

Helly-Type Theorems for Approximate Covering

Julien Demouth · Olivier Devillers · Marc Glisse ·
Xavier Goaoc

Received: 27 August 2008 / Revised: 4 January 2009 / Accepted: 30 January 2009 /
Published online: 21 April 2009
© Springer Science+Business Media, LLC 2009

Abstract Let $\mathcal{F} \cup \{U\}$ be a collection of convex sets in \mathbb{R}^d such that \mathcal{F} covers U . We prove that if the elements of \mathcal{F} and U have comparable size, then a small subset of \mathcal{F} suffices to cover most of the volume of U ; we analyze how small this subset can be depending on the geometry of the elements of \mathcal{F} and show that smooth convex sets and axis parallel squares behave differently. We obtain similar results for surface-to-surface visibility amongst balls in three dimensions for a notion of volume related to form factor. For each of these situations, we give an algorithm that takes \mathcal{F} and U as input and computes in time $O(|\mathcal{F}| * |\mathcal{H}_\epsilon|)$ either a point in U not covered by \mathcal{F} or a subset \mathcal{H}_ϵ covering U up to a measure ϵ , with $|\mathcal{H}_\epsilon|$ meeting our combinatorial bounds.

Keywords Approximate covering · Helly-type theorems · 3D visibility · LP-type problems

The authors acknowledge support from the French–Korean Science and Technology amicable relationship program (STAR) 11844QJ.

J. Demouth · X. Goaoc (✉)
LORIA—Université Nancy 2, Campus Scientifique, B.P. 239, 54506 Villers-lès-Nancy, France
e-mail: Xavier.Goaoc@loria.fr

J. Demouth
e-mail: Julien.Demouth@loria.fr

O. Devillers
INRIA Sophia-Antipolis, B.P. 93, 06902 Sophia-Antipolis, France
e-mail: Olivier.Devillers@sophia.inria.fr

M. Glisse
Gipsa-Lab, CNRS UMR 5216, Domaine Universitaire, B.P. 46, 38402 St Martin d'Hères, France
e-mail: Marc.Glisse@normalesup.org

1 Introduction

A family \mathcal{F} of sets covers a set U if the union of the elements of \mathcal{F} contains U . The classical SETCOVER problem asks, given a covering \mathcal{F} of a finite set U , for the smallest subset of \mathcal{F} that covers U . In the geometric setting, both U and the elements of \mathcal{F} are subsets of a geometric space, for example, points, hyperplanes, or balls in \mathbb{R}^d . The original problem is NP-hard [12], and so are many of its geometric analogues. Therefore, approximation algorithms have been largely investigated, and in general, one looks for a subset of \mathcal{F} that *completely* covers U and whose size is near-optimal; approximation factors better than $\log|U|$ are provably difficult to achieve in the finite case [10, 13], and constant factor approximations were obtained for only a few geometric versions [6] (see also [5]). In this paper, we relax the problem in a different direction: given a covering \mathcal{F} of a set U , we look for a small subset of \mathcal{F} that covers *most* of U . Specifically, in the geometric setting we define an ϵ -covering of U as a collection \mathcal{H} of sets whose union covers U except for a volume of at most ϵ . Although this is a natural question, we are not aware of previous results in this direction.


Results Let \mathcal{F} be a covering of a convex set U by convex sets in \mathbb{R}^d . Let \mathcal{H}_ϵ denote a smallest ϵ -covering of U contained in \mathcal{F} . We let $\tilde{O}(n)$ denote $O(n \log^\beta n)$ for some β . Our main results are the following (the dimension d , as everywhere in this article, is assumed to be fixed):

- If the elements in \mathcal{F} have similar size, i.e., each can be sandwiched between two spheres of fixed radii, then $|\mathcal{H}_\epsilon|$ is bounded polynomially in $1/\epsilon$ and independently of $|\mathcal{F}|$ (Theorem 3).
- $|\mathcal{H}_\epsilon|$ is $O(\frac{1}{\epsilon})$ when \mathcal{F} consists of axis-parallel unit squares in the plane (Theorem 4) and $\tilde{O}(\epsilon^{\frac{1-d}{2}})$ if \mathcal{F} consists of unit balls in \mathbb{R}^d (Theorem 5) or smooth convex sets of bounded curvature (Corollary 11). These bounds are tight in the worst-case (up to the logarithmic factor).
- These results extend to visibility occlusion among disjoint unit balls in \mathbb{R}^3 , where the notion of volume used relates to the form factor (Theorem 12).
- For covering by squares or balls and visibility in 3D, we give algorithms that take \mathcal{F} and U as input and output in $O(|\mathcal{F}| * |\mathcal{H}_\epsilon|)$ -time either a point in U not covered by \mathcal{F} or an ϵ -cover of U contained in \mathcal{F} ; $|\mathcal{H}_\epsilon|$ denotes our bound on the size of the smallest ϵ -covering for that situation (Sect. 6).

Our results imply that there does not exist an arbitrarily large minimal ϵ -cover of a convex set by similar-sized convex sets, which is in sharp contrast with exact covering. The order $\sqrt{\epsilon}$ gap between our bounds in the case of squares and smooth convex sets with bounded curvature in the plane shows that the asymptotic behavior of $|\mathcal{H}_\epsilon|$ as $\epsilon \rightarrow 0$ depends not only on the size but also on the shape of the covering objects.

Geometric problems such as guarding or visibility can be rephrased as covering problems where, given a collection \mathcal{F} and a set U , one has to decide if \mathcal{F} covers U . Such tests can be expensive, e.g., no algorithm with complexity $o(n^4)$ is known for reporting visible pairs among n triangles in \mathbb{R}^3 [15, Problem 7.7.1(f)], so approximation algorithms are often used in practice. Our algorithms are interesting in that they are simple, have complexity linear in $|\mathcal{F}|$, and allow one to control the error a priori.

Fig. 1 Covering of U with no bounded size exact sub-covering



The computation of *form factors* is a classical question in computer graphics [7, 20]. While all form factors of a scene of disjoint convex polygons of complexity n can be approximated within error less than ϵ in time $O(\epsilon^{-2}n \log n + n^2 \log n)$ [16], we are not aware of any result in this direction for nonpolygonal objects.

Helly-Type Theorems Helly's theorem asserts that n convex sets in \mathbb{R}^d have nonempty intersection if any $d + 1$ of them have nonempty intersection. Results of similar flavor—that some property on a set \mathcal{F} can be checked by examining its subsets of bounded size—are known as *Helly-type theorems* and are the object of active research [8, 9, 21]. A collection \mathcal{F} covers U if and only if the intersection of the complements of its elements and U is empty; thus, if \mathcal{F} consists of complements of convex sets in \mathbb{R}^d and covers a convex set U , then $d + 1$ elements in \mathcal{F} suffice to cover U . Cases where such statements are known are, however, rather exceptional as for most classes of objects, there exists arbitrarily large minimal covering families (Fig. 1 illustrates the principle of such a construction for unit disks). Our Theorems 3, 4, 5, and 12 show that the situation is different when *approximate* covering is considered; perhaps the closest related results in this line of research are the quantitative Helly-type theorems of Bárány et al. [3, 4].

