point $x \in O$ *sees* a point $y \in O$ if the line segment $\overline{xy}$ is contained in O.)

Corollary 1. *For every fixed α and β, every point on the boundary of an (α,β)-covered object sees a constant fraction of the length of the boundary. In particular, every (α,β)-covered object is $c\alpha\beta^2\sin^4\alpha$-boundary-good, for some positive constant c.*

Proof. Consider the region I from the proof of Lemma 1. An isosceles triangle with p as one of its vertices, two edges of length $2\mathrm{r}(\alpha,\beta)$ incident to p and the angle $\frac{2\alpha}{5}$ at p, is contained in I and thus it is contained in O. Therefore, $\mathrm{per}(O_p) \geqslant 4\mathrm{r}(\alpha,\beta)\sin\alpha/5$, where O_p denotes the set of all the points of ∂O that p sees. By Theorem 1, $\mathrm{per}(O) \leqslant c'\frac{\mathrm{diam}(O)}{\alpha\beta\sin^2\alpha}$, for some positive constant c'. Therefore, $\mathrm{per}(O_p)/\mathrm{per}(O) \geqslant c\alpha\beta^2\sin^4\alpha$, for some positive constant c. $\square$

The above corollary, coupled with the result of Kirkpatrick [8], implies that the boundary of every (α,β)-covered polygon can be guarded with a constant number (that depends on α and β) of guards. This result had already been shown with constant $8\sqrt{2}\pi\frac{1}{\alpha\beta^2}$ by Aloupis et al. [2]. A recent result by Addario-Berry et al. [1] states that if the boundary of a polygon can be guarded with g guards then its interior can be guarded with an additional $4g-6$ guards. Thus both the boundary and the interior of every (α,β)-covered polygon can be guarded with at most $\frac{40\sqrt{2}\pi}{\alpha\beta^2}$ guards.

4. Perimeter of locally γ-fat objects

We now show that locally γ-fat objects do not have good perimeter-length, by constructing a family of objects that are locally γ-fat with $\gamma = 3\sqrt{3}/(128\pi)$, have diameter at most one, and contain objects of arbitrarily large perimeter length. Our family is comprised of the objects bounded by the curves that converge to the Koch snowflake [11]. Object K_1 is an equilateral triangle, whose circumcircle C_1 has diameter one. We obtain K_{i+1} from K_i by dividing each edge of K_i into three segments of equal length, and attaching an equilateral triangle to the middle segment. Fig. 4 shows the first three objects of this sequence. The perimeter of the objects K_i grows to infinity [11]. Every K_i is contained in the circumcircle of K_1, and so their diameter is bounded by one. It remains to show that all K_i are locally γ-fat for $\gamma = 3\sqrt{3}/(128\pi)$ (a conservative bound).

The Koch construction can be represented as a tree $\mathcal{T}$ as follows: The root of the tree is the triangle K_1. The children of a node are the triangles that are attached to it in a later stage of the construction. (So the root node has three children added in the construction of K_2, six nodes added in the construction of K_3, etc.) It is known that a triangle Δ is contained in the circumcircle of all its ancestor triangles in $\mathcal{T}$. We let $r(\Delta)$ denote the radius of the circumcircle of triangle Δ.

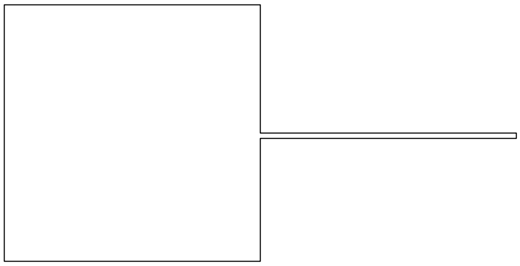

Fig. 5. Illustration for Observation 2.

Consider now a disk D with radius r and center p in K_n, for some n, such that D does not completely contain K_n. Let Δ_p be the first triangle during the construction that contains p. If D does not completely contain Δ_p, then we are done as the equilateral triangle Δ_p is locally γ-fat.

So assume $\Delta_p \subset D$, and let Δ_1 be the smallest ancestor of Δ_p such that $r < 4r(\Delta_1)$. There must be such a triangle as otherwise $K_n \subset D$. We distinguish two cases.

If $\Delta_1 = \Delta_p$, then D's area $|D| = 2\pi r^2 < 2\pi(4r(\Delta_p))^2 = 32\pi(r(\Delta_p))^2$. The area of an equilateral triangle with circumcircle radius ρ is $(3\sqrt{3}/4)\rho^2$, and so

$$\frac{|\Delta_p|}{|D|} > \frac{(3\sqrt{3}/4)(r(\Delta_p))^2}{32\pi(r(\Delta_p))^2} = \frac{3\sqrt{3}}{128\pi} = \gamma.$$

Consider now the case that $\Delta_1 \neq \Delta_p$. All triangles Δ that are ancestors of Δ_p and descendants of Δ_1 have $r \geqslant 4r(\Delta)$, which implies that $\Delta \subset D$. It follows that the union of these triangles is connected, and connects p to $D \cap \Delta_1$. We will complete the proof by showing that $|D \cap \Delta_1| > \gamma|D|$.

