Cuttings for Disks and Axis-Aligned Rectangles in Three-Space*

Eynat Rafalin† Diane L. Souvaine‡ Csaba D. Tóth§

Abstract

We present new asymptotically tight bounds on *cuttings*, a fundamental data structure in computational geometry. For n objects in space and a parameter $r \in \mathbb{N}$, an $\frac{1}{r}$-*cutting* is a covering of the space with simplices such that the interior of each simplex intersects at most n/r objects. For n pairwise disjoint disks in $\mathbb{R}^3$ and a parameter $r \in \mathbb{N}$, we construct a $\frac{1}{r}$-cutting of size $O(r^2)$. For n axis-aligned rectangles in $\mathbb{R}^3$, we construct a $\frac{1}{r}$-cutting of size $O(r^{3/2})$.

As an application related to multi-point location in three-space, we present tight bounds on the cost of spanning trees across barriers. Given n points and a finite set of disjoint disk barriers in $\mathbb{R}^3$, the points can be connected with a straight line spanning tree such that every disk is stabbed by at most $O(\sqrt{n})$ edges of the tree. If the barriers are *axis-aligned rectangles*, then there is a straight line spanning tree such that every rectangle is stabbed by $O(n^{1/3})$ edges. Both bounds are best possible.

1 Introduction

Divide-and-conquer strategies are omnipresent in computer science. In problems involving multi-variable reals, one of the most successful methods over the last two decades has been the *partition technique* (in particular, *cuttings*) of computational geometry. Cuttings are indispensable for optimal data structures that support range searching, point location, motion planning, among others, and they are also used in currently best combinatorial bounds for hard discrete geometry problems [8, 11, 26]. For a set of n objects in $\mathbb{R}^d$ and a parameter $r \in \mathbb{N}$, a $\frac{1}{r}$-*cutting* is a finite collection of simplices (i.e., intersections of $d+1$ halfspaces) that cover $\mathbb{R}^d$ and such that the interior of each simplex intersects at most n/r objects. Even though the definition of cuttings allows overlapping simplices, all cuttings we construct consist of interior-disjoint simplices, which form *subdivisions* of $\mathbb{R}^d$.

The efficiency of applications using this partition technique crucially depend on the size of the $\frac{1}{r}$-cutting, that is, the number of simplices. Optimal $O(r^d)$ size $\frac{1}{r}$-cuttings are known for hyperplanes and $(d-1)$-dimensional simplices in $\mathbb{R}^d$ (note that the size is independent of n). In this paper, we present new tight bounds on the minimum size of cuttings for disjoint 2-dimensional objects in three-space. Our main result is an optimal size cutting for disjoint disks in three-space. For brevity,

*A preliminary version of this paper appeared in the *Proceedings of the 10th Workshop on Algorithms and Data Structures (Halifax, NS, 2007)*, vol. 4619 of LNCS, Springer, pp. 470–482. Work by Rafalin and Souvaine was partially supported by the National Science Foundation under grant CCF-0431027. E. Rafalin's research conducted while at Tufts University. Cs. D. Tóth's research conducted while working at the Massachusetts Institute of Technology.

†Google Inc., Mountain View, CA 94043, eynat@google.com.

‡Department of Computer Science, Tufts University, Medford, MA 02155, dls@cs.tufts.edu

§Department of Mathematics, University of Calgary, AB, T2N 1N4, Canada, cdtoth@ucalgary.ca.

we write *disk* for any planar set of constant description complexity (that is, a set in $\mathbb{R}^3$ that lies in a plane and is described by a constant number of algebraic inequalities of constant degree).

Theorem 1. *For every set of n pairwise disjoint disks in $\mathbb{R}^3$ and every r, $1 \le r \le n+1$, there is a $\frac{1}{r}$-cutting of size $O(r^2)$. This bound is best possible.*

A similar $O(r^2)$ bound was previously known only for disjoint *triangles* in $\mathbb{R}^3$ [29]. A subdivision of $\mathbb{R}^3$ into $O(r^2)$ cells, each bounded by a constant number of algebraic surfaces (rather than simplices bounded by hyperplanes), easily follows from known techniques. In some applications, such *pseudo-cuttings* (with curved boundaries) are satisfactory, others require partitions into *convex* simplices (e.g., Theorems 3 and 4 below). The challenge, we partially resolve in this paper, is to construct optimal size cuttings (with "straight" simplices) for curved objects. We also give a randomized algorithm that computes, for an input of n disjoint disks, an $\frac{1}{r}$-cutting of size $O(r^2)$ in polynomial expected time. It can be derandomized with standard techniques [23].

For disjoint *axis-aligned rectangles*, we can construct substantially smaller cuttings. This is the first sub-quadratic bound on $\frac{1}{r}$-cuttings for a family of 1- or 2-dimensional objects in $\mathbb{R}^3$.

Theorem 2. *For every set of n pairwise disjoint axis-aligned rectangles in $\mathbb{R}^3$ and every r, $1 \le r \le n+1$, there is an $\frac{1}{r}$-cutting of size $O(r^{3/2})$.*

An application: Spanning trees across barriers. A line segment e *stab*s a $(d-1)$-dimensional object b in $\mathbb{R}^d$ if the relative interiors of e and b intersect but the hyperplane spanned by b does not contain e. Chazelle and Welzl [12] showed that a set S of n points (*sites*) in $\mathbb{R}^d$, $d \ge 2$, can be connected by a straight line spanning tree such that at most $O(n^{1-1/d})$ edges of the tree stab every hyperplane. This bound is tight apart from the constant factor (e.g., for any spanning tree on n points in the $\lfloor n^{1/d} \rfloor \times \ldots \times \lfloor n^{1/d} \rfloor$ integer grid in $\mathbb{R}^d$, there is an axis-aligned hyperplane stabbed by $\Omega(n^{1-1/d})$ edges). In particular, for n points in the plane, there is a spanning tree that stabs every line $O(\sqrt{n})$ times. Interestingly, if we replace the lines by a set of disjoint line segments, there is a spanning tree that stabs every segment only a constant number of times. This constant is between 3 and 4 in the worst case by recent results of [3] and [17]. We address analogous problems involving disjoint barriers in $\mathbb{R}^3$, which have applications on multi-point location data structures in three-space [31].

Theorem 3. *Given a set S of n points and a finite set B of pairwise disjoint disk barriers in $\mathbb{R}^3$, there is a straight line spanning tree T on the vertex set S such that every barrier in B is stabbed by at most $O(\sqrt{n})$ edges of T. There are n points and $2\sqrt{n}$ disjoint circular disks in $\mathbb{R}^3$ such that for every spanning tree T on S, on average $\Omega(\sqrt{n})$ edges of T stab a disk.*

Theorem 4. *Given a set S of n points and a finite set B of pairwise disjoint axis-aligned rectangles in $\mathbb{R}^3$, there is a straight line spanning tree T on the vertex set S such that every rectangle in B is stabbed by at most $O(n^{1/3})$ edges of T. There are n points and $O(\sqrt{n})$ disjoint axis-aligned rectangles in $\mathbb{R}^3$ such that for every spanning tree T on S, on average $\Omega(n^{1/3})$ edges of T stab a rectangle.*

If we associate unit *cost* to each stabbing edge-barrier pair, then the total cost of such a spanning tree is bounded by $O(|B| \cdot \sqrt{n})$ for disks and by $O(|B| \cdot n^{1/3})$ for axis-aligned rectangles in the worst case, which are also best possible bounds. Note, however, that for computing a *minimum* cost spanning tree for given sets S and B, one can use a *min-max weight spanning tree* algorithm by Camerini [6]. Compute the number of barriers stabbed by each of the $\binom{n}{2}$ edges of a complete straight line graph on S, which gives an integer weight on each edge. Camerini's algorithm computes a spanning tree, for which our combinatorial bounds apply, in $O(n^2)$ time. A randomized $O(n)$ algorithm of Krager *et al.* [18] can compute the *minimum weight spanning tree* for a given configuration.

Related previous work on cuttings. Cuttings were introduced in the late eighties by Clarkson and Shor [13], and their ideas were later gradually improved and simplified [8, 9, 11, 24, 32]. For an arrangement of hyperplanes in $\mathbb{R}^d$, there is a $\frac{1}{r}$-cutting of size $O(r^d)$ [22]. For disjoint $(d-1)$-dimensional simplices in $\mathbb{R}^d$, there is an $\frac{1}{r}$-cutting of size $O(r^{d-1})$. Pellegrini [30] combined these results and showed that if the arrangement of n $(d-1)$-simplices in $\mathbb{R}^d$ has K vertices, there is an $\frac{1}{r}$-cutting of size $O(r^{d-1} + (K/n^d)r^d)$ for the $(d-1)$-simplies, which is best possible. Similar arguments show that there is a $\frac{1}{r}$-pseudo-cutting of size $O(r^d \text{ polylog } r)$ for $(d-1)$-dimensional semi-algebraic surfaces of constant maximal degree in $\mathbb{R}^d$ [10], however, pseudo-cuttings typically produce a covering with nonconvex regions. Two essentially different methods have been developed for constructing optimal cuttings: One is based on vertical decompositions and the other on sparse ε-nets. We briefly compare them with our strategy.

One method is a delicate construction by Chazelle and Friedman [11, 24] based on the following components: (i) the Clarkson-Shor random sampling technique [13]; (ii) a vertical decomposition algorithm; (iii) and a certain locality property of the decomposition that allows using tail estimates (for more details and examples, see [24]). The bottleneck for this technique is often the *vertical decomposition*, which is the partition of the space into vertical prisms (often called *trapezoids* in the plane) of bounded complexity such that each prism has at most two nonvertical sides. For example, Mulmuley [25] showed that r disjoint triangles in $\mathbb{R}^3$ have a vertical decomposition into $O(r^2)$ triangular prisms (see also [5]). De Berg, Guibas, and Halperin [4] extended this result and showed that n not necessarily disjoint triangles in $\mathbb{R}^3$ have a vertical decomposition of size $O(n^{2+\varepsilon} + K)$ for every $\varepsilon > 0$, where K is the complexity of the arrangement of the triangles. A straightforward application of this method, for instance, gives an $\frac{1}{r}$-cutting of size $O(r^2)$ for m disjoint triangles in $\mathbb{R}^3$.