2 The General Case

In this section we consider the covering of a convex volume U by a family \mathcal{F} of fat convex objects of roughly equal size. The general idea is to decompose U using a grid and show that at least one object in \mathcal{F} completely covers at least one cell. We then bootstrap this idea, decomposing the other cells using a smaller grid. The fact that we do not recurse on one of the cells suffices to obtain the desired bound. We recall that we consider d , r , and R as constants. We start with a simple observation on approximation of a convex set by a grid:

Lemma 1 *Let $A \subset \mathbb{R}^d$ be a convex set of diameter at most R , and Γ a regular grid of step ℓ . The cells of Γ contained in A cover A except for a volume of $O(\ell)$.*

Proof To a cell σ whose interior meets ∂A , we associate the line L_σ through diagonally opposite vertices with direction closest to the normal of some (arbitrary)




Fig. 2 Finding a ball in $U \cap \mathcal{A}$

support hyperplane H to \mathcal{A} in some point interior to σ . At most two cells correspond to the same line since all grid vertices, on one side along L_σ , are separated from \mathcal{A} by H . There are 2^{d-1} pairs of possible directions of such lines L_σ . The projection of the vertices of Γ on a plane orthogonal to one such direction is a lattice whose primitive cell has a volume of $\Theta(\ell^{d-1})$. As a consequence, there are at most $O(R^{d-1}/\ell^{d-1})$ lines with that direction through a vertex of Γ that intersect \mathcal{A} . Thus, there are $O(2^d(R/\ell)^{d-1})$ cells in Γ whose interiors intersect $\partial\mathcal{A}$, and their total volume is $O(2^d R^{d-1} \ell)$. \square

A collection \mathcal{F} of sets has *scale* (r, R) if each element in \mathcal{F} contains a ball of radius r and is contained in one of radius R . We define $\kappa = r/(16R\sqrt{d})$ and prove the following technical lemma:

Lemma 2 *If U is a cube of side length ℓ in \mathbb{R}^d and \mathcal{A} is a convex set of scale (r, R) , with $\ell \leq 2r$, containing the center of U , then $\mathcal{A} \cap U$ contains at least one cell of any regular grid of step at most $\kappa\ell$.*

Proof Let C and C' denote the centers of, respectively, U and a ball B' of radius r contained in \mathcal{A} . We consider the balls B_1 and B_2 of radii $\ell/4$ and $\ell/2$ centered in C .

If B' intersects B_1 , we can find a ball of diameter $\ell/4$ centered on the line segment $[CC']$ contained in $U \cap \mathcal{A}$ (see Fig. 2(a)). If B' does not intersect B_1 , the convex hull of C and B' contains a cone of revolution with apex C , axis (CC') , height $\ell/4$, and half-angle $\theta = \sin^{-1}(r/2R)$ (see Fig. 2(b)), which in turn contains a ball of diameter $\frac{\ell}{4} \times \frac{r}{2R}$ (see Fig. 2(c)). In both cases, $U \cap \mathcal{A}$ contains a ball of radius $(\kappa\sqrt{d})\ell$ and thus a cube of any grid of step at most $\kappa\ell$. \square

We can now state the main result of this section:

Theorem 3 *For any d, r , and R , there exists a polynomial function $H(1/\epsilon) = H_{d,r,R}(\epsilon)$ such that the following holds. Any covering \mathcal{F} of a convex set $U \subset \mathbb{R}^d$ of diameter at most R by a collection of convex sets of scale (r, R) contains an ϵ -covering of U of size at most $H(1/\epsilon)$.*

Proof Let \mathcal{R}_0 be an $\frac{\epsilon}{2}$ -covering of U by $O(\epsilon^{-d})$ cells of a regular grid; Lemma 1 guarantees its existence. We proceed recursively. At step i , we have a subset \mathcal{C}_i of \mathcal{F} and a set \mathcal{R}_i of congruent cubes, each of side length $\ell_i = \kappa^i \ell_0$, that together form an $\epsilon/2$ -cover of U . For each cube $Y \in \mathcal{R}_i$, we select an object in \mathcal{F} that covers its

center and add it to C_{i+1} ; we then subdivide Y using a grid of step $\kappa \ell_i$ and collect the cubes not covered by C_{i+1} into \mathcal{R}_{i+1} . We initialize the recursion with \mathcal{R}_0 and $C_0 = \emptyset$. Lemma 2 implies that in the subdivision of any cube, at least one of the smaller cubes is covered, and thus

$$|\mathcal{R}_{i+1}| \leq |\mathcal{R}_i|(\kappa^{-d} - 1) \quad \text{and} \quad |C_{i+1}| \leq |C_i| + |\mathcal{R}_i|,$$

which resolves in:

$$|\mathcal{R}_i| \leq (\kappa^{-d} - 1)^i |\mathcal{R}_0|$$

and

$$|C_i| \leq \sum_{k=0}^{i-1} |\mathcal{R}_k| \leq \frac{(\kappa^{-d} - 1)^i - 1}{\kappa^{-d} - 2} |\mathcal{R}_0|. \tag{1}$$

As the volume of U not covered at step i is at most $\epsilon/2 + \ell_0^d(1 - \kappa^d)^i |\mathcal{R}_0|$, C_i is an ϵ -cover of U for

$$i \geq \frac{1}{\log \frac{1}{1-\kappa^d}} \log \left(\frac{2\ell_0 |\mathcal{R}_0|}{\epsilon} \right) = \Omega \left(\kappa^{-d} \log \frac{|\mathcal{R}_0|}{\epsilon} \right) = \Omega \left(d\kappa^{-d} \log \frac{1}{\epsilon} \right).$$

Substituting into (1), we get that $|C_i|$ is

$$O\left(\epsilon^{-O(d^2\kappa^{-d} \log \frac{1}{\epsilon})}\right),$$

which concludes the proof. □

If any of the scale or convexity assumptions is dropped, the statement becomes false as \mathcal{F} may consist of objects whose contribution to covering U decrease quickly: e.g., grid cells of decreasing size if r is not bounded, triangles of decreasing fatness (see Fig. 3, left) if R is not bounded, and unions of a fat convex set outside U with triangles of decreasing fatness (see Fig. 3, right).

While a more careful analysis might improve the bound obtained and in particular the dependency of the exponent of $1/\epsilon$ on d , the next sections show that pinning down the precise asymptotic behavior of $H(\epsilon)$ requires taking into account the shape of the objects in \mathcal{F} .