If $\Delta_1 \subset D$, we can argue as in the first case, and so we assume now that D does not contain Δ_1. Let Δ_2 be the child of Δ_1 that is an ancestor of Δ_p. Since p lies in the circumcircle of Δ_2, and Δ_1 intersects this circumcircle, the distance d between p and Δ_1 is at most $2r(\Delta_2)$. On the other hand, we have $r \geqslant 4r(\Delta_2)$, implying that $d \leqslant r/2$. It follows that Δ_1 contains a point in D at distance $r/2$ from the center p, and also intersects the boundary of D. Then $\Delta_1 \cap D$ must contain an equilateral triangle of side length $r/2$, and therefore of area $(\sqrt{3}/4)(r/2)^2$, and so we have

$$\frac{|D \cap \Delta_1|}{|D|} \geqslant \frac{(\sqrt{3}/4)(r/2)^2}{2\pi r^2} = \frac{\sqrt{3}}{32} > \gamma.$$

We summarize this section in the following theorem.

Theorem 2. *For every $L > 0$ there is a locally γ-fat object O of diameter at most one and perimeter larger than L, with $\gamma = 3\sqrt{3}/(128\pi)$.*

5. Conclusions and open problems

Corollary 1 states that (α, β)-covered objects are ε'-boundary good, for some $\varepsilon' := \varepsilon(\alpha, \beta)$. It is simple to see that (α, β)-covered objects are also ε'-area-good:

Observation 1. *For every fixed α and β, every point on the boundary of an (α, β)-covered object O sees a constant fraction of the area of O.*

Proof. Every (α, β)-covered object O contains a good triangle, thus the area of O is at least the area of a good triangle, $\frac{1}{2}(\beta \operatorname{diam}(O))^2 \sin\alpha$. Furthermore, O is contained in a disk of radius $\operatorname{diam}(O)$, and thus the area of O is at most $\pi \operatorname{diam}(O)^2$. Therefore, for $\varepsilon = \frac{\beta^2 \sin\alpha}{2\pi}$, every point on the boundary of an (α, β)-covered object O sees an ε fraction of the area of O, thus (α, β)-covered objects are ε-area good. □

This raises the question if ε-area-good objects are ε-boundary-good for some $\varepsilon' := \varepsilon'(\varepsilon)$, or vice versa. As stated in the following observation, the answer is no.

Observation 2. *There exists no $\varepsilon' := \varepsilon'(\varepsilon)$ such that every ε-good object is ε'-boundary good. Similarly, there exists no $\varepsilon' := \varepsilon'(\varepsilon)$ such that every ε-boundary good object is ε'-area-good.*

To see this, consider the object in Fig. 5. It is the union of a square and a long thin rectangle. If the square and rectangle have the same area, then every point can see at least half the area of the object, however the top right corner of the square

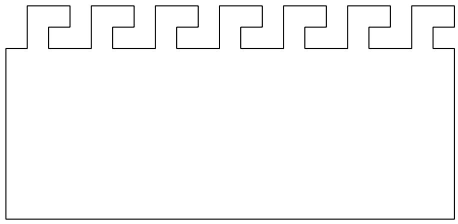

Fig. 6. A locally fat object.

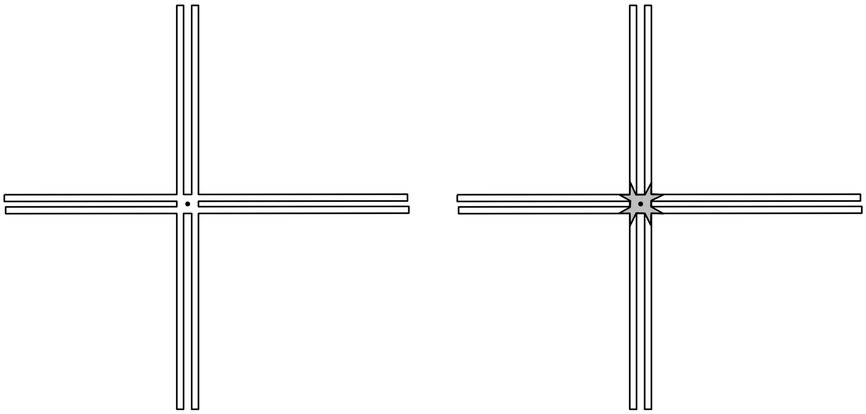

Fig. 7. An ε-area-good object that is neither ε-good nor ε-boundary good.

does not see a constant fraction of the perimeter. Thus every object in this class is 1/2-good, but there exists no constant ε' such that every object in the class is ε'-boundary-good. Similarly, if the square and rectangle have the same perimeter, then every point can see at least half the perimeter, however the top right corner of the rectangle does not see a constant fraction of the area. Thus every object in this class is 1/2-boundary-good, but there exists no constant ε' such that every object in the class is ε'-area-good.

