Helly-type theorems for approximate covering

Julien Demouth, Olivier Devillers, Marc Glisse, Xavier Goaoc

Helly-Type Theorems for Approximate Covering

Julien Demouth* Olivier Devillers† Marc Glisse‡ Xavier Goaoc*

Abstract

Let $\mathcal{F} \cup \{U\}$ be a collection of convex sets in $\mathbb{R}^d$ such that $\mathcal{F}$ covers U. We show that if the elements of $\mathcal{F}$ and U have comparable size, in the sense that each contains a ball of radius r and is contained in a ball of radius R for some fixed r and R, then for any $\epsilon > 0$ there exists $\mathcal{H}_\epsilon \subset \mathcal{F}$, whose size $|\mathcal{H}_\epsilon|$ is polynomial in $1/\epsilon$ and independent of $|\mathcal{F}|$, that covers U except for a volume of at most ϵ. The size of the smallest such subset depends on the geometry of the elements of $\mathcal{F}$; specifically, we prove that it is $O(\frac{1}{\epsilon})$ when $\mathcal{F}$ consists of axis-parallel unit squares in the plane and $\widetilde{O}(\epsilon^{\frac{1-d}{2}})$ when $\mathcal{F}$ consists of unit balls in $\mathbb{R}^d$ (recall that $\widetilde{O}(n)$ means $O(n \log^\beta n)$ for some constant β), and that these bounds are, in the worst-case, tight up to the logarithmic factor .

We extend these results to surface-to-surface visibility in 3 dimensions: if a collection $\mathcal{F}$ of disjoint unit balls occludes visibility between two balls then a subset of $\mathcal{F}$ of size $\widetilde{O}(\epsilon^{-\frac{7}{2}})$ blocks visibility along all but a set of lines of measure ϵ.

Finally, for each of the above situations we give an algorithm that takes $\mathcal{F}$ and U as input and outputs in time $O(|\mathcal{F}| * |\mathcal{H}_\epsilon|)$ either a point in U not covered by $\mathcal{F}$ or a subset $\mathcal{H}_\epsilon$ covering U up to a measure ϵ, with $|\mathcal{H}_\epsilon|$ satisfying the previous bound.

*LORIA - INRIA Grand Est, Univ. Nancy 2, Projet VEGAS. {demouth,goaoc}@loria.fr.

†INRIA Sophia-Antipolis. olivier.devillers@sophia.inria.fr.

‡Gipsa-Lab, CNRS UMR 5216, Grenoble, France. 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 [8] 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 [7, 9] and constant factor approximations were obtained for only a few geometric versions [4] (see also [3]). 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 $\mathcal{H}_\epsilon$ denote a smallest ϵ-covering of U contained in $\mathcal{F}$. Recall that $\widetilde{O}(n)$ means $O(n\log^\beta n)$ for some β. Our main results are the following:

- 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 $O\left(\frac{1}{\epsilon}\right)$ when $\mathcal{F}$ consists of axis-parallel unit squares in the plane (Theorem 4) and $\widetilde{O}(\epsilon^{\frac{1-d}{2}})$ 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 (Section 6).

Our results imply that there do not exist 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 $|\mathcal{H}_\epsilon|$ when $\epsilon \to 0$ 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 $\mathbb{R}^3$ [10, 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 to control the error a priori.

Helly-type theorems. Helly's theorem asserts that n convex sets in $\mathbb{R}^d$ have non-empty intersection if any $d+1$ of them have non-empty 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 [5, 6, 14]. A collection $\mathcal{F}$ covers U if and only if the intersection of the complement of its elements and U is empty; thus, if $\mathcal{F}$ consists of complement 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 (the figure on the left 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.

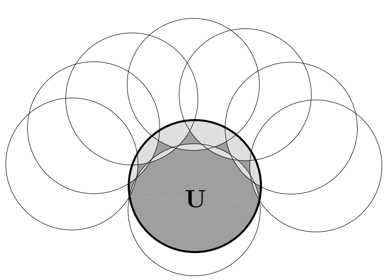

2 The general case

We start with a simple observation on approximation of a convex set by a grid:

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

Proof. To a cell σ whose interior meets $\partial\mathcal{O}$ we associate the line L_σ through diagonally opposite vertices with direction closest to the normal of some (arbitrary) support hyperplane H to $\mathcal{O}$ 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 $\mathcal{O}$ 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 $\mathcal{O}$. Thus, there are $O(2^d(R/\ell)^{d-1})$ cells in Γ whose interior intersect $\partial\mathcal{O}$, 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 $\mathcal{O}$ is a convex set of scale (r, R), such that $\ell \leq 2r$, containing the center of U, then $\mathcal{O} \cap U$ 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 $\mathcal{O}$. We consider the balls B_1 and B_2 of radius $\ell/4$ and $\ell/2$ centered in C.

