

Union of Hypercubes and 3D Minkowski Sums with Random Sizes

Pankaj K. Agarwal¹

Department of Computer Science, Duke University, Durham, NC 27708, USA
pankaj@cs.duke.edu

Haim Kaplan²

School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel
haimk@tau.ac.il

Micha Sharir³

School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel
michas@tau.ac.il

Abstract

Let $T = \{\Delta_1, \dots, \Delta_n\}$ be a set of n pairwise-disjoint triangles in \mathbb{R}^3 , and let B be a convex polytope in \mathbb{R}^3 with a constant number of faces. For each i , let $C_i = \Delta_i \oplus r_i B$ denote the Minkowski sum of Δ_i with a copy of B scaled by $r_i > 0$. We show that if the scaling factors r_1, \dots, r_n are chosen randomly then the expected complexity of the union of C_1, \dots, C_n is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$; the constant of proportionality depends on ε and the complexity of B . The worst-case bound can be $\Theta(n^3)$.

We also consider a special case of this problem in which T is a set of points in \mathbb{R}^3 and B is a unit cube in \mathbb{R}^3 , i.e., each C_i is a cube of side-length $2r_i$. We show that if the scaling factors are chosen randomly then the expected complexity of the union of the cubes is $O(n \log^2 n)$, and it improves to $O(n \log n)$ if the scaling factors are chosen randomly from a “well-behaved” probability density function (pdf). We also extend the latter results to higher dimensions. For any fixed odd value of d , we show that the expected complexity of the union of the hypercubes is $O(n^{\lfloor d/2 \rfloor} \log n)$ and the bound improves to $O(n^{\lfloor d/2 \rfloor})$ if the scaling factors are chosen from a “well-behaved” pdf. The worst-case bounds are $\Theta(n^2)$ in \mathbb{R}^3 , and $\Theta(n^{\lfloor d/2 \rfloor})$ in higher dimensions.

2012 ACM Subject Classification Theory of computation → Computational geometry, Theory of computation → Generating random combinatorial structures

Keywords and phrases Computational geometry, Minkowski sums, Axis-parallel cubes, Union of geometric objects, Objects with random sizes

Digital Object Identifier 10.4230/LIPIcs.ICALP.2018.10

¹ Work on this paper by Pankaj Agarwal and Micha Sharir has been supported by Grant 2012/229 from the U.S.-Israel Binational Science Fund. Work by Pankaj Agarwal has also been supported by NSF under grants CCF-09-40671, CCF-10-12254, and CCF-11-61359, by ARO grants W911NF-07-1-0376 and W911NF-08-1-0452, and by an ERDC contract W9132V-11-C-0003.

² Work by Haim Kaplan has been supported by grant 1841-14 from the Israel Science Fund, and grant 1367-2017 from the German-Israeli Science Fund.

³ Work by Micha Sharir has also been supported by Grant 892-13 from the Israel Science Fund, by the Blavatnik Research Fund in Computer Science at Tel Aviv University, and by the Hermann Minkowski-MINERVA Center for Geometry at Tel Aviv University.



1 Introduction

We advance the state of the art regarding the complexity of the union of combinatorial objects in dimensions $d \geq 3$ beyond worst-case analysis. To that end we study the union complexity of Minkowski sums of pairwise disjoint triangles with a randomly scaled copies of a fixed convex polytope, and Minkowski sums of points with cubes of random sizes. We hope that our techniques will be useful in future studies of such stochastic arrangements.

Specifically, let $T = \{\Delta_1, \dots, \Delta_n\}$ be a collection of n pairwise-disjoint triangles in \mathbb{R}^3 , and let B be a fixed convex polytope in \mathbb{R}^3 with a constant number of faces. We consider the setup where we are given a sequence $\mathbf{r} = \langle r_1, \dots, r_n \rangle$ of non-negative scaling factors (or *sizes*), and we let $C_i = \Delta_i \oplus r_i B$ denote the Minkowski sum of Δ_i with a copy of B scaled by r_i , for $i = 1, \dots, n$. Each C_i is a convex polytope with a constant number of faces. Let $\mathcal{C} = \{C_1, \dots, C_n\}$, and let $\mathcal{U} = \mathcal{U}(T, \mathbf{r}) = \bigcup_{i=1}^n C_i$ denote their union. We also use $\mathcal{U}(\mathcal{C})$ to denote $\mathcal{U}(T, \mathbf{r})$.

The *combinatorial complexity* of \mathcal{U} is the total number of faces of all dimensions on the boundary $\partial\mathcal{U}$ of \mathcal{U} . Each vertex of $\partial\mathcal{U}$ is a vertex of some set C_i , an intersection point between an edge of some C_i and a face of another C_j , or an intersection point of three faces of three distinct C_i 's. By Euler's formula, the overall complexity of $\partial\mathcal{U}$ is proportional to the number of vertices on $\partial\mathcal{U}$. Therefore we measure the combinatorial complexity of \mathcal{U} by the number of its vertices, and denote this quantity by $\psi(T, \mathbf{r})$.

Our goal is to obtain an upper bound on the *expected* value of $\psi(T, \mathbf{r})$, under a suitable stochastic model for choosing the scaling factors \mathbf{r} for the members of \mathcal{C} . We refer to this problem as the *stochastic Minkowski-sum union* problem. Our expected bounds are significantly better than the worst case bounds (for any such union), indicating that inputs which are not adversarial are likely to have lower union complexity.

Another motivation for our analysis is for an efficient computation of the most vulnerable location of an attack in a three-dimensional scene. Concretely, we use the model where we have a scene consisting of objects, modeled as a collection of pairwise disjoint triangles. An attack occurs at some point, and the probability of a triangle to be hit decreases as its distance from the attacking point increases (up to some threshold). We want to compute a point of attack in which we maximize the expected number of triangles that we hit. The same technique to approximately solve this problem, as done in the planar case [1, 2], leads to questions about the complexity of the union of Minkowski sums of the input triangles with balls of random radii. The case studied here replaces the ball by a convex polytope, which we choose to approximate the Euclidean ball. The results obtained in this paper lead to improved solution to the vulnerability problem, by roughly one factor of n .

We consider two stochastic models for choosing the sequence $\mathbf{r} = \langle r_1, \dots, r_n \rangle$ of scaling factors:

The density model. We are given an arbitrary probability density (or mass) function (pdf) π over the non-negative reals; for each $1 \leq i \leq n$, we take r_i to be a random value drawn independently from the distribution determined by π .

The permutation model. We are given a sequence $\Theta = \langle \theta_1, \dots, \theta_n \rangle$ of n arbitrary non-negative real numbers, with $\theta_1 \leq \theta_2 \leq \dots \leq \theta_n$. We draw a random permutation σ on $[1 : n]$, where all permutations are equally likely to be chosen, and assign $r_i := \theta_{\sigma(i)}$ to Δ_i , for each $i = 1, \dots, n$.

For the density model, let $\psi(T, \pi)$ denote the expected value of $\psi(T, \mathbf{r})$, where the expectation is taken over the random choices of \mathbf{r} , drawn from π in the manner specified above. Set $\psi(n) = \max \psi(T, \pi)$, where the maximum is taken over all probability density (mass) functions and over all sets T of n pairwise-disjoint triangles in \mathbb{R}^3 . For the permutation model, in an analogous manner, we let $\psi(T, \Theta)$ denote the expected value of $\psi(T, \mathbf{r})$, where the expectation is taken over the choices of \mathbf{r} , obtained by randomly shuffling the members of Θ . Then, with a slight overloading of the notation, we define $\psi(n) = \max \psi(T, \Theta)$, where the maximum is over all possible choices of the multi-set Θ and over all sets T of n pairwise-disjoint triangles in \mathbb{R}^3 . Our goal is to prove sharp bounds on $\psi(n)$ under these two models.

As noted in [2], the permutation model is more general than the density model, in the sense that an upper bound on $\psi(n)$ in the permutation model immediately implies the same bound on $\psi(n)$ in the density model. Although we mostly focus on the permutation model, we also show that some of our bounds can be improved, by a logarithmic factor, in the density model if the underlying pdf is “well behaved,” in a sense to be made precise below.

Related work. There is extensive work on bounding the complexity of the union of a set of geometric objects, especially in \mathbb{R}^2 and \mathbb{R}^3 , and optimal or near-optimal bounds have been obtained for many interesting cases. We refer the reader to the survey paper by Agarwal *et al.* [18] for a comprehensive summary of most of the known results on this topic.

