



INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Helly-type theorems for approximate covering

Julien Demouth — Olivier Devillers — Marc Glisse — Xavier Goaoc

N° 6342

Octobre 2007

Thèmes COM et SYM



*R*apport
de recherche

Helly-type theorems for approximate covering

Julien Demouth , Olivier Devillers , Marc Glisse , Xavier Goaoc

Thèmes COM et SYM — Systèmes communicants et Systèmes symboliques
Projets Geometrica et Vegas

Rapport de recherche n° 6342 — Octobre 2007 — 12 pages

Abstract: Let \mathcal{F} be a covering of a unit ball $U \subset \mathbb{R}^d$ by unit balls. We prove that for any $\epsilon > 0$, the smallest subset of \mathcal{F} leaving at most a volume ϵ of U uncovered has size $O\left(\epsilon^{\frac{1-d}{2}} \text{polylog} \frac{1}{\epsilon}\right)$. We give an example showing that this bound is tight in the worst-case, up to a logarithmic factor, and deduce an algorithm to compute such a small subset of \mathcal{F} in time $O\left(|\mathcal{F}| \epsilon^{\frac{1-d}{2}} \text{polylog} \frac{1}{\epsilon}\right)$. We then extend these results in several directions, including covering boxes by boxes and visibility among disjoint unit balls in \mathbb{R}^3 .

Key-words: Helly, covering, visibility


Versions approchées de théorèmes de couverture de type Helly

Résumé : Étant donné \mathcal{F} une famille de boules unitaires couvrant une boule unité $U \subset \mathbb{R}^d$, nous démontrons que pour tout $\epsilon > 0$, le plus petit sous-ensemble de \mathcal{F} couvrant U à un volume ϵ près a une taille $O\left(\epsilon^{\frac{1-d}{2}} \text{polylog} \frac{1}{\epsilon}\right)$. Nous donnons également un exemple montrant que cette borne est optimale dans le cas le pire à un facteur logarithmique près, et nous proposons un algorithme calculant un tel ensemble en temps $O\left(|\mathcal{F}| \epsilon^{\frac{1-d}{2}} \text{polylog} \frac{1}{\epsilon}\right)$. Ces résultats se prolongent à d'autres problèmes, en particulier à la couverture de boîte par des boîtes et à la visibilité entre des boules unitaires disjointes de \mathbb{R}^3 .

Mots-clés : Helly, couverture, visibilité

1 Introduction

A family \mathcal{F} of sets covers a set U if the union of the elements of \mathcal{F} contains 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 . Many geometric problems, for example visibility problems, reduce to deciding if a given family \mathcal{F} covers an object U .



A family \mathcal{F} covers a set U if and only if the intersection of U with the complement of the elements of \mathcal{F} is empty. Thus, it follows from Helly's theorem [5] that if U is convex and the family \mathcal{F} consists of complements of convex sets, the smallest subset of \mathcal{F} covering U has size at most $d + 1$. The cases such that this size is bounded are in fact rather exceptional: for most classes of objects there exist arbitrarily large minimal covering families (the figure on the left shows an example of a minimal covering of a unit disk by 10 unit disks in the plane).

Yet, it seems clear that if a family \mathcal{F} of unit disks covers a unit disk U , then a *small* subset of \mathcal{F} is sufficient to cover *most* of U .

Let $\mathcal{G} \cup \{U\}$ be a collection of measurable sets in some measurable space, typically \mathbb{R}^d . Given a real $\epsilon > 0$, we say that \mathcal{G} is an ϵ -covering of U if at most a measure ϵ of U is not covered by \mathcal{G} . In this article, we study the size of the smallest ϵ -covering contained in a covering. We show that, in several cases, for any $\epsilon > 0$ this size can be bounded independently of the size of the covering. In other words, we prove Helly-type theorems for approximate covering.