Vertical decompositions exists for several types of objects in space, although asymptotically tight bounds are known only in very few cases. Chazelle *et al.* [10] showed that there is a vertical decomposition of size $O(r^2)$ for r disjoint disks, and this bound is sharp. An almost sharp upper bound on cuttings for hypersurfaces in four-space was recently proved by Koltun [19, 20]. The cells of such vertical decompositions, however, are bounded by semi-algebraic surfaces, and may not be *convex*. This explains why the technique of Chazelle and Friedman [11] cannot produce optimal size $\frac{1}{r}$-cuttings (with "straight" simplices) for disks in $\mathbb{R}^3$.

The other method known for constructing optimal cuttings for hyperplanes in $\mathbb{R}^d$ is due to Chazelle [7] (following earlier work by Agarwal [1]). It is a hierarchical decomposition based on so-called sparse ε-nets (a multi-purpose random sample related to ε-nets—a combinatorial tool introduced by Haussler and Welzl [16]). We adapt this method to disjoint disks in $\mathbb{R}^3$ by combining it with techniques from binary space partitions. Two recent results also followed an argument reminiscent of [7]: A construction of almost optimal pseudo-cuttings "sensitive" to a collection of algebraic curves (which means that a curve intersects few pseudo-simplices on average [21]); and an algorithm for counting the number of intersecting triples in a set of triangles in 3-space [15]. Prior to our work, Pellegrini [29] has combined binary space partition methods with sparse ε-nets when he constructed optimal cuttings for triangles in $\mathbb{R}^3$.

Binary space partitions. The *binary space partition* (for short, *BSP*) is a data structure produced by a simple hierarchical partition scheme, called *BSP algorithm*: Given a set B of disjoint $(d-1)$-dimensional objects in the interior of a convex cell σ, $\sigma \subset \mathbb{R}^d$, a BSP algorithm partitions σ along a hyperplane into two convex subcells σ_1 and σ_2 (while fragmenting the input objects as well), and recurses on the objects clipped in σ_1 and σ_2, independently, until the interior of each resulting cell is empty of input objects. The *size* of a BSP is the number of fragments of the input objects; intuitively, it measures the fragmentation caused by the partition. Paterson and Yao [27, 28] constructed a BSP of size $O(r^2)$ for r disjoint triangles and a BSP of size $O(r^{3/2})$ for r disjoint axis-aligned rectangles in $\mathbb{R}^3$; they also showed that these bounds are best possible.

Notice that our bounds on the size of $\frac{1}{r}$-cuttings match Paterson and Yao's bounds on the size of the BSPs of r disjoint objects. A BSP naturally constructs a subdivision of the input cell σ into convex subcells. A convex cell in $\mathbb{R}^3$ with k vertices can be partitioned into $O(k)$ simplices, so a BSP leads to a subdivision of σ into simplices whose number is proportional to the *complexity* of the subdivision (which is the total number of faces of dimensions 0, 1, and 2 over all cells). Note, however, that the size of a BSP may be much smaller than the complexity of the resulting subdivision (for example, a BSP of size $O(n^2)$ can produce a subdivision of the space into $O(n^2)$ convex cells of $\Theta(n^3)$ total complexity). For disjoint axis-aligned rectangles, we devise a deterministic BSP-like partition scheme for constructing a $\frac{1}{r}$-cutting of size $O(r^{3/2})$.

Organization. In Section 2, we briefly outline how we build optimal size cuttings for disjoint disks and recall a few basic facts about finite VC-dimensional range spaces and ε-nets. We construct cuttings for a "sparse" collection of disjoint disks in Section 3, and then for arbitrary disjoint disks in Section 4. We present our deterministic algorithm for constructing optimal cuttings for axis-aligned rectangles in Section 5. We prove our bounds on spanning trees in Section 6, and show that our upper bounds are asymptotically optimal in Section 7. We conclude with open problems in Section 8.

2 Preliminaries for Constructing Optimal Cuttings for Disks

Overview of our construction of cuttings. We construct an optimal size cutting for disjoint disks in $\mathbb{R}^3$ by combining several layers of hierarchical space decompositions. We introduce a binary relation between disks (*a avoids b*), and call a configuration *sparse* if the disks mutually avoid each other. For k mutually avoiding disks, we subdivide the space into $k+1$ convex polytopes, the interior of each of which is disjoint from the disks. For mutually avoiding disks, we use this algorithm recursively to build an $\frac{1}{r}$-cutting of size $O(r^{1+\delta})$ for every fixed $\delta > 0$, where the constant of proportionality depends on δ. Our main algorithm for arbitrary disjoint disks is reminiscent of a scheme originally introduced by Agarwal and Chazelle [1, 7] to construct optimal size cuttings for hyperplanes. We partition $\mathbb{R}^3$ into simplices recursively using planes spanned by so-called *sparse ε-nets* of a constant number of disks. The size of the resulting $\frac{1}{r}$-cutting is bounded by $O(r^2)$, which is shown by a charging scheme. The basis for our charging scheme is that the total number of "nonavoiding" pairs cannot increase when the problem is split into subproblems (even though a disk may be fragmented into many pieces, and the fragments of a single disk may occur in several subproblems).

Full and sparse configurations. We define a binary relation between disks in $\mathbb{R}^3$. Consider a set B of n disjoint disks in $\mathbb{R}^3$. Let π_b denote the plane containing disk $b \in B$, and let ∂b denote the (relative) boundary of b. For two disks $a, b \in B$, we say that a *avoids* b with respect to a convex cell σ if $\pi_a \cap \partial b \cap \mathrm{int}(\sigma) = \emptyset$ (Fig. 1). This relation is not symmetric: It is possible that a avoids b but b does not avoid a in a cell σ. Note, also, that even if a avoids b in σ, the plane π_a may intersect $b \cap \sigma$ (but then, π_a intersects ∂b outside of σ).

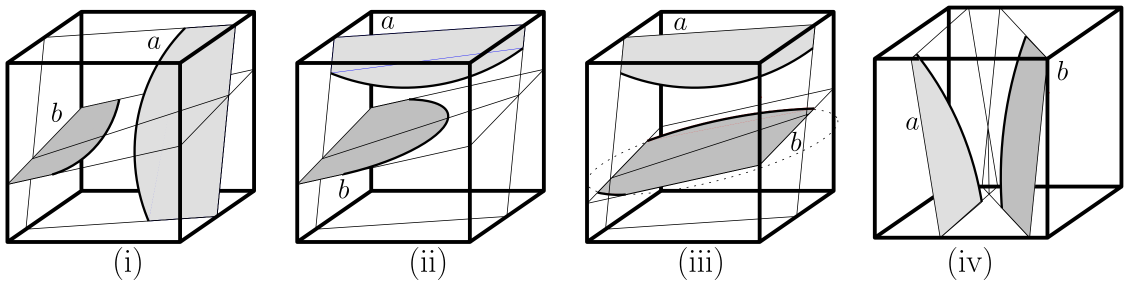

Figure 1: Disks a and b clipped within a cube σ. The (relative) boundaries ∂a and ∂b clipped within σ are bold. In (i) a and b do not avoid each other; in (ii) b avoids a, but a does not avoid b; in (iii) a and b avoid each other but $\pi_a \cap b \cap \sigma \neq \emptyset$; and in (iv) a and b avoid each other and $\pi_a \cap b \cap \sigma = \pi_b \cap a \cap \sigma = \emptyset$.

The *multiplicity* of a nonavoiding pair $(a, b) \in B^2$ with respect to σ is the number of intersection points of $\pi_a \cap \partial b$ that lie in $\mathrm{int}(\sigma)$. We measure the complexity of a set B of disjoint disks relative to the interior of a convex cell σ by $\tau(B, \sigma)$, which is the sum of multiplicities of all nonavoiding pairs with respect to σ.

For a convex cell σ, we define $B_\sigma = \{b \in B : b \cap \mathrm{int}(\sigma) \neq \emptyset\}$ to be the set of disks that intersect the interior of σ, and $\hat{B}_\sigma = \{b \cap \sigma : b \cap \sigma \neq \emptyset\}$ to be the portions of the disks of B_σ clipped within σ. Letting $n_\sigma = |B_\sigma|$, it is clear that $\tau(B, \sigma) = O(n_\sigma^2)$, since the multiplicity of every pair is bounded by a constant (depending on their description complexity). We use the following crucial property of the measure τ.

Proposition 1. *For any subdivision Ξ of a convex cell σ into convex subcells, we have*

$$\sum_{\xi \in \Xi} \tau(B, \xi) \leq \tau(B, \sigma).$$

Proof. Every intersection point $p \in \pi_a \cap \partial b \cap \mathrm{int}(\sigma)$ lies in the interior of at most one subcell $\xi \in \Xi$. Even if $p \in \mathrm{int}(\xi)$ for some $\xi \in \Xi$, it is counted in $\tau(B, \xi)$ only if disk a intersects $\mathrm{int}(\xi)$. $\square$

The technically most difficult part of the proof of Theorem 1 is the construction of a subquadratic $\frac{1}{r}$-cutting for a set of mutually avoiding disks. It uses sparse ε-nets and a randomized incremental subdivision scheme. In particular, it uses two randomized constructions: (1) it chooses a random sample of disks (which is a sparse ε-net) *and* (2) it subdivides a cell incrementally along sample disks in a random order. Our incremental subdivision scheme produces convex cells, but the benefits of ε-nets hold for simplices only: our construction maneuvers between convex and simplicial subdivisions.

ε-nets. We use a few basic facts about ε-nets and range spaces in Sections 3 and 4. (For an in-depth overview, refer to [9] or [24].) A *range space* is a set system (P, Q) on a ground set P and some subsets $Q \subset 2^P$. The *VC-dimension* of a range space is the size of the largest subset $S \subset P$ such that all $2^{|S|}$ subsets of S can be written as $S \cap q$ for some range $q \in Q$. In this paper, we consider range spaces (B, Q) where B is the set of geometric objects in $\mathbb{R}^d$ and every simplex σ defines a range $q_\sigma = \{b \in B : b \cap \text{int}(\sigma) \neq \emptyset\}$. It is not difficult to see that if B is a set of objects of bounded description complexity, then the corresponding range space (B, Q) has bounded VC-dimension [9, 24].

Haussler and Welzl [16] introduced the concept of *ε-net*s in range spaces. An ε-net of a finite range space (P, Q), for a constant $\varepsilon > 0$, is a set $S \subset P$ such that for every $q \in Q$ with $|P \cap q| \geq \varepsilon \cdot |P|$, we have $S \cap q \neq \emptyset$. For every range space of VC-dimension δ, a Bernoulli sample $S \subset P$, that contains every element $p \in P$ independently at random with probability $(\frac{cd}{\varepsilon} \log \frac{d}{\varepsilon})/n$ for an absolute constant $c > 0$, is an ε-net with constant probability [9].