The complexity of the union of n arbitrary shapes of constant complexity in \mathbb{R}^2 is $\Theta(n^2)$, but the bound improves to near linear for a large class of well behaved planar objects; see [18]. Analogously, the complexity of the union of n arbitrary shapes of constant complexity in \mathbb{R}^3 is $\Theta(n^3)$. Over more than a decade, a series of papers have considered geometric objects in \mathbb{R}^3 that have some special properties, and derived *near-quadratic* bounds on the complexity of their union. These cases include: a family of arbitrary convex polytopes⁴ [8], a family of cylinders (of arbitrary radii) [12], (arbitrarily aligned) congruent cubes in three dimensions [19], a family of fat tetrahedra [13], a family of κ -round objects [4], and Minkowski sums of a family of pairwise-disjoint convex polytopes with a fixed convex polytope [7] or with a fixed ball [3]. Quadratic lower bounds are known for all these cases.

The case of the union of axis-parallel cubes is a highly-structured special instance. If the cubes are *isothetic* (that is, they are all congruent), the complexity of their union is⁵ $O(n)$. In dimension d , Boissonnat *et al.* [10] proved that the complexity of the union of n isothetic hypercubes is $\Theta(n^{\lfloor d/2 \rfloor})$. If the hypercubes are (axis-parallel but) of arbitrary sizes, the complexity of their union is $\Theta(n^{\lceil d/2 \rceil})$ in \mathbb{R}^d . These two bounds coincide for even values of d , but there is a gap, by a factor of d , for odd values of d .

There is a rich literature on bounding the complexity of geometric structures under a stochastic model in which the locations of points are drawn randomly from a distribution; see [14, 20, 21] and references therein. The complexity of the union of a set of objects in a semi-stochastic model, in which the locations of the objects were arbitrary but their scaling factors were chosen randomly, was first studied by the authors (with Har-Peled) in [2]. They investigated two planar variants of the stochastic Minkowski-sum union problem. In the first variant, one is given a set S of n pairwise-disjoint line segments in \mathbb{R}^2 , and one replaces each $e \in S$ by the Minkowski sum $e \oplus r_i B$, where B is the unit disk and the scaling factors r_i are

⁴ Here the bound is cubic in the number of polytopes but is only near linear in the number of facets.

⁵ In contrast, the complexity of the union of congruent balls in \mathbb{R}^3 is quadratic in the worst case; see, e.g., [18].

randomly chosen under either of the above two models. It is shown in [2] that the expected complexity of the union of these sums is $O(n \log n)$. In the second variant, S is a collection of arbitrary pairwise-disjoint convex sets (of constant complexity) in the plane, and then the expected complexity of the union of the corresponding random Minkowski sums (again, with randomly scaled copies of the unit disk) is shown in [2] to be $O(n^{1+\varepsilon})$, for any $\varepsilon > 0$. In both cases, the bounds are a significant improvement over the worst case quadratic bound, and almost match the linear upper bound when all scaling factors are equal—the Minkowski sums then form a collection of *pseudo-disks*. Here we study related random arrangements in higher dimensions, a situation that required some new ideas in order to apply some existing techniques.

In a different, but closely related context, Har-Peled and Raichel [16] proved that the expected complexity of the multiplicatively weighted Voronoi diagram of a set of points or line segments in \mathbb{R}^2 is $O(n \text{polylog}(n))$ if the weights are randomly chosen under the permutation model.⁶ Recall that if the weights are arbitrarily chosen, then the worst-case complexity of the weighted Voronoi diagram is quadratic [9]. Chang *et al.* [11] studied various generalizations of multiplicatively weighted Voronoi diagrams and proved sharp bounds on their expected complexity.

Our results. We have two main results in this paper:

Union of hypercubes. We first study (in Section 2) the interesting special case of the stochastic Minkowski-sum union problem in which each triangle is a point in \mathbb{R}^3 and B is an axis-aligned cube in \mathbb{R}^3 of side-length 2. That is, we have a set $P = \{p_1, \dots, p_n\}$ of n points in \mathbb{R}^3 , and $B = \{x \in \mathbb{R}^3 \mid \|x\|_\infty \leq 1\}$, and for each $i \leq n$, $C_i = r_i B + p_i$ is a cube of side-length $2r_i$ centered at p_i . In other words, we study the complexity of the union of n cubes C_i each having a random (according to each of our models) L_∞ diameter and a fixed center. In fact, we study this problem not only in \mathbb{R}^3 but in \mathbb{R}^d for general d . The reader should note that all our definitions (e.g. of $\psi(P, \Theta)$) extend to \mathbb{R}^d . Hypercubes in \mathbb{R}^d arise in many applications and the following result indicates that if their diameters are not chosen in an adversarial manner then the complexity of their union is likely to be smaller than the worst case, by nearly a factor of n . Specifically, we prove:

► **Theorem 1.1.** *Let P be a set of n points in \mathbb{R}^d , let B be the axis-aligned hypercube of side-length 2 in \mathbb{R}^d centered at the origin, and let Θ be a multi-set of scaling factors. Under the permutation model, $\psi(P, \Theta) = O(n \log^2 n)$ for $d = 3$, and $\psi(P, \Theta) = O(n^{\lfloor d/2 \rfloor} \log n)$ for any fixed odd value of $d > 3$. The same bounds hold for the density model.*

Tame distributions. We call a pdf π , with Π as its cdf (cumulative distribution function), *tame* if for every $x > 0$ and for every integer $k \geq 0$,

$$1 - \Pi(kx) \leq (1 - \Pi(x))^k. \quad (1)$$

It is an easy exercise to verify that (1) is satisfied for a large class of pdfs, including uniform, exponential, (one-sided) normal, and log-normal distributions. If the scaling factors are chosen from a tame pdf, then the bounds can be improved by a logarithmic factor:

► **Theorem 1.2.** *Let P be a set of n points in \mathbb{R}^d , and let B be the hypercube of side-length 2 centered at the origin. If the scaling factors are drawn randomly from a tame pdf π , then $\psi(P, \pi) = O(n \log n)$ for $d = 3$ and $\psi(P, \pi) = O(n^{\lfloor d/2 \rfloor})$ for any fixed odd value of $d > 3$.*

⁶ Given a point set $P = \{p_1, \dots, p_n\}$ and a weight $w_i > 0$ for each point p_i , the Voronoi cell of p_i in the multiplicatively weighted Voronoi diagram is $\{x \in \mathbb{R}^d \mid w_i \|x - p_i\| \leq w_j \|x - p_j\| \forall j \leq n\}$.

By Theorems 1.1 and 1.2, the expected complexity of the union of a set of axis-aligned hypercubes, whose sizes are chosen randomly, is almost the same as when their sizes are equal. The overall structure of the proof of both theorems is the same, and consists of two stages. The first stage bounds the expected number of *outer* vertices on the union, i.e., the vertices that lie on a $(d - 2)$ -face or a lower-dimensional face of one of the hypercubes. This is where we had to develop a new technique to exploit the randomness of the input. The second stage then obtains a bound on the number of *inner* vertices, namely vertices formed by the intersection of d facets of d distinct hypercubes, by adapting the argument in [10] and using the bound, derived earlier, on the expected number of outer vertices.

The stochastic Minkowski-sum union problem. Next, we study (in Section 3) the general case in which T is a set of pairwise disjoint triangles in \mathbb{R}^3 and B is a convex polytope in \mathbb{R}^3 with $O(1)$ faces, and prove the following near-quadratic upper bound:

► **Theorem 1.3.** *Let T be a collection of n pairwise-disjoint triangles in \mathbb{R}^3 , let Θ be a multi-set of scaling factors, and let B be a convex polytope with $O(1)$ vertices. Then the value of $\psi(T, \Theta)$ under the permutation model is $O(n^{2+\varepsilon})$, for any fixed constant $\varepsilon > 0$; the constant of proportionality depends on ε and the complexity of B . The same bound holds for the density model.*

If all sizes r_i are equal, the complexity of \mathcal{U} is $O(n^2 \log n)$ [7], where the constant of proportionality depends on the complexity of B . On the other hand, for a bad layout of the triangles of T and a bad (non-random) choice of sizes, the complexity of \mathcal{U} can be $\Theta(n^3)$; see Section 3 for a lower-bound construction. (The bound in Theorem 1.3 is coarse enough so that drawing the scaling factors from a tame pdf does not seem to affect it.)

