On Lines Avoiding Unit Balls in Three Dimensions*

Pankaj K. Agarwal† Boris Aronov‡ Vladlen Koltun§ Micha Sharir¶

Abstract

Let $\mathcal{B}$ be a set of n unit balls in $\mathbb{R}^3$. We show that the combinatorial complexity of the space of lines in $\mathbb{R}^3$ that avoid all the balls of $\mathcal{B}$ is $O(n^{3+\varepsilon})$, for any $\varepsilon > 0$. This result has connections to problems in visibility, ray shooting, motion planning and geometric optimization.

Categories and Subject Descriptors: F.2.2 [Theory of Computation]: Nonnumerical Algorithms and Problems—*geometrical problems and computations*; G.2.1 [Discrete Mathematics]: Combinatorics—*combinatorial complexity*

General Terms: Theory

Keywords: Combinatorial complexity, lines in space, visibility

1. Introduction

Let $\mathcal{B}$ be a set of n unit balls in $\mathbb{R}^3$. A line in $\mathbb{R}^3$ is called *free* (with respect to $\mathcal{B}$) if it does not meet the interior of any ball in $\mathcal{B}$; we also say that the line *avoids* all the balls of $\mathcal{B}$. Let $\mathcal{F} = \mathcal{F}(\mathcal{B})$ denote the space of lines that are free with respect to $\mathcal{B}$. We regard $\mathcal{F}$ as a subset of some 4-dimensional parametric space $\mathbb{L}$ that represents lines in $\mathbb{R}^3$ [12]. In this paper we study the combinatorial complexity of $\mathcal{F}$, which, for the time being, we define as the number of lines in $\mathcal{F}$ that are tangent to four balls in $\mathcal{B}$. (These are the "vertices" of $\mathcal{F}$; see below for a more precise definition.)

The space $\mathcal{F}$ is, in a sense, "antithetical" to the space $\mathcal{T} = \mathcal{T}(\mathcal{B})$ of *line transversals* of $\mathcal{B}$, i.e., the space of lines that intersect all the balls of $\mathcal{B}$. It is known (see, e.g., Koltun and Sharir [19]) that, with an appropriate choice of $\mathbb{L}$, $\mathcal{T}$ can be represented as a *sandwich region* lying between the upper envelope of one family of surfaces in $\mathbb{L}$, consisting of the loci of lines tangent to the balls of $\mathcal{B}$ from below, and the lower envelope of another such family, consisting of the loci of lines tangent to the balls of $\mathcal{B}$ from above. The recent result of [19] immediately yields a near-cubic bound on the complexity of $\mathcal{T}$. However, the space $\mathcal{F}$ of free lines does not appear to admit such a representation, and the best previously known upper bound for its complexity was the trivial bound of $O(n^4)$, obtained by observing that $\mathcal{F}$ is the union of certain cells of the arrangement of the aforementioned tangent-line surfaces in $\mathbb{L}$, and that the complexity of this entire four-dimensional arrangement is at most $O(n^4)$. The following theorem is the main result of this paper.

THEOREM 1.1. *The combinatorial complexity of the space of lines free with respect to a set of n unit balls in $\mathbb{R}^3$ is $O(n^{3+\varepsilon})$, for any $\varepsilon > 0$.*

Motivation. This result is of potential significance in several areas of computational and combinatorial geometry and their applications, including ray shooting and visibility computation, motion planning, and geometric optimization.

The study of the structure and complexity of the space of free lines (with respect to a collection of balls, as well as of other types of objects) has gained considerable attention, in part due to its relevance to the study of global-visibility data structures, such as the visibility skeleton [14] and the visibility complex [15, 24]. The size of these data structures crucially depends on the complexity of the space of free *segments* amid the given objects (defined in analogy to the space of free lines). The currently best known lower and upper bounds on the maximum possible complexity of the space of free segments amid n unit balls are the trivial bounds of $\Omega(n^2)$ and $O(n^4)$, respectively. Motivated by the practical significance of global visibility data structures

*Work on this paper has been supported by a joint grant from the U.S.–Israel Binational Science Foundation. Work by Pankaj Agarwal was also supported by NSF under grants CCR-00-86013 EIA-98-70724, EIA-99-72879, EIA-01-31905, and CCR-02-04118. Work by Boris Aronov was also supported by NSF Grants CCR-99-72568 and ITR CCR-00-81964. Work by Vladlen Koltun was also supported by NSF Grant CCR-01-21555. Work by Micha Sharir was also supported by NSF Grants CCR-97-32101 and CCR-00-98246, by a grant from the Israeli Academy of Sciences for a Center of Excellence in Geometric Computing at Tel Aviv University, and by the Hermann Minkowski–MINERVA Center for Geometry at Tel Aviv University.

†Department of Computer Science, Duke University, Durham, NC 27708-0129, USA; *pankaj@cs.duke.edu*.

‡Department of Computer and Information Science, Polytechnic University, Brooklyn, NY 11201-3840, USA; *aronov@cis.poly.edu*.

§Computer Science Division, University of California, Berkeley, CA 94720-1776, USA; *vladlen@cs.berkeley.edu*.

¶School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel and Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA; *michas@post.tau.ac.il*.

SCG'04, June 8–11, 2004, Brooklyn, New York, USA.

in computer graphics (e.g., to speeding up radiosity computations), researchers have studied the *expected* complexity of the space of free segments and of the space of free lines in a scene composed of n unit balls that are generated randomly according to certain probability distributions. A tight bound of $\Theta(n)$ on these quantities is known for this special case [13]. Better worst-case upper bounds are still desirable, though, and we hope that our new bound on the complexity of the space of free lines will inspire a similar improvement for the complexity of the space of free segments.

Our result is also related to the problem of ray shooting, which is often solved by reduction to the question of determining whether a query line is free [3, 8]. The known algorithms for ray shooting do not compute explicitly the space of all free lines, but instead construct data structures that encode it implicitly. It is therefore interesting to note that the best known algorithms for ray shooting amid balls that answer a query in polylogarithmic time require near-cubic storage even for such an implicit representation [1, 21]; there does not seem to be a more efficient solution for the case of unit balls.

Our result is also closely related to robot motion planning. The space $\mathcal{F}(\mathcal{B})$ is the *free configuration space* [25] of a line robot moving amid unit ball obstacles in $\mathbb{R}^3$, or, alternatively, of a cylindrical robot (of infinite length) moving amid point obstacles or amid congruent balls. Our result is of interest due to the scarcity of results on motion planning with translations and rotations in $\mathbb{R}^3$. The case of a line (or cylinder) robot moving amid congruent balls is the simplest instance of such a problem, since the motion has only four degrees of freedom, as opposed to six for a general rigid robot.

Finally, our result is related to the problem of computing the largest empty cylinder amid n points in $\mathbb{R}^3$, and to several related problems in geometric optimization.

Related work. The complexity of the space of free lines amid various types of *polyhedral* objects has been considered in the past, with mixed success. When the objects constitute a polyhedral region of total complexity n, or even in the special case of a collection of n lines in $\mathbb{R}^3$, the maximum complexity of the space of free lines is easily seen to be $\Theta(n^4)$ [12]. However, the complexity of the space of lines that avoid a polyhedral terrain having n faces is only near-cubic [23]. The $\Omega(n^4)$ lower bound constructions for the cases of polyhedra and lines are rather unrealistic, and the hope is that, when the objects are *fat* or well distributed, the complexity of the space of free lines is reduced. Collections of unit balls are the simplest example of such a well-behaved class of objects, which can still model fully 3-dimensional scenes, and they have indeed already been studied in this context [13, 15]. A related result [9] obtains a near-cubic bound on the complexity of the set of lines free relative to a collection of n homothets of a bounded-complexity convex polyhedron.

To calibrate the result of this paper and to put it in a wider context, we note that lines in $\mathbb{R}^3$ have four degrees of freedom, and can therefore be represented as points in $\mathbb{R}^4$ (the space $\mathbb{L}$ alluded to above).[1] Hence, $\mathcal{F}(\mathcal{B})$ is a subregion of $\mathbb{R}^4$. In fact, if we consider n surfaces in $\mathbb{R}^4$, each being the locus of all (points representing) lines tangent to a fixed ball in $\mathcal{B}$, and form the arrangement of these surfaces, then each cell of the arrangement consists of lines that intersect all the balls in a fixed subset of $\mathcal{B}$ and avoid the remaining balls. This arrangement has $O(n^4)$ cells of all dimensions, and $\mathcal{F}(\mathcal{B})$ is the union of a subset of them. Hence, the complexity of $\mathcal{F}(\mathcal{B})$ is $O(n^4)$, and it can be trivially constructed in near-quartic time. (In fact, these properties hold for any collection $\mathcal{B}$ of n objects of simple shape in $\mathbb{R}^3$.) Hence, a near-cubic upper bound for the case of unit balls is indeed a significant improvement.

Another way of looking at this problem is to define, for each ball $B \in \mathcal{B}$, the region K_B, consisting of all (points representing) lines that intersect B (the boundary of K_B is the aforementioned surface of lines tangent to B). Then $\mathcal{F}(\mathcal{B})$ is the complement of the union of the sets K_B, over $B \in \mathcal{B}$. Thus, analyzing the complexity of $\mathcal{F}(\mathcal{B})$ is an instance of the general study of the complexity of the union of geometric objects, an area that has received considerable attention in recent years; see [4, 6, 7, 11, 16, 17, 18, 20, 22]. Our result is among the very first non-trivial bounds on the complexity of the union of geometric objects in four dimensions (excluding the significantly easier cases of halfspaces, balls, and axis-parallel hypercubes). A companion work by the authors [5] derives a near-cubic bound for the union of κ*-round* objects in four dimensions, extending an earlier result by Koltun and Sharir [19] (see also [4]). These are the only results of this type of which we are aware.