Sparse ε-nets. Chazelle [7] noticed that a random sample drawn from a bounded VC-dimensional range space (P, Q), $|P| = n$, is not only an ε-net, but it may also preserve other properties of P with constant probability. He considered the range space of hyperplanes intersecting a simplex σ in $\mathbb{R}^d$. We state this result in broader terms. Assume that we are given a t-uniform hypergraph G with e hyperedges[1] on the vertex set P. The Bernoulli sample $S \subset P$ that contains every element $p \in P$ independently at random with probability $\varrho = (\frac{cd}{\varepsilon} \log \frac{d}{\varepsilon})/n$ is expected to span $e \cdot \varrho^t$ hyperedges of G. Chazelle showed that (i) S is an ε-net, (ii) it has size $\Theta(n\varrho)$, *and* (iii) it spans at most $2e \cdot \varrho^t$ hyperedges of G with constant probability. We use this observation for $t = 2$ (that is, when G is a graph defined on P) in Section 4.

3 Subdivisions for Mutually Avoiding Disks

In this section, we construct for every $\delta > 0$ an $\frac{1}{r}$-cutting of size $O(r^{1+\delta})$ for mutually avoiding disks in three steps. First, we present a simple algorithm that subdivides a convex cell into $k+1$ convex subcells along k mutually avoiding disks. Second, we extend this scheme to a randomized algorithm that subdivides a simplex into $k+1$ convex subcells along k disks drawn randomly from a set of n mutually avoiding disks, and partitions the n disks into $O(n \log k)$ pieces. Third, we apply this subdivision hierarchically to construct an $\frac{1}{r}$-cutting of size $O(r^{1+\delta})$ for mutually avoiding disks, where $\delta > 0$ is an arbitrary but fixed constant.

Lemma 1. *For k mutually avoiding disks w.r.t. a convex polytope σ in $\mathbb{R}^3$, there is a subdivision of σ into $k+1$ convex cell such that the interior of every cell is disjoint from the disks and every cell is bounded by planes spanned by the disks and sides of σ.*

We will extend incrementally every disk b_i, $i = 1, 2, \ldots, k$, to a planar polygon w_i (*wall*) in the plane π_i spanned by b_i so that the (relative) interior of w_i is disjoint from other disks and previous walls; and its boundary ∂w_i lies on the boundary of σ, other disks, or previous walls.

Let us illustrate first the analogous partition in the plane. The 2-dimensional analogue of Lemma 1 says that a convex polygon Δ containing k disjoint line segments can be subdivided into $k+1$ convex cells using the lines through the line segments. The subdivision in the plane is

[1]A t-uniform hypergraph, $t \in \mathbb{N}$, is a set system on a ground set (*vertex set*) where every set (*hyperedge*) has t elements. A 2-uniform hypergraph, for instance, is a simple graph.

quite straightforward: Extend every segment successively beyond its endpoints until the extension hits the boundary $\partial\Delta$, another segment, or a previously drawn extension. Since we eliminated every reflex vertex, we obtain a convex subdivision of Δ. It is also easy to see that the number of cells is $k+1$: Argue by induction on k. Assume that no segment has two endpoints on $\partial\Delta$ (otherwise, induction completes the proof). Assume, further, that g segments lie entirely in the interior of Δ, and $k-g$ segments each have one endpoint in the interior of Δ. The genus of the set $\Delta \setminus \{p \in \Delta : p \text{ is a point of a segment}\}$ is g; and the extension beyond each segment endpoint either partitions a cell into two cells or decreases the genus of a cell by one. Hence k extensions each split a cell into two cells, and we end up with $k+1$ convex cells. If we have mutually avoiding disks in 3-space, we can use a very similar argument using first Betti numbers instead of the genus of a manifold.

Proof of Lemma 1. We are given a set S of k mutually avoiding disks w.r.t. a convex polytope σ. We proceed by induction on k, the base case being $k=0$. Assume that $\sigma \setminus b$ is connected for every disk b (i.e., no disks partitions σ), otherwise induction on the two components completes the proof. Distinguish two kinds of disks: let $S_1 \subseteq S$ be the set of disks that intersect $\partial\sigma$, and let $S_2 \subseteq S$ be the set of disks that lie entirely in the interior of σ. Let $k_1 = |S_1|$ and $k_2 = |S_2|$, with $k = k_1 + k_2$. Recall that we denote by $\hat{b} = b \cap \mathrm{int}(\sigma)$ the portion of a disk $b \in S$ clipped in the interior of σ.

Consider the 3-manifold $M = \sigma \setminus (\bigcup_{b \in S_1} b)$. We build a subdivision of M by successively erecting planar *walls* in the interior of M. The insertion of a wall may or may not increase the number of components of M. Every wall must satisfy three conditions:

(i) the (relative) interior of every wall is disjoint from ∂M and from previous walls;

(ii) the (relative) boundary of every wall is covered by ∂M and previously erected walls;

(iii) every wall lies in the plane spanned by a disk of S.

Fix a linear order of the disks $b_1, b_2, \ldots, b_k$, in which the first k_1 elements are from S_1 in an arbitrary order, and the last k_2 elements are from S_2 in an arbitrary order. Proceed in k steps. In step i, insert walls in the plane π_i spanned by b_i as follows. Consider the planar arrangement A_i on the convex polygon $\Delta_i = \sigma \cap \pi_i$, formed by the boundary of Δ_i, and all line segments which are either intersections of Δ_i with previous walls or with $\hat{b}_j$, for $b_j \in S_1$ (refer to Fig. 3(i)). Since b_i avoids all other disks w.r.t. σ, π_i does not intersect any disk $b_j \in S_2$, and the endpoints of each segment $\pi_i \cap \hat{b}_j$, $b_j \in S_2$, lie on the boundary of the polygon Δ_i. Property (ii) ensures that the endpoints of the intersection of π_i and each previous wall lie on the boundary of M or on previous walls. The arrangement A_i is a convex subdivision of the polygon Δ_i. Disk $\hat{b}_i$ lies in some convex polygonal face ϕ_i of this subdivision. Let the new walls be the connected components of $\phi_i \setminus M$.

The walls along ϕ_i clearly satisfy properties (i), (ii), and (iii). Since every disk $\hat{b}_i$ was extended with planar walls whose boundaries lie on the boundary of M or on previous walls, the walls subdivide M (and σ) into convex cells. After the first k_1 steps, the extensions of the disks in S_1 already subdivide M into convex cells, and each disk of S_2 lies in the interior of one of those cells.

It remains to determine the total number of cells in the convex subdivision after all k steps. In step $i = 1, 2, \ldots, k_1$, we have $b \in S_1$. For a disk $b_i \in S_1$, let $u(b_i)$ denote the number of components of $\partial b_i \cap \mathrm{int}(\sigma)$, that is, the portion of the (relative) boundary of b_i clipped in the interior of σ. The first Betti number of a 3-manifold is the maximum number of *wall*s that can be erected without disconnecting it [34]. Hence, the first Betti number of M equals $\sum_{i=1}^{k_1}(u(b_i)-1) = (\sum_{i=1}^{k_1} u(b_i)) - k_1$

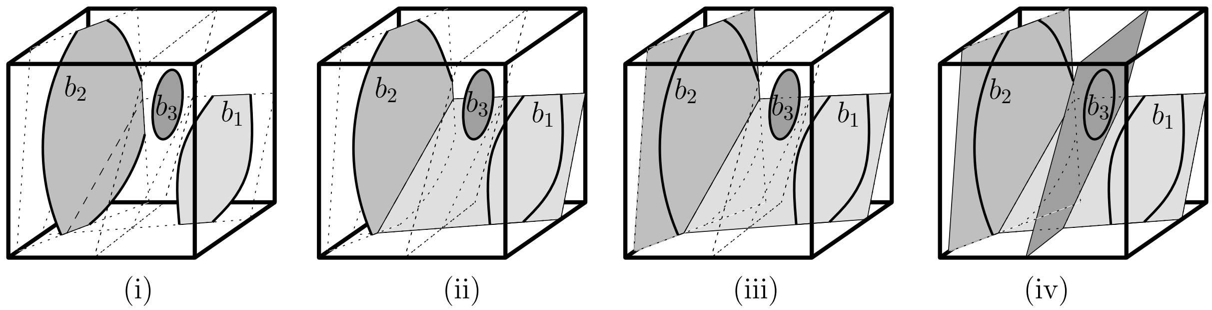

Figure 2: Three disks clipped within a cube σ. Since b_3 lies in the interior of σ, we have $M = \sigma \setminus (b_1 \cup b_2)$. Its first Betti number is 3. In three steps, we erect 6 walls in M and obtain a convex subdivision of M (and σ) into 4 convex cells.

because for each $b_i \in S_1$, independently, we can erect $u(b_i) - 1$ walls in the plane π_i without disconnecting M.

Let $\mathcal{M}(i)$ be the set of connected components (cells) of M after erecting walls in the first i steps (e.g., $\mathcal{M}(0) = \{M\}$), and let g_i denote the sum of the first Betti numbers of these cells. Each wall either partitions one connected component of $\mathcal{M}(i)$ into two components or decreases the Betti number of one component by one. After k_1 steps, each components is convex, and the sum of first Betti numbers is 0. In step $i = 1, 2, \ldots, k_1$, $u(b_i)$ walls are erected in plane π_i. The total increase of connected components together with the total decrease in first Betti numbers is $[|\mathcal{M}(k_1)| - 1] + [(\sum_{i=1}^{k_1} u(b_i)) - k_1] = \sum_{i=1}^{k_1} u(b_i)$, and so the number of cells in $\mathcal{M}(k_1)$ is $k_1 + 1$. Each disk of S_2 lies in the interior of a convex cell of $\mathcal{M}(k_1)$. In step $i = k_1 + 1, \ldots, k$, we have $b_i \in S_2$ and a single wall ϕ_i is erected, which partitions some convex cell into two convex cells. Hence, $k = k_1 + k_2$ steps partition σ into $k + 1$ convex cells. □

Subdivision for a subset of disks. Next we subdivide a simplex σ along the elements of a random sample $S \subset B$ of k disks. If we compute a subdivision described in Lemma 1 for cell σ and the k disks in S, then some of the other disks $b \in B \setminus S$ may be split into several *fragments*. The following lemma gives an upper bound on the number of fragments.