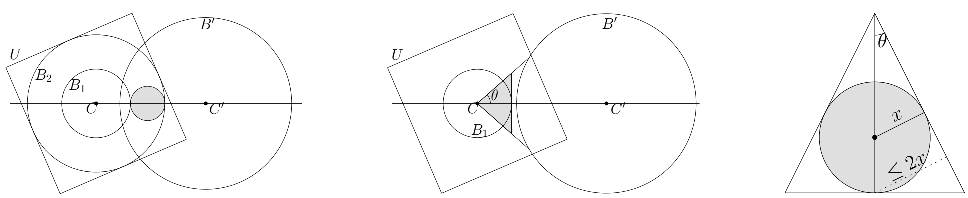

(a) (b) (c)

Figure 1: Finding a ball in $U \cap \mathcal{O}$.

If B' intersects B_1, we can find a ball of diameter $\ell/4$ centered on the line segment $[CC']$ contained in $U \cap \mathcal{O}$ (see Figure 1(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 Figure 1(b)), which in turn contains a ball of diameter $\frac{\ell}{4} \times \frac{r}{2R}$ (see Figure 1(c)). In both cases $U \cap \mathcal{O}$ contains a ball of radius $\left(\kappa\sqrt{d}\right)\ell$ 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 $H(\epsilon) = H_{d,r,R}(\epsilon)$ 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 $H(\epsilon)$.*

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 then 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 $\mathcal{C}_{i+1}$; 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 $\mathcal{C}_0 = \emptyset$. Lemma 2 implies that in the subdivision of any cube, at least one of the smaller cubes is covered, and thus

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

which resolves in:

$$|\mathcal{R}_i| \leq \left(\kappa^{-d} - 1\right)^i |\mathcal{R}_0| \quad \text{and} \quad |\mathcal{C}_i| \leq \sum_{k=0}^{i-1} |\mathcal{R}_k| \leq \frac{\left(\kappa^{-d} - 1\right)^i - 1}{\kappa^{-d} - 2} |\mathcal{R}_0|. \tag{1}$$

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 Equation (1) we get that $|C_i|$ is $O\left(\epsilon^{-O(d^2\kappa^{-d} \log \frac{1}{\kappa})}\right)$, which concludes the proof. □

This result is optimal in the sense that it becomes false if one of the scale or convexity conditions is dropped. 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 $\mathcal{F}$.

3 Covering by squares

For axis parallel boxes in $\mathbb{R}^d$, the analysis of the previous section holds for $\kappa = 1/2$; if, moreover, U is a cube, then $|\mathcal{R}_0|$ is 1 and this bound becomes $O\left(\epsilon^{-O(d2^d)}\right)$. 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 $O\left(\frac{1}{\epsilon}\right)$; this bound is tight in the worst-case.*

Proof. We first prove the lower bound. 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:

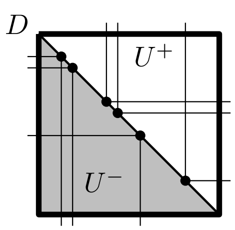

$$x_k - x_1 \geq \frac{1}{2} \quad \Rightarrow \quad \sum_{i=2}^{k} \alpha_i \geq \frac{1}{2}.$$

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:

$$\frac{1}{2} \leq \sum_{i=2}^{k} \alpha_i \leq \left(\sum_{i=2}^{k} \alpha_i^2\right)^{\frac{1}{2}} \left(\sum_{i=2}^{k} 1^2\right)^{\frac{1}{2}} \leq \sqrt{2\epsilon(k-1)} \quad \Rightarrow k \geq \frac{1}{8\epsilon}.$$

Note that this bound 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 $Y \in \mathcal{R}_i$ 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,

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

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 $U \subset \mathbb{R}^d$ 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 $\widetilde{O}\left(\epsilon^{\frac{1-d}{2}}\right)$. 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 (Section 4.1) before discussing the general case (Section 4.2). We then show that Theorem 5 extends to covering by other smooth objects if their curvature has bounded norm (Section 4.3).