2 Union of Hypercubes

In this section we prove Theorems 1.1 and 1.2. For a point $x \in \mathbb{R}^d$, we use $\|x\|_\infty$ to denote the L_∞ -norm of x . For a compact set $S \subset \mathbb{R}^d$, let $d_\infty(x, S) = \min_{y \in S} \|x - y\|_\infty$. A hypercube C of L_∞ -radius (or *radius* for brevity) r centered at c is the set $C = \{x \in \mathbb{R}^d \mid \|x - c\|_\infty \leq r\}$; the *side-length* of this hypercube is $2r$. Let $P = \{p_1, \dots, p_n\}$ be a set of n points in \mathbb{R}^d , and let B be the (axis-aligned) hypercube of radius 1 centered at the origin. Given a sequence $\mathbf{r} = \langle r_1, \dots, r_n \rangle$ of non-negative scaling factors, we define, for each $i \leq n$, $C_i = r_i B + p_i$ to be a hypercube of radius r_i centered at p_i . Let $\mathcal{C} = \{C_1, \dots, C_n\}$, and let $\mathcal{U} = \mathcal{U}(\mathcal{C}) = \bigcup_{i=1}^n C_i$ denote the union of the hypercubes.

For simplicity, we assume that the hypercubes of \mathcal{C} are in *general position*, in the sense that no two facets of any pair of hypercubes lie in a common hyperplane. This assumption is made only to simplify the analysis, and the general, possibly degenerate case can be reduced to the general-position setup by means of a suitable small perturbation.

A vertex v in the arrangement $\mathcal{A}(\mathcal{C})$ of \mathcal{C} is the intersection of d facets of the hypercubes in \mathcal{C} and each of these facets is orthogonal to a different coordinate axis. These facets need not belong to distinct hypercubes. We refer to v as an *inner* vertex if the d facets on which it lies belong to d distinct hypercubes, and as an *outer* vertex if at least two of these facets belong to the same hypercube. The *level* of v is j if it lies in the interior of exactly j hypercubes of \mathcal{C} ; the level of every vertex on $\partial\mathcal{U}$ is 0. For $j \geq 0$, let $\mathcal{J}_j(\mathcal{C})$ (resp. $\mathcal{O}_j(\mathcal{C})$) denote the number of inner (resp. outer) vertices at level j in $\mathcal{A}(\mathcal{C})$.

If the sequence \mathbf{r} of scaling factors is chosen randomly, using either the permutation model or the density model, $\mathcal{J}_j(\mathcal{C})$ and $\mathcal{O}_j(\mathcal{C})$ are random variables. For a sequence Θ of scaling factors, let $\bar{\mathcal{J}}_j(P, \Theta) = \mathbb{E}[\mathcal{J}_j(\mathcal{C})]$ denote the expected value of $\mathcal{J}_j(\mathcal{C})$ in the permutation

model, over a random permutation of Θ . For any $n > 0$ and for any $j \leq n - d$, let $\bar{J}_j(n) = \max \bar{J}_j(P, \Theta)$ denote the expected number of inner vertices of level j under the permutation model, where the maximum is taken over all point sets of size n . We define $\bar{O}_j(P, \Theta)$ and $\bar{O}_j(n)$ similarly. By definition (recall the terminology introduced earlier), $\psi(P, \Theta) = \bar{J}_0(P, \Theta) + \bar{O}_0(P, \Theta)$ and $\psi(n) \leq \bar{J}_0(n) + \bar{O}_0(n)$.

Finally, in the density model, which we use explicitly only when the scaling factors are chosen randomly from a tame pdf π , we define $\bar{J}_j(P, \pi)$ and $\bar{O}_j(P, \pi)$ as the expected number of inner and outer level- j vertices, respectively, and we also define $\bar{J}_j(n)$ and $\bar{O}_j(n)$ as the maximum value of these respective quantities, where the maximum is taken over all sets P of n points and all tame pdf's.⁷

Our goal is to bound $\bar{J}_0(n)$ and $\bar{O}_0(n)$. Our strategy is first to derive an upper bound for $\bar{O}_0(n)$, either for the permutation model or for the density model for a tame pdf, and then use either of these bounds in a charging scheme that leads to a corresponding bound on $\bar{J}_0(n)$. The following lemma will be crucial in proving the bound on $\bar{J}_0(n)$.

► **Lemma 2.1.** $\bar{O}_1(n) = O(\bar{O}_0(n))$.

To keep the presentation simple, we first bound $\bar{O}_0(n)$ in \mathbb{R}^3 (Section 2.1) and then extend the argument to higher dimensions (Section 2.2). Finally, we prove an upper bound on $\bar{J}_0(n)$ in \mathbb{R}^d , for any odd dimension $d \geq 3$ (Section 2.3).

2.1 Outer vertices in 3D

An outer vertex of $\mathcal{U}(\mathcal{C})$ in \mathbb{R}^3 is either a vertex of a cube in \mathcal{C} or the intersection point of an edge of a cube C_i with a face of another cube C_j . There are $O(n)$ vertices of the first type, so it suffices to bound the expected number of the second type of outer vertices.

We fix a point, say p_1 , of P and an edge e of the cube $C_1 = p_1 + r_1 B$ centered at p_1 . We bound the expected number of outer vertices of \mathcal{U} lying on e . Note that e is not a fixed segment in \mathbb{R}^3 —its physical location in \mathbb{R}^3 as well as its length depend on the value of r_1 . Nevertheless, we can combinatorially define e so that it refers to a fixed edge (one of the 12 possible edges) of C_1 . For simplicity, we assume that e is parallel to the z -axis. We bound the expected number of vertices lying on e , first for the permutation model, and then for the density model for a tame pdf.

The permutation model. We are given a sequence $\Theta = \langle 0 \leq \theta_1 \leq \dots \leq \theta_n \rangle$ of n arbitrary non-negative scaling factors. We choose a random permutation σ on $[1 : n]$, and assign $r_i := \theta_{\sigma(i)}$.

Let $\mu(e, P, \Theta)$ denote the expected number of connected components of $\partial\mathcal{U} \cap e$. The expected number of outer vertices on e is at most $2\mu(e, P, \Theta)$. For a fixed value $\theta_i \in \Theta$ for r_1 , let $\mu(e, P, \Theta \mid r_1 = \theta_i)$ denote the conditional expected number of connected components of $\partial\mathcal{U} \cap e$ assuming that $r_1 = \theta_i$. With this fixed choice of r_1 , e is a segment of length $2\theta_i$, lying at a fixed location in \mathbb{R}^3 . Observe that

$$\mu(e, P, \Theta) = \frac{1}{n} \sum_{i=1}^n \mu(e, P, \Theta \mid r_1 = \theta_i). \quad (2)$$

The following probabilistic lemma is the main technical tool used in the analysis in \mathbb{R}^3 .

⁷ One can also take the maximum over all pdfs, but then the various expectations are at most the corresponding quantities under the permutation model, and no better bounds are known so far. We will therefore be interested only in tame pdfs.

► **Lemma 2.2.** *For any $1 \leq k \leq n$, $\mu(e, P, \Theta \mid r_1 = \theta_{n-k+1}) \leq \frac{4n}{k}$.*

Proof. Let $r = \theta_{n-k+1}$. Let e^+ (resp. e^-) denote the top (resp. bottom) half of e , i.e., the portion lying above (resp. below) p_1 . Further partition e^+ into two parts e_1^+ and e_2^+ , each of length $r/2$, and similarly partition e^- . We bound the expected number of connected components of \mathcal{U} lying on e_1^+ , e_2^+ , e_1^- and e_2^- separately. Let $e' = e_1^+$, the bounds for the other parts of e are derived analogously.

Let $P(e^*) = \langle p_{i_1}, p_{i_2}, \dots, p_{i_{n-1}} \rangle$ denote the sorted sequence of the points of $P \setminus \{p_1\}$ by increasing values of $\delta_i = d_\infty(p_i, e')$. We fix an assignment r_2, \dots, r_n of scaling factors (which is a permutation of $\Theta \setminus \{r\}$) to the points of $P \setminus \{p_1\}$, so that p_i is assigned the scaling factor r_i , for $i = 2, \dots, n$. Let p_{i_j} be the first point in the sequence $P(e^*)$ for which $r_{i_j} \geq r$. If $\delta_{i_j} \leq r/2$ then, since $|e'| = r/2$, it follows by the triangle inequality that $e' \subseteq C_{i_j}$ and there are no connected components of \mathcal{U} on e' .