Fig. 3 Optimality of the assumptions of Theorem 3

3 Covering by Squares

For axis-parallel boxes in \mathbb{R}^d , the analysis of the previous section holds for $\kappa = 1/2$; since $|\mathcal{R}_0|$ is less than $(R/\epsilon)^d$, the bound becomes $O(\epsilon^{-O(d2^d)})$. We improve this bound in the planar case:

Theorem 4 *Let $U \subset \mathbb{R}^2$ be an axis-parallel square of side r covered by a finite collection \mathcal{F} of larger axis-aligned squares. For $\epsilon > 0$ sufficiently small, the smallest ϵ -covering of U contained in \mathcal{F} has size $O(\frac{1}{\epsilon})$; this bound is tight in the worst-case.*

Proof We first prove the lower bound illustrated by Fig. 4. Let U be a unit square, \mathcal{F} the (infinite) family of unit squares tangent to one of the diagonals of U , and $\mathcal{G} \subset \mathcal{F}$ an ϵ -cover of U . Consider the subset $\mathcal{G}^+ \subset \mathcal{G}$ of squares lying above the diagonal, and let x_1, \dots, x_k denote the abscissae of the tangency points of the squares in \mathcal{G}^+ , sorted increasingly. Let $\alpha_i = x_i - x_{i-1}$. For ϵ small enough, since \mathcal{G} is an ϵ -cover, we have $x_k - x_1 \geq \frac{1}{2}$, or equivalently,

$$\sum_{i=2}^k \alpha_i \geq \frac{1}{2}.$$

The uncovered area of U above the diagonal and between the $(i - 1)$ th and the i th squares is $\frac{1}{2}\alpha_i^2$. Thus, $\sum_{i=2}^k \alpha_i^2 \leq 2\epsilon$, and Hölder’s inequality yields that

$$\frac{1}{2} \leq \sum_{i=2}^k \alpha_i \leq \left(\sum_{i=2}^k \alpha_i^2 \right)^{\frac{1}{2}} \left(\sum_{i=2}^k 1^2 \right)^{\frac{1}{2}} \leq \sqrt{2\epsilon(k - 1)}$$


and, finally,

$$k \geq \frac{1}{8\epsilon}.$$

Note that this construction can be modified so that \mathcal{F} is finite.

We now turn our attention to the upper bound. Half of any rectangle Y contained in U can be covered by a pair $\{X_1, X_2\} \subset \mathcal{F}$: choose X_1 maximal for the inclusion among the squares in \mathcal{F} that contain the center of Y and, if a corner of X_1 lies inside Y , X_2 covering that point (otherwise X_1 suffices). We set $\mathcal{R}_0 = \{U\}$ and $\mathcal{C}_0 = \emptyset$ and iterate as follows: \mathcal{C}_{i+1} consists of \mathcal{C}_i augmented by all pairs $\{X_1, X_2\}$ for $Y \in \mathcal{R}_i$,

Fig. 4 Lower bound in the case of squares in the plane



and \mathcal{R}_{i+1} collects all rectangular pieces remaining uncovered (at most two pieces per element $Y \in \mathcal{R}_i$). Since the area not covered by \mathcal{C}_i is halved at every iteration, we get that \mathcal{C}_i is a 2^{-i} -covering of U . Furthermore,

$$|\mathcal{C}_i| \leq 2 \sum_{k=0}^{i-1} |\mathcal{R}_k| \leq 2 \sum_{k=0}^{i-1} 2^k |\mathcal{R}_0| = (2^{i+1} - 2) |\mathcal{R}_0|,$$

and the upper bound follows. □

4 Covering by Balls

When the objects of \mathcal{F} are balls in \mathbb{R}^d , we can prove the following, almost tight, bound:

Theorem 5 *Let \mathcal{F} be a covering of a convex set $U \subset \mathbb{R}^d$ of diameter at most R by finitely many balls, each of radius at least r . For any $\epsilon > 0$, the smallest ϵ -covering of U contained in \mathcal{F} has size $\tilde{O}(\epsilon^{\frac{1-d}{2}})$. This bound is tight up to the logarithmic factor in the worst-case.*

For the clarity of the exposition, we prove the result in two dimensions (Sect. 4.1) before discussing the general case (Sect. 4.2). We then discuss the extension of Theorem 5 to coverings by smooth objects with curvature of bounded norm (Sect. 4.3).

4.1 The Planar Case

Upper Bound The general idea is that the area to cover (the unit disk U) will be split in several disks Y_\star having the same radius that will be covered separately. One of these disks Y_i is covered by taking locally¹ an approximation of the covering disks in \mathcal{F} by an half-plane. At this point, we get an approximate covering; the uncovered part due to the approximation of disks by half-planes is split in a new generation of smaller disks $Y_{\star\star}$. Then we can iterate on this new family of disks.

We start the formal proof by the technical lemma below. For two disks X and Y , we denote by X^Y the half-plane containing X and bounded by the tangent to X at the projection² of the center of Y on the boundary of X (see Fig. 5). We denote by \mathcal{F}^Y the collection $\{X^Y \mid X \in \mathcal{F}\}$.

Lemma 6 *Let Y be a disk of radius $0 < r < 1$, and \mathcal{F} a covering of a unit disk U by larger disks. Then, $U \cap Y$ can be covered by $C(Y) \cup R(Y)$, where $C(Y)$ contains three elements of \mathcal{F} , and $R(Y)$ is a collection of at most $\frac{3}{r}$ disks of radius $4r^2$.*

¹Locally here means inside Y_i

²More formally, this point is defined as the intersection between ∂X and the ray from the center of X directed toward the center of Y . If the two disks have the same center, we can choose any tangent to X .




Fig. 5 Left: $(X^Y \cap Y) \setminus (X \cap Y)$ is inscribed in a rectangle (the thick rectangle) of sides at most $2r$ and $1 - \cos(\sin^{-1} r) \leq 4r^2$. Right: the covering of $(X^Y \cap Y) \setminus (X \cap Y)$ by disks of radius $4r^2$

Proof Since the collection \mathcal{F}^Y covers U , it also covers $U \cap Y$, and, by Helly’s theorem, three of these half-planes must cover $U \cap Y$ (because complements of half-planes are convex sets). We denote by $C(Y)$ the corresponding disks in \mathcal{F} . The portion of Y covered by X^Y but not by X forms a crescent shape inscribed in a rectangle limited by the boundary line of X^Y , the chord joining the intersection points of the circles ∂X and ∂Y , and the tangents to Y orthogonal to X^Y (see Fig. 5, left). The sides of this rectangle are smaller than $2r$, the diameter of Y , and $4r^2$. This rectangle can thus be covered by overlapping disks of radius $4r^2$ centered on its larger axis (Fig. 5 (right)). By choosing the disks so that the height covered at the intersection between two disks is $4r^2$, we need only $\frac{1}{r}$ disks. \square

We can now prove Theorem 5 for the case $d = 2$:

Proof of Theorem 5 for $d = 2$ We fix some $0 < r < 1$ and start by covering U by a collection \mathcal{R}_0 of μr^{-2} disks of radius r , for some constant μ , and let \mathcal{C}_0 denote the empty set. We then iterate as follows: \mathcal{R}_{i+1} collects the balls $R(Y)$, and \mathcal{C}_{i+1} consists of \mathcal{C}_i augmented by all $C(Y)$ for $Y \in \mathcal{R}_i$, where $C(\cdot)$ and $R(\cdot)$ denote the sets defined in Lemma 6. By induction, for any $i \geq 0$, $\mathcal{C}_i \cup \mathcal{R}_i$ covers U . Let α_i denote the area of U that is covered by disks in \mathcal{R}_i ; \mathcal{C}_i is an α_i -cover of U . The disks in \mathcal{R}_i have radius r_i satisfying the recurrence relation

$$r_i = 4r_i^2 \quad \text{with } r_0 = r,$$


and thus $r_i = \frac{1}{4}(4r)^{2^i}$. The number of disks in \mathcal{R}_i is governed by the relation

$$|\mathcal{R}_i| \leq \frac{3}{r_{i-1}} |\mathcal{R}_{i-1}|,$$

which gives

$$|\mathcal{R}_i| \leq \left(\prod_{k=0}^{i-1} 12(4r)^{-2^k} \right) |\mathcal{R}_0| \leq 12^i (4r)^{1-2^i} \mu r^{-2}$$