4.1 The planar case

Upper bound. 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 projection[1] of the center of Y on the boundary of X (see Figure 2). We denote by $\mathcal{F}^Y$ the collection $\left\{X^Y \mid X \in \mathcal{F}\right\}$. We first start by a technical lemma:

Lemma 6. *Let Y be a disk of radius $r < 1$ and $\mathcal{F}$ a covering of a unit disk U by larger disks. Then, $U \cap Y$ can be covered by a triple $C(Y) \subset \mathcal{F}$ and a collection $R(Y)$, of at most $\frac{3}{r}$ 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 $U \cap Y$. We denote by $C(Y)$ the corresponding balls in $\mathcal{F}$. For any disk $X \in \mathcal{F}$, the area $\left(X^Y \cap Y\right) \setminus (X \cap Y)$ is inscribed in a rectangle (see Figure 2) with sides respectively smaller than $2r$ and $4r^2$. This rectangle can thus be covered by overlapping disks of radius $4r^2$ centered on its larger axis

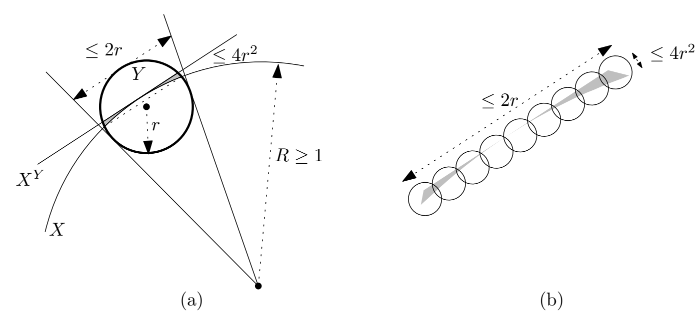

Figure 2: (a) $\left(X^Y \cap Y\right) \setminus (X \cap Y)$ is inscribed in a rectangle of sides at most $2r$ and $1 - \cos(\sin^{-1} r) \leq 4r^2$. (b) A covering with disks of radius $4r^2$ of $\left(X^Y \cap Y\right) \setminus (X \cap Y)$.

(Figure 2(b)). By choosing the disks so that the height covered at the intersection between two disks is, at least, $4r^2$, we need only $\frac{1}{r}$ disks. □

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

Proof of Theorem 5 for $d = 2$. We fix some $r < 1$ 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 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

$$r_i = 4r_{i-1}^2, \quad \text{with} \quad r_0 = r,$$

[1] If the two disks have the same center, we can choose any tangent to X.

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:

$$|\mathcal{R}_i| \leq \left(\prod_{k=0}^{i-1} 12(4r)^{-2^k}\right)|\mathcal{R}_0| \leq 12^i(4r)^{1-2^i}\mu r^{-2}, \quad \text{and} \quad \alpha_i \leq \pi r_i^2|\mathcal{R}_i| = 12^i\mu\pi(4r)^{2^i-1}.$$

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 $O\left(\epsilon^{-\frac{1}{2}}\log^6\frac{1}{\epsilon}\right)$. ☐

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 figure on the left) 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 $x_1, \ldots, x_k$ denote the abscissae of the tangency points of the disks in $\mathcal{G}^+$, sorted increasingly.

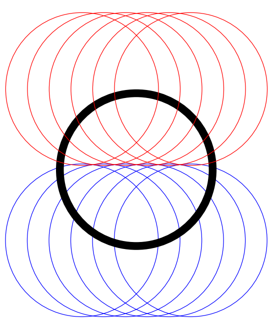

Let $\alpha_i = x_i - x_{i-1}$. For ϵ small enough, since $\mathcal{G}$ is an ϵ-cover we have:

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

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

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

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 if we let the disks intersect the x-axis on arbitrarily small lengths, the same can easily be achieved with a finite family.

4.2 Arbitrary dimension

Upper bound. We start with a generalization of Lemma 6.