Proof outline. Our analysis of the complexity of $\mathcal{F}(\mathcal{B})$ proceeds through a number of steps, each constituting a further reduction of the problem. We begin, in Section 2, by defining the space $\mathcal{F}$ of free lines as a subset of an arrangement of surfaces in $\mathbb{R}^4$, and show that in order to bound the complexity of $\mathcal{F}$ it is sufficient to bound the number of *vertices* of this arrangement that lie on the boundary of $\mathcal{F}$. These vertices correspond to free lines in $\mathbb{R}^3$ that are tangent to quadruples of balls in $\mathcal{B}$.

We then define, still in Section 2, the notion of vertices that are *deep* with respect to a reference direction in $\mathbb{S}^2$. Roughly speaking, a vertex that represents a line ℓ is deep with respect to a direction u if, when we enter each of the four balls B tangent to ℓ in the direction orthogonal to both u and ℓ, we traverse a "long" chord of B. It is shown that there exists a reference direction for which at least half of the vertices are deep. Hence it suffices to bound the maximal number of vertices that are deep with respect to some reference direction, a task that is undertaken in Section 3. We bound the number of deep vertices by deriving a recurrence relation on this quantity. We project the centers of the balls in $\mathcal{B}$ onto the xy-plane, and partition the projected centers using a (nonuniform) planar grid, such that every row and every column respectively contain n/r centers, for a certain r. We concentrate on vertices defined by quadruples of balls whose four centers project into four distinct rows and four distinct columns, and use recursion to bound the number of other types of vertices. We fix a quadruple of grid cells that lie in distinct rows and columns, and argue that the four corresponding sets of projected ball centers are

[1] Actually, lines should be represented in real projective 4-space. Since not all of them can be mapped to points in the finite portion of projective space, the real representation leaves out a set of lower dimension of unrepresentable lines. We ignore this issue for the sake of convenience, since, with a sufficiently generic choice of the coordinate system, all vertices of $\mathcal{F}$ will be represented by finite points.

doubly well-separated with respect to some "middle" point o, meaning intuitively that o lies "in between" one pair of cells and also in between the complementary pair of cells.

The main technical step of the proof, given in Section 4, concentrates on proving a near-cubic bound on the number of deep vertices defined by a doubly well-separated quadruple of sets of balls. By carefully defining a coordinate frame for the four-dimensional space of lines, we are able to express, in Section 4.1, each such vertex as a vertex in one of a number of *sandwich regions* of envelopes of four-dimensional arrangements that we define. This is where we exploit the nature of deep vertices. Despite of a good (near-cubic) existing bound on the complexity of sandwich regions in four dimensions [19], this is not the end of the story yet, since the number of sandwich regions that we define, as well as the overall size of all the sets of surfaces that form these regions, may be unbounded. In Section 4.2 we show that the functions defining the sandwich regions have *sparse domains*, in a certain sense defined in that section. This is where we exploit the nature of doubly well-separated sets and the fact that the balls are congruent. In Section 5 we establish a refined bound on the complexity of sandwich regions of envelopes of arrangement of trivariate functions with sparse domains, a result that we believe to be of independent interest. Using this result, we are able to prove the desired near-cubic bound on the number of deep vertices defined by doubly well-separated sets of balls.

Plugging this bound into the analysis of Section 3, we obtain the desired recurrence for the number of deep vertices, which solves to a near-cubic bound. Due to the reductions mentioned above, this bound leads directly to the main result of the paper.

2. Reduction to Deep Vertices

Let $\mathcal{B} = \{B_1, \dots, B_n\}$ be a set of n unit balls in $\mathbb{R}^3$. We assume that the balls in $\mathcal{B}$ are in *general position*, meaning that no line is tangent to any five distinct balls and that only a finite number of lines are tangent to any four balls. This involves no loss of generality because we can apply a sufficiently small random perturbation to the balls, putting them in general position, without losing more than an (expected) constant fraction of the number of faces of $\partial\mathcal{F}(\mathcal{B})$; see [26, Section 7.3.1] for similar arguments.

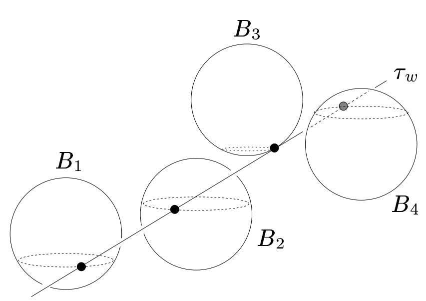

Figure 1. The illustrated line τ_w is represented by a vertex w of $\mathcal{F}(\mathcal{B})$.

Let $\mathbb{L}$ be the space of all lines in $\mathbb{R}^3$. Since lines in $\mathbb{R}^3$ can be parameterized by four parameters, $\mathbb{L}$ can be regarded as a 4-dimensional (real) parametric space. We recall the notation introduced above: For a ball $B \in \mathcal{B}$, let $K_B \subseteq \mathbb{L}$ denote the set of lines that intersect B; ∂K_B is the set of all lines that are tangent to B. By our general-position assumption, $\mathcal{F} = \mathcal{F}(\mathcal{B}) = \mathrm{cl}(\mathbb{L} \setminus \bigcup_{B\in\mathcal{B}} K_B)$, where $\mathrm{cl}(\cdot)$ denotes closure. The boundary of $\mathcal{F}$, denoted as $\partial\mathcal{F}$, thus consists of portions of the boundaries ∂K_B, for $B \in \mathcal{B}$. Any such portion is in fact a face of the arrangement $\mathcal{A}$ of the boundaries ∂K_B, for $B \in \mathcal{B}$. The *combinatorial complexity* of $\mathcal{F}$ is the number of faces of $\mathcal{A}$ of all dimensions that appear on $\partial\mathcal{F}$. A *vertex* w of $\mathcal{F}$ represents a free line τ_w in $\mathbb{R}^3$ that is tangent to four balls of $\mathcal{B}$; see Figure 1. We denote by $v_w \in \mathbb{S}^2$ the orientation[2] of τ_w, and by $\mathcal{B}_w \subseteq \mathcal{B}$ the set of four balls tangent to τ_w. Let $\mathcal{W} = \mathcal{W}(\mathcal{B})$ denote the set of vertices of $\mathcal{F}$.

Lemma 2.1. *For a set $\mathcal{B}$ of n unit balls in $\mathbb{R}^3$, the combinatorial complexity of $\mathcal{F}$ is $O(|\mathcal{W}| + n^3)$.*

Proof. Any face of $\partial\mathcal{F}$ that is incident upon a vertex of $\mathcal{F}$ can be charged to any of its vertices. The general-position assumption implies that each vertex is charged in this manner at most a constant number of times, and the number of such faces is thus $O(|\mathcal{W}|)$. A face f that is not incident upon a vertex of $\mathcal{F}$ is defined by at most three balls $B_1, \dots, B_l \in \mathcal{B}$, where $l \in \{1, 2, 3\}$, and either coincides with, or is incident upon, an entire connected component of $\bigcap_{i=1}^{l} \partial K_{B_i}$. We can thus charge f to such a component, and observe that, as above, the general position assumption implies that no component is charged more than $O(1)$ times. Since any intersection $\bigcap_{i=1}^{l} \partial K_{B_i}$ has only $O(1)$ connected components, the total number of faces that are not incident upon a vertex is $O(n^3)$. □

Fix a *reference direction* u in $\mathbb{S}^2$. Let w be a vertex in $\mathcal{W}$, and let $v_w \neq u$ be the orientation of τ_w. Let B_1, B_2, B_3, B_4 be the balls in $\mathcal{B}_w$, and let ℓ_i be the line passing through the tangency point $\tau_w \cap B_i$ and orthogonal to both u and v_w; these four lines are parallel. We call w *deep* (with respect to u) if for each $i \in \{1, 2, 3, 4\}$, the length of the chord $\ell_i \cap B_i$ is at least $2 \sin\alpha$ (or, equivalently, the central angle that the chord subtends is at least 2α), where α is a constant to be fixed below. Project τ_w and the balls in $\mathcal{B}_w$ onto a plane π_w normal to τ_w. Then τ_w projects to a point w^* and the four balls $B_1, \dots, B_4$ project to four unit disks $B_1^*, \dots, B_4^*$ whose boundaries all pass through w^*. The four lines ℓ_i project to

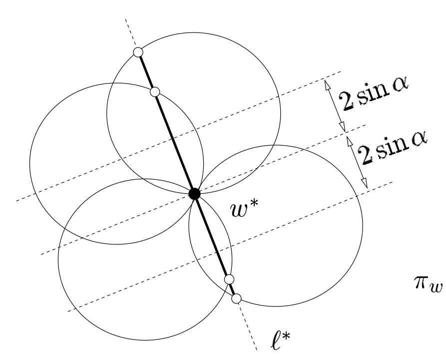

Figure 2. A deep vertex w of $\mathcal{F}$ and the associated projected configuration in the plane π_w normal to τ_w.

a common line ℓ^* that passes through w^*. If w is deep, then

[2]Strictly speaking, the lines in $\mathbb{L}$ are not considered to be oriented, so v_w is encoded by *two* antipodal points on $\mathbb{S}^2$. We ignore this technicality for the sake of simplicity.

the length of each of the chords $\ell^* \cap B_i^*$ is at least $2\sin\alpha$; see Figure 2.

