On the Union of Fat Tetrahedra in Three Dimensions*

Esther Ezra[†] Micha Sharir[‡]

Abstract

We show that the combinatorial complexity of the union of n "fat" tetrahedra in 3-space (i.e., tetrahedra all of whose solid angles are at least some fixed constant) of arbitrary sizes, is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$; the bound is almost tight in the worst case, thus almost settling a conjecture of Pach *et al.* [24]. Our result extends, in a significant way, the result of Pach *et al.* [24] for the restricted case of *nearly congruent cubes*. The analysis uses cuttings, combined with the Dobkin-Kirkpatrick hierarchical decomposition of convex polytopes, in order to partition space into subcells, so that, on average, the overwhelming majority of the tetrahedra intersecting a subcell Δ behave as fat *dihedral* wedges in Δ. As an immediate corollary, we obtain that the combinatorial complexity of the union of n cubes in $\mathbb{R}^3$, having arbitrary side lengths, is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$ (again, significantly extending the result of [24]). Our analysis can easily be extended to yield a nearly-quadratic bound on the complexity of the union of arbitrarily oriented fat triangular prisms (whose cross-sections have arbitrary sizes) in $\mathbb{R}^3$. Finally, we show that a simple variant of our analysis implies a nearly-linear bound on the complexity of the union of fat triangles in the plane.

*Work on this paper has been supported by NSF Grant CCF-05-14079, by a grant from the U.S.-Israeli Binational Science Foundation, by Grant 155/05 from the Israel Science Fund, and by the Hermann Minkowski–MINERVA Center for Geometry at Tel Aviv University. Work on this paper by the first author has also been supported by an IBM Doctoral Fellowship. A preliminary version of this paper will appear in *Proc.* 48^{nd} *Annu. IEEE Sympos. Found. Comput. Sci.* 2007.

[†]School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel; `estere@tau.ac.il`.

[‡]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@tau.ac.il`.

1 Introduction

Let $\mathcal{T}$ be a collection of n (arbitrarily oriented) tetrahedra of arbitrary sizes in 3-space. Let $\mathcal{A}(\mathcal{T})$ denote the three-dimensional arrangement induced by the facets of the tetrahedra in $\mathcal{T}$, i.e., the decomposition of 3-space into vertices, edges, faces, and three-dimensional cells, each being a maximal connected set contained in the intersection of a fixed subcollection of facets of the tetrahedra of $\mathcal{T}$ and not meeting any other facet. The *combinatorial complexity* of the union of the tetrahedra in $\mathcal{T}$ is the number of vertices, edges and faces of the arrangement appearing on the union boundary. The problem studied in this paper is to obtain a nearly-quadratic upper bound on the combinatorial complexity of the union, in the special case where all the given tetrahedra are *fat*, meaning that the solid angles at their vertices are all at least some fixed constant $\alpha > 0$.

Previous results. The problem of determining the combinatorial complexity of the union of geometric objects has received considerable attention in the past twenty years, although most of the earlier work has concentrated on the planar case. See [1] for a recent comprehensive survey of the area.

The case involving *pseudodiscs* (that is, a collection of simply connected planar regions, where the boundaries of any two distinct objects intersect at most twice), arises for Minkowski sums of a fixed convex object with a set of pairwise disjoint convex objects (which is the problem one faces in translational motion planning of a convex robot), and has been studied by Kedem *et al.* [18]. In this case, the union has only linear complexity. Matoušek *et al.* [21, 22] proved that the union of n α-*fat* triangles (where each of their angles is at least α) in the plane has only $O(n)$ holes, and its combinatorial complexity is $O(n \log \log n)$. The constant of proportionality, which depends on the fatness factor α, has later been improved by Pach and Tardos [25]. Extending the study to the realm of curved objects, Efrat and Sharir [15] studied the union of planar convex fat objects. Here we say that a planar convex object c is α-*fat*, for some fixed $\alpha > 1$, if there exist two concentric disks, $D \subseteq c \subseteq D'$, such that the ratio between the radii of D' and D is at most α. In this case, the combinatorial complexity of the union of n such objects, such that the boundaries of each pair of objects intersect in a constant number of points, is $O(n^{1+\varepsilon})$, for any $\varepsilon > 0$. See also Efrat and Katz [13] and Efrat [12] for related (and slightly sharper) nearly-linear bounds.