De Berg [4] introduced the class of locally γ-fat objects and proved that, for some $\gamma := \gamma(\alpha, \beta)$, every (α, β)-covered object is locally γ-fat; and that there exists no $\alpha := \alpha(\gamma)$ and $\beta := \beta(\gamma)$ such that every locally γ-fat object is (α, β)-covered. Thus locally γ-fat objects are generalizations of (α, β)-covered objects. Locally γ-fat objects are not ε-area-good or ε-boundary-good for any $\varepsilon := \varepsilon(\gamma)$. To see this, consider the object in Fig. 6 (which is an adaptation of a similar example by de Berg).

The class of objects depicted in Fig. 6 also shows that locally γ-fat objects cannot be guarded by a constant number of guards for any constant that depends on γ only.

We end with two open problems, the question marks in Table 2:

Open Problem 1.

(a) Does there exist a $c := c(\varepsilon)$ such that for every ε-area-good object O, $\text{per}(O) \leqslant c \cdot \text{diam}(O)$?
(b) Does there exist an $c := c(\varepsilon)$ such that for every ε-boundary-good object O, $\text{per}(O) \leqslant c \cdot \text{diam}(O)$?

Note that these questions are easy for convex objects, since every convex object has its perimeter bounded by π times its diameter. A much stronger property is known for locally γ-fat convex objects O. Namely, Chew et al. [3] proved that for any two points p and q on ∂O, there is a path on ∂O from p to q whose length is bounded by the length of the segment $\overline{pq}$ times a constant $\gamma' := \gamma'(\gamma)$.

We conclude the paper with the following figure, which demonstrates that a property that applies to every point on the boundary of an object does not necessarily extend to the points in its interior. In particular, ε-area-good objects are not necessarily ε-good. Consider the object in Fig. 7. It is a polygon that is the union of four rectangles, with the hole in the middle filled in. Each point on the boundary can see at least a quarter of the area and perimeter. However, the point in the center can only see a negligible fraction of the boundary and area.

Acknowledgements

This research was initiated at the Carleton–Eindhoven Workshop on Computational Geometry, July 18–22, 2005 organized by Mark de Berg and Prosenjit Bose. The authors are grateful to Boris Aronov for suggesting the main problem studied in this paper and to Greg Aloupis, Chris Gray, Stefan Langerman and Bettina Speckmann for many discussions. We would like to thank the two anonymous referees for their helpful comments and for suggesting to study the perimeter of locally γ-fat objects.

References

[1] Louigi Addario-Berry, Omid Amini, Jean-Sébastien Sereni, Stéphan Thomassé, Guarding art galleries: The extra cost for sculptures is linear, in: Proc. of the 11th Scandinavian Workshop on Algorithm Theory (SWAT'08), 2008, pp. 41–52.
[2] Greg Aloupis, Prosenjit Bose, Vida Dujmović, Chris Gray, Stefan Langerman, Bettina Speckmann, Guarding fat polygons and triangulating guarded polygons, in: Proc. 20th Canadian Conference on Computational Geometry (CCCG'08), 2008, pp. 107–111.
[3] L. Paul Chew, Haggai David, Matthew J. Katz, Klara Kedem, Walking around fat obstacles, Inform. Process. Lett. 83 (3) (2002) 135–140.
[4] Mark de Berg, Improved bounds on the union complexity of fat objects, Discrete Comput. Geom. 40 (1) (2008) 127–140.
[5] Mark de Berg, Matthew Katz, Frank van der Stappen, Jules Vleugels, Realistic input models for geometric algorithms, Algorithmica 34 (2002) 81–97.
[6] Alon Efrat, The complexity of the union of (α, β)-covered objects, SIAM J. Comput. 34 (2005) 775–787.
[7] Klara Kedem, Ron Livne, Janos Pach, Micha Sharir, On the union of Jordan regions and collision-free translational motion amidst polygonal obstacles, Discrete Comput. Geom. 1 (1986) 59–71.
[8] David Kirkpatrick, Guarding galleries with no nooks, in: Proc. of the 12th Canadian Conference on Computational Geometry (CCCG'00), 2000, pp. 43–46.
[9] Jiři Matoušek, Janos Pach, Micha Sharir, Shmuel Sifrony, Emo Welzl, Fat triangles determine linearly many holes, SIAM J. Comput. 23 (1994) 154–169.
[10] Pavel Valtr, Guarding galleries where no point sees a small area, Israel J. Math. 104 (1998) 1–16.
[11] Helge von Koch, Sur une courbe continue sans tangente, obtenue par une construction géométrique élémentaire, Arch. Mat. Astron. Fys. 1 (1904) 681–702.