So assume that $\delta_{i_j} > r/2$. Then, since the sequence $P(e^*)$ is sorted by distance to e' , it follows that $\delta_{i_\ell} > r/2$ for every $\ell \geq j$. Hence, any cube C_{i_ℓ} , for $\ell \geq j$, will intersect e' only if $r_{i_\ell} \geq r/2$. But then $e' \cap C_{i_\ell}$ must contain an endpoint of e' . It follows that all the cubes C_{i_ℓ} , for $\ell \geq j$, can contribute at most one connected component to $\partial\mathcal{U} \cap e'$. The cubes C_{i_ℓ} , for $\ell < j$, can increase the number of components by at most $j - 1$, so in total we get at most j connected components of \mathcal{U} on e' .

As argued in [2, Lemma 3.3] (see also [15]), the expected value of the index j of the first appearance of one of $k - 1$ distinguished elements in a random permutation of $n - 1$ elements is n/k , so the expected value of j is n/k . This gives an average of at most n/k connected components of $\partial\mathcal{U} \cap e'$. From this the lemma follows. ◀

Plugging Lemma 2.2 into (2), we obtain $\mu(e, P, \Theta) = O(\log n)$. Summing this bound over all $O(n)$ edges of cubes in \mathcal{C} and using Lemma 2.1, we obtain the following.

► **Corollary 2.3.** *Let P be a set of n points in \mathbb{R}^3 , and let Θ be a multi-set of n non-negative scaling factors. Then $\bar{\Theta}_0(P, \Theta), \bar{\Theta}_1(P, \Theta) = O(n \log n)$.*

The density model for tame distributions. Next, we show that if the scaling factors of the cubes in \mathcal{C} are chosen randomly and independently from a tame pdf then the expected number of outer vertices on a fixed edge g of any cube, say for concreteness the cube C_1 centered at p_1 , is only $O(1)$. For simplicity, we assume that p_1 is at the origin. Let π be a tame pdf and Π its cdf. By definition, $1 - \Pi(kx) \leq (1 - \Pi(x))^k$ for all $x > 0$ and for each $k \geq 1$. Adapting the previous notation, let $\mu(g, P, \pi)$ denote the expected number of connected components on $\partial\mathcal{U} \cap g$.

► **Lemma 2.4.** $\mu(g, P, \pi) \leq 2e^3/(e - 1)$.

Proof. Assume, as before, that g is parallel to the z -axis, and let g^+ (resp. g^-) denote the top (resp. bottom) half of g , i.e., the portion of g lying above (resp. below) p_1 . Recall that g^+ is a random segment whose length and location depend on the value of r_1 , the scaling factor assigned to p_1 . We bound $\mu(g^+, P, \pi)$, the expected number of connected components of $\partial\mathcal{U} \cap g^+$. A symmetric argument will bound $\mu(g^-, P, \pi)$, and the sum of these two quantities will bound $\mu(g, P, \pi)$.

Set $\alpha = \Pi^{-1}(1 - 1/e)$ and let $I_k = [k\alpha, (k + 1)\alpha)$, for $k = 0, 1, \dots$. For a fixed value of r_1 we partition g^+ into $\ell + 1$ intervals, $\Delta_0, \dots, \Delta_\ell$, such that, for $k < \ell$, the projection of Δ_k on the z -axis is I_k and the projection of Δ_ℓ is contained in I_ℓ . For convenience we define Δ_k , for $k > \ell$, to be an empty interval. Note that ℓ is a random variable whose value depends

on r_1 . Let $\mu(\Delta_k) := \mu(\Delta_k, P, \pi)$ denote the expected number of connected components of $\partial\mathcal{U} \cap \Delta_k$. Then

$$\mu(g^+, P, \pi) \leq \sum_{k=0}^{\infty} \mu(\Delta_k). \quad (3)$$

The probability that Δ_k is not empty is equal to $\Pr[r_1 \geq k\alpha]$, which is

$$\Pr[r_1 \geq k\alpha] = 1 - \Pi(k\alpha) \leq (1 - \Pi(\alpha))^k = e^{-k}. \quad (4)$$

Fix a value of r_1 for which Δ_k is not empty. Let $P(\Delta_k) = \langle p_{i_1}, p_{i_2}, \dots, p_{i_{n-1}} \rangle$ denote the sorted sequence of the points of $P \setminus \{p_1\}$ by increasing values of $\delta_i = d_\infty(p_i, \Delta_k)$. Fix an assignment r_2, \dots, r_n of scaling factors to the points of $P \setminus \{p_1\}$, so that p_i is assigned the scaling factor r_i , drawn from π , for $i = 2, \dots, n$.

Let p_{i_j} be the first point in the sequence $P(\Delta_k)$ for which $r_{i_j} \geq 2\alpha$. If $\delta_{i_j} \leq \alpha$ then since $|\Delta_k| \leq \alpha$ it follows by the triangle inequality that $\Delta_k \subseteq C_{i_j}$, and there are no connected components of \mathcal{U} on Δ_k .

So let us assume that $\delta_{i_j} > \alpha$. Then, since the sequence $P(\Delta_k)$ is sorted by distance to Δ_k , it follows that $\delta_{i_\ell} \geq \alpha$ for every $\ell \geq j$. Thus, for a cube C_{i_ℓ} , for $\ell \geq j$, to intersect Δ_k , we must have that $r_{i_\ell} \geq \alpha$. But then $\Delta_k \cap C_{i_\ell}$ must contain an endpoint of Δ_k . It follows that all the cubes C_{i_ℓ} , for $\ell \geq j$, can contribute at most one connected component to $\partial\mathcal{U} \cap \Delta_k$. The cubes C_{i_ℓ} , for $\ell < j$, can increase the number of components by at most $j - 1$, so in total we get at most j connected components of \mathcal{U} on Δ_k .

Notice that j is a geometric random variable, where the success probability of each trial is $p = 1 - \Pi(2\alpha) \leq (1 - \Pi(\alpha))^2 = 1/e^2$. The expectation of j is thus $1/p = e^2$. Therefore $\mu(\Delta_k \mid \Delta_k \neq \emptyset) \leq e^2$.

Combining this with the bound on the probability that Δ_k is not empty, in Equation (4), we get that $\mu(\Delta_k) = \Pr[\Delta_k \neq \emptyset] \cdot \mu(\Delta_k \mid \Delta_k \neq \emptyset) \leq e^{-k+2}$. Substituting these bounds into Equation (3), the lemma follows. \blacktriangleleft

Putting everything together and using Lemma 2.1, we obtain the following:

► **Lemma 2.5.** *Let P be a set of n points in \mathbb{R}^3 and π a tame pdf. Then $\bar{\Theta}_0(P, \pi), \bar{\Theta}_1(P, \pi) = O(n)$.*

2.2 Outer vertices in higher odd dimensions

We now bound $\bar{\Theta}_0(n)$ in \mathbb{R}^d , for any fixed odd value $d > 3$. We fix a point of P , say, p_1 , and a $(d - 2)$ -dimensional face f of C_1 . We bound the expected number of outer vertices lying on f . As in the previous case, we first obtain the bound for the permutation model and then for the density model (with a tame pdf).

The permutation model. Let Θ denote a fixed set of scaling factors, and let $\mu(f, P, \Theta)$ denote the expected number of vertices of \mathcal{U} lying on f . We bound the conditional expectation $\mu(f, P, \Theta \mid r_1 = \theta_i)$, assuming that the scaling factor of p_1 is θ_i , and then use (2) to bound $\mu(f, P, \Theta)$.

Recall that once the value r_1 is fixed to θ_i , the hypercube $C_1 = p_1 + r_1 B$ is also fixed, and so is its face f . The problem is therefore to bound the expected number of outer vertices lying on a fixed $(d - 2)$ -dimensional hypercube f . A naive approach, based on induction on d within f does not work. We use instead a different approach.

The following lemma is analogous to Lemma 2.2.