In three and higher dimensions, it was shown by Aronov *et al.* [5] that the complexity of the union of k convex polyhedra with a total of n facets in $\mathbb{R}^3$ is $O(k^3 + nk \log k)$, and it can be $\Omega(k^3 + nk\alpha(k))$ in the worst case. The bound was improved by Aronov and Sharir [4] to $O(nk \log k)$ (and $\Omega(nk\alpha(k))$) when the given polyhedra are Minkowski sums of a fixed convex polyhedron with k pairwise-disjoint convex polyhedra. (This problem arises in the case of a translating convex polyhedral robot in $\mathbb{R}^3$ amid a collection of polyhedral obstacles.) Boissonnat *et al.* [6] proved that the maximum complexity of the union of n axis-parallel hypercubes in $\mathbb{R}^d$ is $\Theta\left(n^{\lceil d/2 \rceil}\right)$, and that the bound improves to $\Theta\left(n^{\lfloor d/2 \rfloor}\right)$ if all hypercubes have the same size. Pach *et al.* [24] showed that the combinatorial complexity of the union of n nearly congruent arbitrarily oriented cubes in three dimensions is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$ (see also [23] for a subcubic bound on the complexity of the union of fat wedges in 3-space). Agarwal and Sharir [2] have shown that the complexity of the union of n congruent infinite cylinders is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$. In fact, the more general problem studied in [2] involves the union of the Minkowski sums of n pairwise disjoint triangles with a ball (where congruent infinite cylinders are obtained when the triangles become lines), and the nearly quadratic bound is extended in [2] to this case as well. Finally, Aronov *et al.* [3] showed that the union complexity of n κ-*round* objects in $\mathbb{R}^3$ is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, where an object c is κ-round if for each $p \in \partial c$ there exists a ball $B \subset c$ that touches p and its radius is at least $\kappa \cdot diam(c)$. The bound is $O(n^{3+\varepsilon})$, for any $\varepsilon > 0$, for κ-round objects in $\mathbb{R}^4$. Each of the above known nearly-quadratic bounds (for the three-dimensional case) is almost tight in the worst case.

To recap, all of the above results indicate that the combinatorial complexity of the union of fat objects is roughly "one order of magnitude" smaller than the complexity of the arrangement that they induce. While considerable progress has been made on the analysis of unions in three dimensions, the case of the union of fat polyhedra has so far been lagging behind, where only very few nearly-quadratic bounds are known.

Our results. In this paper we make significant progress on the problem of bounding the complexity of the union of fat polyhedra, by deriving a nearly-quadratic bound on the combinatorial complexity of the union of fat tetrahedra. Our bound, which is the first known subcubic bound for this general problem, is almost tight in the worst case.

Specifically, a tetrahedron is called α-*fat* if each of its four solid angles (at its four respective apices) is at least α. We show that, for any fixed $\alpha > 0$, the complexity of the union of n α-fat tetrahedra is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, where the constant of proportionality depends on ε and on α. Our proof technique relies only on the nearly-quadratic bound of the union of α-fat dihedral wedges, established by Pach *et al.* [24]; a dihedral wedge is called α-fat if its dihedral angle is at least α. An immediate application of our result is a nearly-quadratic bound on the complexity of the union of arbitrary cubes. In particular, the second part of the analysis of Pach *et al.* [24], for the specific case of nearly congruent cubes, is not needed any more, since it is subsumed by our analysis, which does not use that part, and applies in a much wider context.

The analysis is based on cuttings, which incorporate the Dobkin-Kirkpatrick hierarchical decomposition scheme for convex polytopes [11], in order to partition space into subcells (simplices), so that, on average, the overwhelming majority of the tetrahedra intersecting a subcell Δ behave as α'-fat *dihedral* wedges within Δ, where α' is another constant that depends on α. Since, as shown in [24], the complexity of the union of α'-fat dihedral wedges is nearly quadratic, it only remains to analyze the number of other types of vertices (incident to some of the few "bad" tetrahedra that cross Δ), a task which is handled by the cutting-based divide-and-conquer mechanism (see below for details).

Our analysis can also be applied when the given objects are arbitrarily oriented α-*fat* triangular prisms (that is, all the dihedral angles in each prism are at least α) having cross-sections of arbitrary sizes. In this case, the complexity of the union is nearly-quadratic as well, and the bound is nearly worst-case tight. We are not aware of any previous known subcubic bound in this case, except for the nearly-quadratic bound of Aronov and Sharir [4], for the special case where the prisms are Minkowski sums of lines in 3-space with a fixed (not necessarily fat) polyhedron.