Lemma 2. *We are given a set B of n mutually avoiding disks w.r.t. a simplex σ and a random sample $S \subset B$ of size k. With probability at least $3/4$ there is a subdivision of σ into $k+1$ convex cells such that the interior of every cell is disjoint from all disks in S, every cell is bounded by planes spanned by disks of S and sides of σ, and the total number of fragments of disks of B is $O(n \log k)$.*

Proof. Let (B, R) be the range space over B where every planar quadrilateral domain q defines a range $\{b \in B : b \cap q \neq \emptyset\}$, containing all disks in B that intersect q. The range space (B, R) has finite VC-dimension. This implies that with probability at least 3/4, the sample S is an ε-net for (B, R) with some $\varepsilon = \Omega((\log k)/k)$. In what follows, we assume that S is an ε-net for (B, R).

Fix a linear order $b_1, b_2, \ldots, b_k$ of the sample disks of S, in which the disks intersecting the boundary of σ come first in a random permutation, followed by the disks lying in the interior of σ in an arbitrary order. Apply the subdivision algorithm described in Lemma 1. It subdivides σ along polygonal walls, each lying in a plane spanned by a disk of S, into $k + 1$ convex cells such

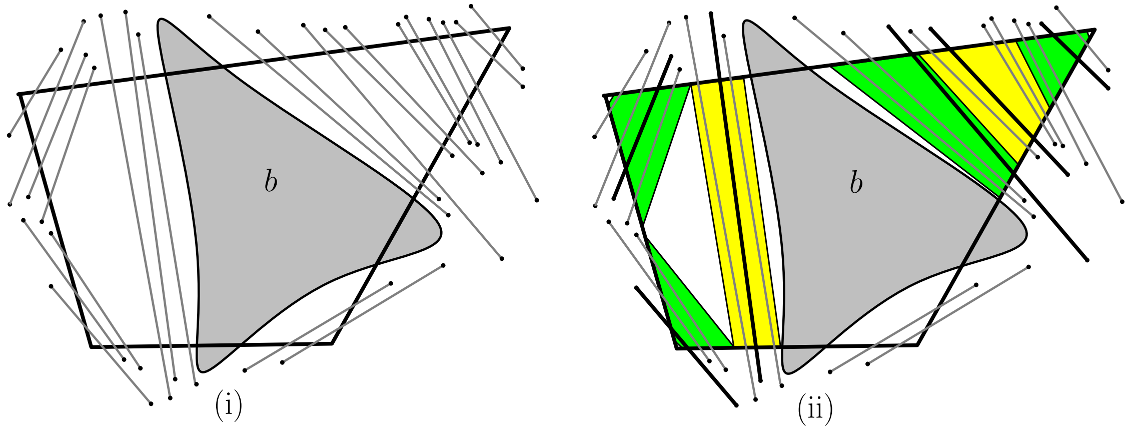

Figure 3: (i) The cross section of a disk b and a simplex σ in the plane π_b. If b avoids any other disk b', the endpoints of the segment $b' \cap \pi_b$ cannot be in the interior of σ. (ii) The segments clipped into $\sigma \cap \pi_b$ can be partitioned into groups, each covered by a quadrilateral. An $\frac{1}{3}$-net must contain a segment from each group that contains at least three segments.

that the interior of each cell is disjoint from the all disks in S. It remains to bound the expected number of fragments of disks of $B \setminus S$.

We show that the wall erected in plane π_i in step i of the incremental subdivision algorithm is expected to cut at most $O(n/i)$ disks. This implies that the expected number of fragments over k steps is $O(n \log k)$, and so there is a permutation where the number of fragments is $O(n \log k)$.

Consider the set $B_i \subset B$ of disks that intersect the plane π_i spanned by $b_i \in S$. Let $L_i = \{b \cap \pi_i : b \in B_i\}$ be the set of intersection segments of these disks with the polygon $\Delta_i = \sigma \cap \pi_i$. Since b_i avoids every disk in B_i, the disks of B_i must intersect the boundary of σ and every segment in L_i connects two points on the boundary of the polygon $\Delta_i = \sigma \cap \pi_i$ (see Fig. 3(i)). We introduce a partial order on the segments of L_i: We say that $\ell_1 \prec \ell_2$ if ℓ_1 separates b_i and ℓ_2 in Δ_i. Note that Δ_i is a triangle or a quadrilateral (since σ is a simplex), and each segment in L connects two sides of Δ_i. Select $O(1/\varepsilon) = O(k/\log k)$ quadrilaterals in π_i, each containing at least εn segments of L_i (Fig. 3(ii)). Since S is an ε-net for (B, R), every quadrilateral contains a segment $b_j \cap \Delta_i$, $b_j \in S$.

At step i, a disk $b \in B \setminus S$ is split into two fragments if the segment $b \cap \Delta_i$ lies in a quadrilateral that is not separated from b_i by any previous sample disk b_j, $j < i$. Each of the $O(1/\varepsilon) = O(k/\log k)$ quadrilaterals contains a sample disk b_j, $j < i$, with probability at least $(i-1)/k$. The expected number of quadrilaterals not separated from b_i is $O((1/\varepsilon)/i)) = O(k/(i \log k))$. Hence, the expected number of disks fragmented in step i is $O(1/(\varepsilon i)) \cdot O(\varepsilon n) = O(n/i)$. □

Cuttings for mutually avoiding disks in a simplex. We are now ready to prove the main result of this section.

Lemma 3. *For every $\delta > 0$, there is a constant $c(\delta)$ with the following property. Given a set B of n mutually avoiding disks w.r.t. a simplex σ_0 in $\mathbb{R}^3$, there is a subdivision of σ_0 into $c(\delta)n^{1+\delta}$ simplices such that the interior of every simplex is disjoint from any disk of B.*

Remark. Applying Lemma 2 with $k = n$, we would obtain a subdivision of σ into $n + 1$ convex cells. A triangulation of these cells would give a subdivision claimed in Lemma 3, only if the total complexity of the convex cells is $O(n^{1+\delta})$. However, the there are instances where our algorithm in Lemma 2 produces $n+1$ cells with $O(n^2)$ total complexity. In the proof of Lemma 3 below, we recursively invoke Lemma 2, each time with a constant parameter k, and partially triangulate some cells at each level of the recursion. We will distinguish heavy and light cells at each level of the recursion, and recurse on the heavy cells only. This allows controlling the total complexity of the resulting convex subdivision, and triangulation.

Proof of Lemma 3. We show that the following algorithm computes a required subdivision. The input is a simplex σ_0 and a set B of mutually avoiding disks w.r.t. σ_0. Let $k>0$ be a constant to be specified later. Put $i:=0$ and $\mathcal{C}_0 := \{\sigma_0\}$. For every $i\in\mathbb{N}$, $\mathcal{C}_i$ will be a subdivision of σ_0 into convex cells (not necessarily simplices). Repeat the following step until all cells in C_i are disjoint from any disk in B for some $i\in\mathbb{N}$.

Compute a subdivision $\mathcal{C}_{i+1}$, which is a refinement of $\mathcal{C}_i$. Cells of $\mathcal{C}_i$ that are already disjoint from the disks are not refined. Refine every convex cell $\sigma\in\mathcal{C}_i$ where $B_\sigma\neq\emptyset$ as follows. With probability at least 3/4, a random sample S_σ of size $c_1 k\log k$ drawn from B_σ is an $\frac{1}{k}$-net for the range space (B_σ, Q) defined in Section 2, where $c_1>0$ is an absolute constant (which depends only on the VC-dimension of the range space). Apply Lemma 2 for the disks B_σ and the *initial simplex* σ_0. With probability at least 3/4, we obtain a subdivision $\mathcal{D}_\sigma$ of σ_0 into at least $|S_\sigma| = c_1 k\log k+1$ convex cells such that the interior of every cell is disjoint of the sample disks, and the disks of $\hat{B}_\sigma$ are fragmented into at most $c_2 n_\sigma \log k$ pieces, with another absolute constant $c_2>0$.

We say that a cell $\sigma'\in\mathcal{D}_0$, $\sigma'\subset\sigma_0$, is *heavy* if it intersects more than $n_\sigma/(c_2\log k)^\alpha$ disks of $\hat{B}_\sigma$, where α satisfies $(c_2\log k)^\alpha\leq k$ (or $\alpha\leq\log k/\log(c_2\log k)$). At most $(c_2\log k)^{\alpha+1}$ cells of $\mathcal{D}_0$ are heavy, since there are at most $c_2 n_\sigma\log k$ fragments of $\hat{B}_\sigma$ over all cells of $\mathcal{D}_\sigma$. Triangulate every heavy cell $\sigma'\in\mathcal{D}_0$. Since S_σ is an $\frac{1}{k}$-net for B_σ, each simplex in the triangulation of σ' intersect at most $n_\sigma/k\leq n_\sigma/(c_2\log k)^\alpha$ disks of B_σ. Since every cell of $\mathcal{D}_\sigma$ is bounded by at most $c_1 k\log k+4$ faces (by Lemma 2), its triangulation consists of at most $4c_1 k\log k$ simplices. (Here we use Euler's polyhedron theorem pertaining to the surface of a polytope in 3-space.) We obtain a subdivision of σ_0 into at most $c_1 k\log k + (c_2\log k)^{\alpha+1}(4c_1 k\log k)\leq c_1(4c_2^{\alpha+1}+1)k\log^{\alpha+1}k = O(k\log^{\alpha+2}k)$ convex cells, each containing at most $n_\sigma/(c_2\log k)^\alpha$ fragments of $\hat{B}_\sigma$. By clipping these cells in σ, we have a subdivision $\mathcal{E}_\sigma=\{\sigma\cap\sigma' : \sigma\in\mathcal{D}_\sigma\}$ of σ into $O(k\log^{\alpha+2}k)$ convex cells. This completes the description of the refinement of a cell $\sigma\in\mathcal{C}_i$ with $B_\sigma\neq\emptyset$.

The recursion terminates with a subdivision $\mathcal{C}=\mathcal{C}_i$ of σ_0 into convex cells, each of which is empty of disks. Triangulate each convex cell of $\mathcal{C}$. Return the resulting subdivision $\mathcal{F}$ of σ_0. This completes the description of our algorithm.