Fig. 6 Lower bound in the case of disks in the plane



and

$$\alpha_i \leq \pi r_i^2 |\mathcal{R}_i| = 12^i \mu \pi (4r)^{2^i - 1}.$$

Moreover, for each element in \mathcal{R}_{i-1} , we add three disks from \mathcal{F} to \mathcal{C}_i . Thus, the size of \mathcal{C}_i is given by

$$|\mathcal{C}_i| \leq 3 \sum_{k=0}^{i-1} |\mathcal{R}_k| \leq 3\mu r^{-2} 12^i \sum_{k=0}^{i-1} (4r)^{1-2^k} = O(12^i i (4r)^{1-2^{i-1}}).$$

Let $\epsilon > 0$ and k be such that

$$12^{k-1} \mu \pi (4r)^{2^{k-1}-1} \leq \epsilon.$$

The previous inequalities then bound α_i by ϵ and $|\mathcal{C}_k|$ by $O(\epsilon^{-\frac{1}{2}} \log^6 \frac{1}{\epsilon})$. □

Lower Bound The following construction shows that the upper bound in Theorem 5 is optimal for $d = 2$ up to the logarithmic factor.

Lemma 7 *There exists a family \mathcal{F} of unit disks in \mathbb{R}^2 covering a unit disk $U \subset \mathbb{R}^2$ such that, for arbitrary small $\epsilon > 0$, any ϵ -covering of U contained in \mathcal{F} has size $\Omega(\epsilon^{-\frac{1}{2}})$.*

Proof We equip the plane \mathbb{R}^2 with a frame (O, x, y) where O denotes the center of U . Let \mathcal{F} be the (infinite) family of all unit disks tangent to the x -axis inside U (see the Fig. 6) and let \mathcal{G} be a finite subset of \mathcal{F} that covers U except for an area of at most ϵ . Consider the subset $\mathcal{G}^+ \subset \mathcal{G}$ of disks whose centers are above the x -axis and let x_1, \dots, x_k denote the abscissae of the tangency points of the disks in \mathcal{G}^+ , sorted increasingly. Clearly \mathcal{G}^+ must cover the part of U above the x -axis up to an area of at most ϵ .

Let $\alpha_i = x_i - x_{i-1}$. For ϵ small enough, since \mathcal{G} is an ϵ -cover, we have

$$x_k - x_1 \geq 1 \quad \Rightarrow \quad \sum_{i=2}^k \alpha_i \geq 1.$$

The uncovered area of U above the x -axis and between the $(i - 1)$ th disk and the i th disk is at least $\frac{\alpha_i^3}{24}$ since this area is bounded from below by

$$2 \int_0^{\frac{\alpha_i}{2}} \frac{1}{2} x^2 dx = \frac{\alpha_i^3}{24}.$$

Thus, $\sum_{i=2}^k \alpha_i^3 \leq 24\epsilon$, and Hölder’s inequality yields

$$1 \leq \sum_{i=2}^k \alpha_i \leq \left(\sum_{i=2}^k \alpha_i^3 \right)^{\frac{1}{3}} \left(\sum_{i=2}^k 1^{\frac{3}{2}} \right)^{\frac{2}{3}} \leq (24\epsilon)^{\frac{1}{3}} (k - 1)^{\frac{2}{3}}.$$

The statement follows. □

Remark This example involves an infinite covering family, but the same can easily be achieved with a finite family by letting the disks intersect the x -axis on arbitrarily small lengths.

4.2 Arbitrary Dimension

Upper Bound We start with a generalization of Lemma 6.


Lemma 8 *Let Y be a ball of radius $0 < \ell < 1$, and \mathcal{F} a covering of a unit ball $U \subset \mathbb{R}^d$ by larger balls. Then $U \cap Y$ can be covered by $C(Y) \cup R(Y)$, where $C(Y)$ contains $(d + 1)$ elements of \mathcal{F} , and $R(Y)$ is a collection of $O(\ell^{1-d})$ balls of radius $\rho(\ell) = O(\ell^2)$.*

Proof Given two balls X and Y , we denote by X^Y the half-space containing X and bounded by the hyperplane tangent to X at the projection of the center of Y on ∂X (this tangency point is the intersection between ∂X and the oriented line from the center of X to the center of Y). Notice that this is well defined whenever X and Y have distinct centers. We denote by \mathcal{F}^Y the collection of all X^Y for X in \mathcal{F} .

Let Y be some ball. If a ball of \mathcal{F} has the same center as Y , then it covers Y , and we are done. We can then assume that this is not the case. Since \mathcal{F} covers U , \mathcal{F}^Y also covers U , and in particular it covers $Y \cap U$, and by Helly’s theorem there are $d + 1$ elements in \mathcal{F}^Y that cover $Y \cap U$; we denote by $C(Y)$ the corresponding $d + 1$ balls in \mathcal{F} . If the radius of Y is $\ell < 1$, then the region $(X^Y \cap Y) \setminus (X \cap Y)$ is included in a cylinder defined by a $(d - 1)$ -dimensional ball of radius ℓ and an orthogonal segment of length $O(\ell^2)$. This region can thus be covered by a collection $R_X(Y)$ of $O(\ell^{1-d})$ balls of radius $\rho(\ell) = O(\ell^2)$. Covering the $d + 1$ regions corresponding to the $d + 1$ balls $X \in C(Y)$ gives a collection $R(Y) = \bigcup_{X \in C(Y)} R_X(Y)$, which concludes the proof. □

The proof of Theorem 5 for a general d follows the same approach as in the case $d = 2$, so we omit the details of the computations.

Fig. 7 For the proof of Lemma 9



Proof of Theorem 5 We fix some constant $r_0 \in (0, r)$ small enough such that for some constant $K > 0$, the function ρ introduced in Lemma 8 satisfies $\rho(t) \leq Kt^2$ for any $0 < t \leq r_0$. Denote $C = Kr_0$. We further assume that r_0 is small enough so that $0 < C < 1$. Again, we construct a small ϵ -covering from \mathcal{F} by starting with a covering \mathcal{R}_0 of U by $O((R/r_0)^d)$ balls of radius r_0 , setting $\mathcal{C}_0 = \emptyset$ and iterating:

$$\mathcal{C}_{i+1} \leftarrow \mathcal{C}_i \cup \left(\bigcup_{Y \in \mathcal{R}_i} C(Y) \right) \quad \text{and} \quad \mathcal{R}_{i+1} \leftarrow \bigcup_{Y \in \mathcal{R}_i} R(Y).$$

After k iterations, \mathcal{C}_k has size at most $O(C^{(1-d)2^k-1} D^{k-1})$ (where D is a positive constant) and covers U except for the region covered by the balls in \mathcal{R}_k , which consists of $O(C^{(1-d)2^k} D^k)$ balls of radius $O(C^{2^k})$. The volume possibly not covered by the balls in \mathcal{C}_k is thus $O(C^{(1-d)2^k+d2^k} D^k) = O(C^{2^k} D^k)$. By choosing k such that $\epsilon = C^{2^k} D^k$, we get an $O(\epsilon)$ -cover of U of size $O(\epsilon^{\frac{1-d}{2}} \text{polylog}(\frac{1}{\epsilon}))$. \square

Note that the constant hidden in the $O()$ notation depends on d .