An immediate consequence of our results is a bound $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, on the complexity of the union of any family of polyhedral objects, so that each of them is the union of some number of α-fat tetrahedra (or triangular prisms), and the total number of these tetrahedra or prisms is n.

The problem studied in this paper is a natural extension of the two-dimensional problem of bounding the complexity of the union of fat triangles, which was studied by Matoušek *et al.* [21, 22] and by Pach and Tardos [25]. We show that a simple specialization of our analysis to the two-dimensional case yields an upper bound $O(n^{1+\varepsilon})$, for any $\varepsilon > 0$, on this complexity. The analysis, based on the new approach, is almost immediate (albeit yielding a slightly suboptimal bound), and is significantly simpler than the analysis in [21, 22, 25].

By now, the arsenal of techniques for analyzing the complexity of the union of geometric objects in 3-space is quite rich: It includes, for example, the technique of Aronov *et al.* [3], for bounding the combinatorial complexity of the union of n κ-round objects in $\mathbb{R}^3$, by reducing the problem to subproblems involving sandwich regions between upper and lower envelopes (see also [2]), and the technique of Pach *et al.* [24] for bounding the complexity of the union of n cubes in 3-space by bounding the number of "special cubes" in the arrangement of these cubes (see also [5]). However, we were unable to extend any of these alternative techniques to our context, and had to develop new machinery. We believe it to be of independent interest, and hope that it will find additional applications to related problems.

2 The Union of Fat Tetrahedra

2.1 Preliminaries and overview

We borrow the following notation from Pach *et al.* [24] (some of which has already been mentioned in the introduction). A *dihedral* (resp., *trihedral*) *wedge* is the intersection of two (resp., three) halfspaces. A dihedral (resp., trihedral) wedge is *α-fat* if its dihedral (resp., solid) angle is at least α. A trihedral wedge is also associated with the three dihedral angles at its edges. It is easily verified that there exists constant $\alpha' > 0$, which depends only on α, such that, for any α-fat trihedral wedge, each of its three dihedral angles is at least α' [1].

Similar definitions and observations apply to α-fat tetrahedra, namely, tetrahedra all of whose solid angles are at least α. In particular, there exist (the same) constant $\alpha' > 0$, such that, for any α-fat tetrahedron, each of its six dihedral angles is at least α'.

Let $\mathcal{T} = \{T_1, \ldots, T_n\}$ be a collection of n α-fat tetrahedra in 3-space, and let $\mathcal{U} = \bigcup \mathcal{T}$ denote their union. For simplicity of the analysis, we assume that the given tetrahedra are in *general position* (see [5] for an argument that this involves no loss of generality). This general position assumption implies that each vertex of the arrangement $\mathcal{A}(\mathcal{T})$ of the (facets of the) tetrahedra of $\mathcal{T}$ lies on exactly three tetrahedra facets, and is thus incident upon only a constant number of edges and faces. This is easily seen to imply that the combinatorial complexity of $\mathcal{U}$ is $O(|V(\mathcal{T})|)$, where $V(\mathcal{T})$ is the set of vertices of $\mathcal{A}(\mathcal{T})$ that appear on the boundary of the union.

[1] By intersecting the wedge with a sphere centered at it apex, this amounts to asserting that if the (normalized) area of a spherical triangle is at least α then each of its angles is at least α'.

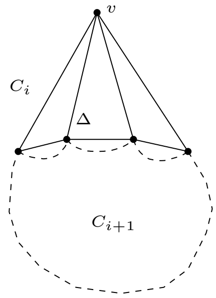

Figure 1: Illustrating a step in the construction of the Dobkin-Kirkpatrick hierarchical decomposition: The vertex v is peeled off C_i, the convex hull of the "hole" that it leaves is constructed, and the new facets are connected to v by tetrahedra that fill up this portion of $C_i \setminus C_{i+1}$.