► **Lemma 2.6.** For any $1 \leq k \leq n$, $\mu(f, P, \Theta \mid r_1 = \theta_{n-k+1}) = O\left(\frac{n^{\lfloor d/2 \rfloor}}{k}\right)$.

Proof. The proof is similar to that of Lemma 2.2. Let $r = \theta_{n-k+1}$. We partition f into 4^{d-2} hypercubes (subfaces) each of radius $r/4$. Fix such a smaller hypercube g . Let $\mu(g, P, \Theta \mid r_1 = r)$ be the expected number of outer vertices of \mathcal{U} incident to g , conditioned on the choice of r_1 .

Let $\delta_i = d_\infty(p_i, f)$ and let $P(g) = \langle p_{i_1}, p_{i_2}, \dots, p_{i_{n-1}} \rangle$ denote the sorted sequence of the points of $P \setminus \{p_1\}$ by increasing values of $\delta_i = d_\infty(p_i, g)$. We fix an assignment r_2, \dots, r_n of scaling factors from $\Theta \setminus \{r\}$ to the points of $P \setminus \{p_1\}$, so that p_i is assigned the scaling factor r_i , for $i = 2, \dots, n$. Let $\hat{\mathcal{C}} = \{C_i := p_i + r_i B \mid p_i \in P \setminus \{p_1\}\}$ be the resulting set of hypercubes. Let $U_g = \mathcal{U}(\hat{\mathcal{C}}) \cap g$. Our goal is to bound the number of vertices of U_g . Let F be the $(d-2)$ -dimensional flat spanned by g . For a hypercube $C_i \in \hat{\mathcal{C}}$, let $K_i = C_i \cap F$. If nonempty, K_i is a $(d-2)$ -dimensional hypercube of radius r_i .

Let p_{i_j} be the first point in the sequence $P(g)$ for which $r_{i_j} \geq r$. Set $\mathcal{K}_< = \{K_{i_\ell} \mid \ell < j \wedge K_{i_\ell} \neq \emptyset\}$ and $\mathcal{K}_\geq = \{K_{i_\ell} \mid \ell \geq j \wedge K_{i_\ell} \neq \emptyset\}$. By definition, $|\mathcal{K}_<| < j$ and $U_g = \mathcal{U}(\mathcal{K}_< \cup \mathcal{K}_\geq) \cap g$.

If $\delta_{i_j} \leq r/2$ then, since the side-length of g is $r/2$, it follows by the triangle inequality that $f \subseteq C_{i_j}$ and the number of vertices in U_g is 0. So assume that $\delta_{i_j} > r/2$. Then, since the p_i 's are sorted by their distance to g , it follows that $\delta_{i_\ell} > r/2$ for every $\ell \geq j$. Hence, for a cube C_{i_ℓ} with $\ell \geq j$ to intersect g we must have $r_{i_\ell} > r/2$.

► **Lemma 2.7.** Let g be a $(d-2)$ -dimensional hypercube of radius $r/4$. For any $(d-2)$ -dimensional hypercube K of radius at least $r/4$, there is another hypercube \tilde{K} of radius exactly $r/4$ such that $K \cap g = \tilde{K} \cap g$.

Since all hypercubes in \mathcal{K}_\geq have radius at least $r/4$, we obtain, by applying Lemma 2.7 to all of them, a collection $\tilde{\mathcal{K}}_\geq = \{\tilde{K}_i \mid K_i \in \mathcal{K}_\geq\}$ such that $\mathcal{U}(\mathcal{K}_\geq) \cap g = \mathcal{U}(\tilde{\mathcal{K}}_\geq) \cap g$. It follows that $U_g = \mathcal{U}(\mathcal{K}_< \cup \tilde{\mathcal{K}}_\geq) \cap g$. It therefore suffices to bound the number of vertices of $\tilde{U}_g = \mathcal{U}(\mathcal{K}_< \cup \tilde{\mathcal{K}}_\geq)$, and also the number of vertices of $\tilde{U}_g \cup g$ (in order to upper bound the number of vertices in U_g which are on the boundary of g). We bound the number of vertices of \tilde{U}_g , the analysis for $\tilde{U}_g \cup g$ is similar.

We call a vertex v of \tilde{U}_g *pure* if it does not lie on the boundary of any hypercube in $\mathcal{K}_<$, otherwise we call it *mixed*. A pure vertex is also a vertex of $\mathcal{U}(\tilde{\mathcal{K}}_\geq)$. Since $\tilde{\mathcal{K}}_\geq$ is a set of at most n congruent hypercubes in \mathbb{R}^{d-2} , the number of vertices in $\mathcal{U}(\tilde{\mathcal{K}}_\geq)$ is $O(n^{\lfloor (d-2)/2 \rfloor}) = O(n^{\lfloor d/2 \rfloor - 1})$.

Each mixed vertex of \tilde{U}_g is incident to at least one facet of a hypercube in $\mathcal{K}_<$. Fix such a facet ϕ , and let Φ be the $(d-3)$ -dimensional flat spanned by ϕ . For each hypercube $K \in \mathcal{K}_< \cup \tilde{\mathcal{K}}_\geq$, let $K^* = K \cap \Phi$ be the (possibly empty) $(d-3)$ -dimensional hypercube contained in Φ . Set $\mathcal{K}^* = \{K^* \mid K \in \mathcal{K}_< \cup \tilde{\mathcal{K}}_\geq\}$. A mixed vertex incident on ϕ is also a vertex of $\mathcal{U}(\mathcal{K}^*)$. Since $d-3$ is even, the result in [10] implies the number of such vertices is (always) $O(n^{(d-3)/2}) = O(n^{\lfloor d/2 \rfloor - 1})$. Multiplying the bound by the number of facets in the hypercubes of $\mathcal{K}_<$, the total number of mixed vertices is $O(|\mathcal{K}_<| n^{\lfloor d/2 \rfloor - 1})$. Hence, the total number of vertices in \tilde{U}_g is $O(|\mathcal{K}_<| n^{\lfloor d/2 \rfloor - 1} + n^{\lfloor d/2 \rfloor - 1}) = O(j n^{\lfloor d/2 \rfloor - 1})$. As in the proof of Lemma 2.2, the expected value of j (assuming that $r_1 = \theta_{n-k+1}$) is $O(n/k)$. Hence, $\mu(g, P, \Theta \mid r_1 = \theta_{n-k+1}) = O\left(\frac{n^{\lfloor d/2 \rfloor}}{k}\right)$. The lemma follows by summing this over all 4^{d-2} subfaces g in the subdivision of f . ◀

Plugging Lemma 2.6 in (2), we obtain that $\mu(f, P, \Theta) = \frac{1}{n} \sum_{k=1}^n O\left(\frac{n^{\lfloor d/2 \rfloor}}{k}\right) = O(n^{\lfloor d/2 \rfloor - 1} \log n)$. Finally, summing this bound over all points of P and using Lemma 2.1, we obtain the main result of this section:

► **Lemma 2.8.** *Let P be a set of n points in \mathbb{R}^d for some fixed odd value of $d > 3$, and let Θ be a multi-set of scaling factors. Then $\bar{\mathcal{O}}_0(P, \Theta), \bar{\mathcal{O}}_1(P, \Theta) = O(n^{\lfloor d/2 \rfloor} \log n)$.*

The density model for tame distributions. Next, we show (see the full version for details) that if the scaling factors of the cubes in \mathcal{C} are chosen randomly and independently from a tame pdf π (whose cdf is denoted by Π) then the expected number of outer vertices on a fixed $(d-2)$ -dimensional face f of any cube, say of C_1 , is $O(n^{\lfloor d/2 \rfloor - 1})$. That is,

► **Lemma 2.9.** $\mu(f, P, \pi) = O(n^{\lfloor d/2 \rfloor - 1})$.