Lemma 2.2. *There exists a reference direction $u \in \mathbb{S}^2$ such that the number of vertices in $\mathcal{W}$ that are deep with respect to u is at least $|\mathcal{W}|/2$, provided that $\alpha \le \pi/16$.*

Proof. We will show that if $\alpha \le \pi/16$ and u is chosen randomly and uniformly from $\mathbb{S}^2$, then the expected number of vertices in $\mathcal{W}$ that are deep with respect to u is at least $|\mathcal{W}|/2$. This implies in particular that there must exist a direction u for which the property described in the lemma holds.

Let $w \in \mathcal{W}$, let $v := v_w$ be the orientation of $\tau := \tau_w$, and let $\mathcal{B}_w = \{B_1, B_2, B_3, B_4\}$ be the set of balls tangent to τ. Fix one of these balls, say $B := B_1$. Without loss of generality, assume that τ is the z-axis, and that B is the unit ball centered at $(1, 0, 0)$. Regard the direction u as *bad* (for B) if the horizontal line ℓ in direction $u \times v$ through the origin intersects B in a chord whose length is smaller than $2\sin\alpha$. Clearly, if u is bad then ℓ meets (the horizontal cross section of) B in a chord that subtends a central angle of less than 2α, or equivalently, the direction $u \times v$ of ℓ forms an angle less than α with the y-axis. See Figures 3(i) and (ii). Therefore, $u \in \mathbb{S}^2$, being orthogonal to $u \times v$, must lie in a spherical double wedge of opening angle 2α centered around the great circle of $\mathbb{S}^2$ normal to the y-axis. See Figure 3(iii). Thus a random choice of u would land in this double wedge with probability $2\alpha/\pi$.

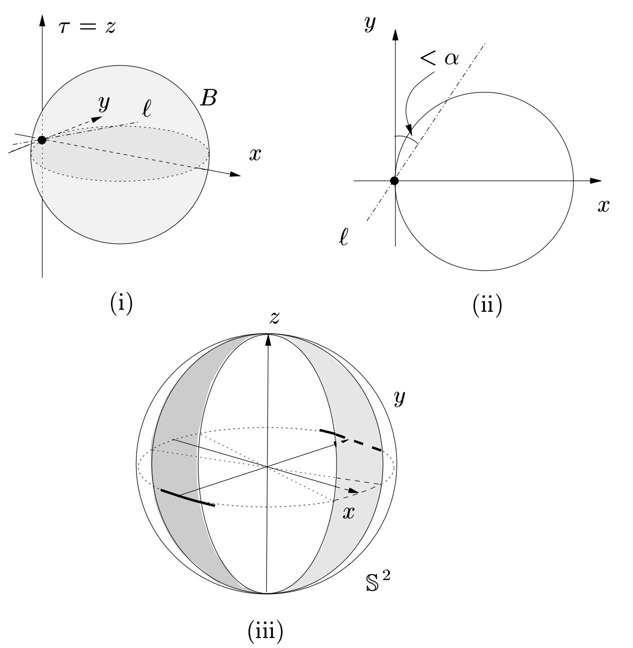

Figure 3. Illustrating the proof of Lemma 2.2. (i) The line ℓ and the ball B. (ii) The cross section of (i) by the xy-plane. (iii) The double wedge of bad directions on $\mathbb{S}^2$.

Applying this analysis to all four balls $B_1, \ldots, B_4$, we conclude that the probability that a random direction u is bad relative to at least one ball is no more than $8\alpha/\pi$, and setting α to any value not exceeding $\pi/16$ assures that a random choice of u is good for w with probability at least $1/2$, i.e., with probability at least $1/2$, w is a deep vertex with respect to u. This completes the proof. $\square$

3. Counting Deep Vertices

In view of Lemma 2.2, it suffices to bound the number of vertices in $\mathcal{W}$ that are deep with respect to some fixed direction u, which, without loss of generality, we may assume to be the $(+z)$-direction. For a subset $\mathcal{R} \subseteq \mathcal{B}$, let $\mathcal{D}(\mathcal{R}) \subseteq \mathcal{W}(\mathcal{R})$ denote the set of vertices that are deep with respect to the vertical direction; set $\psi(\mathcal{R}) := |\mathcal{D}(\mathcal{R})|$ and

$$\psi(m) := \max_{\substack{\mathcal{R} \subset \mathcal{B} \\ |\mathcal{R}| = m}} \psi(\mathcal{R}).$$

We derive a recurrence relation for $\psi(m)$.

Let c_i^* be the xy-projection of the center c_i of B_i, and let $C^* = \{c_1^*, \ldots, c_n^*\}$ be the set of the n projected centers in $\mathbb{R}^2$. Fix a parameter r, and partition C^* into r subsets of equal cardinality n/r by $r-1$ lines $x = a_1, \ldots, x = a_{r-1}$ lying in the xy-plane parallel to the y-axis. We assume for simplicity and without loss of generality that n is integrally divisible by r. Construct a similar partition of C^* into r subsets of equal cardinality by $r-1$ lines $y = b_1, \ldots, y = b_{r-1}$ lying in the xy-plane parallel to the x-axis. These $2(r-1)$ lines partition the xy-plane into the $r \times r$ non-uniform grid

$$\mathcal{G} := \{\kappa_{ij} := [a_{i-1}, a_i] \times [b_{j-1}, b_j] \mid i, j = 1, \ldots, r\},$$

where we put $a_0, b_0 := -\infty$ and $a_r, b_r := +\infty$. For $i, j = 1, \ldots, r$, let $\mathcal{B}_{ij} := \{B_l \mid c_l^* \in \kappa_{ij}\}$; set $n_{ij} := |\mathcal{B}_{ij}|$. We have $\sum_i n_{ij} = n/r$ for any fixed j and $\sum_j n_{ij} = n/r$ for any fixed i. Let w be a deep vertex in $\mathcal{D}(\mathcal{B})$, let c_1^*, c_2^*, c_3^*, and c_4^* be the projections of the centers of the four respective balls $B_1, B_2, B_3, B_4 \in \mathcal{B}_w$, and let $\kappa_{i_1j_1}, \ldots, \kappa_{i_4j_4}$ denote the four (not necessarily distinct) cells of $\mathcal{G}$ that contain $c_1^*, \ldots, c_4^*$, respectively.

Suppose first that at least two of these cells lie in the same row or in the same column of $\mathcal{G}$. We estimate the number of such vertices w recursively, by solving $2\binom{r}{3}$ subproblems, each involving the balls whose projected centers lie in any fixed *triple* of rows or columns of $\mathcal{G}$. The number of balls in each subproblem is at most $3n/r$. The contribution of such vertices to $\psi(n)$ is at most $2\binom{r}{3}\psi(3n/r)$.

Otherwise, let $\mathcal{D}(\mathcal{B}_{i_1j_1}, \mathcal{B}_{i_2j_2}, \mathcal{B}_{i_3j_3}, \mathcal{B}_{i_4j_4}) \subseteq \mathcal{D}(\mathcal{B})$ denote the set of vertices w for which the four centers of the balls in $\mathcal{B}_w$ project into cells $\kappa_{i_1j_1}, \kappa_{i_2j_2}, \kappa_{i_3j_3}$, and $\kappa_{i_4j_4}$ of $\mathcal{G}$, lying in four distinct rows i_1, i_2, i_3, i_4 and in four distinct columns j_1, j_2, j_3, j_4, respectively. Then

$$\psi(n) \le 2\binom{r}{3}\psi\left(\frac{3n}{r}\right) + \sum |\mathcal{D}(\mathcal{B}_{i_1j_1}, \mathcal{B}_{i_2j_2}, \mathcal{B}_{i_3j_3}, \mathcal{B}_{i_4j_4})|, \quad (1)$$

where the summation is over all (unordered) quadruples of cells of $\mathcal{G}$ that lie in distinct rows and in distinct columns.

We next introduce the notion of well separated pairs of balls. We call two balls B_1 and B_2, with respective projected centers c_1^*, c_2^*, *well separated* with respect to a point $o \in \mathbb{R}^2$ if

$$|oc_1^*|, |oc_2^*| \le |c_1^* c_2^*|. \quad (2)$$

We call a set of four balls *doubly well separated* with respect to o if it is the union of two complementary pairs, both of which are well separated with respect to o. We extend this definition by saying that four sets $\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3, \mathcal{B}_4$ of balls are doubly well separated with respect to o if there exist two complementary pairs of sets, say $\mathcal{B}_1, \mathcal{B}_2$ and $\mathcal{B}_3, \mathcal{B}_4$, such that any pair of balls in $\mathcal{B}_1 \times \mathcal{B}_2$ is well separated with respect to o, and so is any pair of balls in $\mathcal{B}_3 \times \mathcal{B}_4$.

Lemma 3.1 (Separation Lemma). *Given four distinct rows i_1, i_2, i_3, i_4 and four distinct columns j_1, j_2, j_3, j_4, there exists a point o so that the four sets $\mathcal{B}_{i_1j_1}, \mathcal{B}_{i_2j_2}, \mathcal{B}_{i_3j_3}$, and $\mathcal{B}_{i_4j_4}$ are doubly well separated with respect to o.*

Proof. Draw an x-parallel line λ_1 so that two of the cells $\kappa_{i_1j_1}, \kappa_{i_2j_2}, \kappa_{i_3j_3}, \kappa_{i_4j_4}$ lie below it and two lie above. Similarly, draw a y-parallel line λ_2 that has two of the cells on either side. Let o be the intersection point of λ_1 and λ_2. Clearly, either each of the four quadrants determined by λ_1 and λ_2 contains exactly one of the four cells, or two opposite quadrants each contain two of these cells and the two other quadrants do not contain any cell. In either case, there exist two complementary pairs of cells, so that the two cells in each pair lie in opposite quadrants; refer to Figure 4.