We classify the vertices of $\mathcal{A}(\mathcal{T})$ as in [5]: An intersection vertex v of $\mathcal{A}(\mathcal{T})$ (i.e., not a vertex of one of the tetrahedra of $\mathcal{T}$) is said to be an *outer* vertex if it is the intersection of an edge of one tetrahedron and the relative interior of a facet of another tetrahedron, or an *inner* vertex, if v is the intersection of the relative interiors of three facets of three distinct tetrahedra. Trivially, the number of outer vertices in the entire arrangement $\mathcal{A}(\mathcal{T})$ is $O(n^2)$, so our main goal is to bound the number of inner vertices that appear on $\partial\mathcal{U}$. The main result of the paper is:

Theorem 2.1 *The complexity of the union of n α-fat tetrahedra in $\mathbb{R}^3$ is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, where the constant of proportionality depends on ε and α. The bound is almost tight in the worst case.*

It is relatively easy (using standard techniques; see, e.g., [26]) to construct a set of n α-fat tetrahedra that yield $\Omega(n^2\alpha(n))$ vertices on the boundary of their union (see also [23, 24] for further details). We thus devote the remainder of this section to deriving the upper bound stated in Theorem 2.1.

Curve-sensitive cuttings. We use a divide-and-conquer approach, based on a simple variant of "curve-sensitive" cuttings [19]. Specifically, let $\mathcal{F}$ be the set of all facets of the tetrahedra in $\mathcal{T}$. For any $r \le n$ there exists a $(1/r)$-cutting Ξ for $\mathcal{F}$, which is a partition of $\mathbb{R}^3$ into $O(r^3 \log^3 r)$ simplices, such that every simplex (also referred to as a *cell* of Ξ) is crossed by at most n/r facets of $\mathcal{F}$, with the additional property that any edge of a tetrahedron in $\mathcal{T}$ crosses at most $O(r \log^2 r)$ cells of Ξ.

One can obtain such a cutting using the following (simple) construction[2]. We first draw a random sample R of $O(r \log r)$ of the planes containing the facets of $\mathcal{F}$, and add to that collection four additional planes that define a sufficiently large simplex σ_0 that encloses all the vertices of $\mathcal{A}(\mathcal{F})$. We form the arrangement $\mathcal{A}(R)$ of R, consider only its portion within σ_0,

[2]For simplicity, we do not use the refined technique of [7, 20], which improves the size of the cutting down to $O(r^3)$, since it does not affect the asymptotic bound that we obtain on the complexity of the union, and since the analysis is cleaner without this refinement.

and triangulate each of its cells C contained in σ_0, using the Dobkin-Kirkpatrick hierarchical decomposition of convex polytopes [11].

The number of simplices is proportional to the overall complexity of $\mathcal{A}(R)$, and is thus $O(r^3 \log^3 r)$. The ε-net theory [8, 17] implies that, with high probability, each simplex of the resulting decomposition is crossed by at most n/r (planes containing) facets of $\mathcal{F}$. We pick one sample R for which this property holds, and fix it in the foregoing analysis.

So far, the use of the Dobkin-Kirkpatrick hierarchy is not essential—many other triangulation schemes for the cells of $\mathcal{A}(R)$ (e.g., bottom-vertex triangulation) would do equally well. However, the Dobkin-Kirkpatrick hierarchy *is* crucial for our divide-and-conquer approach in a manner that will be described later on.

Here is a brief review of the technique, given for the sake of completeness, and also because we will exploit several features of the construction in our analysis. Let $C \subseteq \sigma_0$ be a fixed (bounded) cell of $\mathcal{A}(R)$, which is a convex polytope. The hierarchical decomposition of C is a sequence $C_1, \ldots, C_k$ of $k \geq 1$ convex polytopes, such that (i) $C_1 = C$ and C_k is a simplex, (ii) $C_{i+1} \subset C_i$, for $1 \leq i < k$, (iii) $V(C_{i+1}) \subset V(C_i)$, for $1 \leq i < k$, where $V(P)$ is the set of the vertices of a polytope P, and (iv) the vertices in $V(C_i) \setminus V(C_{i+1})$ form an independent set in the planar skeleton graph of ∂C_i, for $1 \leq i < k$.

