On the perimeter of fat objects

Prosenjit Bose[*]   Otfried Cheong[†]   Vida Dujmović[*]

Abstract

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, $\mathsf{per}(O) \leq c \frac{\mathsf{diam}(O)}{\alpha \beta \sin^2 \alpha}$, 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.

[*]School of Computer Science, Carleton University, Ottawa, Canada {jit,vida}@cg.scs.carleton.ca. Research supported in part by NSERC.

[†]Department of Computer Science, Korea Advanced Institute of Science and Technology, Daejeon, South Korea otfried@kaist.edu. Research supported by the Korea Science and Engineering Foundation Grant R01-2008-000-11607-0.

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). 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 table can be interpreted 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 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.

object class	(α', β')-covered	locally γ'-fat
(α, β)-covered	YES	YES [4]
locally γ-fat	NO [4] (Fig. 2(b))	YES
ε-area-good	NO	NO
ε-boundary-good	NO	NO

object class	ε'-area-good	ε'-boundary-good
(α, β)-covered	YES (Obs. 1)	YES (Cor. 1)
locally γ-fat	NO (Fig. 2(b))	NO (Fig. 2(b))
ε-area-good	YES	NO (Obs. 2)
ε-boundary-good	NO (Obs. 2)	YES

Table 1: Pairwise relationships between four classes of objects.

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. 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.

object class	good union-complexity	
(α, β)-covered	YES [6]	
locally γ-fat	YES [4]	
ε-area-good	NO	
ε-boundary–good	NO	
object class	good guardability	good perim.-length
(α, β)-covered	YES [2, 1], also [8] & Cor. 1	YES (Thm. 1)
locally γ-fat	NO (Fig. 2(b))	NO (Thm. 2)
ε-area-good	YES [10, 1]	?
ε-bound.–good	YES [8, 1]	?

Table 2: Properties of the four classes of objects.

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   Perimeter of (α, β)–covered objects

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 $\alpha \leq \pi/3$. 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.

Consider a parallelogram having as one of its vertices a point p in the plane and having angle $\frac{\alpha}{2}$ at p, $0 < \alpha \leq \pi/3$. 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*. Let Γ denote the set of directions, $\Gamma = \{d_i := i\frac{\alpha}{5} \mid i \in \mathbb{Z}, 0 \leq i < \frac{10\pi}{\alpha}\}$. Let $\mathsf{r}(\beta) := \frac{1}{2}\beta \sin \alpha \cdot \mathsf{diam}(O)$.

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 $\mathsf{r}(\beta)$-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 $2\mathsf{r}(\beta)$ centered at p and the region of the plane bounded by two rays at p with directions d_i and d_{i+2}. The height of $T(p)$ is greater than $2\mathsf{r}(\beta)$. Thus $I \subset \{T(p) \setminus \partial T(p)\} \cup \{p\}$. This completes the

---

*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[†] 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.

proof, since the d_{i+1}-directional $r(\alpha, \beta)$-long tile at p is contained in I. □

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 $r(\beta)$-long tile at p is contained in $\{O \setminus \partial O\} \cup \{p\}$. Since O is (α, β)–covered object, Lemma 1 implies that $\bigcup_{d \in \Gamma} O_d = \partial O$.

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 $r(\beta)$-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. *$\text{per}(O_d \cap L) \leq 2r(\beta)$.*

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 $r(\beta)$-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 Figure 1(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 $r(\beta)$-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 $X(p)$ be the intersection point of the line of slope d_2 that contains p, and the ray at ℓ with direction d_1. Similarly, $Y(p)$ 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 Figure 1 (a). We now show that the sum of the lengths of these Jordan curves in $O_d \cap L$ is at most $2r(\beta)$.

Consider any set of $n \geq 2$ 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
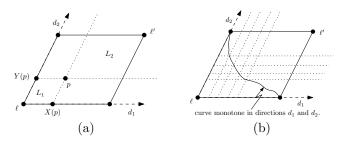

Figure 1: Illustration for the proof of Lemma 2.

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 $\overline{v_i v_{i+1}}$, $1 \leq i < n$, as illustrated in Figure 1(b). In what follows, we prove that for all integers $n \geq 2$ the length of $\overline{C}_n$ is at most $2r(\beta)$. 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 $2r(\beta)$ on the length of the curves in $O_d \cap L$ as well.