Our results. Our main results are the following:

- We prove that the size of the smallest ϵ -covering contained in a covering of a unit ball in \mathbb{R}^d by n unit balls can be bounded independently of n (Theorem 1). We also give examples of covering in which our bound is tight up to a polylogarithmic factor (Lemma 6).
- We extend this result to two other covering problems: the covering of a unit box by unit boxes (Theorems 10 and 11) and visibility occlusion among disjoint unit balls in \mathbb{R}^3 (Theorem 13). In the particular case of the covering of a unit square by unit squares in the plane, we also give examples where our bound is tight. In all cases, our bounds are independent of the size of the covering family.
- For each of the above situations, we obtain algorithms (Theorems 7, 12 and 16) that, given a family of objects $\mathcal{F} \cup \{U\}$, decide in $O(|\mathcal{F}|f(\epsilon))$ -time that \mathcal{F} does not cover U or that \mathcal{F} is an ϵ -cover of U – with $f(\epsilon)$ denoting our bound on the size of the smallest ϵ -covering subset for that situation. More precisely, our algorithms return either a point in U not covered by \mathcal{F} or an ϵ -cover of U of size $f(\epsilon)$ contained in \mathcal{F} .

Related work. The classical problem of set cover asks, given a covering \mathcal{F} of a set U , for the smallest subset of \mathcal{F} that covers U . Since most versions are **NP**-hard, approximation algorithms have been largely investigated; namely, one looks for a subset of \mathcal{F} that covers U and whose size is within some multiplicative factor from the optimal. Approximation in this setting is also difficult as in many cases, no polynomial-time algorithm yielding constant factor approximations exists unless **P=NP** [6]. In the geometric setting, the situation is essentially the same [3, 4]. In this paper, we relax the covering problem in a different direction by allowing the covering to be imperfect. We are not aware of previous results in this direction.

2 Covering a unit ball by unit balls

In this section, we study the case where all objects are unit balls in \mathbb{R}^d . We prove:

Theorem 1. *Let U be a unit ball in \mathbb{R}^d and \mathcal{F} a covering of U by unit balls. For any $\epsilon > 0$, the smallest subset of \mathcal{F} , that is an ϵ -covering of U , has size $O\left(\epsilon^{\frac{1-d}{2}} \text{polylog} \frac{1}{\epsilon}\right)$.*

We also give examples for which this bound is tight up to the logarithmic factor (Theorem 6). For the clarity of the exposition, we prove the result in two dimensions (Section 2.1) before discussing the general case (Section 2.2). Our proof leads to an algorithm (discussed in Section 2.3) for computing, given \mathcal{F} , U and ϵ , either a point in U not covered by \mathcal{F} or a small ϵ -covering of U contained in \mathcal{F} .

2.1 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 projection¹ of the center of Y on the boundary of X (see Figure 1). We denote by \mathcal{F}^Y the collection $\{X^Y \mid X \in \mathcal{F}\}$. We first start by a technical lemma:

Lemma 2. *Let Y be a disk of radius $r < 1$ and \mathcal{F} a covering of a unit disk U by unit 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 $(X^Y \cap Y) \setminus (X \cap Y)$ is inscribed in a rectangle (see Figure 1) 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 1(b)).




Figure 1: (a) $(X^Y \cap Y) \setminus (X \cap Y)$ 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 $(X^Y \cap Y) \setminus (X \cap Y)$.

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. \square

We can now prove Theorem 1 for the case $d = 2$:

Proof of Theorem 1 for $d = 2$. Let r be a real in the open interval $(\frac{1}{100}, \frac{1}{10})$, let \mathcal{R}_1 be a covering of U by μr^{-2} disks of radius r , for some constant μ , and let \mathcal{C}_1 denote the empty set. We construct an ϵ -cover by iterating the following operation:

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

where $C(\cdot)$ and $R(\cdot)$ denote the sets defined in Lemma 2. By induction, we have that, 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

$$r_i = 4r^2_{i-1}, \quad \text{with} \quad r_1 = r,$$

and thus $r_i = \frac{1}{4}(4r)^{2^{i-1}}$. 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:

$$|\mathcal{R}_i| \leq \left(\prod_{k=1}^{i-1} 12(4r)^{-2^{k-1}} \right) |\mathcal{R}_1| \leq 12^{i-1} (4r)^{1-2^{i-1}} \mu r^{-2}, \quad \text{and} \quad \alpha_i \leq \pi r_i^2 |\mathcal{R}_i| = 12^{i-1} \mu \pi (4r)^{2^{i-1}-1}.$$

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

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

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


for $r \in (\frac{1}{100}, \frac{1}{10})$. Any sufficiently small $\epsilon > 0$ can be expressed as $\epsilon = 12^{k-1} \mu \pi (4r)^{2^{k-1}-1}$ for some integer k and real $r \in (\frac{1}{100}, \frac{1}{10})$. The previous inequalities then give that $\alpha_i \leq \epsilon$ and $|\mathcal{C}_k| = O\left(\epsilon^{-\frac{1}{2}} \log^6 \frac{1}{\epsilon}\right)$. \square

Lower bound. The following construction shows that the upper bound in Theorem 1 is optimal for $d = 2$ up to the logarithmic factor.