Putting everything together, using Lemma 2.1, we obtain the following:

► **Lemma 2.10.** *Let P be a set of n points in \mathbb{R}^d for some fixed odd value of $d > 3$, and let π be a tame pdf. Then $\bar{\mathcal{O}}_0(P, \Theta), \bar{\mathcal{O}}_1(P, \Theta) = O(n^{\lfloor d/2 \rfloor})$.*

2.3 Inner vertices

We now bound the expected number of inner vertices in $\mathcal{U}(\mathcal{C})$ in \mathbb{R}^d under both the permutation model and the density model for tame distributions. We first consider a fixed assignment of scaling factors to the points of P , so $\mathcal{C} = \{C_1, \dots, C_n\}$ is a fixed set of n hypercubes in \mathbb{R}^d . We obtain a recurrence for $\mathcal{J}_0(\mathcal{C})$ for this fixed \mathcal{C} , by closely following the charging scheme proposed in Boissonnat *et al.* [10], which shows that there are many distinct outer vertices of level 0 plus vertices (inner or outer) of level 1 in the neighborhood of any inner vertex of $\mathcal{U}(\mathcal{C})$. This gives an upper bound on $\mathcal{J}_0(\mathcal{C})$ in terms of $\mathcal{J}_1(\mathcal{C})$, $\mathcal{O}_0(\mathcal{C})$, and $\mathcal{O}_1(\mathcal{C})$. Next, we consider a random assignment of scaling factors and obtain a recurrence for $\bar{\mathcal{J}}_0(n)$ in terms of $\bar{\mathcal{J}}_1(n)$, $\bar{\mathcal{O}}_0(n)$, and $\bar{\mathcal{O}}_1(n)$, from the recurrence for $\mathcal{J}_0(\mathcal{C})$. We use results from Sections 2.1 and 2.2 to bound the last two terms, and we use a counting argument to bound $\bar{\mathcal{J}}_1(n)$ in terms of $\bar{\mathcal{J}}_0(n-1)$.

The charging scheme. Let v be an inner vertex of $\partial\mathcal{U}$, lying on d facets $\mathbf{f} = \{f_1, \dots, f_d\}$ of a set $\mathbf{C} = \{C_1, \dots, C_d\}$ of d distinct hypercubes, such that f_i is a facet of C_i orthogonal to the x_i -axis. For each $i \leq d$, let γ_i denote the intersection segment of the $d-1$ facets in $\mathbf{f} \setminus \{f_i\}$. The segment γ_i is parallel to the x_i -axis. The vertex v is incident to all d segments $\gamma_1, \dots, \gamma_d$. For each i , we trace the segment γ_i from v entering the hypercube C_i (recall that v lies on the x_i -orthogonal facet f_i of C_i) until we encounter another vertex, w_i , of the arrangement $\mathcal{A}(\mathcal{C})$ of \mathcal{C} . Three cases can arise:

- (i) The vertex w_i lies on the facet of C_i opposite to f_i (i.e., the other facet of C_i orthogonal to the x_i -axis). This event can happen only if C_i is smaller than all the other $d-1$ hypercubes in \mathbf{C} , i.e., it can happen for at most one index i .
- (ii) The vertex w_i is an outer vertex at level 1 (it is contained in C_i), lying on a $(d-2)$ -face of one of the hypercubes $C_j \in \mathbf{C} \setminus \{C_i\}$.
- (iii) The vertex w_i is an inner vertex at level 1 (it is contained in C_i), lying on the boundary of another hypercube $C_0 \in \mathcal{C} \setminus \mathbf{C}$.

In the first case, we simply ignore this segment γ_i and the vertex w_i —at most one of the d segments will be ignored. In cases (ii) and (iii), v charges w_i one unit. This way, each inner vertex v of \mathcal{U} receives at least $d-1$ units of charge.

An outer vertex w at level 1, lying on a $(d-2)$ -face ϕ of C_j , can be reached from an inner vertex only along one of the two corresponding facets of C_j (in a direction normal to the other facet containing ϕ), so w is charged at most twice by a type (ii) event.

An inner vertex w at level 1 can be charged at most d times, once along each of its d incident intersection segments. Suppose w is charged s times. If $s \leq 1$, then w pays one unit of charge for its unique charging inner vertex (or does not pay at all). If $s > 1$, we distribute $s - 1$ of the s units of charge to nearby outer vertices at level 0 or 1 so that each outer vertex is charged at most once (at most one inner vertex can pass a charge to it), and w also pays only one unit of charge. The charge distribution scheme culminates in the following recurrence.

$$(d - 1)\mathcal{J}_0(\mathcal{C}) \leq \mathcal{J}_1(\mathcal{C}) + 3\mathcal{O}_1(\mathcal{C}) + \mathcal{O}_0(\mathcal{C}). \quad (5)$$

Recurrence for a fixed assignment. We now obtain a recurrence for $\mathcal{J}_0(\mathcal{C})$ by using the following lemma; see [10] for the easy proof.

► **Lemma 2.11.** $\sum_{i=1}^n \mathcal{J}_0(\mathcal{C} \setminus \{C_i\}) = (n - d)\mathcal{J}_0(\mathcal{C}) + \mathcal{J}_1(\mathcal{C})$.

Using Lemma 2.11 and (5), we obtain the following recurrence.

$$\begin{aligned} (n - 1)\mathcal{J}_0(\mathcal{C}) &= (n - d)\mathcal{J}_0(\mathcal{C}) + (d - 1)\mathcal{J}_0(\mathcal{C}) \\ &\leq (n - d)\mathcal{J}_0(\mathcal{C}) + \mathcal{J}_1(\mathcal{C}) + 3\mathcal{O}_1(\mathcal{C}) + \mathcal{O}_0(\mathcal{C}) && \text{(Eq. (5))} \\ &= \sum_{i=1}^n \mathcal{J}_0(\mathcal{C} \setminus \{C_i\}) + 3\mathcal{O}_1(\mathcal{C}) + \mathcal{O}_0(\mathcal{C}) && \text{(Lemma 2.11).} \end{aligned} \quad (6)$$

Recurrence for a random assignment. We remind the reader that so far the analysis has been applied to a fixed assignment of scaling factors to the points of P , that is, to a fixed set \mathcal{C} of hypercubes. We now want to replace (6) by the expectations of the terms appearing there, with respect to a random assignment of the scaling factors. By linearity of expectation,

$$(n - 1)\mathbb{E}[\mathcal{J}_0(\mathcal{C})] \leq \sum_{i=1}^n \mathbb{E}[\mathcal{J}_0(\mathcal{C} \setminus \{C_i\})] + 3\mathbb{E}[\mathcal{O}_1(\mathcal{C})] + \mathbb{E}[\mathcal{O}_0(\mathcal{C})]. \quad (7)$$

For a fixed choice of the scaling factor r_i of C_i , the assignment of the remaining scaling factors to the elements of $\mathcal{C} \setminus \{C_i\}$ is a random permutation (of the $n - 1$ elements in $\Theta \setminus \{r_i\}$) under the permutation model, and a random assignment under the density model. The conditional expected value of $\mathcal{J}_0(\mathcal{C} \setminus \{C_i\})$ over any fixed choice of r_1 is thus at most $\bar{\mathcal{J}}_0(n - 1)$, and therefore the same also holds for the unconditional expected value of $\mathcal{J}_0(\mathcal{C} \setminus \{C_i\})$. Hence, we obtain the following:

$$(n - 1)\bar{\mathcal{J}}_0(n) \leq n\bar{\mathcal{J}}_0(n - 1) + 3\bar{\mathcal{O}}_1(n) + \bar{\mathcal{O}}_0(n). \quad (8)$$

Dividing both sides by $n(n - 1)$ and setting $\chi(m) = \frac{1}{m}\bar{\mathcal{J}}_0(m)$, we obtain

$$\chi(n) \leq \chi(n - 1) + \frac{1}{n(n - 1)}(3\bar{\mathcal{O}}_1(n) + \bar{\mathcal{O}}_0(n)). \quad (9)$$

We now substitute the bounds on $\bar{\mathcal{O}}_0(n)$ and $\bar{\mathcal{O}}_1(n)$ from Section 2.1 and Section 2.2 for both the permutation model and the density model for tame distributions. For example, for the permutation model in \mathbb{R}^3 , $\bar{\mathcal{O}}_0(n), \bar{\mathcal{O}}_1(n) = O(n \log n)$. Therefore, we obtain $\chi(n) \leq \chi(n - 1) + \frac{b \log n}{n}$, where $b > 0$ is a constant. Hence, $\chi(n) = O(\log^2 n)$ and $\bar{\mathcal{J}}_0(n) = O(n \log^2 n)$ in this case. For any fixed odd $d > 3$ and still under the permutation model, we obtain

$\chi(n) \leq \chi(n-1) + b'n^{\lfloor d/2 \rfloor - 2} \log n$, where b' is another constant. Since $d > 3$, the solution of the above recurrence is $\chi(n) = O(n^{\lfloor d/2 \rfloor - 1} \log n)$ and thus $\bar{J}_0(n) = O(n^{\lfloor d/2 \rfloor} \log n)$. This completes the proof of Theorem 1.1.