Lower Bound To generalize the lower bound we use the following lemma:

Lemma 9 *Let p be a point, and U a convex set of volume v in \mathbb{R}^d . Let δ be the distance from p to a point in U farthest from it. The part of U at distance greater than $\delta/2$ from p has volume $\Omega(v)$.*

Proof We refer to Fig. 7. We call $q \in U$ the farthest point from p (or one of them). Let H be the hyperplane of points equidistant from p and q , and let H' be the hyperplane parallel to H at distance 2δ from q and δ from p . H intersects U in a convex set U^0 . We draw the half-cone C centered at q that intersects H in U^0 . The part of U at distance larger than $\delta/2$ from p contains the part U^+ of U that is on the same side of H as q . Furthermore, U^+ contains the part C^+ of C on the same side of H as q . The part U^- of U on the other side of H is contained in the region C^- delimited by C , H , and H' . Since the volume of C^- is equal to $4^d - 1$ times that of C^+ , the statement follows. \square

We can now prove that the bound of Theorem 5 is optimal up to the logarithmic factor.

Fig. 8 For the proof of Theorem 10



Theorem 10 *There exists a family \mathcal{F} of unit balls in \mathbb{R}^d covering a unit ball $U \subset \mathbb{R}^d$ such that for arbitrarily small $\epsilon > 0$, any subset of \mathcal{F} that is an ϵ -covering of U has size $\Omega(\epsilon^{-\frac{d-1}{2}})$.*

Proof Let H be a hyperplane through the center of U , let B be the $(d - 1)$ -dimensional unit ball $U \cap H$, and let \mathcal{F} denote the set of all unit balls tangent to H in a point of B . Observe that \mathcal{F} covers U . We assume that H is given by $x_d = 0$ and, to simplify the description, consider it to be horizontal.

The portion of U on one side of H is covered by the balls of \mathcal{F} that are on that side of H . We thus only argue about the portion U^+ of U above H and covered by the set \mathcal{F}^+ of balls in \mathcal{F} above H . We denote by ∂U^+ the part of the boundary of U above H .

Let $\mathcal{G}^+ \subset \mathcal{F}^+$ be a family of k balls. For each ball $X \in \mathcal{G}^+$, let P_X denote the parabola with equation $2 * x_d = \sum_{i=1}^{d-1} (x_i - t_i)^2$, where $(t_1, \dots, t_{d-1}, 0)$ is the tangency point of X with H . Since X is completely above the parabola P_X (see Fig. 8 on the left), the volume of U not covered by \mathcal{G}^+ is bounded from below by the volume of the region above B and under the parabolas and ∂U^+ .

Let \mathcal{T}^+ denote the set of tangency points of \mathcal{G}^+ on H . The height of the lowest parabola above a point p in B is proportional to the square of the distance from p to the closest point in \mathcal{T}^+ . Let C be a cell of the Voronoï diagram of \mathcal{T}^+ restricted to B , and let v denote its volume. The diameter of C is $\Omega(v^{\frac{1}{d-1}})$, and, by Lemma 9, a subset of C of volume $\Omega(v)$ is at distance $\Omega(v^{\frac{1}{d-1}})$ from its center in \mathcal{T}^+ . The volume between this cell and the parabola above it is thus $\Omega(v^{1+\frac{2}{d-1}})$. Since the cells partition B , the sum of their volumes is $\Omega(1)$, and Hölder’s inequality yields:

$$\Omega(1) = \sum_{i=1}^k v_i \leq \left(\sum_{i=1}^k v_i^{1+\frac{2}{d-1}} \right)^{\frac{d-1}{d+1}} \left(\sum_{i=1}^k 1^{\frac{d+1}{2}} \right)^{\frac{2}{d+1}} = \left(\sum_{i=1}^k v_i^{1+\frac{2}{d-1}} \right)^{\frac{d-1}{d+1}} k^{\frac{2}{d+1}}.$$

Hence, the volume below the parabolas is $\Omega(k^{-\frac{2}{d-1}})$. To take ∂U^+ into account, we consider the ball B' obtained by scaling B by a factor $\frac{1}{2}$. The previous argument still yields that the volume between B' and the parabolas is $\Omega(k^{-\frac{2}{d-1}})$. Also, above any point in B' , the ratio of the height of the lowest parabola to that of ∂U^+ is bounded. Thus, the volume above B' and below the parabolas and ∂U^+ is $\Omega(k^{-\frac{2}{d-1}})$. It follows that the volume not covered by \mathcal{G}^+ is $\Omega(k^{-\frac{2}{d-1}})$. Equivalently, any subset of \mathcal{F}^+ leaving a volume at most ϵ of U^+ uncovered has size $\Omega(\epsilon^{\frac{1-d}{2}})$. □

4.3 Smooth Convex Sets

Lemma 8 requires that (i) given a ball Y , the set $U \cap Y$ be convex and that (ii) the difference between $X^Y \cap Y$ and $X \cap Y$ can be covered by $O(\frac{1}{r})$ balls of radius $O(r^2)$. If an object is convex and its boundary has a curvature of bounded norm, then for any point M on this boundary, the object contains a ball (of radius bounded away from 0) and is contained in a half-space delimited by a hyperplane tangent to both the object and the ball in M ; this means that covering the region between the ball and the hyperplane is enough to cover the region between the object and the hyperplane. Theorem 5 thus extends to:

Corollary 11 *Let $U \subset \mathbb{R}^d$ be a convex set of diameter at most R , and \mathcal{F} a covering of U by smooth convex sets whose curvatures have a norm at most γ . For any $\epsilon > 0$, the smallest subset of \mathcal{F} that is an ϵ -covering of U has size $\tilde{O}(\epsilon^{\frac{1-d}{2}})$.*

5 Visibility Among 3D Unit Balls

Two among n objects are *visible* if they support the endpoint of a segment that intersects no other object, and such a segment is called a *visibility segment*. Visibility between objects can be recast as a covering problem by observing that two objects are mutually visible if and only if the set of segments they support is not covered by the set of segments supported by these two objects and intersecting some other object. Yet, it is not clear whether Theorem 3 applies in this setting. In this section we show that Theorem 5 yields a similar result for visibility among balls.