Lemma 3. *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 below) 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.

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:

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

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. \square

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.

2.2 Arbitrary dimension

Upper bound. We start with a generalization of Lemma 2.

Lemma 4. *Let Y be a ball of radius $r < 1$ and \mathcal{F} a covering of a unit ball $U \subset \mathbb{R}^d$ by unit balls. Then $U \cap Y$ can be covered by a $(d+1)$ -tuple $C(Y) \subset \mathcal{F}$ and a collection $R(Y)$ of $O(r^{1-d})$ balls of radius $\rho(r) = O(r^2)$.*

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 $r < 1$ 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 r and an orthogonal segment of length $O(r^2)$. This region can thus be covered by a collection $R_X(Y)$ of $O(r^{1-d})$ balls of radius $\rho(r) = O(r^2)$. 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 1 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 1. Let \mathcal{F} be a covering of a unit ball U by unit balls in \mathbb{R}^d . Let $0 < r_0 < 1$ be small enough so that there exists a constant $K > 0$ such that for any $0 < r \leq r_0$, $\rho(r) \leq K \cdot r^2$ for the function ρ introduced in Lemma 4. 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 $O(r_0^{-d})$ 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 5. *Let p be a point and R a convex region of volume v of \mathbb{R}^d . Let δ be the distance from p to its furthest point in R . The part of R at distance larger than $\delta/2$ from p has volume $\Omega(v)$.*

Proof. We refer to the figure below. We call $q \in R$ 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 R in a convex set U . We draw the half-cone C centered at q and whose intersection with H is U . The part of R at distance larger than $\delta/2$ from p contains the part R^+ of R that is on the same side of H as q . Furthermore, R^+ contains the part C^+ of C on the same side of H as q . The part R^- of R on the other side of H is contained in the region C^- delimited by C , H and H' . Since the volumes 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 1 is optimal up to the logarithmic factor.

Theorem 6. *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 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 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 \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 5, 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 $\Omega(\epsilon^{\frac{1-d}{2}})$. \square

2.3 Algorithm

To turn the proof of Theorem 1 into an algorithm, it suffices to describe how, given a ball Y , the sets $C(Y)$ and $R(Y)$ defined in Lemmas 2 and 4 can be computed. We consider d as a constant.

Recall that \mathcal{F} is a covering of a unit ball U by unit balls, and that Y is some ball with radius $r < 1$. After computing \mathcal{F}^Y , in $O(|\mathcal{F}|)$ time, the problem becomes to find a $(d+1)$ -tuple in \mathcal{F}^Y that covers $U \cap Y$. This can be done by turning it into a LP-type problem, a special class of optimization problems introduced by Sharir and Welzl [9].

Specifically, let $\phi : 2^{\mathcal{F}^Y} \rightarrow \mathbb{R}$ be the map that associates to $\mathcal{G} \subset \mathcal{F}^Y$ the real

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

where \oplus and $D(t)$ denote respectively the Minkowski sum operator and the disk of radius t centered at the origin. (\mathcal{F}^Y, ϕ) is a LP-type problem of dimension $d+1$, and thus there exists a subset $\mathcal{H} \subset \mathcal{F}^Y$ of size $d+1$ such that $\phi(\mathcal{H}) = \phi(\mathcal{F}^Y)$; thus, if \mathcal{F}^Y covers $U \cap Y$ so does \mathcal{H} . Moreover, \mathcal{H} can be computed in time $O(|\mathcal{F}^Y|)$ (using e.g. the algorithm of Seidel [8]) and if \mathcal{F}^Y does not cover $U \cap Y$, then the hyperplanes bounding the half-spaces

$$\{x \oplus D(\phi(\mathcal{H})) \mid x \in \mathcal{H}\}$$

intersect in a point not covered by \mathcal{F}^Y . As a result, we obtain:

Theorem 7. *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 $O\left(\epsilon^{\frac{1-d}{2}} \text{polylog}\left(\frac{1}{\epsilon}\right)\right)$ contained in \mathcal{F} in time $O\left(|\mathcal{F}| \epsilon^{\frac{1-d}{2}} \text{polylog}\left(\frac{1}{\epsilon}\right)\right)$.*

This (standard) formulation was obtained using Amenta's technique [1, 2] for relating Helly-type theorems and LP-type problems.