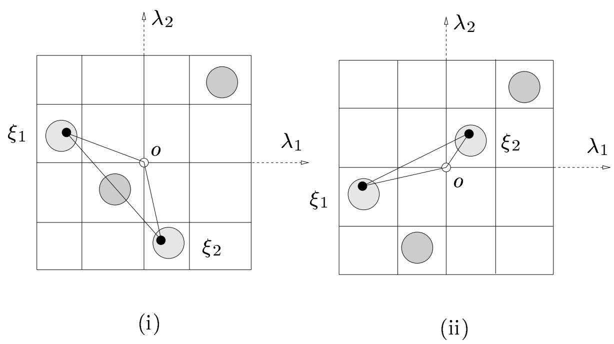

Figure 4. Illustrating the proof of the Separation Lemma: (i) The cells lie in four distinct quadrants. (ii) The cells lie in two opposite quadrants. The shaded disks represent the sets $\mathcal{B}_{i_kj_k}$ of projected centers. The lighter pairs of sets are well separated with respect to o, and so are the darker pairs.

Hence, if ξ_1 and ξ_2 are any two points, one from each of the two opposite cells in a pair, then the angle $\angle \xi_1 o \xi_2$ is at least $\pi/2$, implying that $|\xi_1\xi_2| \geq |\xi_1 o|, |\xi_2 o|$, and thus showing that any pair of balls whose centers project into these two cells is well separated with respect to o. This completes the proof of the lemma. $\square$

Let $\mathcal{B}_1$, $\mathcal{B}_2$, $\mathcal{B}_3$, and $\mathcal{B}_4$ be four sets of unit balls that are doubly well separated with respect to some point $o \in \mathbb{R}^2$. In the following section (see Lemma 4.5) we prove that the number of "rainbow" vertices in $\mathcal{D}(\mathcal{B}_1 \cup \mathcal{B}_2 \cup \mathcal{B}_3 \cup \mathcal{B}_4)$, i.e., vertices w such that $\mathcal{B}_w$ contains one ball of each of the sets $\mathcal{B}_1, \ldots, \mathcal{B}_4$, is $O(|\mathcal{B}_1 \cup \mathcal{B}_2 \cup \mathcal{B}_3 \cup \mathcal{B}_4|^{3+\varepsilon})$, for any $\varepsilon > 0$. We apply this bound to each of the $O(r^8)$ quadruples of cells of $\mathcal{G}$ that lie in four distinct rows and in four distinct columns, and observe that for each such quadruple the corresponding size $|\mathcal{B}_1 \cup \mathcal{B}_2 \cup \mathcal{B}_3 \cup \mathcal{B}_4|$ is at most $4n/r$, to conclude that these cells contribute a total of

$$O\left(r^8 \cdot \left(\frac{4n}{r}\right)^{3+\varepsilon}\right) = O(r^5 n^{3+\varepsilon})$$

to $\psi(n)$. Thus the recurrence (1) for $\psi(\cdot)$ becomes

$$\psi(n) \leq \frac{r^3}{3}\psi(3n/r) + Cr^5 n^{3+\varepsilon},$$

for any $\varepsilon > 0$ and for a constant $C = C(\varepsilon)$. The solution of the recurrence is $\psi(n) \leq An^{3+\varepsilon}$, where the constant A depends on ε and C. Indeed, substituting into the recurrence we obtain

$$\begin{aligned}\psi(n) &\leq \frac{r^3}{3}A(3n/r)^{3+\varepsilon} + Cr^5 n^{3+\varepsilon} \\ &\leq An^{3+\varepsilon}\left(\frac{3^{2+\varepsilon}}{r^\varepsilon} + \frac{Cr^5}{A}\right) \leq An^{3+\varepsilon},\end{aligned}$$

provided that r is chosen to be greater than $3^{(2+\varepsilon)/\varepsilon}$ and A to be sufficiently large. Together with Lemmas 2.1 and 2.2, this proves Theorem 1.1, the main result of the paper.

4. Counting Rainbow Vertices in Doubly Well Separated Sets

4.1 Reduction to Sandwich Regions

Let $\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3, \mathcal{B}_4$ be four pairwise-disjoint sets of unit balls in $\mathbb{R}^3$ that are doubly well separated with respect to the origin, say. Set $\mathcal{B} := \mathcal{B}_1 \cup \mathcal{B}_2 \cup \mathcal{B}_3 \cup \mathcal{B}_4$. We call a subset $\mathcal{R} \subseteq \mathcal{B}$ of size at most four *rainbow* if each ball in $\mathcal{R}$ belongs to a different set $\mathcal{B}_j$. A vertex w of $\mathcal{F}(\mathcal{B})$ is *rainbow* if $\mathcal{B}_w$ is rainbow. We wish to bound the number of rainbow vertices in $\mathcal{D}(\mathcal{B})$; this is the missing step in the proof of Theorem 1.1.

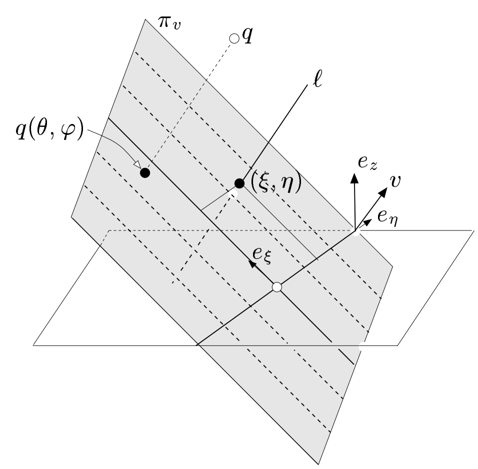

Figure 5. Representation of lines in space; ℓ is represented by the quadruple $(\theta, \varphi, \xi, \eta)$.

We begin by specifying our choice of the four parameters $(\theta, \varphi, \xi, \eta)$ that represent a line ℓ in $\mathbb{R}^3$. The first two parameters (θ, φ) are the spherical coordinates of the direction v of ℓ, where φ is its *azimuth*, i.e., the angle it forms with the positive z-direction, and θ is the horizontal orientation of the xy-projection of v. Let π_v denote the plane orthogonal to v and passing through the origin o. We define below a 2-dimensional coordinate frame (ξ, η) in π_v with o as its origin. The last two parameters (ξ, η) that represent ℓ are the coordinates in this frame of the point $\pi_v \cap \ell$.

Let $e_z := (0, 0, 1)$ be the vertical unit vector. The unit vectors e_ξ, e_η that define the ξ- and η-axes are constructed as follows. We set

$$e_\eta := (\cos(\theta + \pi/2), \sin(\theta + \pi/2), 0) = (-\sin\theta, \cos\theta, 0).$$

We observe that e_η is orthogonal to both e_z and v. (Note that θ is undefined when v is vertical. We exclude lines with $\varphi = 0$ from our analysis; with a generic choice of u, no vertex of $\mathcal{W}(\mathcal{B})$ will correspond to a vertical line.) We then choose $e_\xi := v \times e_\eta$, i.e., the ξ-axis is the intersection of π_v with the vertical plane through o spanned by v and e_z. See Figure 5. For a point $q \in \mathbb{R}^3$ and a direction $v = (\theta, \varphi) \in \mathbb{S}^2$, let $q(v) = q(\theta, \varphi)$ denote the orthogonal projection of q onto the plane π_v, whose coordinates are represented in the (ξ, η)-system. Note that the η-coordinate of $q(\theta, \varphi)$ does not depend on the value of φ. Similarly, for a ball $B \in \mathcal{B}$, we denote by $B(v)$ its orthogonal projection onto π_v (which is a unit disk), represented in the (ξ, η)-system.

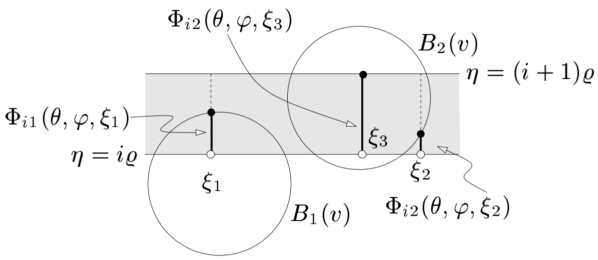

Figure 6. Definition of the function $\Phi_{ij}(\theta, \varphi, \xi)$; all three cases are illustrated.

We choose a sufficiently small parameter $\varrho < 2\sin\alpha$, and partition the plane π_v into ξ-vertical strips of the form

$$\sigma_i^v = \sigma_i^{(\theta,\varphi)} \colon \varrho i \leq \eta \leq \varrho(i+1), \text{ for } i \in \mathbb{Z}.$$

We partition $\mathcal{D}(\mathcal{B})$ into subsets as follows. For $i \in \mathbb{Z}$, set

$$\mathcal{D}_i := \{w = (\theta_w, \varphi_w, \xi_w, \eta_w) \mid (\xi_w, \eta_w) \in \sigma_i^{(\theta_w, \varphi_w)}\}.$$

We bound the size of each $\mathcal{D}_i$ separately, as follows. Fix i and define a family $\mathfrak{F}_i = \{\Phi_{i1}, \ldots, \Phi_{in}\}$ of n partially defined trivariate functions, one for each ball in $\mathcal{B}$, with θ, φ, ξ being the independent variables. A function Φ_{ij} is defined for all (θ, ϕ, ξ) such that one of the points $(\xi, \varrho i)$, $(\xi, \varrho(i+1))$ in the plane π_v is contained in $B_j(v)$, where $v := (\theta, \varphi)$.