Lemma 8. *Let Y be a ball of radius $\ell < 1$ 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 a $(d+1)$-tuple $C(Y) \subset \mathcal{F}$ and a collection $R(Y)$ 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. Notice that this is well defined whenever X and Y have distinct centers. We call $\mathcal{F}^Y$ 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 it 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 $\left(X^Y \cap Y\right) \setminus (X \cap Y)$ 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. $\square$

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.

Proof of Theorem 5. We fix some constant $r_0 \in (0, r)$ small enough so that for some constant $K > 0$, the function ρ introduced in Lemma 8 satisfies $\rho(t) \leq Kt^2$ for any $0 < t \leq r_0$. Call $C = Kr_0$. 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 $O(\epsilon^{\frac{1-d}{2}} \operatorname{polylog}(\frac{1}{\epsilon}))$. $\square$

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 region of volume v of $\mathbb{R}^d$. Let δ be the distance from p to its furthest point in U. The part of U at distance larger than $\delta/2$ from p has volume $\Omega(v)$.*

Proof. We refer to the figure below. We call $q \in U$ the furthest point from p (or one of them). Let H be the hyperplane that consists 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 $\mathcal{C}^-$ delimited by $\mathcal{C}$, H and H'. Since the volume of C^+ is equal to $4^d - 1$ time that of C^-, the statement follows. $\square$

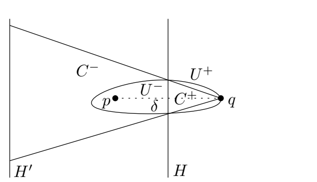

We can now prove that the bound of Theorem 5 is optimal up to the logarithmic factor.

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 an 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.

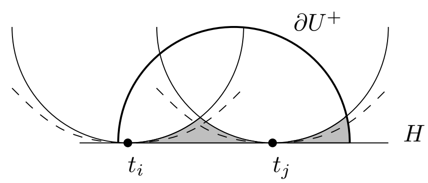

Let $\mathcal{G}^+ \subset \mathcal{F}^+$ be a family of k balls. For each ball $X \in \mathcal{G}^+$, let P_X denote the parabola with equation $2 * x_d = \sum_{i=1}^{d-1}(x_i - t_i)^2$ where $(t_1, ..., t_{d-1}, 0)$ is the tangency point of X with H. Since X is completely above the parabola P_X (see the figure 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 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 $\widetilde{O}\left(\epsilon^{\frac{1-d}{2}}\right)$.*

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 $\widetilde{O}\left(\epsilon^{-\frac{7}{2}}\right)$, such that the amount of visibility between A and B in $\mathcal{G}_\epsilon \cup \{A, B\}$ is $O(\epsilon)$.*

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 [11]. 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

$$dG = dxdy \sin\theta d\theta d\phi$$

and the statement follows from 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

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

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

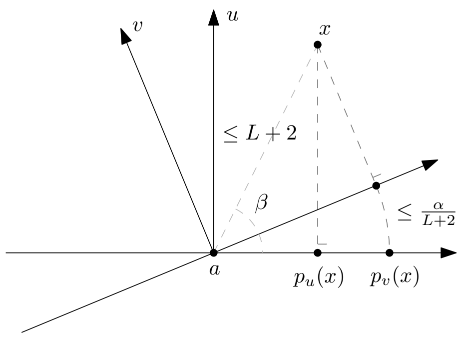

in any direction v close enough to u. Let $\alpha > 0$ be some constant and $v \in \mathbb{S}^2$ be a vector making, with u, an angle of at most $\frac{\alpha}{L+2}$ where L is the distance between a and b. For any ball $X \in \mathcal{F}_u$, with center x, we have (see Figure on the left)

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

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 $p_u(A) \cap p_u(B)$ but an area of at most $\epsilon_0 + 2\pi\alpha|\mathcal{H}_u|$.

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

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

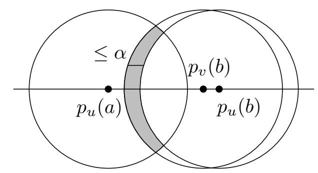

is bounded from above by 2α (see the figure on the right). Hence, $p_v(\mathcal{H}_u)$ covers $p_v(A) \cap p_v(B)$ but an area of at most:

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

Note that for a ball $X \in \mathcal{H}_u$, having a non-empty 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 $O(\alpha)$ and refer to Appendix B for the details.