Next, we bound the number of simplices in $\mathcal{F}$. We can represent the convex cells in all subdivisions $\mathcal{C}_i$ by a rooted tree: The root corresponds to σ_0; and if our algorithm subdivides a cell $\sigma\in\mathcal{C}_i$, then the children of σ correspond to the cells of $\mathcal{E}_\sigma$. Notice that for every $\sigma\in\mathcal{C}_i$, we have $n_\sigma\leq n/(c_2\log k)^{i\alpha}$. This implies that the depth of the recursion tree is at most $\log n/\log(c_2\log k)^\alpha$. Each convex cell in $\mathcal{C}$ is bounded by at most $c_1 k\log k\cdot\log n/\log(c_2\log k)^\alpha = O(\log n)$ planes. The triangulation of a cell of $\mathcal{C}$ creates at most $O(\log n)$ simplices, and so $\mathcal{F}$ consists of $O(|\mathcal{C}|\log n)$ simplices. It remains to bound the size of $\mathcal{C}$. Since we subdivide cells hierarchically until every cell is empty of disks of B, the size of $\mathcal{C}$ is proportional to the total number of fragments of B produced during the algorithm.

Consider a recursion step. A subdivision along disks of an $\frac{1}{k}$-net S_σ of a cell $\sigma\in\mathcal{C}_i$ increases the number of fragments by a factor of at most $c_2\log k$ (c.f., Lemma 2). The triangulation of heavy cells also increases the number of fragments: There are up to $(c_2\log k)^{\alpha+1}$ heavy cells, each heavy cell is

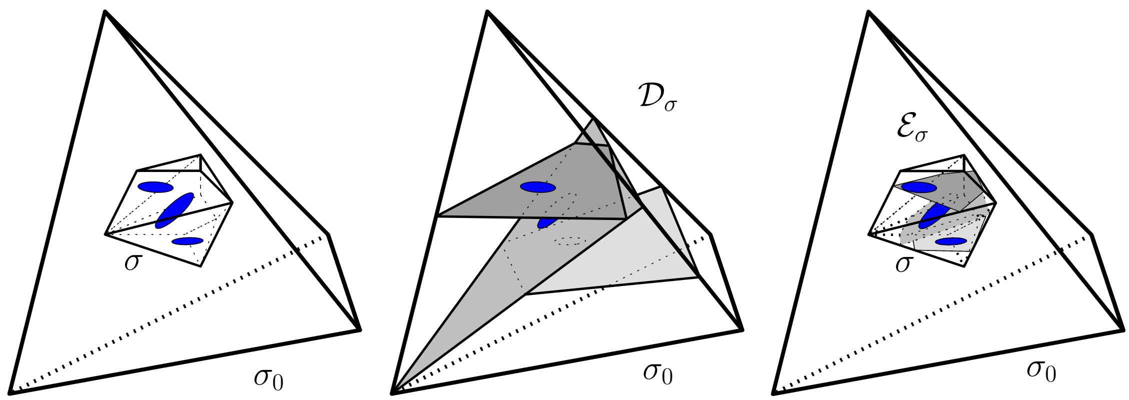

Figure 4: A convex cell σ in a simplex σ_0, containing three disks (left); a convex subdivision of σ_0 induced by the three disks (middle); this subdivision restricted to σ is a convex subdivision of σ.

triangulated into at most $4c_1 k \log k$ simplices, and each simplex contains at most n_σ/k fragments. This argument produces an upper bound of

$$(c_2 \log k)^{\alpha+1} \cdot 4c_1 k \log k \cdot \frac{n_\sigma}{k} = n_\sigma \cdot 4c_1 c_2^{\alpha+1} \log^{\alpha+2} k$$

on the number of fragments in the resulting simplices. That is, the refinement of $\mathcal{C}_i$ into $\mathcal{C}_{i+1}$ increases the number of fragments by a factor of at most $c_3 \log^{\alpha+2} k$, where $c_3 > 0$ is an absolute constant. Throughout the algorithm, the total number of fragments produced in $\log n / \log(c_2 \log k)^\alpha$ recursive steps is at most

$$n \cdot (c_3 \log^{\alpha+2} k)^{\frac{\log n}{\log(c_2 \log k)^\alpha}} = n \cdot n^{\frac{\log(c_3 \log^{\alpha+2} k)}{\log(c_1 \log k)^\alpha}} \leq n^{1+\frac{\log c_3 + (\alpha+2)\log\log k}{\log c_1 + \alpha \log\log k}} \leq n^{1+\delta},$$

if α is sufficiently large. We can set α, $\alpha \leq \log k / \log(c_2 \log k)$, to be arbitrarily large, if k is a sufficiently large integer. This completes the proof of Lemma 3. □

Corollary 1. *For every $\delta > 0$, there is a constant $c(\delta)$ with the following property. For every set B of n mutually avoiding disks and every $r \in \mathbb{N}$, there is a $\frac{1}{r}$-cutting of size $c(\delta) r^{1+\delta}$.*

Proof. Perform the recursive algorithm in the proof of Lemma 3 until the interior of every cell intersects at most n/r disks. The result follows by an analogous argument. □

4 Optimal Cuttings for Disjoint Disks

In this section, we prove Theorem 1 and construct an $\frac{1}{r}$-cutting of size $O(r^2)$ for a set B of n mutually disjoint disks in $\mathbb{R}^3$ and an $r \in \mathbb{N}$. We use a hierarchical partition scheme of Agarwal [1] and Chazelle [7], originally designed for hyperplane arrangements.

Let σ_0 be a bounding simplex of all input disks in $\mathbb{R}^3$. We construct a subdivision recursively, in stages. In stage $k \in \mathbb{N}$, we have a subdivision $\mathcal{D}_k$ of a bounding volume into simplices, initially $\mathcal{D}_0 = \{\sigma_0\}$. Recall that in Section 2, we defined $\tau(B, \sigma) = \sum_{a,b \in B} |\pi_a \cap \partial b \cap \mathrm{int}(\sigma)|$, the total number of points in the interior of σ that lie on a boundary of one disk and in the plane spanned by another disk. We have noted that it is at most quadratic in the number n_σ of disks that intersect σ, that is, $\tau(B,\sigma) = O(n_\sigma^2)$. For a small constant $c_0 > 0$, to be specified later we distinguish three types of cells:

- $\sigma \in \mathcal{D}_k$ is *dormant* if $n_\sigma < n/r_0^{k+1}$ (that is, it intersects less than n/r_0^{k+1} disks of B);
- $\sigma \in \mathcal{D}_k$ is *full* if $n_\sigma \geq n/r_0^{k+1}$ and $\tau(B,\sigma) \geq c_0 \cdot n_\sigma^2$;
- $\sigma \in \mathcal{D}_k$ is *sparse* if $n_\sigma \geq n/r_0^{k+1}$ and $\tau(B,\sigma) < c_0 \cdot n_\sigma^2$.

Our algorithm proceeds as follows. In stage k, we leave every dormant cell intact and refine the full and sparse cells of $\mathcal{D}_k$ to obtain the next subdivision $\mathcal{D}_{k+1}$. Consider every full or sparse simplex $\sigma \in \mathcal{D}_k$, and define the (undirected) graph G_σ on B_σ that has an edge between two disks $a, b \in B_\sigma$ if they do not avoid each other. By our remarks on sparse ε-nets in Section 2, if we pick every element of B_σ independently at random with probability $\varrho = cdr_0 \log(dr_0)/n_\sigma$, then with positive probability we obtain a $\frac{1}{r_0}$-net $S_\sigma \subset B_\sigma$ of size $\Theta(dr_0 \log(dr_0))$ such that $\tau(S_\sigma, \sigma) = \Theta(\varrho^2 \tau(B_\sigma, \sigma))$.

If σ is full, then compute the full arrangement of the planes spanned by S_σ within σ and triangulate the resulting convex cells. We obtain a subdivision of σ into $O(r_0^3 \log^3 r_0)$ simplices. If σ is sparse, we can apply Lemma 2 for B_σ and σ. Indeed, if $c_0 < 1/r_0^2$, then $\tau(S_\sigma, \sigma) = 0$ with constant probability. By Lemma 3, we obtain a subdivision of σ into $O((r_0 \log r_0)^{1+\delta})$ simplices, for some $\delta < 1$. Note that $O((r_0 \log r_0)^{1+\delta})$ is less than $r_0^2/2$ if r_0 is a sufficiently large constant. This is the refinement of σ into simplices in $\mathcal{D}_{k+1}$. Iterate until the resulting subdivision is a $\frac{1}{r_0}$-cutting. The analysis of this algorithm is based on a charging scheme, reminiscent of [7, 8, 9].

Lemma 4. *$\mathcal{D}_k$ is an $O(r_0^{2k})$-size $1/r_0^k$-cutting for B.*

Proof. Consider the subdivision $\mathcal{D}_k$, for $k \in \mathbb{N}$. It is easy to show (by induction on k) that every simplex $\sigma \in \mathcal{D}_k$ intersects at most n/r_0^k and at least n/r_0^{k+1} disks in B, and so $\mathcal{D}_k$ is an $1/r_0^k$-cutting for B. It remains to estimate the size of $\mathcal{D}_k$.

First, consider the full simplices of $\mathcal{D}_k$ and those dormant simplices of $\mathcal{D}_k$ produced by subdividing a full simplex of $\mathcal{D}_j$, for some $1 \leq j \leq k$. A full simplex $\sigma \in \mathcal{D}_j$ intersects at least n/r_0^{j+1} disks, and since it is full, we have $\tau(B,\sigma) \geq c_0(n/r_0^{j+1})^2$. Since we initially have $\tau(B,\sigma_0) = O(n^2)$, at stage j at most $n^2/(c_0(n/r_0^{j+1})^2) = r_0^{2(j+1)}/c_0$ full simplices of $\mathcal{D}_j$ are partitioned, each of them producing $O(r_0^3 \log^3 r_0)$ dormant simplices. In the first k stages, the total number of dormant simplices produced by full simplices is at most $\sum_{j=1}^{k} (r_0^{2(j+1)}/c_0) O(r_0^3 \log^3 r_0) = O(r_0^{2k+5} \log^3 r_0)$.