Specifically, the domain Δ_{ij} of Φ_{ij} is defined as follows. For a triple (θ, φ, ξ), let $v := (\theta, \varphi)$ and consider the strip σ_i^v. If $K_{ij} := B_j(v) \cap \sigma_i^v = \emptyset$ then $\Phi_{ij}(\theta, \varphi, \xi)$ is undefined for all ξ. If K_{ij} is nonempty, let $I_{ij} = I_{ij}(\theta, \varphi)$ denote the set of all ξ for which at least one of $(\xi, \varrho i), (\xi, \varrho(i+1))$ is in K_{ij}; it is easily seen that I_{ij} is a nonempty interval. Then $\Phi_{ij}(\theta, \varphi, \xi)$ is defined if and only if $\xi \in I_{ij}$. Its value depends on which of the points $(\xi, \varrho i), (\xi, \varrho(i+1))$ lie in K_{ij}. Let λ be the length of the intersection of K_{ij} with the η-parallel line through $(\xi, \varrho i)$. If only $(\xi, \varrho i) \in K_{ij}$, then $\Phi_{ij}(\theta, \varphi, \xi) = \lambda$ (refer to $\Phi_{i1}(\theta, \varphi, \xi_1)$ in Figure 6); if only $(\xi, \varrho(i+1)) \in K$, then $\Phi_{ij}(\theta, \varphi, \xi) = \varrho - \lambda$ ($\Phi_{i2}(\theta, \varphi, \xi_2)$ in the figure); if both $(\xi, \varrho i), (\xi, \varrho(i+1)) \in K_{ij}$, then $\Phi_{ij}(\theta, \varphi, \xi) = \varrho$ ($\Phi_{i2}(\theta, \varphi, \xi_3)$ in the figure). Note that Φ_{ij} may be discontinuous at the intersection point of $\partial B_j(v)$ and $\eta = \varrho i$, for example, see Φ_{i2} in Figure 6.

Next, we partition the domain of Φ_{ij} into two subdomains Δ_{ij}^+, Δ_{ij}^-, as follows: $(\theta, \varphi, \xi) \in \Phi_{ij}$ belongs to Δ_{ij}^+ if the point $(\xi, \varrho i) \in B_j(v)$ and to Δ_{ij}^- otherwise (i.e., $(\xi, \varrho i) \notin B_j(v)$ and $(\xi, \varrho(i+1)) \in B_j(v)$). We denote the restriction of Φ_{ij} to Δ_{ij}^+ (resp., to Δ_{ij}^-) as Φ_{ij}^+ (resp., as Φ_{ij}^-). Each function Φ_{ij}^+, Φ_{ij}^- has *constant description complexity*, meaning that it can be defined in terms of a constant number of polynomial equalities and inequalities of constant maximum degree (using the standard re-parametrization of θ, φ by $\tan(\theta/2)$ and $\tan(\varphi/2)$, respectively).

For each i, set $\mathfrak{F}_i^+ := \{\Phi_{ij}^+\}_{j=1}^n$, and $\mathfrak{F}_i^- := \{\Phi_{ij}^-\}_{j=1}^n$, where only functions with nonempty domains of definition are included in these collections. Let

$$\Xi_i := \{(\theta, \varphi, \xi, \eta) \in \mathbb{S}^2 \times \mathbb{R}^2 \mid \max_j \Phi_{ij}^+(\theta, \varphi, \xi) \leq \eta \leq \min_j \Phi_{ij}^-(\theta, \varphi, \xi)\}.$$

Ξ_i is a "sandwich region," i.e., the region lying between the upper envelope of $\mathfrak{F}_i^+$ and the lower envelope of $\mathfrak{F}_i^-$. Set $n_i = |\mathfrak{F}_i^-| + |\mathfrak{F}_i^+|$. We call a vertex of Ξ_i *rainbow* if the four function graphs incident upon the vertex correspond to balls from different families $\mathcal{B}_j$.

Lemma 4.1 (Sandwich Lemma). *Every rainbow vertex of $\mathcal{D}_i$ corresponds to a rainbow vertex of Ξ_i.*

Proof. This is a consequence of the definition of deep vertices. Let $w = (\theta_w, \varphi_w, \xi_w, \eta_w)$ be a deep vertex. The projection of each of the four balls $B \in \mathcal{B}_w$ is a unit disk in π_{v_w} and the line τ_w intersects π_{v_w} at a point $w(v_w)$ lying on the boundaries of all four disks, where $v_w = (\theta_w, \varphi_w)$. Since w is deep, the η-parallel line through $w(v_w)$, whose direction is orthogonal to both v_w and the z-direction, intersects each of these disks in a chord of length at least $2\sin\alpha > \varrho$; see Figure 2. This implies that each of the four corresponding functions Φ_{ij} (where i is the index of the strip that contains $w(v_w)$) is defined at $(\theta_w, \varphi_w, \xi_w)$, and they all have the same value there. Moreover, since no other ball projects to a disk that contains $w(v_w)$, it follows that the intersection of the graphs of the four Φ_{ij}'s is indeed a (rainbow) vertex of Ξ_i. □

By the result of Koltun and Sharir [19], the number of vertices of Ξ_i is $O(n_i^{3+\varepsilon})$, for any $\varepsilon > 0$. Summing over all strips σ_i^v, the overall number of deep vertices under consideration is $\sum_{i\in\mathbb{Z}} O(n_i^{3+\varepsilon})$, for any $\varepsilon > 0$.

Unfortunately, this bound is too weak—in fact, it can be arbitrarily large, because, if the center of a ball B_j lies at distance r from o, B_j contributes a function Φ_{ij} with nonempty domain of definition to $\Theta(r/\varrho)$ strips, so if the balls lie arbitrarily far from the origin, the number of indices i with $n_i \approx n$ may be arbitrarily large. We handle this problem below by counting only the number of rainbow vertices, exploiting the separation property of $\mathcal{B}$, and refining the bound of [19].

4.2 Exploiting Separation

Let J_{ij} be the θ-*projection* of Δ_{ij}, and let χ_i be the number of rainbow triples $B_1, B_2, B_3 \in \mathcal{B}$ so that $J_{i1} \cap J_{i2} \cap J_{i3} \neq \emptyset$, i.e., there exist θ_0, φ_j, ξ_j, for $1 \leq j \leq 3$, so that $(\theta_0, \varphi_j, \xi_j) \in \Delta_{ij}$ for each j. Our refinement, obtained in the following section (Corollary 5.4), shows that the number of rainbow vertices in Ξ_i is $O(n^\varepsilon \chi_i)$. We will show (see Lemma 4.4) that $\sum_i \chi_i = O(n^3)$, which will then imply that the total number of deep rainbow vertices in $\mathcal{F}(\mathcal{B})$ is $O(n^{3+\varepsilon})$, and thus complete the proof of Theorem 1.1.

The following is a key technical lemma that encapsulates the significance of separation for our analysis.

Lemma 4.2 (Sparseness Lemma). *There exists a constant $c = c(\varrho)$ with the following property. For any pair $\{B_1, B_2\}$*

of distinct unit balls that are well-separated with respect to the origin, there are at most c integers k with the property that $J_{k1} \cap J_{k2} \neq \emptyset$, i.e., the θ-projections of Δ_{k1} and Δ_{k2} overlap,

Proof. Let $\mathcal{I}$ denote the set of integers k for which $J_{k1} \cap J_{k2} \neq \emptyset$. Let $k \in \mathcal{I}$, and let $\theta_0 \in J_{k1} \cap J_{k2}$. Recall that the η-direction in π_v depends only on θ and not on φ. Consequently, if we keep θ_0 fixed and vary φ, the plane π_v rotates about the fixed horizontal η-axis. For any fixed point $q \in \mathbb{R}^3$, its projection onto π_v traces a vertical circular arc in $\mathbb{R}^3$, which appears within the rotating plane π_v as a line segment orthogonal to the η-axis. Note also that the k-th strip σ_k^v rotates with π_v about the fixed η-axis, but its representation in the moving (ξ, η)-system does not change.

Let c_1, c_2 be the respective centers of B_1, B_2. Since $\theta_0 \in J_{k1} \cap J_{k2}$, the two (parallel) segments traced within π_v by the projections of c_1 and c_2 lie at η-distance at most 1 from the (fixed) strip σ_k^v. In particular, when π_v becomes horizontal, both xy-projections $q_1 = c_1^*$, $q_2 = c_2^*$ of these centers lie in the xy-plane, each at η-distance at most 1 from $\sigma_k^{(\theta_0,0)}$ and thus at η-distance at most $2+\varrho$ from each other. Technically, we need to regard the horizontal position of π_v as a limit position, obtained as $\varphi \to 0$, because the (ξ, η)-system is undefined when $\varphi = 0$. We interpret $\sigma_k^{(\theta_0,0)}$ simply as strips orthogonal to the line at orientation $(-\sin\theta_0, \cos\theta_0, 0)$.) Our goal is thus to bound the number of integers k for which there exists θ so that both points q_1 and q_2 lie in the extended strip $\hat{\sigma}_k^{(\theta,0)}$ of width $2+\varrho$, obtained by expanding $\sigma_k^{(\theta,0)}$ by distance 1 in the positive and negative η-directions. Note that $\hat{\sigma}_k^{(\theta,0)}$ rotates about the origin as θ varies. We distinguish between two cases:

Case (i): $|q_1q_2| \le 2 + \varrho$. Since B_1, B_2 are well separated, $|oq_1| \le |q_1q_2| \le 2+\varrho$. Hence any strip $\hat{\sigma}_k^{(\theta,0)}$ that contains q_1 must lie at distance at most $2+\varrho$ from the origin, and the number of such strips is at most $2\lceil (3+\varrho)/\varrho \rceil$. This clearly also serves as an upper bound on the number of extended strips that contain (for the same θ) both q_1 and q_2.