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 $\arcsin\left(\frac{2}{L}\right)$ on $\mathbb{S}^2$. We can thus choose a collection D of $O\left(\alpha^{-2}\right)$ 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:

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

$\mathcal{H}$ has size $O\left(\alpha^{-2}h\right)$ 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:

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

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

$$|\mathcal{H}| = O\left(\epsilon_0^{-\frac{7}{2}} \operatorname{polylog} \frac{1}{\epsilon_0}\right) \quad \text{and} \quad V = O\left(\epsilon_0 + \alpha h\right) = O\left(\epsilon_0 \operatorname{polylog} \frac{1}{\epsilon\, 0}\right).$$

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 $O\left(\epsilon^{-\frac{7}{2}} \operatorname{polylog} \frac{1}{\epsilon}\right)$, such that the amount of visibility between A and B in $\mathcal{H} \cup \{A, B\}$ is $O(\epsilon)$. $\square$

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 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 $O\left(\frac{|\mathcal{F}|}{\epsilon}\right)$-time a point in U not covered by $\mathcal{F}$ or an ϵ-cover of U of size $O\left(\frac{1}{\epsilon}\right)$ contained in $\mathcal{F}$.*

Covering by balls. In the case of covering by balls, the main difficulty is to compute $C(Y)$, $R(Y)$ following immediately. Helly's theorem yields that given a collection $\mathcal{F}^Y$ of n halfspaces and a ball $Y \subset \mathbb{R}^d$, either there are $d+1$ halfspaces in $\mathcal{F}^Y$ that cover Y or there is a point in Y not covered by any half-space in $\mathcal{F}^Y$. In the case of covering of a ball Y by balls $\mathcal{F}$, finding $C(Y)$ reduces in $O(|\mathcal{F}|)$ time into solving the associated computational problem: finding such $d+1$ half-spaces or such a point.

Recall that *LP-type problems* are a special class of optimization problems [13]. Using a technique introduced by Amenta [1, 2], we can formulate the above problem as a LP-type problem. Specifically, let $\phi : 2^{\mathcal{F}^Y} \to \mathbb{R}$ be the map that associates to $\mathcal{G} \subset \mathcal{F}^Y$ the real

$$\phi(\mathcal{G}) = \min\left\{t \in [0, +\infty) \middle| \cup_{x\in\mathcal{G}}\, x \oplus D(t) \text{ covers } Y \cap U\right\}$$

where $\oplus$ and $D(t)$ denote respectively the Minkowski sum operator and the disk of radius t centered at the origin. The pair $(\mathcal{F}^Y, \phi)$ a LP-type problem of dimension $d+1$. This implies that $C(Y)$ can be computed in $O(|\mathcal{F}|)$ time using e.g. the algorithm of Seidel [12]. As a consequence, we obtain:

Corollary 15. *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 $\widetilde{O}\left(\epsilon^{\frac{1-d}{2}}\right)$ contained in $\mathcal{F}$ in time $\widetilde{O}\left(|\mathcal{F}|\epsilon^{\frac{1-d}{2}}\right)$.*

For the sake of completeness, the details are presented in Appendix A.

Visibility among unit balls. Corollary 15 makes the proof of Theorem 12 constructive and we get:

Corollary 16. *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 $\widetilde{O}\left(|\mathcal{F}|\epsilon^{-\frac{7}{2}}\right)$-time a visibility segment between A and B or a subset $\mathcal{G}_\epsilon \subset \mathcal{F}$, of size $\widetilde{O}\left(\epsilon^{-\frac{7}{2}}\right)$, 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. 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.

References

[1] N. Amenta. *Helly Theorems and Generalized Linear Programming*. Ph.D. thesis, U.C. Berkeley, 1993.

[2] N. Amenta. Helly-type theorems and generalized linear programming. *Discrete and Computational Geometry*, 12, 1994.

[3] H. Bronnimann and M.T. Goodrich. Almost optimal set covers in finite vc-dimension. *Discrete and Computational Geometry*, 14:463–479, 1995.

[4] K.L. Clarkson and K. Varadarajan. Improved approximation algorithms for geometric set cover. *Discrete and Computational Geometry*, 37:43–58, 2007.

[5] L. Danzer, B. Grünbaum, and V. Klee. Helly's theorem and its relatives. In V. Klee, editor, *Convexity*, Proc. of Symposia in Pure Math., pages 101–180. Amer. Math. Soc., 1963.