Before proving the theorem, 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 $\phi : 2^{\mathcal{H}} \rightarrow \Omega$ from the family of subsets of \mathcal{H} to some completely ordered set Ω . The pair (\mathcal{H}, ϕ) 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 7. We start by a simple lemma:

Lemma 8. *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 $\mathcal{G} \subset \mathcal{H}$ the real

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

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 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 $\cup_{x \in \mathcal{H}} x$. \square

From there, the proof of Theorem 7 is almost immediate:

Proof of Theorem 7. We construct the sets \mathcal{C}_i and \mathcal{R}_i by repeating, as indicated in the proof of Theorem 1, 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 8 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 1, $|\mathcal{C}_k|$ is $O(\epsilon^{\frac{1-d}{2}} \text{polylog}(\frac{1}{\epsilon}))$ and $|\mathcal{R}_{k-1}|$ is $O(\epsilon^{\frac{1-d}{2}})$ so the time complexity of the algorithm is $O(|\mathcal{F}| \epsilon^{\frac{1-d}{2}} \text{polylog}(\frac{1}{\epsilon}))$. □

3 Extensions

3.1 Covering by other shapes

3.1.1 Smooth convex sets

Lemma 2 requires that (i) given a disk 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 1 thus extends to:

Corollary 9. *Let U be a bounded convex set in \mathbb{R}^d 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 $O(\epsilon^{\frac{1-d}{2}} \text{polylog} \frac{1}{\epsilon})$.*

Note that the constant in the $O()$ depends on γ and d .

3.1.2 Boxes

While our proof technique requires the covering objects to be smooth, similar results hold in the non-smooth case. We illustrate this with the example of the covering of a unit box (i.e. a cube with axis-aligned edges) by unit boxes in \mathbb{R}^d .

Upper bounds. First, let us observe that a result similar to Theorem 1 holds in arbitrary dimension.

Theorem 10. *Let U be a unit box in \mathbb{R}^d and \mathcal{F} a covering of U by unit boxes. For any $\epsilon > 0$, the smallest subset of \mathcal{F} that is an ϵ -covering of U has size $O(\epsilon^{-O(d2^d)})$.*

Sketch of proof. Let \mathcal{F} be a family of unit boxes that covers a unit box U , in \mathbb{R}^d . Some box $X \in \mathcal{F}$ covers the center c of U and thus completely covers one of the 2^d half-sized boxes of \mathbb{R}^d defined by U and the axis-parallel hyperplanes through c . By applying this idea recursively, we get after k iterations a set of m elements of \mathcal{F} that cover all but possibly s boxes of volume $(2^{-k})^d$, with

$$m = \frac{(2^d - 1)^k - 1}{2^d - 2} \quad \text{and} \quad s = (2^d - 1)^k.$$

Thus, for

$$k \geq \log_{\frac{2^d}{2^d-1}} \frac{1}{\epsilon},$$

we obtain a family of size $O(\epsilon^{-O(d2^d)})$ that leaves uncovered a volume at most ϵ of U . □

In the planar case, a better bound can be obtained as follows:

Theorem 11. *Let U be a unit square in the plane and \mathcal{F} a finite covering of U by unit squares. For any $\epsilon > 0$, the smallest subset of \mathcal{F} that is an ϵ -covering of U has size $O\left(\frac{1}{\epsilon}\right)$.*

Proof. Let $Y \subset U$ be an axis-aligned rectangle. Since \mathcal{F} covers Y , there is a box X_1 that covers the center of Y . Since \mathcal{F} is finite, we can choose X_1 such that $Y \cap X_1$ is maximal for the inclusion. If X_1 does not have a corner inside Y , X_1 covers at least half of Y . Otherwise, let p be the corner of X_1 that is inside Y . Since \mathcal{F} is finite and completely covers Y , some other square X_2 covers p and together, X_1 and X_2 cover at least half of Y . Let $C(Y)$ denote these two squares (or just X_1 in the first case) and $R(Y)$ consist of the 2 rectangular pieces (at most) of Y not covered by $C(Y)$.