If the scaling factors of all points are chosen from a tame pdf, the bounds for $\bar{O}_0(n), \bar{O}_1(n)$ can be obtained using Lemmas 2.5 and 2.10. This implies that in this case $\bar{J}_0(n) = O(n \log n)$ for $d = 3$ and $\bar{J}_0(n) = O(n^{\lfloor d/2 \rfloor})$ for any fixed $d > 3$. This completes the proof of Theorem 1.2.

3 Union of Stochastic Minkowski Sums

We now consider the general (three-dimensional) stochastic Minkowski-sum union problem defined in the introduction, i.e., we now have a set $T = \{\Delta_1, \dots, \Delta_n\}$ of n pairwise-disjoint triangles in \mathbb{R}^3 and a convex polytope B with $O(1)$ vertices. Let $\mathbf{r} = \langle r_1, \dots, r_n \rangle$ be a sequence of n non-negative scaling factors. $C_i = \Delta_i \oplus r_i B$ is a convex polytope with $O(1)$ vertices. Set $\mathcal{C} = \{C_i \mid 1 \leq i \leq n\}$. Let $\Theta = \langle \theta_1, \dots, \theta_n \rangle$ be a sequence of n scaling factors. Our goal is to bound the expected number of vertices in $\mathcal{U}(T, \mathbf{r})$, denoted by $\psi(T, \Theta)$, when \mathbf{r} is chosen from Θ using the permutation model, i.e., we choose a random permutation σ of $[1 : n]$ and set $r_i = \theta_{\sigma(i)}$. For simplicity, we assume that T, B, Θ are such that \mathcal{C} is in general position for any permutation of Θ . Let $\mathcal{F}(T, \mathbf{r})$ denote the complement of $\mathcal{U}(T, \mathbf{r})$.

As in Section 2, we first analyze a certain conditional expectation and then bound the overall expected value $\psi(T, \Theta)$. A crucial ingredient in our analysis is the following technical lemma, adapted from Aronov and Sharir [5, Slicing Theorem] (see also [6, Chopping Theorem]).

► **Lemma 3.1.** *Let $R \subseteq T$ be a subset of t triangles and let $\mathbf{r} = \langle r, r, \dots, r \rangle$. Then (a) the complexity of $\mathcal{U}(T, \mathbf{r})$ is $O(t^2 \log t)$, and (b) there exists a decomposition \mathcal{F}^∇ of $\mathcal{F} := \mathcal{F}(R, \mathbf{r})$ into $O(t^2 \log t)$ tetrahedra with pairwise-disjoint interiors.*

We note that each tetrahedron in \mathcal{F}^∇ is defined by at most s triangles of T , where s is a constant that specifies the (maximum) number of input triangles that are needed to define the features (vertices, edges, and facets) of the corresponding tetrahedron. Moreover, any s or fewer triangles define at most $O(1)$ tetrahedra. We thus obtain the following lemma:

► **Lemma 3.2.** *The overall number of tetrahedra that can ever arise in the decomposition $\mathcal{F}^\nabla(R, \mathbf{r})$ for any subset $R \subseteq T$, where all the scaling factors in \mathbf{r} are equal, is $O(n^s)$.*

The recursive partition scheme. We now follow, or rather adapt, the recursive scheme used in [2]. Specifically, fix a parameter t , whose value will be determined later, and put $\rho = \theta_{n-t}$, the $(t+1)$ -st largest scaling factor in Θ . Let $\Theta^<$ be the sequence of the $n-t$ smallest values in Θ . We fix a subset $T^> \subseteq T$ of t triangles, and set $T^< = T \setminus T^>$. Let $\psi(T, \Theta \mid T^>)$ denote the conditional expected complexity of \mathcal{U} conditioned over the subset of those permutations σ that assign the t largest scaling factors in Θ to the triangles of $T^>$; the restriction of a random permutation from this subset to $[1 : n-t]$ (i.e., the part of the permutation that assigns the scaling factors of $\Theta^<$ to the triangles in $T^<$) is a random permutation, i.e., each permutation of $[1 : n-t]$ is equally likely to arise.

We first obtain a bound on $\psi(T, \Theta \mid T^>)$. For each $\Delta_i \in T$, put $\bar{C}_i = \Delta_i \oplus \rho B$. Set $\bar{\mathcal{C}}^> = \{\bar{C}_i \mid \Delta_i \in T^>\}$ and $\bar{\mathcal{U}} := \mathcal{U}(\bar{\mathcal{C}}^>)$. Let $\bar{\mathcal{F}}$ denote the complement of $\bar{\mathcal{U}}$. Since the Δ_i 's are pairwise disjoint and we now add the same convex polytope ρB to each of them, Lemma 3.1(a) implies that $\bar{\mathcal{U}}$ has $O(t^2 \log t)$ complexity. Moreover, by Lemma 3.1(b), there exists a decomposition, denoted as $\bar{\mathcal{F}}^\nabla$, of $\bar{\mathcal{F}}$ into $O(t^2 \log t)$ tetrahedra with pairwise-disjoint interiors. For each tetrahedron τ of $\bar{\mathcal{F}}^\nabla$, let $T_\tau^< = \{\Delta_i \in T^< \mid \bar{C}_i \cap \tau \neq \emptyset\}$; set $n_\tau = |T_\tau^<|$.

Note that $\bar{\mathcal{U}}$, $\bar{\mathcal{F}}$, $\bar{\mathcal{F}}^\nabla$, and the sets $T_\tau^<$ are independent of the choice of permutation; they only depend on ρ and the set $T^>$.⁸

We now fix a permutation σ , and thus an assignment of scaling factors, such that the largest t scaling factors are assigned to the triangles of $T^>$. Thus the scaling factors of all triangles in $T^<$ (resp. $T^>$) are at most (resp. at least) ρ . Therefore $C_i \subseteq \bar{C}_i$ for each $\Delta_i \in T^<$ and $C_i \supseteq \bar{C}_i$ for each $\Delta_i \in T^>$. Consequently, $\bar{\mathcal{U}} \subseteq \mathcal{U}(\mathcal{C})$.

For a cell $\tau \in \bar{\mathcal{F}}^\nabla$, set $\mathcal{C}_\tau^< = \{C_i \mid T_i \in T_\tau^<\}$. Clearly, no other triangle $\Delta_i \in T^<$ can have its real expansion C_i meet τ . We define $\mathcal{C}_\tau = \{C_i \cap \tau \mid \Delta_i \in T_\tau^< \cup T^>\}$, and set $\mathcal{U}_\tau = \mathcal{U}(\mathcal{C}_\tau)$. By construction,

$$\mathcal{U} = \mathcal{U}(\mathcal{C}) = \bar{\mathcal{U}} \cup \bigcup_{\tau \in \bar{\mathcal{F}}^\nabla} \mathcal{U}_\tau. \tag{10}$$

We call a vertex v of \mathcal{U} *internal* if v does not lie on the boundary of any polytope C_i with $T_i \in T^>$, and we call v *external* otherwise. The number of external vertices is trivially $O(tn^2)$, so it suffices to bound the number of internal vertices of \mathcal{U} . Since $\bar{\mathcal{U}} \subseteq \mathcal{U}$, every internal vertex lies in $\bar{\mathcal{F}}$. Suppose v lies in the tetrahedron τ of $\bar{\mathcal{F}}^\nabla$. Then by (10), v is a vertex of \mathcal{U}_τ . Furthermore, v is an internal vertex, so it is not incident on any C_i for $T_i \in T^>$, and thus v is a vertex of $\mathcal{U}(\mathcal{C}_\tau^<)$. The total number of internal vertices in τ is bounded by the combinatorial complexity of $\mathcal{U}(\mathcal{C}_\tau^<)$, denoted by $\psi(T_\tau^<, \mathbf{r}_\tau)$ (as defined in the introduction), where \mathbf{r}_τ is the set of scaling factors of the triangles in T_τ . Hence, the total number of internal vertices on \mathcal{U} is $\sum_{\tau \in \bar{\mathcal{F}}^\nabla} \psi(T_\tau^<, \mathbf{r}_\tau)$, and (10) implies that the total number vertices $\psi(T, \mathbf{r})$ on \mathcal{U} satisfies $\psi(T, \mathbf{r}) \leq \sum_{\tau \in \bar{\mathcal{F}}^\nabla} \psi(T_\tau^<, \mathbf{r}_\tau) + O(tn^2)$.