A natural “volume” to quantify approximate visibility between two objects—similarly to the ϵ -coverings discussed so far—is given by the measure of the set of lines supporting visibility segments between these two objects. In fact, this corresponds, up to normalization, to the *form factor* used in computer graphics (when constant basis functions are used) to quantify visibility for simulating illumination. We call this measure the *amount of visibility* between the two objects. Building on Theorem 5, we prove:

Theorem 12 *Let $\mathcal{F} \cup \{A, B\}$ be a collection of disjoint unit balls in \mathbb{R}^3 such that A and B are mutually invisible. For any $\epsilon > 0$, there exists a subset $\mathcal{G}_\epsilon \subset \mathcal{F}$, of size $\tilde{O}(\epsilon^{-\frac{7}{2}})$, such that the amount of visibility between A and B allowed by $\mathcal{G}_\epsilon \cup \{A, B\}$ is $O(\epsilon M_{A,B})$, where $M_{A,B}$ denotes the measure of the set of lines intersecting both A and B .*

Measure in Line Space Recall that there exists, up to scaling by some constant, a unique measure over lines in \mathbb{R}^3 that is invariant under rigid motions [18]. We choose the constant such that the set of lines intersecting a unit ball has measure $4\pi^2$.

Let S be a measurable set of lines, let \vec{S} denote its set of directions, and, for $u \in \mathbb{S}^2$, let $S(u)$ be the set of lines in S with direction u . Finally, let $|\vec{S}|$ denote the area of \vec{S} (on the unit sphere of directions), and let $|S(u)|$ be the measure of $S(u)$, i.e., the area of the intercept of $S(u)$ with a plane orthogonal to u .

Lemma 13 *The measure of a set of lines S is bounded from above by $|\vec{S}| \times \max_u |S(u)|$.*

Proof Let us represent a line by its direction, given in spherical coordinates $(\theta, \phi) \in [0, 2\pi) \times [0, \pi]$, and a point (x, y) in the plane orthogonal to its direction through the origin. With our choice of constant, the density of the measure on the space of lines is then

$$dG = dx dy \sin \theta d\theta d\phi,$$

and the statement follows by integrating separately along the couples (x, y) and (θ, ϕ) . □

We now prove Theorem 12:

Proof of Theorem 12 Let us fix $\epsilon_0 > 0$. Let a and b be the respective centers of balls A and B . Given $u = (\theta, \phi) \in \mathbb{S}^2$, we denote by $p_u(\cdot)$ the projection on the plane through a with normal u , equipped with a frame with origin at a and with $p_u([0, 2\pi), \phi)$ as x -axis (in the sequel, points $p_u(\cdot)$ are considered in the two-dimensional affine space). The proof consists of four steps:

Step 1 We first find a small subset of \mathcal{F} that blocks visibility between A and B for some given direction $u \in \mathbb{S}^2$. Let \mathcal{F}_u denote the collection of balls that block visibility between A and B along u (i.e., a ball X belongs to \mathcal{F}_u if some oriented line with direction u intersects X in-between A and B). Since A and B are mutually invisible, $p_u(\mathcal{F}_u)$ is a collection of unit discs that covers $p_u(A) \cap p_u(B)$. Furthermore, $p_u(A) \cap p_u(B)$ is a bounded convex set. Hence, Corollary 11 yields that for any $\epsilon_0 > 0$, there exists a subset $\mathcal{H}_u \subset \mathcal{F}_u$ of size at most

$$|\mathcal{H}_u| = O\left(\epsilon_0^{-\frac{1}{2}} \text{polylog} \frac{1}{\epsilon_0}\right)$$

such that $p_u(\mathcal{H}_u)$ is an ϵ_0 -covering of $p_u(A) \cap p_u(B)$.

Step 2 We now argue that a subset that almost blocks visibility in direction u still almost blocks visibility in any direction v close enough to u . Let $\alpha > 0$ be some constant, and $v \in \mathbb{S}^2$ be a vector making, with u , an angle of at most $\frac{\alpha}{L+2}$, where L is the distance between a and b . For any ball $X \in \mathcal{F}_u$ with center x , we have (see Fig. 9)

$$\begin{aligned} p_u(x)p_v(x) &\leq (L+2) \left(\cos\left(\beta - \frac{\alpha}{L+2}\right) - \cos \beta \right) \\ &\leq 2(L+2) \sin\left(\beta - \frac{\alpha}{2(L+2)}\right) \sin\left(\frac{\alpha}{2(L+2)}\right) \\ &\leq \alpha, \end{aligned}$$

since $\sin x \leq x$ for $x \leq 1$. So, the disk with center $p_u(x)$ and radius $1 - \alpha$ is contained in $p_v(X)$. It follows that, for any vector v making angle at most $\frac{\alpha}{L+2}$ with u , $p_v(\mathcal{H}_u)$ covers $p_u(A) \cap p_u(B)$ except an area of at most $\epsilon_0 + 2\pi\alpha|\mathcal{H}_u|$.

Fig. 9 Blocking visibility at Step 2





Fig. 10 Bounding areas at Step 2



By definition of p_u , we have $p_u(A) = p_v(A)$ and, for the same reason as above, $p_u(b)p_v(b) \leq \alpha$. Thus, the area of the difference

$$(p_v(A) \cap p_v(B)) \setminus (p_u(A) \cap p_u(B))$$

is bounded from above by 2α (see Fig. 10). Hence, $p_v(\mathcal{H}_u)$ covers $p_v(A) \cap p_v(B)$ except an area of at most


$$\epsilon_0 + 2\alpha + 2\pi\alpha|\mathcal{H}_u|.$$

Note that for a ball $X \in \mathcal{H}_u$, having a nonempty intersection $p_v(A) \cap p_v(B) \cap p_v(X)$ does not guarantee that X blocks visibility between A and B : lines with directions u and v may intersect the three balls in different orders. It thus remains to remove the area covered by $p_v(\mathcal{H}_u \setminus \mathcal{F}_v)$; we claim that this area is $O(\alpha)$.

Indeed, first, observe that if a ball X is in $\mathcal{F}_u \setminus \mathcal{F}_v$ and is such that $p_v(X) \cap p_v(A) \cap p_v(B) \neq \emptyset$, then the balls $\{A, B, X\}$ have two distinct geometric permutations (along direction u we have AXB , whereas along direction v the permutation is ABX or XAB). Since these are disjoint unit balls, the centers of two of them are at most a distance of $2\sqrt{2}$ apart (see Fig. 11). If these two balls are A and B , then the theorem holds since they have at most a constant number of blockers. Otherwise, an immediate packing argument yields that at most a constant number, say c_1 , of balls in $\mathcal{F}_u \setminus \mathcal{F}_v$ contribute to cover $p_v(A) \cap p_v(B)$. Also, there is some direction w in the interval $[u, v]$ such that $p_w(X)$ is tangent to $p_w(Y)$ with $Y \in \{A, B\}$. Since

$$p_v(x)p_v(y) \geq p_w(x)p_w(y) - p_w(x)p_v(x) - p_w(y)p_v(y) \geq 2 - 2\alpha,$$

Fig. 11 Two different geometric permutations



the area $p_v(X) \cap p_v(Y)$ is bounded from above by

$$2 \int_{1-\alpha}^1 \sqrt{1-x^2} dx,$$

which is, at most, 2α (since $\sqrt{1-x^2} \leq 1$ on $[1-\alpha, 1]$). This also bounds the contribution of $p_v(X)$ in covering $p_v(A) \cap p_v(B)$, and the claim follows.