It is shown in [11] that there always exists a hierarchical decomposition for C that satisfies $k = O(\log |V(C_i)|)$, $\sum_{i=1}^{k} |V(C_i)| = O(|V(C)|)$, and $\max_i \max_{v \in V(C_i) \setminus V(C_{i+1})} deg(v, C_i) \leq c$, for some absolute constant $c \geq 3$, where $deg(v, C_i)$ is the degree of v in the skeleton graph of ∂C_i. Specifically, we obtain C_{i+1} from C_i by the following steps: (a) Find an independent subset $V_i^* \subseteq V(C_i)$ of vertices of degree at most c, whose size is $\Theta(|V(C_i)|)$ (e.g., an independent set of size $|V(C_i)|/24$ whose vertices have degree at most 11 can be shown to exist, as a simple consequence of Euler's polyhedral formula). (b) For each $v \in V_i^*$, remove v and its adjacent edges and facets from C_i. (c) The removal of such a vertex v leaves a "hole" in C_i. The convex hull of the set of neighboring vertices of v is constructed, and its outer (triangular) facets are added as new facets of C_{i+1}, thereby closing the hole that v has left [3]. (d) Finally, the gap between ∂C_i and C_{i+1} in the neighborhood of v (formally, the connected component of $\mathrm{int}(C_i) \setminus C_{i+1}$ whose closure contains v) is triangulated into $O(1)$ simplices by connecting v with each of the new facets of C_{i+1} that bound the gap; see Figure 1 for an illustration.

A simplicial subcell Δ is said to be generated at step i if it has a vertex v that is removed from C_i; that is, Δ is one of the simplices that fill up the gap formed by the removal of v. Note that the three other vertices of Δ belong to C_{i+1}.

The Dobkin-Kirkpatrick decomposition has several useful properties that we will exploit. One of these properties is that a line that crosses a cell τ of $\mathcal{A}(R)$ crosses only $O(\log r)$ of its simplices (it can visit at most two gaps of $C_i \setminus C_{i+1}$, for each of the logarithmically many indices i). Since a line (or, rather, an edge of a tetrahedron in $\mathcal{T}$) crosses at most $O(r \log r)$ cells of $\mathcal{A}(R)$ (it has to cross a plane of R to move from one cell to another), it crosses at most $O(r \log^2 r)$ simplices, as claimed.

[3] The independence of V_i^* guarantees that the holes, and their hulls, are openly pairwise disjoint.

The problem decomposition—an overview. We construct the cutting Ξ, as just described, with a sufficiently large constant value of r, and bound the number of inner vertices of the union in each cell of Ξ separately. Fix a cell Δ of Ξ. We classify each facet $F \in \mathcal{F}$ that intersects Δ as being either *long* in Δ, if $\partial F \cap \Delta = \emptyset$, or *short*, otherwise. As just discussed, the number of cells Δ in which F is short is $O(r \log^2 r)$.

Let us fix a tetrahedron $T \in \mathcal{T}$. For each cell Δ of Ξ, either (i) Δ is disjoint from T, or (ii) Δ is fully contained in T, or (iii) Δ intersects only one or two facets of T, or (iv) Δ intersects at least three facets of T. In case (i) T has no effect on the union within Δ. In case (ii) Δ is fully covered and does not contain any portion of the boundary of the union. In case (iii) we say that T meets Δ as a dihedral wedge (which can also be a halfspace), and call T a *D-tetrahedron* in Δ, and in case (iv) we say that T meets Δ as a tetrahedron or a trihedral wedge, and call T a *T-tetrahedron* in Δ.

If Δ meets only one facet F of T, we replace T by the halfspace bounded by that facet and containing T. Similarly, if Δ meets two facets of T, we replace T by the dihedral wedge formed by the planes supporting these facets and containing T. Clearly, these replacements do not affect the union of the tetrahedra within Δ. The case where at least three facets of T meet Δ (case (iv)) is more involved—this is after all the situation we started with. What saves us is the property that the number of T-tetrahedra is small on average. This is one of the main technical insights in our analysis, and is established below in Lemmas 2.3 and 2.4.

Each inner intersection vertex v of the union that appears in Δ is consequently classified as either *DDD*, if all three facets that are incident to v belong to three respective D-tetrahedra in Δ, *DDT*, if two of these facets belong to two respective D-tetrahedra and one belongs to a T-tetrahedron, *DTT*, if one of these facets belongs to a D-tetrahedron and two belong to two respective T-tetrahedra, or *TTT*, if all three facets belong to three respective T-tetrahedra. In all four cases, the three relevant tetrahedra are distinct.