Case (ii): $|q_1q_2| > 2+\varrho$. Since the situation is symmetric with respect to rotation about the origin, we may assume, without loss of generality, that q_1 and q_2 have the same y-coordinate h. Since B_1 and B_2 are well-separated, q_1q_2 is the longest edge in the triangle oq_1q_2, which implies that q_1 and q_2 lie in different quadrants, and that $h \le |q_1q_2|$. See Figure 7. A point $x \in \hat{\sigma}_k^{(\theta,0)}$ must satisfy

$$k\varrho - 1 \le x \cdot e_\eta \le (k+1)\varrho + 1. \tag{3}$$

In particular, as the extended strip $\hat{\sigma}_k^{(\theta,0)}$ rotates about the

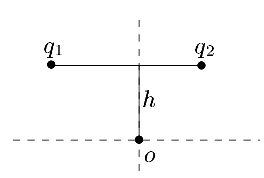

Figure 7. The triangle oq_1q_2.

origin, it contains both q_1 and q_2 when the corresponding e_η satisfies

$$\begin{aligned} k\varrho - 1 \le q_1 \cdot e_\eta \le (k+1)\varrho + 1 \text{ and} \\ k\varrho - 1 \le q_2 \cdot e_\eta \le (k+1)\varrho + 1. \end{aligned} \tag{4}$$

Subtracting the two inequalities, we obtain

$$|q_1q_2 \cdot e_\eta| \le (k+1)\varrho + 1 - (k\varrho - 1) = 2 + \varrho.$$

On the other hand,

$$|q_1q_2 \cdot e_\eta| = |q_1q_2 \cdot (-\sin\theta, \cos\theta, 0)| = |q_1q_2||\sin\theta|$$

Therefore

$$|\sin\theta| \le \frac{2+\varrho}{|q_1q_2|}. \tag{5}$$

In addition, since both $q_1, q_2 \in \hat{\sigma}_k^{(\theta,0)}$, so is the point $(0,h,0)$, which must therefore also satisfy (3), namely,

$$k\varrho - 1 \le h\cos\theta \le (k+1)\varrho + 1. \tag{6}$$

Let $\theta_0 \in (0, \pi/2)$ satisfy $\sin\theta_0 = (2+\varrho)/|q_1q_2|$. The range of θ where (5) holds is thus contained in the union of the two angular intervals $(-\theta_0, \theta_0)$ and $(\pi - \theta_0, \pi + \theta_0)$. Any k that satisfies (4) and thus also (6) must therefore be in one of the two ranges

$$\left[\min_{\theta \in (-\theta_0, \theta_0)} \frac{h\cos\theta - 1}{\varrho} - 1,\ \max_{\theta \in (-\theta_0, \theta_0)} \frac{h\cos\theta + 1}{\varrho} \right],$$

and

$$\left[\min_{\theta \in (\pi-\theta_0, \pi+\theta_0)} \frac{h\cos\theta - 1}{\varrho} - 1,\ \max_{\theta \in (\pi-\theta_0, \pi+\theta_0)} \frac{h\cos\theta + 1}{\varrho} \right].$$

The number of integers k in the former range is at most

$$\left\lceil \frac{h}{\varrho}\left(\max_{\theta \in (-\theta_0,\theta_0)} \cos\theta - \min_{\theta \in (-\theta_0,\theta_0)} \cos\theta \right) + \frac{2}{\varrho} + 1 \right\rceil =$$
$$\left\lceil \frac{h(1-\cos\theta_0)}{\varrho} + \frac{2}{\varrho} + 1 \right\rceil < \left\lceil \frac{h\sin\theta_0}{\varrho} + \frac{2}{\varrho} + 1 \right\rceil \le \left\lceil \frac{4}{\varrho} + 2 \right\rceil,$$

where we have used the facts that $h \le |q_1q_2|$ and that $|q_1q_2|\sin\theta_0 = 2+\varrho$. The analysis for the second angular interval yields the same estimate, for a total of no more than $2(4/\varrho + 3)$ values of k. This completes the proof of the lemma. □

Corollary 4.3. *There exists a constant $c = c(\varrho)$ with the following property. For any triple $\{B_1, B_2, B_3\}$ of distinct unit balls such that B_1 and B_2 are well-separated with respect to the origin, there are at most c integers k with the property that $J_{k1} \cap J_{k2} \cap J_{k3} \neq \emptyset$.*

Lemma 4.4. $\sum_i \chi_i = O(n^3)$.

Proof.

$$\begin{aligned} \sum_i \chi_i &= \sum_i \left| \{(a,b,c) \mid a < b < c \text{ and } J_{ia} \cap J_{ib} \cap J_{ic} \neq \emptyset\} \right| \\ &= \sum_{a<b<c} \left| \{ i \mid J_{ia} \cap J_{ib} \cap J_{ic} \neq \emptyset \} \right| \\ &\le \binom{n}{3} c(\varrho) = O(n^3). \quad \square \end{aligned}$$

As already mentioned, in Corollary 5.4 in the following section we prove that $|\Xi_i| = O(n^\varepsilon \chi_i)$, making crucial use of the separation property—see below for details. This implies the following result, which serves as the missing link in the proof of Theorem 1.1, as presented in Sections 2 and 3.

Lemma 4.5. *Let $\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3, \mathcal{B}_4$ be four sets of unit balls in $\mathbb{R}^3$ of total size n which are doubly well separated with respect to the origin. Then the number of deep rainbow vertices in $\mathcal{F}(\mathcal{B}_1 \cup \mathcal{B}_2 \cup \mathcal{B}_3 \cup \mathcal{B}_4)$ is $O(n^{3+\varepsilon})$, for any $\varepsilon > 0$.*

5. Functions with Sparse Domains

This section provides refined upper bounds on the complexity of the "sandwich" region of two collections of trivariate *partially-defined* functions. We believe that this refinement is of independent interest and expect it to be useful for other applications as well. Let $\mathcal{F} = \{f_1, \ldots, f_n\}$ be a family of partially-defined trivariate functions of *constant description complexity* (as defined in the preceding section) and in general position, which is the disjoint union $\mathcal{F}_1 \cup \mathcal{F}_2 \cup \mathcal{F}_3 \cup \mathcal{F}_4$ of four subfamilies. In what follows we will not distinguish between a function and its graph. We call a subset $\mathcal{R} \subseteq \mathcal{F}$ of size at most four *rainbow* if each function in $\mathcal{R}$ belongs to a different set $\mathcal{F}_j$. A vertex in the arrangement $\mathcal{A}(\mathcal{F})$ of (the graphs of) the functions in $\mathcal{F}$ is called *rainbow* if the subset of four functions whose graphs are incident upon v is rainbow. Let $\mathcal{F}^+$ and $\mathcal{F}^-$ be another partition of $\mathcal{F}$ into two subsets, and let $\Xi(\mathcal{F}^+, \mathcal{F}^-)$ denote the sandwich region lying between the upper envelope of $\mathcal{F}^+$ and the lower envelope of $\mathcal{F}^-$, i.e.,

$$\Xi(\mathcal{F}^+, \mathcal{F}^-) := \{(x, y, z, w) \in \mathbb{R}^4 \mid \max_{f \in \mathcal{F}^+} f(x, y, z) \le w \le \min_{f \in \mathcal{F}^-} f(x, y, z)\}.$$

Let Π be the set of rainbow vertices of $\Xi(\mathcal{F}^+, \mathcal{F}^-)$. We wish to bound the size of Π. As already noted, the recent result of [19] implies that the size of Π is $O(n^{3+\varepsilon})$, for any $\varepsilon > 0$. The goal of this section is to refine this bound in terms of a parameter that counts overlaps between the domains of the functions in $\mathcal{F}$, in a manner already alluded to in the previous section and defined more precisely below. We begin with the following simple but weak bound, which will be central in the derivation of our main result.

Lemma 5.1. *If $|\mathcal{F}_1| \ge |\mathcal{F}_2| \ge |\mathcal{F}_3| \ge |\mathcal{F}_4|$, then*

$$|\Pi| = O(|\mathcal{F}_1||\mathcal{F}_2||\mathcal{F}_3|^{1+\varepsilon}),$$

for any $\varepsilon > 0$.

Proof. Put $n_i = |\mathcal{F}_i|$, for $i = 1, \ldots, 4$. Partition $\mathcal{F}_1$ into $t := \lceil n_1/n_3 \rceil$ pairwise-disjoint subsets $X_1, \ldots, X_t$, each of size at most n_3, and partition similarly $\mathcal{F}_2$ into $s := \lceil n_2/n_3 \rceil$ pairwise-disjoint subsets $Y_1, \ldots, Y_s$, each also of size at most n_3. Let $\mathcal{F}^+_{ij} := \mathcal{F}^+ \cap (X_i \cup Y_j \cup \mathcal{F}_3 \cup \mathcal{F}_4)$ and $\mathcal{F}^-_{ij} := \mathcal{F}^- \cap (X_i \cup Y_j \cup \mathcal{F}_3 \cup \mathcal{F}_4)$, for $i = 1, \ldots, t$ and $j = 1, \ldots, s$. Then any vertex of Π is a vertex of $\Xi(\mathcal{F}^+_{ij}, \mathcal{F}^-_{ij})$ for some i and j. Since $|\mathcal{F}^+_{ij} \cup \mathcal{F}^-_{ij}| \le 4n_3$, the aforementioned bound of [19] implies that $\Xi(\mathcal{F}^+_{ij}, \mathcal{F}^-_{ij})$ has $O(n_3^{3+\varepsilon})$ vertices, for any $\varepsilon > 0$. Hence, the overall number of vertices in Π is at most $ts \cdot O(n_3^{3+\varepsilon}) = O(n_1 n_2 n_3^{1+\varepsilon})$, for any $\varepsilon > 0$, as asserted. $\square$

For each $i = 1, \ldots, n$, let Δ_i be the domain of f_i, and let J_i denote the z-projection (i.e., the orthogonal projection onto the z-axis) of Δ_i. Without loss of generality, we assume that J_i is a connected interval; otherwise we decompose f_i into $O(1)$ further partially-defined functions so that each of them satisfies this property. Let E denote the set of endpoints of intervals J_i, for $1 \le i \le n$. For the sake of simplicity, we assume that all the elements of E are distinct, but the analysis can be easily adapted to handle degenerate cases as well.