Step 3 We now almost block visibility between A and B by applying the previous construction to a sample of \mathbb{S}^2 . The directions T of common line transversals to A and B make up a disc of radius $\arcsin(\frac{2}{L})$ on \mathbb{S}^2 . We can thus choose a collection D of $O(\alpha^{-2})$ directions such that the discs of radii $\frac{\alpha}{L+2}$ centered on these directions completely cover T . Let \mathcal{H} and h denote respectively

$$\mathcal{H} = \bigcup_{u \in D} \mathcal{H}_u \quad \text{and} \quad h = \max_{u \in \mathbb{S}^2} |\mathcal{H}_u| = O\left(\epsilon_0^{-\frac{1}{2}} \text{polylog} \frac{1}{\epsilon_0}\right).$$

\mathcal{H} has size $O(\alpha^{-2}h)$ and, for any $u \in \mathbb{S}^2$, $p_u(\mathcal{H} \cap \mathcal{F}_u)$ covers $p_u(A) \cap p_u(B)$ except an area of at most

$$\epsilon_0 + O(\alpha) + 2\pi\alpha h.$$

Let V denote the set of lines intersecting A and B and no ball in \mathcal{H} between A and B . Lemma 13 yields that the measure of V is bounded from above by

$$(\epsilon_0 + O(\alpha) + 2\pi\alpha h)\pi \arcsin^2\left(\frac{2}{L}\right) = O((\epsilon_0 + \alpha h)M_{A,B}),$$

where $M_{A,B}$ denotes the measure of the set of lines through A and B .

Step 4 We now have a set \mathcal{H} of size $O(\alpha^{-2}\epsilon_0^{-\frac{1}{2}} \text{polylog} \frac{1}{\epsilon_0})$ that blocks visibility between A and B up to a set of lines of measure $V = O((\epsilon_0 + \alpha h)M_{A,B})$. By choosing $\alpha = \epsilon_0^{\frac{3}{2}}$ we get

$$|\mathcal{H}| = O\left(\epsilon_0^{-\frac{7}{2}} \text{polylog} \frac{1}{\epsilon_0}\right)$$

and

$$V = O((\epsilon_0 + \alpha h)M_{A,B}) = O\left(\epsilon_0 \operatorname{polylog} \frac{1}{\epsilon_0} M_{A,B}\right).$$

Finally, setting ϵ_0 such that $\epsilon = \epsilon_0 \operatorname{polylog} \frac{1}{\epsilon_0}$, \mathcal{H} is a subset of \mathcal{F} of size $O(\epsilon^{-\frac{7}{2}} \operatorname{polylog} \frac{1}{\epsilon})$, such that the amount of visibility between A and B in $\mathcal{H} \cup \{A, B\}$ is $O(\epsilon M_{A,B})$. □

6 Algorithms

The proofs of Theorems 4, 5, and 12 are constructive, provided that $C(Y)$ and $R(Y)$ can be effectively computed. As in the previous sections, we consider here d as a constant.

Covering by Squares In the case of covering by squares, the sets $C(Y)$ and $R(Y)$ can be computed trivially in $O(|\mathcal{F}|)$ time. We thus have the following consequence:

Corollary 14 *Given a covering \mathcal{F} of a unit square U by unit squares, we can compute in $O(\frac{|\mathcal{F}|}{\epsilon})$ -time a point in U not covered by \mathcal{F} or an ϵ -cover of U of size $O(\frac{1}{\epsilon})$ contained in \mathcal{F} .*

Covering by Balls In the case of covering by balls, the main difficulty is to compute $C(Y)$ ($R(Y)$ follows immediately). We recast this task as an LP-type problem.

We first recall some classical facts on the class of optimization problems called LP-type problems (or *generalized linear programming*) [19]. Let \mathcal{H} be a set, and ϕ a map $\phi : 2^{\mathcal{H}} \rightarrow \Omega$ from the family of subsets of \mathcal{H} to some completely ordered set Ω . The pair (\mathcal{H}, ϕ) is an LP-type problem if it satisfies two properties:

Monotonicity: if $F \subseteq G \subseteq \mathcal{H}$, then $\phi(F) \leq \phi(G)$.

Locality: if $F \subseteq G \subseteq \mathcal{H}$ and $\phi(F) = \phi(G)$, then for any $x \in \mathcal{H}$,

$$\phi(F \cup \{x\}) \neq \phi(F) \iff \phi(G \cup \{x\}) \neq \phi(G).$$

A subset $B \subseteq F$ such that $\phi(B) = \phi(F)$, which is minimal for this property, is a *basis* of F . The *combinatorial dimension* of an LP-type problem is the maximal cardinality of a basis. For an LP-type problem with constant combinatorial dimension, a basis B of \mathcal{H} can be computed in $O(|\mathcal{H}|)$ time, using, e.g., the algorithm of Clarkson [11].

Lemma 15 *Let \mathcal{H} be a family of half-spaces in \mathbb{R}^d , and Y a ball. We can compute, in $O(|\mathcal{H}|)$ -time, either a $(d + 1)$ -tuple in \mathcal{H} that covers Y or a point in Y not covered by \mathcal{H} .*

Proof Let $\phi : 2^{\mathcal{H}} \rightarrow \mathbb{R}$ be the map that associates to $\mathcal{G} \subseteq \mathcal{H}$ the real

$$\phi(\mathcal{G}) = \min \left\{ t \in [0, +\infty) \mid \bigcup_{x \in \mathcal{G}} x \oplus D(t) \text{ covers } Y \right\},$$

where \oplus and $D(t)$ denote respectively the Minkowski sum operator and the disk of radius t centered at the origin. The problem (\mathcal{H}, ϕ) is an LP-type problem [1, 2]. Furthermore, Helly’s theorem implies that its combinatorial dimension is bounded by $d + 1$. Thus, a basis B of \mathcal{H} can be computed in $O(|\mathcal{H}|)$ time. If $\phi(B) = 0$, then B is a $(d + 1)$ -tuple in \mathcal{H} that covers Y , otherwise \mathcal{H} does not cover Y . In the latter case, observe that the boundaries of the half-spaces $x \oplus D(\phi(B))$ for $x \in B$ intersect in a point that is not covered by $\bigcup_{x \in \mathcal{H}} x$. \square

As a consequence, we obtain:

Corollary 16 *Let \mathcal{F} be a covering of a unit ball $U \subset \mathbb{R}^d$ by unit balls. We can compute a point in U not covered by \mathcal{F} or an ϵ -cover of U of size $\tilde{O}(\epsilon^{\frac{1-d}{2}})$ contained in \mathcal{F} in time $\tilde{O}(|\mathcal{F}| \epsilon^{\frac{1-d}{2}})$.*

Proof We construct the sets \mathcal{C}_i and \mathcal{R}_i by repeating, as indicated in the proof of Theorem 5, the operation

$$\mathcal{C}_{i+1} \leftarrow \mathcal{C}_i \cup \left(\bigcup_{Y \in \mathcal{R}_i} C(Y) \right) \quad \text{and} \quad \mathcal{R}_{i+1} \leftarrow \bigcup_{Y \in \mathcal{R}_i} R(Y).$$