Next, consider the sparse simplices of $\mathcal{D}_k$ and those dormant simplices of $\mathcal{D}_k$ produced by subdividing some sparse ones. We charge every sparse simplex to its closest full ancestor or to σ_0 (even if σ_0 is sparse). A full simplex $\sigma \in \mathcal{D}_j$ (and σ_0) may produce $O(r_0^3 \log^3 r_0)$ offsprings, but every sparse simplex produces at most $r_0^2/2$ offsprings. Consequently, each sparse simplex $\sigma \in \mathcal{D}_j$ (resp., σ_0) leads to at most $O(r_0^3 \log^3 r_0) \cdot \sum_{i=0}^{k-j-1} (r_0^2/2)^i = O((r_0^{2(k-j)+1} \log^3 r_0)/2^{k-j})$ sparse simplices and dormant simplices produced by some sparse ones. In the first k stages, the total number of sparse simplices and their descendants is at most $\sum_{j=1}^{k} O(r_0^{2j+5} \log^3 r_0) \cdot O((r_0^{2(k-j)+1} \log^3 r_0)/2^{k-j}) = O(r_0^{2k+6} \log^6 r_0)$. Hence $\mathcal{D}_k$ is a $1/r_0^k$-cutting and it consists of $r_0^{2k} \cdot O(r_0^6 \log^6 r_0)$ simplices. Since r_0 and the constant of proportionality do not depend on k, we can consider $O(r_0^6 \log^6 r_0)$ as a constant. We conclude that $\mathcal{D}_k$ is a $\frac{1}{r}$-cutting of size $O(r^2)$ for $k = \lceil \log r / \log r_0 \rceil$. We provide a matching lower bound in Section 7. $\square$

5 Optimal Cuttings for Disjoint Axis-Aligned Rectangles

In this section, we prove Theorem 2 and present a deterministic algorithm that constructs a $\frac{1}{r}$-cutting of size $O(r^{3/2})$ for a finite set of disjoint axis-aligned rectangles in $\mathbb{R}^3$. An *axis-aligned ε-cutting* for n objects in $\mathbb{R}^d$ and a parameter $\varepsilon > 0$ is a covering of $\mathbb{R}^d$ with axis-aligned boxes such that the interior of every box intersects at most εn objects. Since every full-dimensional box in $\mathbb{R}^3$ can be partitioned into $3! = 6$ tetrahedra, an axis-aligned ε-cutting gives an ε-cutting of roughly the same size. The following theorem implies Theorem 2.

Theorem 5. *For every set of n disjoint axis-aligned rectangles in $\mathbb{R}^3$ and every $r \in \mathbb{N}$, there is an axis-aligned $\frac{1}{r}$-cutting of size $O(r^{3/2})$. This bound is best possible.*

Let σ_0 be the axis-aligned open bounding box of the set B of m input rectangles. We recursively partition σ_0 into axis-aligned boxes. For an open box $\sigma \subset \sigma_0$, let $B_\sigma = \{b \in B : \sigma \cap b \neq \emptyset\}$ denote the set of rectangles that intersect σ; and let $\hat{B}_\sigma = \{b \cap \sigma : b \cap \sigma \neq \emptyset\}$ be the portions of the disks of B_σ clipped within σ.

We distinguish a few types of special rectangles with respect to a box σ. A rectangle $b \in B$ is called *x-long* for σ if the x-extent of b contains the x-extent of σ; we can similarly define *y-long* and *z-long* rectangles. If a rectangle $b \in B$ is long for σ in two axis dimensions, then we say that b is *free* for σ. (This terminology follows that of [2, 14, 33], originally introduced for describing BSPs for axis-aligned fat rectangles.)

For an axis-aligned box $\sigma \subset \sigma_0$, let V_σ denote the set of vertices of B_σ lying in the interior of σ; let F_σ denote the set of vertices of $\hat{B}_\sigma$ lying in the relative interior of a side of the box σ; and let M_σ denote the set of midpoints of free rectangles of $\hat{B}_\sigma$. Finally, we denote by W_σ the weighted set $V_\sigma \cup F_\sigma$ where every point of V_σ has weight 8 and every point of F_σ has weight 1. For $q = x, y, z$, we say that the *q-median* of a weighted point set is the plane orthogonal to the q-axis such that the points lying on either side of this plane represent at most half of the total weight.

Initialize the algorithm with a trivial subdivision $\mathcal{D} = \{\sigma_0\}$ and apply the following partition step for every $\sigma \in D$ that intersects more than n/r input rectangles while such a cell σ exists.

1. If there is no free rectangle for σ then partition σ into 8 cells as follows.
   - Let $\mathcal{E}_\sigma = \{\sigma\}$. For $q = x, y, z$, do:
     - For every $\sigma' \in \mathcal{E}_\sigma$, partition σ' by the q-median of the weighted set $W_{\sigma'}$ into two subcells σ'_1 and σ'_2; and put $\mathcal{E}_\sigma := \mathcal{E}_\sigma - \{\sigma'\} + \{\sigma'_1, \sigma'_2\}$.
   - Put $\mathcal{D} := (\mathcal{D} \setminus \{\sigma\}) \cup \mathcal{E}_\sigma$.
2. If there is a free rectangle orthogonal to the q-axis, for some $q \in \{x, y, z\}$, then do:
   - Partition σ by two parallel planes into three subcells σ_1, σ_2, and σ_3; where the two splitting planes are the q-median of the weighted point set W_σ and the q-median of M_σ; and let $\mathcal{D} = (\mathcal{D} \setminus \{\sigma\}) \cup \{\sigma_1, \sigma_2, \sigma_3\}$.

Analysis. Since the rectangles are disjoint, all free rectangles for a box σ are parallel. Notice that if $b \in B$ is free for a box $\sigma \in \mathcal{D}$, then step 2 is applied, which cuts σ along a plane parallel to b, and so the fragment $b \cap \sigma$ is not partitioned anymore in the algorithm.

Consider step 1 to be three consecutive levels of binary cuts (it partitions a box into 8 subboxes), and step 2 as two consecutive levels of cuts (even though, it produces only 3 subboxes). First

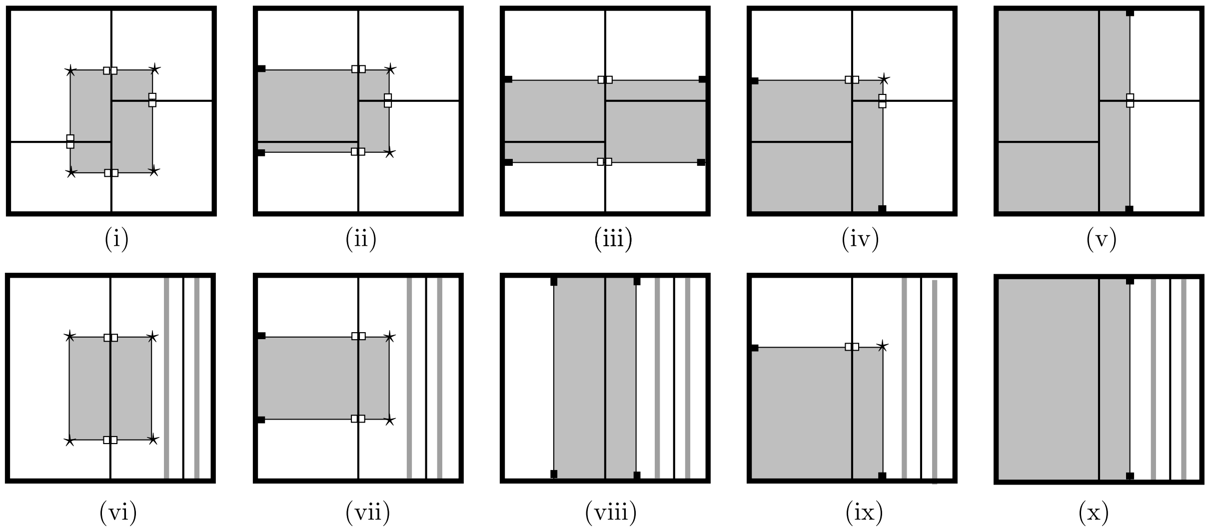

Figure 5: Different ways how a rectangle $b \cap \sigma$ is partitioned. First row: no free rectangles. Second row: two free rectangles. Vertices of V_σ are marked with stars, points of F_σ are marked with full squares, and vertices of $F_{\sigma'}$ for the resulting subboxes $\sigma' \subset \sigma$ are marked with empty squares.

consider step 1. The total weight of W_σ at most doubles, *i.e.* $\sum_{\sigma' \in \mathcal{E}_\sigma} |W_{\sigma'}| \leq 2|W_\sigma|$, because in every rectangle $b \cap \sigma \in \hat{B}_\sigma$, the weight associated to points in $b \cap \sigma$ at most doubles (see Fig. 5(i–v)). Observe that this bound is tight—the weight actually doubles in the examples in Fig. 5(iii) and (v). The number of vertices lying in a cell cannot increase (that is, $\sum_{\sigma' \in \mathcal{E}_\sigma} |V_{\sigma'}| \leq |V_\sigma|$). If we split a cell $\sigma \in \mathcal{D}$ along a plane and split an edge of a rectangle $b \in \hat{B}_\sigma$, then the resulting two fragments of b each have a new vertex in F_{σ_1} and F_{σ_2} along the splitting plane. In fact, all newly created vertices of $F_{\sigma'}$, $\sigma' \in \mathcal{E}_\sigma$, come in pairs, with one new vertex in each side of the latest cutting plane. Even though each binary cut along a q-median may increase the total weight of $\bigcup_{\sigma' \in \mathcal{E}_\sigma} W_{\sigma'}$, it remains a q-median of the increased weighted set. It follows that after all three levels of cuts in step 1, we have $|W_{\sigma'}| \leq |W_\sigma|/4$ for every $\sigma' \in \mathcal{E}_\sigma$. So one level of binary cuts in step 1 decreases the average weight $|W_\sigma|$ of a single cell by a factor of $4^{1/3} = 2^{2/3}$. Observe that this phenomenon holds for step 2, as well: the weight associated to a rectangle $b \cap \sigma \in B_\sigma$ increases by a factor of at most $\frac{11}{9}$ (see Fig. 5(vi–x); the ratio $\frac{11}{9}$ is attained in Fig. 5(vii)). So one level of cutting in step 2 decreases the average weight $|W_\sigma|$ of a cell by a factor of $\sqrt{11/9} < 2^{2/3}$.

If step 1 is applied to a box σ, then $M_\sigma = \emptyset$, and each rectangle in B_σ gives rise to at most two free rectangles in the eight resulting subcells. Hence $\sum_{\sigma' \in \mathcal{E}_\sigma} |M_{\sigma'}| \leq |W_\sigma|$. Let us consider a box σ that contains a free rectangle (that is, $|M_\sigma| \geq 1$, and so step 2 is applied). If a subcell σ_i, $i \in \{1, 2, 3\}$ contains one of the free rectangles of σ, then it cannot contain any *new* free rectangles. Due to the median cut for M_σ, the number of free rectangles in such subcells drops by a factor of at least 2. If a subcell σ_i, $i \in \{1, 2, 3\}$ is disjoint from all free rectangles of σ, then it may still have newly created free rectangles, (e.g., in Fig. 5(x)) and their number is at most $|W_\sigma|$.