2.2 T-tetrahedra are scarce

Our next goal is to show that, for each tetrahedron $T \in \mathcal{T}$, the overall number of simplices Δ of Ξ, such that T crosses Δ and is a T-tetrahedron in Δ, is only $O(r \log^2 r)$. We emphasize that this part of the analysis does not use the fatness of T — the bound holds for *any* tetrahedron T.

Let us fix a tetrahedron T of $\mathcal{T}$, and consider the set of simplices Δ in Ξ that meet at least three facets of T. It suffices to consider only simplices Δ in which all facets of T are long: The edge-sensitivity of the cutting implies that the overall number of simplices Δ that are crossed by an edge of T is $O(r \log^2 r)$.

We establish the above bound in two steps, in the respective Lemmas 2.3 and 2.4. In the first step (Lemma 2.3) we bound the number of cells of the untriangulated arrangement $\mathcal{A}(R)$ that meet at least three facets of T; a crucial ingredient of the analysis is established in Lemma 2.2. Then we fix such a cell C, and bound (in Lemma 2.4) the number of sub-simplices of Ξ in C that have this property. We first prove the following geometric property:

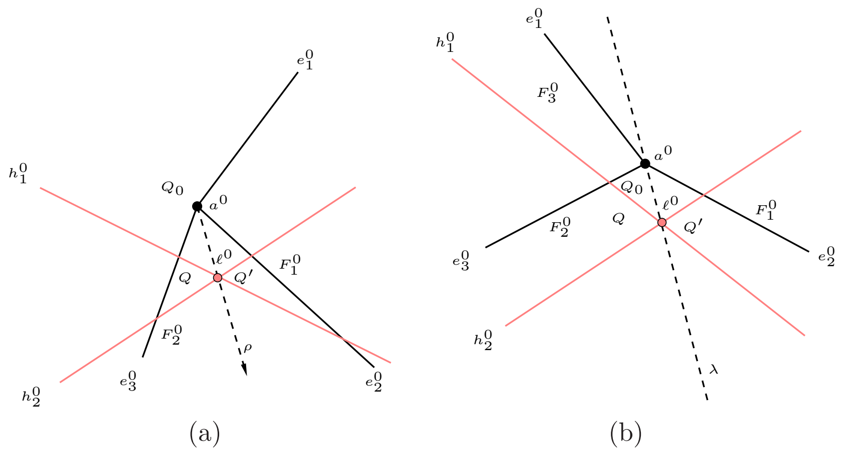

Figure 2: The proof of Lemma 2.2.

Lemma 2.2 *Let W be a trihedral wedge with apex a, let h_1, h_2 be two planes, whose intersection line crosses W. For $i = 1, 2$, denote by h_i^+ the halfspace bounded by h_i and containing a, and by h_i^- the complementary halfspace. Then at least one of the wedges $C = h_1^+ \cap h_2^-$, $C' = h_1^- \cap h_2^+$ is crossed by at most two facets of W.*

Proof: Project everything onto some plane h_0 orthogonal to both h_1, h_2, and denote the projection of object u by u^0. The line $\ell = h_1 \cap h_2$ projects to a point ℓ^0, and h_1^0, h_2^0 are two lines passing through ℓ^0 and partitioning h_0 into four quadrants, so that one of them, Q_0, contains a^0, and the two quadrants Q, Q' adjacent to Q_0 are the projections of C, C', respectively.

Let F_1, F_2, F_3 be the facets of W, and let e_i denote the edge incident to F_i and F_{i+1}, for $i = 1, 2, 3$ (where e_3 is incident to F_3 and F_1, and we also denote it by e_0). The edges e_i project to three respective rays e_i^0 that emanate from a^0. Note that, due to the assumption that ℓ crosses W, ℓ^0 must be contained in the projection of W. We next consider the following two possibilities:

(a) e_1^0, e_2^0, e_3^0 are contained in a common halfplane, bounded by a line λ through a^0. Assume, without loss of generality, that e_2^0 lies between e_1^0 and e_3^0. Then F_1^0 and F_2^0 are openly disjoint, so at least one of them, say F_1^0, does not contain ℓ^0; see Figure 2(a). In this case the ray ρ from a^0 through ℓ^0 is disjoint from F_1^0 (except for its apex). We next claim that F_1^0 lies fully in one of the halfplanes bounded by (the line containing) ρ. It will then easily follow that F_1^0 cannot meet both Q and Q', because each of them is fully contained on a different side of the line containing ρ. Indeed, if the ray opposite to ρ is contained in F_1^0, then it implies that ℓ^0 and F_1^0 lie on different sides of λ, but then ℓ^0 is disjoint from the projection of W, contradicting our assumption.