Assume that we are given \mathcal{C}_i and \mathcal{R}_i . For every ball $Y \in \mathcal{R}_i$, we run the algorithm described in Lemma 15 and obtain either a point in $Y \cap U$ not covered by \mathcal{F} or a family $C(Y)$; in the former case, we stop and return that U is not covered, and in the latter, we compute $R(Y)$. Overall, the times spent for computing \mathcal{C}_i and \mathcal{R}_i are respectively $O(|\mathcal{C}_i| * |\mathcal{F}|)$ and $O(|\mathcal{R}_i|)$. Let k denote the number of iterations performed. Since we need not compute \mathcal{R}_k , the complexity of the algorithm is $O(|\mathcal{C}_k| * |\mathcal{F}| + |\mathcal{R}_{k-1}|)$; with the same convention as in the proof of Theorem 5, $|\mathcal{C}_k|$ is $O(\epsilon^{\frac{1-d}{2}} \text{polylog}(\frac{1}{\epsilon}))$, and $|\mathcal{R}_{k-1}|$ is $O(\epsilon^{\frac{1-d}{2}})$, and so the time complexity of the algorithm is $O(|\mathcal{F}| \epsilon^{\frac{1-d}{2}} \text{polylog}(\frac{1}{\epsilon}))$. \square

Visibility Among Unit Balls Corollary 16 makes the proof of Theorem 12 constructive and we get:

Corollary 17 *Let \mathcal{F} be a collection of disjoint unit balls in \mathbb{R}^3 , and let A and B be two unit balls. We can compute in $\tilde{O}(|\mathcal{F}| \epsilon^{-\frac{7}{2}})$ -time a visibility segment between A and B or a subset $\mathcal{G}_\epsilon \subset \mathcal{F}$, of size $\tilde{O}(\epsilon^{-\frac{7}{2}})$, such that the amount of visibility between A and B in $\mathcal{G}_\epsilon \cup \{A, B\}$ is $O(\epsilon)$.*

7 Conclusion

We showed that the size of the smallest ϵ -covering contained in a covering \mathcal{F} of a set U can be bounded polynomially in $1/\epsilon$ and independently of $|\mathcal{F}|$ when all sets are convex and the size of the sets in \mathcal{F} are comparable with that of U . The order $\sqrt{\epsilon}$ gap between our bounds for smooth sets and squares indicate that the asymptotic

behavior of the size of the smallest ϵ -covering depends on the shape of the objects. Is this gap characteristic of the presence of sharp corners, as is the case in other contexts such as approximate range searching [17] or the size of the convex hull of a random sample [14], or do other simple shapes lead to different bounds?

These bounds yield simple and efficient algorithms for, given a family \mathcal{F} and a set U , certifying either that \mathcal{F} does not cover U or that \mathcal{F} misses at most a volume ϵ of U . We gave an application to approximate 3D visibility with an algorithm to decide in linear time if two balls are visible or if their form factor is at most ϵ . A natural continuation would be to compare these results to the provable bounds on the error provided by methods for approximating visibility queries used in application areas, e.g., sampling and point-to-point visibility in computer graphics.

Acknowledgements We thank Otfried Cheong and Andreas Holmsen for helpful discussions in the early stage of this work. We are also grateful to Ken Clarkson, Mira Lee, and Antoine Vigneron for pointing out several relevant references.

References

1. Amenta, N.: Helly theorems and generalized linear programming. PhD thesis, U.C. Berkeley (1993)
2. Amenta, N.: Helly-type theorems and generalized linear programming. *Discrete Comput. Geom.* **12** (1994)
3. Barany, I., Katchalski, M., Pach, J.: Quantitative Helly-type theorems. *Proc. Am. Math. Soc.* **86**(1), 109–114 (1982)
4. Barany, I., Katchalski, M., Pach, J.: Helly's theorem with volumes. *Am. Math. Mon.* **91**(6), 362–365 (1984)
5. Bronnimann, H., Goodrich, M.T.: Almost optimal set covers in finite VC-dimension. *Discrete Comput. Geom.* **14**, 463–479 (1995)
6. Clarkson, K.L., Varadarajan, K.: Improved approximation algorithms for geometric set cover. *Discrete Comput. Geom.* **37**, 43–58 (2007)
7. Cohen, M.F., Wallace, J.R.: *Radiosity and Realistic Image Synthesis*. Academic Press, New York (1993)
8. Danzer, L., Grünbaum, B., Klee, V.: Helly's theorem and its relatives. In: Klee, V. (ed.) *Convexity, Proceedings of Symposia in Pure Mathematics*, pp. 101–180. Amer. Math. Soc., Reading (1963)
9. Eckhoff, J.: Helly, Radon and Carathéodory type theorems. In: Gruber, P.M., Wills, J.M. (eds.) *Handbook of Convex Geometry*, pp. 389–448. North Holland, Amsterdam (1993)
10. Feige, U.: A threshold of $\ln n$ for approximating set cover. *J. ACM* **45**(4), 634–652 (1998)
11. Gärtner, B., Welzl, E.: Linear programming—randomization and abstract frameworks. In: 13th Annual Symposium on Theoretical Aspects of Computer Science. *Lecture Notes in Computer Science*, vol. 1046, pp. 669–687. Springer, Berlin (1996)
12. Karp, R.: Reducibility among combinatorial problems. In: *Complexity of Computer Computations, Proc. Sympos. IBM Thomas J. Watson Res. Center*, pp. 85–103 (1972)
13. Lund, C., Yannakakis, M.: On the hardness of approximating minimization problems. *J. ACM* **41**(5), 960–981 (1994)
14. Golin, M.: How many maxima can there be? *Comput. Geom. Theory Appl.* **2**, 335–353 (1993)
15. Pach, J., Sharir, M.: *Combinatorial Geometry with Algorithmic Applications*, The Alcalá Lectures, Alcalá, Spain, August 31–September 5, 2006
16. Pellegrini, M.: Monte Carlo approximation of form factors with error bounded a priori. In: *Symposium on Computational Geometry*, pp. 287–296 (1995)
17. Arya, S., Malamatos, T., Mount, D.M.: The effect of corners on the complexity of approximate range searching. In: *Symposium on Computational Geometry*, pp. 11–20 (2006)
18. Santalo, L.A.: *Integral Geometry and Geometric Probability*, 2nd edn. Cambridge University Press, New York (2004)

19. Sharir, M., Welzl, E.: A combinatorial bound for linear programming and related problems. In: STACS '92: Proceedings of the 9th Annual Symposium on Theoretical Aspects of Computer Science, pp. 569–579. Springer, Berlin (1992)
20. Sillion, F., Puech, C.: Radiosity and Global Illumination. Morgan Kaufmann, San Mateo (1994)
21. Wenger, R.: Helly-type theorems and geometric transversals. In: Goodman, J.E., O'Rourke, J., (eds.) Handbook of Discrete & Computational Geometry, 2nd edn., pp. 73–96. CRC Press, Boca Raton (2004)