Note that $|W_\sigma| + |M_\sigma|$ is an upper bound on the number $|B_\sigma|$ of rectangles intersecting the box σ. Initially, $|W_{\sigma_0}| + |M_{\sigma_0}| \leq 32n$. After $\frac{3}{2} \log_2 32r$ binary partition steps, every cell intersects at most n/r rectangles of B. When our recursive partition algorithm terminates, the number of cells in $\mathcal{D}$ is at most $2^{\frac{3}{2}\log 32r} = O(r^{3/2})$. This completes the proof of the upper bound in Theorem 2. We show in Section 7 that this bound is best possible.

6 From Cuttings to Spanning Trees

Chazelle and Welzl [12, 9] constructed spanning trees for n points in $\mathbb{R}^d$ that stab any hyperplane $O(n^{1-1/d})$ times. A key step in their argument relies on cuttings.

Lemma 5. **[12]** *For a set of n points and m hyperplanes in $\mathbb{R}^d$, there are two points, p and q, such that the line segment pq stabs at most $C\lceil m/n^{1/d}\rceil$ hyperplanes with some absolute constant C.*

Intuitively, Lemma 5 says that there are two points "close" to each other, in the sense that they are separated by few hyperplanes. One important distinction between hyperplanes and disjoint objects in $\mathbb{R}^d$ is that the number of hyperplanes separating two points in a hyperplane arrangement is a *pseudo-distance* (satisfies the triangle inequality). Unfortunately, the number of disjoint objects stabbed by a segment between two points is not a pseudo-distance. Nevertheless, analogous results hold for disjoint disks and axis-aligned rectangles which, by an iterative reweighing algorithm [12], immediately imply Theorem 3 and 4, respectively.

Lemma 6. *Given n points and m mutually disjoint 2-dimensional sets of bounded description complexity (barriers) in $\mathbb{R}^3$, there are two points, p and q, such that the straight line segment pq stabs $O(m/\sqrt{n})$ barriers.*

Lemma 7. *Given n points and m mutually disjoint axis-aligned rectangles in $\mathbb{R}^3$, there are two points, p and q, such that the straight line segment pq stabs $O(m/n^{2/3})$ rectangles.*

These bounds follow immediately from Theorems 1 and 2, respectively, by setting the upper bound on the number of cells in the cuttings equal to $n-1$. The above mentioned iterative reweighing algorithm of Chazelle and Welzl [12] carries over *verbatim* to our problem if we use Lemma 6 and 7 instead of Lemma 5. We briefly review the algorithm. We are given a set S of n points and m pairwise disjoint barriers in $\mathbb{R}^3$. Assume that there is a pair of points in S separated by $C\lceil m/n^{\alpha}\rceil$ barriers, for some exponent α, $0 \le \alpha \le 1$. The reweighing algorithm assigns integer weights to the m input barriers: a barrier of weight $k \in \mathbb{N}$ is modeled as k disjoint barriers in general position infinitesimally close to the original barrier. Straight line edges between points are added one by one, in $n-1$ steps. If the total weight of all barriers is m_i in step i, then the algorithm inserts an edge pq that stabs at most $C\lceil m_i/n^{\alpha}\rceil$ barriers, deletes point p from S, and doubles the weight of all hyperplanes stabbed by pq, thus raising the total weight to $m_{i+1} \le m_i + C\lceil m_i/n^{\alpha}\rceil$. After $n-1$ steps, the algorithm has inserted $n-1$ edges that form a spanning tree over S, and a simple calculation shows that every barrier is stabbed by at most $O(n^{1-\alpha})$ edges.

7 Lower Bound Constructions

In this section we present constructions that establish the lower bounds in Theorems 3 and 4. These are reminiscent to constructions that Paterson and Yao [27, 28] gave for n disjoint triangles (resp., axis-aligned rectangles) whose BSP complexity is $\Omega(n^2)$ (resp., $\Omega(n^{3/2})$); these constructions also match our upper bounds on the size of ε-cuttings for disjoint triangles, disks, and axis-aligned rectangles, respectively.

Lemma 8. *For every $n \in \mathbb{N}$, there is a set S of n points and a set of disjoint triangles (or disks) in $\mathbb{R}^3$ such that the edges of any straight line spanning tree over S stab some triangle (or disk) $\Omega(\sqrt{n})$ times.*

Lemma 9. *For every $n \in \mathbb{N}$, and $1 \le r \le n+1$, there is a set of n disjoint triangles (or disks) in $\mathbb{R}^3$ such that in any covering of $\mathbb{R}^3$ with t simplices there is a simplex whose interior intersects at least $n/\sqrt{t}$ triangles (or disks).*

Proof of Lemmas 8 and 9. Assume that $n = k^2$ for some integer $k \in \mathbb{N}$. We arrange n points on the hyperbolic paraboloid $z = xy$. The barriers are vertical trapezoids (each trapezoid can be represented by two triangles); $k-1$ trapezoids are parallel to the xz-plane and lie above the surface $z = xy$; another $k-1$ trapezoids are parallel to the yz-plane and lie below the surface $z = xy$. $T = T_a \cup T_b$, where

$$T_a = \{t_{a,i} = ([0, 2k^2] \times [2ki - k] \times [-2k^4, 2k^4]) \cap (xy < z) : i = 1, 2, \ldots, k-1\},$$

and

$$T_b = \{t_{b,i} = ([2ki - k] \times [0, 2k^2] \times [-2k^4, 2k^4]) \cap (z < xy) : i = 1, 2, \ldots, k-1\}.$$

Define a point set on the manifold $z = xy$ as

$$S = \{p_{i,j} = (2ki + j, 2kj - i, (2ki + j)(2k - i)) : i = 1, 2, \ldots, k;\ j = 1, 2, \ldots, k\} \subset \mathbb{R}^3.$$

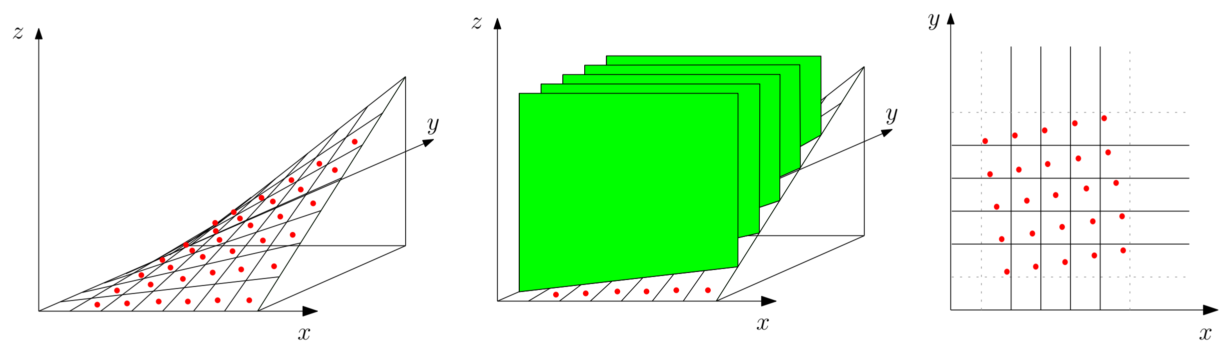

Figure 6: The 25 points of S on the surface $z = xy$ for $k = 5$ (left); trapezoids lying above the surface $z = xy$ (middle); the 25 points seen from above (right).

To prove Lemma 8, we show that every edge of a straight line spanning tree over S stabs at least one barrier. Note that every line stabs $z = xy$ in at most two points, and so the relative interior of every edge lies completely below or above $z = xy$. First consider an edge $e_1 = p_{i,j}p_{i',j'}$ with $i \neq i'$ and $j \neq j'$. If e_1 lies above (below) $z = xy$ then it stabs the trapezoid $t_{a,i}$ ($t_{b,j}$). Next consider an edge $e_2 = p_{i,j}p_{i,j'}$ with $j < j'$ ($j > j'$). Since e_2 lies above (below) $z = xy$, it stabs the trapezoid $t_{a,i}$ ($t_{b,j}$). The total number of stabbing edge-triangle pairs is at least $n - 1 = k^2 - 1$, and so one of $2(k-1)$ barriers is stabbed by at least $(k+1)/2 = \Omega(\sqrt{n})$ edges.

To prove Lemma 9, consider a cover of $\mathbb{R}^3$ with t simplices. There is a simplex σ whose closure contains at least k^2/t points of S. It is enough to find two points $p, p' \in \sigma \cap S$ such that the line segment pp' stabs $\Omega(n/\sqrt{t})$ barriers; since σ is convex, every barrier stabbed by pp' intersects the interior of σ. By the previous argument, each row $\{p_{i,j} : i = 1, 2, \ldots, k\}$ and each column

$\{p_{i,j} : j = 1, 2, \ldots, k\}$ contains at most $k/(4\sqrt{t})$ points of $\sigma \cap S$, otherwise the segment between the two extremal points in a row or column stabs $k/(4\sqrt{t}) = \Omega(n/\sqrt{t})$ barriers. Therefore, $\sigma \cap S$ must contain two points, $p_{i,j}$ and $p_{i',j'}$, such that $|i - i'| > k/(4\sqrt{t})$ and $|j - j'| > k/(4\sqrt{t})$. If the segment $p_{i,j}p_{i',j'}$ has positive slope, then it lies below $z = xy$ and it stabs at least $k/(4\sqrt{t}) = \Omega(n/\sqrt{t})$ barriers in T_b. Otherwise it has negative slope, it lies above $z = xy$, and it stabs at least $k/(4\sqrt{t}) = \Omega(n/\sqrt{t})$ barriers in T_a. □

Lemma 10. *For every $n \in \mathbb{N}$, there is a set S of n points and a set of disjoint axis-aligned rectangles in three-space such that the edges of any straight line spanning tree over S stab some rectangle $\Omega(n^{1/3})$ times.*

Lemma 11. *For every $n \in \mathbb{N}$, and $1 \leq r \leq n+1$, there is a set of n disjoint axis-aligned rectangles in $\mathbb{R}^3$ such that, in any covering of $\mathbb{R}^3$ with t axis-aligned boxes, there is a box whose interior intersects at least $\Omega(n/t^{2/3})$ rectangles.*

Proof of Lemmas 10 and 11. Assume that $n = k^3$ for some integer $k \in \mathbb{N}$. We arrange a set $S = \{(2i_1, 2i_2, 2i_3) : 1 \leq i_1, i_2, i_3 \leq k\}$ of $n = k^3$ points in the lattice $2\mathbb{Z} \times 2\mathbb{Z} \times 2\mathbb{Z}$. Three families of axis-aligned rectangles, each of size $k(k-1)$ partition the cube $[0, k]^3$ into k^3 unit cubes. Our barriers are (relatively) open rectangles, each one is represented as a cross product of two open intervals and a point.

$$R_x = \{[2i+1] \times (2j-1, 2j+1) \times (1, 2k+1) : i = 1, 2, \ldots, k-1; j = 1, 2, \ldots, k\}.$$

$$R_y = \{(1, 2k+1) \times [2i+1] \times (2j-1, 2j+1) : i = 1, 2, \ldots, k-1; j = 1, 2, \ldots, k\}.$$

$$R_x = \{(2j-1, 2j+1) \times (1, 2k+1) \times [2i+1] : i = 1, 2, \ldots, k-1; j = 1, 2, \ldots, k\}.$$

Every edge of a spanning tree over S stabs at least one barrier. The total number of stabbing edge-barrier pairs is at least $n - 1 = k^3 - 1$, and so one of the $3(k-1)^2$ barriers is cuts at least $(k^3 - 1)/2k(k-1) = \Omega(n^{1/3})$ edges. This proves Lemma 10.