(b) e_1^0, e_2^0, e_3^0 are not contained in a common halfplane. In this case, all three projections F_1^0, F_2^0, F_3^0 are openly disjoint and cover h_0, so ℓ^0 is contained in exactly one of them, say F_2^0. See Figure 2(b). The line λ through a^0 and ℓ^0 fully contains one of the two other facets, say F_3^0, on one side. As in (a), each of the quadrants Q, Q' lies fully on one (distinct) side of this line.

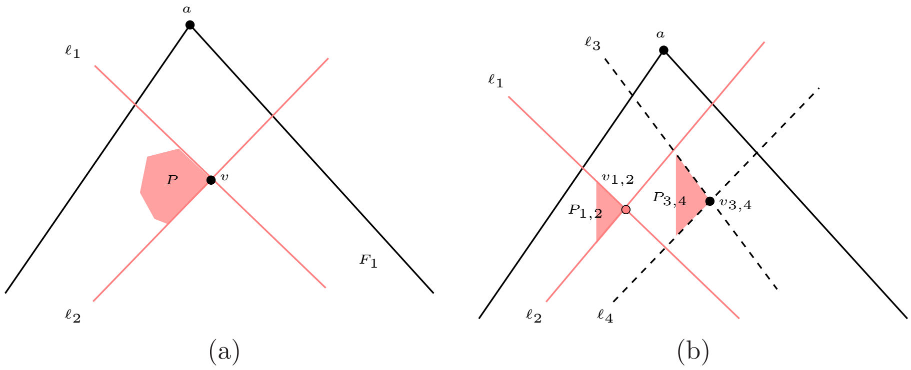

Figure 3: (a) The cross section of a cell of $\mathcal{A}(\mathcal{H})$ on F_1, and the pair of lines ℓ_1, ℓ_2 that it charges. (b) The vertices $v_{1,2}$ and $v_{3,4}$ cannot both be edges of G, because the polygon $P_{3,4}$ (partially shaded) is fully contained in the wedge spanned by the respective planes containing ℓ_1, ℓ_2 (and opposite to the wedge that contains $P_{1,2}$), which meets only two facets of W.

Hence, one of these quadrants cannot meet F_3^0, a contradiction that completes the proof. $\square$

Lemma 2.3 *Let T be an arbitrary tetrahedron. The overall number of cells C of $\mathcal{A}(R)$, for which at least three facets of T meet C, each as a long facet in C, is $O(r \log r)$.*

Proof: Let F_1, F_2, F_3 be a triple of facets of T, let W be the trihedral wedge induced by these facets, and let a denote its apex (which is a vertex of T). We show below that the overall number of cells C of $\mathcal{A}(R)$, for which all three facets of W meet C, each as a long facet in C, is $O(r \log r)$ (note that if F_i is long in C then the extended facet of W is also long). By repeating this argument for each triple of facets of T, the lemma follows.

Let $\mathcal{H}_0$ denote the set of all the planes in R that intersect W. Each such plane intersects W in either a wedge or a triangle (which might be unbounded). We first dispose of all planes that intersect W in a wedge. Each such plane h is disjoint from one of the facets of W, and thus one of the halfspaces that it induces, say, the positive side h^+ of h, meets only two facets of W. Thus all the cells of $\mathcal{A}(R)$ under consideration are contained in h^-. Hence all these cells lie in the convex polyhedron K, which is the intersection of the respective halfspaces h^- induced by the above "good" planes h.

Let $\mathcal{H}$ denote the set of "bad" planes in $\mathcal{H}_0$; each of them intersects W in a (possibly unbounded) triangle. Let $\mathcal{C}$ denote the collection of all cells of $\mathcal{A}(\mathcal{H})$ that meet all three facets of W but do not meet any edge of W. Fix a facet F_1 of W, form the intersections $C \cap F_1$, over all cells $C \in \mathcal{C}$, and denote by $\mathcal{C}_1$ the resulting collection of polygons. Note that $\mathcal{C}_1$ is a collection of cells of the 2-dimensional arrangement, within the plane h_{F_1} containing F_1, of the set $\mathcal{L}_1$ of the intersection lines between the planes of $\mathcal{H}$ and h_{F_1}. The number of unbounded polygons in $\mathcal{C}_1$ is thus $O(|\mathcal{H}|) = O(r \log r)$, so we focus on the bounded elements of this collection. Fix such a polygon P, and let v be its vertex which is the most counterclockwise as seen from the apex a (from some fixed side of h_{F_1}). Denote the two intersection lines that are incident to v and bound P by ℓ_1, ℓ_2, where ℓ_1 separates P and a within h_{F_1}, and ℓ_2 does not separate them;