We define $\mathcal{R}_0 = \{U\}$ and $\mathcal{C}_0 = \emptyset$ and iterate:

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

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,

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

The statement follows. □

Lower bound. Applying the same technique as in Lemma 6 to the (infinite) family of squares tangent to one of the diagonals of U (see Figure 2) shows that the bound of Theorem 11 is tight. This construction can be generalized to dimension d using the same arguments as for balls, and yields a lower bound of $\Omega(\epsilon^{1-d})$ which does not match our upper bound for $d > 2$.




Figure 2: Left: For the proof of Theorem 11. Right: Lower bound of $\Omega\left(\frac{1}{\epsilon}\right)$ for the size of the smallest ϵ -cover of U in \mathcal{F}

Algorithms. The above proof, combined with the technique described in Section 2.3, leads to the following algorithm:

Theorem 12. *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} .*

3.2 Approximated 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.

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 1, we prove:




Figure 3: Two different geometric permutations.

Theorem 13. *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 $O\left(\epsilon^{-\frac{7}{2}} \text{polylog} \frac{1}{\epsilon}\right)$, 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 [7]. 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 14. *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 from integrating separately along the couples (x, y) and (θ, ϕ) . \square

Lemma 15. *The area covered by $p_v(\mathcal{H}_u \setminus \mathcal{F}_v)$ is $O(\alpha)$.*

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

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

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

We now prove Theorem 13:

Proof of Theorem 13. 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 9 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 \$\epsilon_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)

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

since \$\sin x \le x\$ for \$x \le 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) \le \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 \$2\alpha\$ (see the figure on the left). 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)\$; area is \$O(\alpha)\$ (see Lemma 15).

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 14 yields that the measure of \$V\$ is bounded from above by:

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

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

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

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

Algorithm. Combining the proofs of Theorem 7 and Theorem 13, we obtain:

Theorem 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 $O\left(|\mathcal{F}|\epsilon^{-\frac{7}{2}}\text{polylog}\left(\frac{1}{\epsilon}\right)\right)$ -time a visibility segment between A and B or a subset $\mathcal{G}_\epsilon \subset \mathcal{F}$, of size $O\left(\epsilon^{-\frac{7}{2}}\text{polylog}\frac{1}{\epsilon}\right)$, such that the amount of visibility between A and B in $\mathcal{G}_\epsilon \cup \{A, B\}$ is $O(\epsilon)$.*

4 Conclusion

In this paper, we considered the problem of, given a covering \mathcal{F} of a set U , bounding the size of the smallest ϵ -covering $\mathcal{G} \subset \mathcal{F}$ of U . We showed that this size can be bounded independently of $|\mathcal{F}|$ for coverings of a unit ball by unit balls, for coverings of a unit box by unit boxes and for 3-dimensional visibility among disjoint unit balls. We also presented simple algorithms that, given a family \mathcal{F} and a set U , certify either that \mathcal{F} does not cover U or that \mathcal{F} misses at most a volume ϵ of U .

Our bounds show that the size of the smallest ϵ -covering depends on the geometry of the covering sets: our (almost) tight bounds for disks and squares differ by $\Theta(\sqrt{\epsilon})$. Do other simple shapes lead to different bounds? Another obvious open problem is closing the gap between our lower and upper bounds for covering by boxes in dimension $d \geq 3$.

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] E. Helly. Über Mengen konvexer Körper mit gemeinschaftlichen Punkten. *Jahresbericht Deutsch. Math. Verein.*, 32:175–176, 1923.
- [6] C. Lund and M. Yannakakis. On the hardness of approximating minimization problems. *Journal of the ACM*, 41(5):960–981, 1994.
- [7] L. A. Santalo. *Integral Geometry and Geometric Probability*. Cambridge University Press, New York, NY, 2004. Second edition.
- [8] R. Seidel. Small-dimensional linear programming and convex hulls made easy. *Discrete and Computational Geometry*, 6(5):423–434, 1991.
- [9] 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.



Unité de recherche INRIA Sophia Antipolis
2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex (France)

Unité de recherche INRIA Futurs : Parc Club Orsay Université - ZAC des Vignes
4, rue Jacques Monod - 91893 ORSAY Cedex (France)

Unité de recherche INRIA Lorraine : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex (France)

Unité de recherche INRIA Rennes : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex (France)

Unité de recherche INRIA Rhône-Alpes : 655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier (France)

Unité de recherche INRIA Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex (France)

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)
<http://www.inria.fr>
ISSN 0249-6399