For Lemma 11, consider a cover of $\mathbb{R}^3$ with t axis-aligned boxes. There is a box σ whose closure contains at least k^3/t points of S. If the orthogonal projection of σ in direction $q = x, y, z$ contains h_q points of the planar lattice $2\mathbb{Z} \times 2\mathbb{Z}$, then $k^3/t = \sqrt{h_x h_y h_z}$. There is a direction, say $q = x$, in which the orthogonal projection of σ contains at least $(k^3/t)^{2/3} = k^2/t^{2/3}$ points of $2\mathbb{Z} \times 2\mathbb{Z}$. This implies that the interior of σ intersects at least $\Omega(k^2/t^{2/3}) = \Omega(n/t^{2/3})$ rectangles of R_x. □

8 Conclusions and Open Problems

We have shown that for any finite set of pairwise disjoint planar objects of constant description complexity (for short, disks) in 3-space and a parameter r, $1 \leq r \leq n+1$, one can construct an $\frac{1}{r}$-cutting of size $O(r^2)$. That is, there is a partition of $\mathbb{R}^3$ into $O(r^2)$ tetrahedra such that the interior of each intersects at most a $\frac{1}{r}$ fraction of the objects. Previous results on nonpolygonal objects established only a pseudo-cutting, a partition of $\mathbb{R}^3$ into (typically nonconvex) regions bounded by $O(1)$ semi-algebraic surfaces. Our $O(r^2)$ bound on the size of the cutting is asymptotically tight. We also construct a $\frac{1}{r}$-cutting of optimal size $O(r^{3/2})$ for disjoint axis-aligned rectangles.

In our proof, we assumed that the disks and rectangles are pairwise disjoint. For possibly intersecting planar objects in $\mathbb{R}^3$, the smallest $\frac{1}{r}$-cutting may have $\Theta(r^3)$ size, the size of a $\frac{1}{r}$-cutting for a plane arrangement in $\mathbb{R}^3$. However, if just a few pairs of disks intersect, our results are likely to extend in the sense of Pellegrini [30]: We conjecture that for n disks in $\mathbb{R}^3$ and a parameter r, there is an $\frac{1}{r}$-cutting of size $O(r^2 + (K/n^3)r^3)$, where K is the number of vertices of the arrangement of the disks.

Optimal bounds are known on the minimum size of $\frac{1}{r}$-cuttings for hyperplane arrangements in $\mathbb{R}^d$ for any fixed $d \in \mathbb{N}$. Our proof techniques, however, do not seem to generalize to higher dimensions. One of our key tools is a BSP-like hierarchical partition, but no tight bound is known on the combinatorial complexity of BSPs for disjoint $(d-1)$-dimensional objects in $\mathbb{R}^d$, $d \geq 4$. We leave higher dimensional generalizations as an open problem: What is the minimum value $f_d(r)$ such that any finite set of disjoint objects in $\mathbb{R}^d$, each having bounded description complexity and lying in a hyperplane, (e.g., disjoint $(d-1)$-balls in $\mathbb{R}^d$) admits a $\frac{1}{r}$-cuttings of size $f_d(r)$?

We do not know if our results hold if we drop the condition that each object in $\mathbb{R}^3$ is planar. Is there an $\frac{1}{r}$-cutting of size $O(r^2)$ for any finite set of disjoint objects in $\mathbb{R}^3$, each having constant description complexity and lying in a 2-dimensional algebraic variety (e.g., disjoint spherical caps)? We do not even know if there is an $\frac{1}{r}$-cutting of size $O(r^3)$ for any finite set of 2-dimensional algebraic varieties of constant description complexity (it is known only that a *pseudo-cutting* of size $O(r^3 \text{ polylog } r)$ always exists).

References

[1] P. K. Agarwal. Partitioning arrangements of lines I: an efficient deterministic algorithm. *Discrete Comput. Geom.* 5:449–483, 1990.

[2] P. K. Agarwal, E. F. Grove, T. M. Murali, and J. S. Vitter. Binary space partitions for fat rectangles. *SIAM J. Comput.* 29(5):1422–1448, 2000.

[3] T. Asano, M. de Berg, O. Cheong, L. J. Guibas, J. Snoeyink, and H. Tamaki, Spanning trees crossing few barriers, *Discrete and Comput. Geom.*, 30(4):591–606, 2003.

[4] M. de Berg, L. J. Guibas, and D. Halperin. Vertical decompositions for triangles in 3-space. *Discrete Comput. Geom.* 15(1):35–61, 1996.

[5] J-D. Boissonnat, and M. Yvinec. Algorithmic geometry. Cambridge University Press, Cambridge, 1998

[6] P. M. Camerini. The min-max spanning tree problem and some extensions. *Inf. Process. Lett.* 7(1):10–14, 1978.

[7] B. Chazelle. Cutting hyperplanes for divide-and-conquer. *Discrete Comput. Geom.* 9:145–158, 1993.

[8] B. Chazelle. Cuttings. In *Handbook of Data Structures and Applications*, pp. 501–511, CRC Press, Boca Raton, FL, 2005.

[9] B. Chazelle, The Discrepancy Method, Cambridge University Press, Cambridge, 2000.

[10] B. Chazelle, H. Edelsbrunner, L. Guibas, M. Sharir. A singly-exponential stratification scheme for real semi-algebraic varieties. *Theoret. Comp. Sci.* 84:77–105, 1991.

[11] B Chazelle and J. Friedman. A deterministic view of random sampling and its use in geometry. *Combinatorica* 10(3):229-249, 1990.

[12] B. Chazelle and E. Welzl. Quasi-optimal range searching in spaces of finite VC-dimension. *Discrete Comput. Geom.*, 4(5):467–489, 1989.

[13] K. L. Clarkson and P. W. Shor. Application of random sampling in computational geometry, II. *Discrete Comput. Geom.* 4:387–421, 1989.

[14] A. Dumitrescu, J. S. B. Mitchell, and M. Sharir. Binary space partitions for axis-parallel segments, rectangles, and hyperrectangles. *Discrete Comput. Geom.* 31(2):207–227, 2004

[15] E. Ezra and M. Sharir. Counting and representing intersections among triangles in three dimensions. *Comput. Geom. Theory Appl.* 32(3):196–215, 2005.

[16] D. Haussler and E. Welzl. ε-nets and simplex range queries. *Discrete Comput. Geom.* 2:127–151, 1987.

[17] M. Hoffmann and Cs. D. Tóth. Connecting points in the presence of obstacles in the plane. In *Proc. 14th Canadian Conference on Comput. Geom. (Lethbridge, AB, 2002)*, pp. 63–67.

[18] D. R. Karger, P. N. Klein, and R. E. Tarjan. A randomized linear-time algorithm to find minimum spanning trees. *J. ACM*, 42:321–328, 1995.

[19] V. Koltun. Almost tight upper bounds for vertical decompositions in four dimensions. *J. ACM* 51(5):699–730, 2004.

[20] V. Koltun. Sharp bounds for vertical decompositions of linear arrangements in four dimensions. *Discrete Comput. Geom.* 31(3):435–460, 2004.

[21] V. Koltun and M. Sharir, Curve-sensitive cuttings, *SIAM J. Comput.* 34(4):863–878, 2005.

[22] J. Matoušek. Cutting hyperplane arrangements. *Discrete Comput. Geom.* 6:385–406, 1991

[23] J. Matoušek. Derandomization in computational geometry. In *Handbook of Computational Geometry (J.-R. Sack and J. Urrutia, eds.)*, pp. 559-595, North-Holland, Amsterdam, 2000.

[24] J. Matoušek, Lectures on Discrete Geometry, Springer-Verlag, Berlin, 2002.

[25] K. Mulmuley. Hidden surface removal with respect to a moving view point. In *Proc. Sympos. Theory of Computing*, pp. 512–522, ACM Press, 1991.

[26] J. Pach and M. Sharir, Geometric incidences, in *Towards a theory of geometric graphs*, AMS, vol. 342 of *Contemp. Math.*, 2004, pp. 185–223.

[27] M. Paterson and F. F. Yao. Efficient binary space partitions for hidden-surface removal and solid modeling. *Discrete Comput. Geom.* 5:485–503, 1990.

[28] M. Paterson and F. F. Yao. Optimal binary space partitions for orthogonal objects. *J. Algorithms* 13(1):99–113, 1992.

[29] M. Pellegrini. On point location and motion planning among simplices. *SIAM J. Comput.* 25(5):1061–1081, 1996.

[30] M. Pellegrini. On counting pairs of intersecting segments and off-line triangle range searching. *Algorithmica* 17(4):380–398, 1997.

[31] N. Sarnak and R.E. Tarjan. Planar point location using persistent search trees. *Commun. of the ACM*, 29(7):669–679, 1986.

[32] M. Sharir, The Clarkson-Shor technique revisited and extended, *Combinatorics Probability and Computation*, 12(2):191–201, 2003.

[33] Cs. D. Tóth. Binary space partition for orthogonal fat rectangles. In *Proc. 11th European Sympos. Algorithms*, pp. 494–505, vol. 2832 of LNCS, Springer-Verlag, Berlin, 2003.

[34] A. J. Zomorodian. Topology for Computing. Vol. 16 of *Cambridge Monographs on Applied and Computational Mathematics*, Cambridge University Press, 2005.