By the triangle inequality, for each $1 \leq i < n$, $\text{per}(\overline{v_i v_{i+1}}) \leq \text{per}(X(v_i)X(v_{i+1})) + \text{per}(Y(v_i)Y(v_{i+1}))$. 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, $\forall_{s,t \in \overline{C}_n : s \neq t} X(s) \neq X(t)$ and $\forall_{s,t \in \overline{C}_n : s \neq t} Y(s) \neq Y(t)$. Therefore, since the sides of L have length $r(\beta)$, $\text{per}(\overline{C}_n) \leq \sum_{i=1}^{n-1} \left(\text{per}(\overline{X(v_i)X(v_{i+1})}) + \text{per}(\overline{Y(v_i)Y(v_{i+1})}) \right) \leq 2r(\beta)$. □

The proof of the following lemma is omitted.

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

Lemma 1 and Lemma 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 $\text{per}(O) \leq c \frac{\text{diam}(O)}{\alpha \beta \sin^2 \alpha}$, for some positive constant c.*

This theorem in turn implies the following corollary, the proof of which is omitted due to space constraints.

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.*

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

---

[‡]A point $x \in O$ *sees* a point $y \in O$ if the line segment $\overline{xy}$ is contained in O.

(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].

3 Perimeter of locally γ-fat objects

The proof of the following theorem is omitted due to space constraints.

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 $\gamma = 3\sqrt{3}/(128\pi)$.*

Our family is comprised of the objects bounded by the curves that converge to the Koch snowflake [11].

4 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.*

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 Figure 2(a). If the square and the long thin 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 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 the rectangle have the same perimeter, then 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.
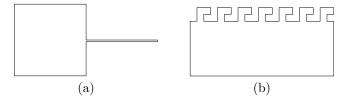

Figure 2: (a) Observation 2 (b) A locally fat object.

De Berg [4] 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 Figure 2(b) (which is an adaptation of a similar example by de Berg). This class of objects 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, $\mathsf{per}(O) \leq c \cdot \mathsf{diam}(O)$?*

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

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 $\gamma' := \gamma'(\gamma)$.

References

[1] L. Addario-Berry, O. Amini, J.-S. Sereni, and S. Thomassé. Guarding art galleries: The extra cost for sculptures is linear. In *Proc. of the 11th Scandinavian Workshop on Algorithm Theory (SWAT'08)*, pp. 41–52, 2008.

[2] G. Aloupis, P. Bose, V. Dujmović, C. Gray, S. Langerman, and B. Speckmann. Guarding fat polygons and triangulating guarded polygons. In *Proc. 20th Canadian Conference on Comput. Geom. (CCCG'08)*, pp. 107–111, 2008.

[3] L. P. Chew, H. David, M. J. Katz, and K. Kedem. Walking around fat obstacles. *Inf. Process. Lett.*, 83(3):135–140, 2002.

[4] M. de Berg. Improved bounds on the union complexity of fat objects. *Discrete Comput. Geom.*, 40(1):127–140, 2008.

[5] M. de Berg, M. Katz, F. van der Stappen, and J. Vleugels. Realistic input models for geometric algorithms. *Algorithmica*, 34:81–97, 2002.

[6] A. Efrat. The complexity of the union of (α,β)-covered objects. *SIAM J. Comput.*, 34:775–787, 2005.

[7] K. Kedem, R. Livne, J. Pach, and M. Sharir. On the union of Jordan regions and collision-free translational motion amidst polygonal obstacles. *Discrete Comput. Geom.*, 1:59–71, 1986.

[8] D. Kirkpatrick. Guarding galleries with no nooks. In *Proc. of the 12th Canadian Conf. on Comput. Geom. (CCCG'00)*, pp. 43–46, 2000.

[9] J. Matoušek, J. Pach, M. Sharir, S. Sifrony, and E. Welzl. Fat triangles determine linearly many holes. *SIAM J. Comput.*, 23:154–169, 1994.

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

[11] H. von Koch. Sur une courbe continue sans tangente, obtenue par une construction géométrique élémentaire. *Archiv för Matemat., Astron. och Fys.*, 1:681–702, 1904.