Let χ be the number of rainbow triples (f_i, f_l, f_k) so that $J_i \cap J_l \cap J_k \ne \emptyset$. This is our (rather weak) measure of the amount of overlap between the domains of the given functions. Our goal is to bound $|\Pi|$ in terms of χ.

Let $\Sigma \subseteq E \times \mathcal{F} \times \mathcal{F}$ be the set of triples (p, f_l, f_k) so that $p \in J_l \cap J_k$, and the function f_i for which p is an endpoint of J_i forms a rainbow triple with f_l and f_k. Set $\sigma_0 := |\Sigma| = 2\chi$.

We use a recursive scheme to bound $|\Pi|$. At each recursive step, we have an open z-interval Z. Let $\mathcal{F}_Z := \{f_i \in \mathcal{F} \mid J_i \cap Z \ne \emptyset\}$. A function $f_i \in \mathcal{F}_Z$ is called *short* in Z if at least one of the endpoints of J_i lies in Z; otherwise, f_i is *long*. Let $\mathcal{L}_Z$ (resp., $\mathcal{S}_Z$) denote the set of long (resp., short) functions in Z, and set $E_Z := E \cap Z$, $\Sigma_Z := \Sigma \cap (E_Z \times \mathcal{F}_Z \times \mathcal{F}_Z)$, and $\sigma_Z := |\Sigma_Z|$. Let $\Pi_Z \subseteq \Pi$ denote the subset of vertices $w \in \Pi$ whose z-coordinates lie in Z and are such that *at least two* of the four functions whose graphs are incident upon w are in $\mathcal{S}_Z$.

We bound the size of Π_Z recursively. Initially, we have $Z = (-\infty, \infty)$, $S_Z = \mathcal{F}$, $\mathcal{L}_Z = \emptyset$, $E_Z = E$, and $\Pi_Z = \Pi$. Let $\psi(m, \sigma) := \max |\Pi_Z|$, where the maximum is taken over all z-intervals Z with $|E_Z| \le m$ and $|\Sigma_Z| \le \sigma$.

We first describe the intuition behind our analysis. We wish to decompose the problem into subproblems, each involving certain subsets $\mathcal{F}'_1 \subseteq \mathcal{F}_1$, $\mathcal{F}'_2 \subseteq \mathcal{F}_2$, $\mathcal{F}'_3 \subseteq \mathcal{F}_3$, $\mathcal{F}'_4 \subseteq \mathcal{F}_4$, and apply the weak bound of Lemma 5.1 to bound the number of vertices in Π that are incident upon one function graph in each of $\mathcal{F}'_1, \mathcal{F}'_2, \mathcal{F}'_3, \mathcal{F}'_4$. The decomposition should satisfy the following properties:
(i) Every vertex in Π is counted in one of these subproblems.
(ii) The subsets $\mathcal{F}'_1, \mathcal{F}'_2, \mathcal{F}'_3, \mathcal{F}'_4$ should be such that the z-projections of the domains of every rainbow triple of functions in $\mathcal{F}'_1 \cup \mathcal{F}'_2 \cup \mathcal{F}'_3 \cup \mathcal{F}'_4$ overlap.
(iii) No triple counted in χ Boris says: ***** $\longrightarrow$Awkward wording, but was worse before. participates in more than a small number of subproblems (in the subsequent analysis, this number is polylogarithmic). Such a problem decomposition will be generated as we recursively analyze the sets Π_Z.

Let Z be a z-interval as above, with $|E_Z| \le m$ and $|\Sigma_Z| \le \sigma$. If $\mathcal{S}_Z = \emptyset$, then $|\Pi_Z| = 0$. Otherwise we partition Z into two intervals Z_1, Z_2 so that each subinterval contains half of the endpoints in E_Z. For $i = 1, 2$, set $\mathcal{F}_i := \mathcal{F}_{Z_i}$, $\mathcal{L}_i := \mathcal{L}_{Z_i}$, $\mathcal{S}_i := \mathcal{S}_{Z_i}$, $E_i := E_{Z_i}$, $\Sigma_i := \Sigma_{Z_i}$, and $\sigma_i := |\Sigma_i|$. We also denote by $\mathcal{N}_i \subseteq \mathcal{L}_i$ the set of functions that are long in Z_i but short in Z. Note that $\mathcal{S}_i \cup \mathcal{N}_i$ is the set of all functions $f_l \in \mathcal{S}_Z$ such that $J_l \cap Z_i \ne \emptyset$. Let $w \in \Pi_Z$ be a vertex whose z-coordinate lies in Z_1, say. If w is incident upon at least two functions in $\mathcal{S}_1$, then $w \in \Pi_{Z_1}$ and is counted recursively in the subproblem associated with Z_1. Otherwise, since w is assumed to be incident upon at least two short functions in $\mathcal{S}_Z$, w is incident upon *at least* one function in $\mathcal{N}_1$, *at most* one function in $\mathcal{S}_1$, and *at most* two functions in $\mathcal{L}_1 \setminus \mathcal{N}_1$. In fact, enumerating all possible cases, w must be incident upon one function in $\mathcal{N}_1$, another function in $\mathcal{S}_1 \cup \mathcal{N}_1$, and two functions in $\mathcal{L}_1$ (any other combination would pass w to the recursive subproblem at Z_1). Let $\Pi'_1 \subseteq \Pi_Z$ be the set of such vertices. We define Π'_2 for the other subinterval Z_2 in a fully analogous manner. We thus obtain the following recurrence for $\psi(m, \sigma)$.

$$\psi(m, \sigma) \le \begin{cases} \psi(m/2, \sigma_1) + \psi(m/2, \sigma_2) + |\Pi'_1| + |\Pi'_2|, & \sigma > 0, \\ 0, & \sigma = 0. \end{cases}$$

Since $E_1 \cap E_2 = \emptyset$, $\sigma_1 + \sigma_2 \le \sigma$. We now establish the following bound on $|\Pi'_1|$ and $|\Pi'_2|$.

LEMMA 5.2. *For* $i = 1, 2$, $|\Pi'_i| = O(\sigma_Z n^\varepsilon)$, *for any* $\varepsilon > 0$.

PROOF. Let $w \in \Pi'_1$ be a vertex incident upon $f_1 \in \mathcal{F}_1$, $f_2 \in \mathcal{F}_2$, $f_3 \in \mathcal{F}_3$, and $f_4 \in \mathcal{F}_4$. By definition, we may assume that, up to a permutation, $f_1, f_2 \in \mathcal{L}_1$, $f_3 \in \mathcal{N}_1$, and $f_4 \in \mathcal{S}_1 \cup \mathcal{N}_1$. Set

$$X_1 := \mathcal{F}_1 \cap \mathcal{L}_1, \qquad X_3 := \mathcal{F}_3 \cap \mathcal{N}_1,$$
$$X_2 := \mathcal{F}_2 \cap \mathcal{L}_1, \qquad X_4 := \mathcal{F}_4 \cap (\mathcal{S}_1 \cup \mathcal{N}_1).$$

By Lemma 5.1, the number of those rainbow vertices $w \in \Xi(\mathcal{F}^+, \mathcal{F}^-)$ that are incident upon one function in each of X_1, X_2, X_3, X_4 is $O(n^\varepsilon)$ times the product of the three largest sizes of the sets X_i, for any $\varepsilon > 0$. Hence this number is definitely upper bounded by $O(n^\varepsilon \sum_{i<j<k} |X_i||X_j||X_k|)$.

Let $f_i \in X_i$, $f_j \in X_j$, $f_k \in X_k$. Since at least two of these functions are long in Z_1, it follows that $J := J_i \cap J_j \cap J_k \neq \emptyset$. Moreover, since at least one of these functions is short in Z, J must have an endpoint p in Z, which is an endpoint of, say, J_k. We then charge (f_i, f_j, f_k) to the triple (p, f_i, f_j), which is counted in Σ_Z. It is easily seen that each triple in Σ_Z is charged at most $O(1)$ times in this manner. Hence we have:

$$\sum_{i<j<k} |X_i||X_j||X_k| = O(\sigma_Z).$$

Repeating this argument for all other $O(1)$ types of vertices in Π'_1 (obtained by permuting the sets $\mathcal{F}_1, \ldots, \mathcal{F}_4$, and/or the short/long classification of the four functions incident upon the vertex), we can conclude that $|\Pi'_1| = O(n^\varepsilon \sigma_Z)$, for any $\varepsilon > 0$. Using a fully symmetric argument, it follows that $|\Pi'_2| = O(n^\varepsilon \sigma_Z)$, for any $\varepsilon > 0$. This completes the proof of the lemma. □

Plugging the bounds on $|\Pi'_1|$ and $|\Pi'_2|$ into the recurrence, we obtain that

$$\psi(m, \sigma) \leq \begin{cases} \psi(m/2, \sigma_1) + \psi(m/2, \sigma_2) + O(n^\varepsilon \sigma), & \sigma > 0, \\ 0, & \sigma = 0, \end{cases}$$

where $\sigma = \sigma_1 + \sigma_2$. The solution of this recurrence is

$$\psi(m, \sigma) = O(n^\varepsilon \sigma \log m).$$

Initially, $Z = (-\infty, +\infty)$, $m = |E| = 2n$, and $\sigma_0 = |\Sigma| = 2\chi$. Therefore

$$|\Pi| \leq \Psi(2n, 2\chi) = O(n^\varepsilon \chi \log n) = O(\chi n^{\varepsilon'}),$$

for any $\varepsilon' > \varepsilon$. Putting everything together, we obtain the following result.