[6] J. Eckhoff. Helly, radon and caratheodory type theorems. In Jacob E. Goodman and Joseph O'Rourke, editors, *Handbook of Convex Geometry*, pages 389–448. North Holland, 1993.

[7] U. Feige. A threshold of lnn for approximating set cover. *Journal of the ACM*, 45(4):634–652, 1998.

[8] R. Karp. Reducibility among combinatorial problems. In *In Complexity of Computer Computations, Proc. Sympos. IBM Thomas J. Watson Res. Center*, pages 85–103, 1972.

[9] C. Lund and M. Yannakakis. On the hardness of approximating minimization problems. *Journal of the ACM*, 41(5):960–981, 1994.

[10] J. Pach and M. Sharir. Combinatorial Geometry with Algorithmic Applications – The Alcala Lectures. Alcala (Spain), August 31 - September 5, 2006.

[11] L. A. Santalo. *Integral Geometry and Geometric Probability*. Cambridge University Press, New York, NY, 2004. Second edition.

[12] R. Seidel. Small-dimensional linear programming and convex hulls made easy. *Discrete and Computational Geometry*, 6(5):423–434, 1991.

[13] M. Sharir and E. Welzl. A combinatorial bound for linear programming and related problems. In *STACS '92: Proceedings of the 9th Annual Symposium on Theoretical Aspects of Computer Science*, pages 569–579. Springer-Verlag, 1992.

[14] R. Wenger. Helly-type theorems and geometric transversals. In Jacob E. Goodman and Joseph O'Rourke, editors, *Handbook of Discrete & Computational Geometry*, chapter 4, pages 73–96. CRC Press LLC, Boca Raton, FL, 2nd edition, 2004.

A For the proof of Corollary 15

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

Monotonicity: if $F \subset G \subset \mathcal{H}$ then $\phi(F) \geq \phi(G)$.

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

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

A subset $B \subset F$, such that $\phi(B) = \phi(F)$, which is minimal for this property is a *basis* of F. The *combinatorial dimension* of a LP-type problem is the maximal cardinality of a basis, possibly $+\infty$. Recall that for LP-type problem with constant combinatorial dimension, a basis B of $\mathcal{H}$ can be computed in $O(|\mathcal{H}|)$-time (see [1, Chapter 7]).

We now prove the case $d = 2$ of Theorem 15. We start by a simple lemma:

Lemma 17. *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 $\phi : 2^{\mathcal{H}} \to \mathbb{R}$ be the map that associates to $\mathcal{G} \subset \mathcal{H}$ the real

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

where $\oplus$ and $D(t)$ denote respectively the Minkowski sum operator and the disk of radius t centered at the origin. The problem $(\mathcal{H}, \phi)$ is clearly a LP-type problem. Furthermore, Helly's theorem implies that its combinatorial dimension is bounded, more precisely by $d+1$, and 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 $x \oplus D(\phi(B))$ intersect in a point that is not covered by $\cup_{x \in \mathcal{H}} x$. □

From there, the proof of Theorem 15 is almost immediate:

Proof of Theorem 15. 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 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 17 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 time spent for computing $\mathcal{C}_i$ and $\mathcal{R}_i$ is 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 $O(\epsilon^{\frac{1-d}{2}} \operatorname{polylog}(\frac{1}{\epsilon}))$ and $|\mathcal{R}_{k-1}|$ is $O\left(\epsilon^{\frac{1-d}{2}}\right)$ so the time complexity of the algorithm is $O\left(|\mathcal{F}|\epsilon^{\frac{1-d}{2}} \operatorname{polylog}(\frac{1}{\epsilon})\right)$. □

B For the proof of Theorem 12

In this appendix we prove the following claim used in the proof of Theorem 12.

Claim. The area covered by $p_v(\mathcal{H}_u \setminus \mathcal{F}_v)$ is $O(\alpha)$.

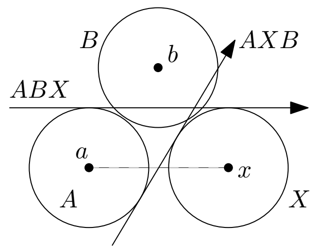

Figure 3: Two different geometric permutations.

Proof: 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 separated by a distance of at most $2\sqrt{2}$ (see Figure 3). 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

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

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. $\square$