Expected number of vertices. We now bound the expected complexity of \mathcal{U} , conditioned on a fixed choice of $T^>$. That is, we condition the analysis on the subset of those permutations σ that assign the t largest scaling factors in Θ to the triangles of $T^>$. The set $\Theta_\tau^<$ of scaling factors assigned to the triangles in $T_\tau^<$ is not fixed, but, since $T_\tau^<$ is a fixed set, conditioned only on the choice of $T^>$, the set $\Theta_\tau^<$ is a *random subset* of $\Theta^< = \{\theta_1, \dots, \theta_{n-t}\}$ of the fixed size n_τ . Moreover, the assignment (under the original random permutation σ) of these scaling factors to the triangles in $T_\tau^<$ is a *random* permutation of $\Theta_\tau^<$. Hence, conditioning further on the choice of $\Theta_\tau^<$, the expected value of $\psi(T_\tau^<, \mathbf{r}_\tau)$ is $\psi(T_\tau^<, \Theta_\tau^<) \leq \psi(n_\tau)$. Hence, the last expression also bounds the unconditional expected complexity of $\mathcal{U}(\mathcal{C}_\tau^<)$, albeit still conditioned on a fixed choice of $T^>$. Summing this over all tetrahedra of $\bar{\mathcal{F}}^\nabla$, we obtain that $\psi(T, \Theta \mid T^>) \leq \sum_{\tau \in \bar{\mathcal{F}}^\nabla} \psi(n_\tau) + O(tn^2)$.

To bound the unconditional expected value $\psi(T, \Theta)$ we notice that the subset $T^>$ of the triangles that are assigned the t largest scaling factors is a random subset of T . Since $T^>$ is a random sample of T of size t and since, by Lemma 3.2, there are only $O(n^s)$ tetrahedra that can appear in the decomposition $\bar{\mathcal{F}}^\nabla$, the following lemma is a consequence of a standard random-sampling argument; see [17, Section 4.6] for a proof.

► **Lemma 3.3.** *For any constant $c > 0$, with probability $1 - O(\frac{1}{n^c})$, every (open) tetrahedron τ of $\bar{\mathcal{F}}^\nabla$ intersects at most $\frac{c'n}{t} \log n$ of the sets of $\bar{\mathcal{C}}^<$, where c' is a constant that depends on both c .*

Choosing $c = 3$ in Lemma 3.3, with probability $1 - O(\frac{1}{n^3})$ our $T^>$ is such that $n_\tau \leq \frac{c'n}{t} \log n$ for every $\tau \in \bar{\mathcal{F}}^\nabla$. With probability $O(\frac{1}{n^3})$, $T^>$ may fail to satisfy this property,

⁸ Clearly, $\bar{\mathcal{U}}$ and $\bar{\mathcal{F}}$ are uniquely determined. For $\bar{\mathcal{F}}^\nabla$, the statement means that if we follow an agreed-upon implementation of the construction in the proof of Lemma 3.1, $\bar{\mathcal{F}}^\nabla$ is also uniquely determined.

but, since the complexity of $\mathcal{U}(\mathcal{C})$ is always $O(n^3)$ (again, with a constant of proportionality that depends on the complexity of B), the contribution of these “bad” choices of T to $\psi(T, \Theta)$ is $O(1)$. If n is below some appropriate constant n_0 , we can use a trivial bound of $O(n^3)$ for the complexity of \mathcal{U} . Altogether, we obtain the following recurrence,

$$\psi(n) \leq \begin{cases} a_0 n^3 & \text{for } n \leq n_0, \\ a_1 t^2 \log t \cdot \psi\left(\frac{c'n}{t} \log n\right) + a_2 n^2 t & \text{for } n > n_0, \end{cases} \quad (11)$$

where n_0, a_0, a_1, a_2, c' are suitable constants. With appropriate choice of parameters, the solution of this recurrence is $\psi(n) \leq An^{2+\varepsilon}$, for any $\varepsilon > 0$, where A depends on ε and on the other constants appearing in the recurrence. This proves Theorem 1.3.

References

- 1 P. K. Agarwal, A. Efrat, S. K. Ganjugunte, D. Hay, S. Sankararaman, and G. Zussman. The resilience of WDM networks to probabilistic geographical failures. *IEEE/ACM Trans. on Netw.*, 21(5):1525–1538, 2013.
- 2 P. K. Agarwal, S. Har-Peled, H. Kaplan, and M. Sharir. Union of random minkowski sums and network vulnerability analysis. *Discrete Comput. Geom.*, 52(3):551–582, 2014.
- 3 P.K. Agarwal and M. Sharir. Pipes, cigars, and kreplach: the union of minkowski sums in three dimensions. *Discrete Comput. Geom.*, 24(4):645–657, Jan 2000.
- 4 B. Aronov, A. Efrat, V. Koltun, and M. Sharir. On the union of κ -round objects in three and four dimensions. *Discrete Comput. Geom.*, 36(4):511–526, 2006.
- 5 B. Aronov and M. Sharir. Triangles in space or building (and analyzing) castles in the air. *Combinatorica*, 10(2):137–173, 1990.
- 6 B. Aronov and M. Sharir. Castles in the air revisited. *Discrete Comput. Geom.*, 12(2):119–150, 1994.
- 7 B. Aronov and M. Sharir. On translational motion planning of a convex polyhedron in 3-space. *SIAM J. Comput.*, 26(6):1785–1803, 1997.
- 8 B. Aronov, M. Sharir, and B. Tagansky. The union of convex polyhedra in three dimensions. *SIAM J. Comput.*, 26(6):1670–1688, 1997.
- 9 F. Aurenhammer, R. Klein, and D. T. Lee. *Voronoi diagrams and Delaunay triangulations*. World Scientific, 2013.
- 10 J.-D. Boissonnat, M. Sharir, B. Tagansky, and M. Yvinec. Voronoi diagrams in higher dimensions under certain polyhedral distance functions. *Discrete Comput. Geom.*, 19(4):485–519, 1998.
- 11 H.-C. Chang, S. Har-Peled, and B. Raichel. From proximity to utility: A voronoi partition of pareto optima. *Discrete Comput. Geom.*, 56(3):631–656, 2016.
- 12 E. Ezra. On the union of cylinders in three dimensions. *Discrete Comput. Geom.*, 45(1):45–64, 2011.
- 13 E. Ezra and M. Sharir. On the union of fat tetrahedra in three dimensions. *J. ACM*, 57(1):2:1–2:23, 2009.
- 14 M. J. Golin and H.-S. Na. On the average complexity of 3d-voronoi diagrams of random points on convex polytopes. *Comput. Geom: Theory Appls.*, 25(3):197–231, 2003.
- 15 R. L. Graham, D. E. Knuth, and O. Patashnik. *Concrete Mathematics: A Foundation for Computer Science*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 2nd edition, 1994.
- 16 S. Har-Peled and B. Raichel. On the complexity of randomly weighted multiplicative voronoi diagrams. *Discrete Comput. Geom.*, 53(3):547–568, 2015.
- 17 J. Matoušek. *Lectures on Discrete Geometry*. Springer-Verlag, Berlin, Heidelberg, 2002.

- 18 J. Pach P. K. Agarwal and M. Sharir. State of the union (of geometric objects). In J. Pach J. Goodman and R. Pollack, editors, *Surveys on Discrete and Computational Geometry*, pages 9–48. Amer. Math. Soc., Providence, RI, 2008.
- 19 J. Pach, I. Safruti, and M. Sharir. The union of congruent cubes in three dimensions. *Discrete Comput. Geom.*, 30(1):133–160, 2003.
- 20 R. Schneider and J. A. Wieacker. Integral geometry. In P. M. Gruber and J. M. Wills, editors, *Handbook of Convex Geometry*, volume B, pages 1349–1390. North-Holland, Amsterdam, 1993.
- 21 W. Weil and J. A. Wieacker. Stochastic geometry. In P. M. Gruber and J. M. Wills, editors, *Handbook of Convex Geometry*, volume B, pages 1393–1438. North-Holland, Amsterdam, 1993.