THEOREM 5.3. *Let* $\mathcal{F}$ *be a family of* n *partially defined trivariate functions of constant description complexity, and let* $\{\mathcal{F}_1, \mathcal{F}_2, \mathcal{F}_3, \mathcal{F}_4\}$ *be a partition of* $\mathcal{F}$. *Let* χ *denote the number of rainbow triples of functions in* $\mathcal{F}$ *so that the* z*-projections of their domains have a common intersection. Then, for any partition* $\mathcal{F}^+, \mathcal{F}^-$ *of* $\mathcal{F}$ *into two subsets, the number of rainbow vertices in the sandwich region* $\Xi(\mathcal{F}^+, \mathcal{F}^-)$ *between the upper envelope of* $\mathcal{F}^+$ *and the lower envelope of* $\mathcal{F}^-$ *is* $O(\chi n^\varepsilon)$, *for any* $\varepsilon > 0$.

Let $\mathfrak{F}_i^+, \mathfrak{F}_i^-$ be the sets of partial trivariate functions defined for the i-th strip in the previous section. Letting the θ-axis play the role of the z-axis, we obtain the following last ingredient in the proof of Theorem 1.1.

COROLLARY 5.4. *For any* $i \in \mathbb{Z}$, *let* $\mathfrak{F}_i^+, \mathfrak{F}_i^-$, *and* χ_i *be as defined in the previous section. The number of (deep) rainbow vertices in* $\Xi(\mathfrak{F}_i^+, \mathfrak{F}_i^-)$ *is* $O(n^\varepsilon \chi_i)$, *for any* $\varepsilon > 0$.

It is easy to extend Theorem 5.3 to the following "uncolored" version which is not needed for our analysis, but may be more appropriate for future applications.

THEOREM 5.5. *Let* $\mathcal{F}$ *be a set of* n *partially defined trivariate functions of constant description complexity. Let* $\mathcal{F}^+, \mathcal{F}^-$ *be a partition of* $\mathcal{F}$ *into two subsets. Let* χ *be the number of triples of functions in* $\mathcal{F}$ *so that the* z*-projections of their domains have a common intersection point. Then the number of vertices in* $\Xi(\mathcal{F}^+, \mathcal{F}^-)$ *is* $O(\chi n^\varepsilon)$.

PROOF. Partition $\mathcal{F}$ into four subfamilies $\mathcal{F}_1, \ldots, \mathcal{F}_4$ at random, assigning each function to each color class independently with probability $1/4$, argue that the expected number of rainbow vertices of $\Xi(\mathcal{F}^+, \mathcal{F}^-)$ is at least a constant fraction of the size of $\Xi(\mathcal{F}^+, \mathcal{F}^-)$ and invoke Theorem 5.3. The number χ of course also bounds the number of rainbow triples with overlapping domains. □

Remark 5.6 An interesting open problem is whether the theorems also hold when χ counts the potentially smaller number of triples of functions whose actual domains (as opposed to their z-projections) have a nonempty intersection. In fact, even the analogous problem for *bivariate* functions is still open: Can one express the complexity of the lower envelope of n partial bivariate functions of constant description complexity (or of the sandwich region between two envelopes) in terms of the number χ of pairs of functions whose domains have nonempty intersection? This can be done if χ counts the number of pairs of functions for which the y-projections of their domains overlap, using a considerably simpler variant of the preceding analysis.

6. Open Problems

The paper raises many interesting and challenging open problems, some of which have already been noted above:

(i) Devise an algorithm that constructs $\mathcal{F}(\mathcal{B})$ in near-cubic time. This would have immediate algorithmic implications for the ray shooting and motion planning problems mentioned in the introduction.

(ii) We note that our bound is not known to be nearly worst-case tight. The best (and trivial) construction of which we are aware yields a lower bound of $\Omega(n^2)$. It would therefore be interesting to close the gap between the upper and lower bounds.

(iii) It would be interesting to extend the analysis to families $\mathcal{B}$ of other classes of objects. The simplest such extensions would be to families of balls of arbitrary radii. Another extension would be to families of "fat" objects; see for example [10]. Our analysis extends to the case where the balls in $\mathcal{B}$ are "nearly congruent," in the sense that the ratio between the largest and the smallest radii is bounded by some fixed constant, but we do not know how to extend it to the case where there is no such restriction.

(iv) Finally, as mentioned in the introduction, a useful extension of our analysis would be to derive a near-cubic bound on the complexity of the space of free *segments* amid n unit balls (or other classes of objects).

Acknowledgments

The authors would like to express their appreciation of the city of Heidelberg, whose magic ambiance was the fertile soil in which, in March 2003, the initial seeds of this work were planted and began to grow.

References

[1] P. K. Agarwal, B. Aronov, and M. Sharir, Computing envelopes in four dimensions with applications. *SIAM J. Comput.* 26 (1997), 1714–1732.

[2] P. K. Agarwal, B. Aronov, and M. Sharir, Line transversals of balls and smallest enclosing cylinders in three dimensions. *Discrete Comput. Geom.* 21 (1999), 373–388.

[3] P. K. Agarwal and J. Matoušek, Ray shooting and parametric search. *SIAM J. Comput.* 22 (1993), 794–806.

[4] P. K. Agarwal and M. Sharir, Pipes, cigars, and kreplach: The union of Minkowski sums in three dimensions. *Discrete Comput. Geom.* 24 (2000), 645–685.

[5] B. Aronov, A. Efrat, V. Koltun and M. Sharir, On the union of κ-round objects in three and four dimensions, these proceedings.

[6] B. Aronov and M. Sharir, On translational motion planning of a convex polyhedron in 3-space, *SIAM J. Comput.* 26 (1997), 1785–1803.

[7] B. Aronov, M. Sharir, and B. Tagansky, The union of convex polyhedra in three dimensions, *SIAM J. Comput.* 26 (1997), 1670–1688.

[8] M. de Berg, *Efficient Algorithms for Ray Shooting and Hidden Surface Removal,* Ph.D. Dissertation, Dept. Comput. Sci., Utrecht Univ., 1992.

[9] M. de Berg, H. Everett, and L.J. Guibas, The union of moving polygonal pseudodiscs—combinatorial bounds and applications. *Computational Geometry: Theory and Applications* 11 (1998), 69–82.

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

[11] J.-D. Boissonnat, M. Sharir, B. Tagansky, and M. Yvinec, Voronoi diagrams in higher dimensions under certain polyhedral distance functions, *Discrete Comput. Geom.* 19 (1998), 485–519.

[12] B. Chazelle, H. Edelsbrunner, L. Guibas, M. Sharir, and J. Stolfi, Lines in space: Combinatorics and algorithms. *Algorithmica* 15 (1996), 428–447.

[13] O. Devillers, V. Dujmović, H. Everett, X. Goaoc, S. Lazard, H.-S. Na, and S. Petitjean, The expected number of 3D visibility events is linear. Rapport de recherche no. 4671, INRIA, 2002.

[14] F. Durand, G. Drettakis, and C. Puech, The visibility skeleton: A powerful and efficient multi-purpose global visibility tool. In *Proc. SIGGRAPH 97*, 89–100, 1997.

[15] F. Durand, G. Drettakis, and C. Puech, The 3D visibility complex. *ACM Transactions on Graphics*, 21 (2002), 176–206.

[16] A. Efrat and M. Sharir, The complexity of the union of fat objects in the plane, *Discrete Comput. Geom.* 23 (2000), 171–189.

[17] A. Efrat, The complexity of the union of (α, β)-covered objects, In *Proc. 15th Annu. ACM Sympos. Comput. Geom.*, 1999, pp. 134–142.

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

[19] V. Koltun and M. Sharir, The partition technique for overlays of envelopes, *SIAM J. Comput.* 32 (2003), 841–863.

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

[21] S. Mohaban and M. Sharir, Ray shooting amidst spheres in three dimensions and related problems, *SIAM J. Comput.* 26 (1997), 654–674.

[22] J. Pach, I. Safruti, and M. Sharir, The union of congruent cubes in three dimensions, *Discrete Comput. Geom.* 30 (2003), 133–160.

[23] M. Pellegrini, On lines missing polyhedral sets in 3-space. *Discrete Comput. Geom.* 12 (1994), 203–221.

[24] M. Pocchiola and G. Vegter, The visibility complex. *Internat. J. Comput. Geom. Appl.* 6 (1996), 279–308.

[25] M. Sharir, Algorithmic motion planning, In *Handbook of Discrete and Computational Geometry,* CRC Press LLC, Boca Raton, FL, 733–754, 1997. (A second edition is in press.)

[26] M. Sharir and P.K. Agarwal, *Davenport-Schinzel Sequences and Their Geometric Applications,* Cambridge University Press, Cambridge-New York-Melbourne, 1995.