see Figure 3(a). We then charge P to the pair (ℓ_1, ℓ_2); clearly, the charge is unique (the two respective planes h_1, h_2, which intersect h_{F_1} in ℓ_1, ℓ_2, can intersect only once on F_1).

Let G be the graph whose vertices are the intersection lines $\ell \equiv h \cap h_{F_1}$, for $h \in \mathcal{H}$, and whose edges are all the charged pairs (ℓ_1, ℓ_2) just defined. We claim that G is planar. It will then follow that the number of edges of G is at most $3|\mathcal{H}| - 6 = O(r \log r)$, which thus also bounds the number of cells of $\mathcal{A}(\mathcal{H})$ that meet all three facets of W. Any such cell C induces at most one cell of $\mathcal{A}(R)$ that can touch all three facets of W, namely the intersection $C \cap K$. Hence the number of such cells of $\mathcal{A}(R)$ is also $O(r \log r)$.

To establish the claim, assume, without loss of generality, that a is the origin in h_{F_1}, and apply the standard duality transform that maps points (u, v) to the respective lines $ux+vy+1 = 0$ and vice versa (where lines through a are ignored). This duality maps the lines ℓ in $\mathcal{L}_1$ to points ℓ^*, and each of the above pairs (ℓ_1, ℓ_2) is mapped to the segment connecting the points ℓ_1^*, ℓ_2^* dual to the respective lines ℓ_1, ℓ_2. By construction, and by the properties of this duality, any point q within the polygon $P \in \mathcal{C}_1$ which is represented by (ℓ_1, ℓ_2), is mapped to a line q^* that separates the origin o and ℓ_1^*, and has ℓ_2^* on the same side as the origin. That is, q^* intersects the segment $\ell_1^*\ell_2^*$. Conversely, for any line q^* that separates ℓ_1^* and ℓ_2^* as above, its primal point q must lie in the wedge between ℓ_1 and ℓ_2 that contains P. Moreover, if q^* separates ℓ_1^* and ℓ_2^* in the opposite way (i.e., so that the origin and ℓ_1^* lie on the same side of q^*), q lies in the opposite wedge between ℓ_1 and ℓ_2.

The collection of dual segments, as constructed above, defines a straight-line embedding of G in the dual plane, and we claim that this drawing is crossing-free. Indeed, suppose to the contrary that two edges $\ell_1^*\ell_2^*$, $\ell_3^*\ell_4^*$ of the drawing cross each other. The preceding discussion then implies that, back in the primal plane h_{F_1}, each of the resulting vertices $v_{1,2} = \ell_1 \cap \ell_2$, $v_{3,4} = \ell_3 \cap \ell_4$ lies in the double wedge of the other vertex that does not contain a. Denote by $P_{1,2}$ (resp., $P_{3,4}$) the polygon of $\mathcal{C}_1$ whose most counterclockwise vertex is $v_{1,2}$ (resp., $v_{3,4}$). In particular, one of these vertices, say $v_{1,2}$ lies clockwise to the other vertex $v_{3,4}$, in which case $v_{1,2}$ must lie in the wedge of ℓ_3, ℓ_4 that contains $P_{3,4}$, and $v_{3,4}$ must lie in the wedge of ℓ_1, ℓ_2 opposite to the one containing $P_{1,2}$. Let $C_{1,2}$ (resp., $C_{3,4}$) denote the cell of $\mathcal{A}(\mathcal{H})$ that contains $P_{1,2}$ (resp., $P_{3,4}$); also, for $i = 1, \ldots, 4$, let h_i denote the plane of $\mathcal{H}$ containing ℓ_i. Then, since $C_{1,2}$ meets all three facets of W, it follows by Lemma 2.2 that the wedge spanned by h_1, h_2, and opposite to the wedge containing $P_{1,2}$ (and $C_{1,2}$), meets only two facets of W, but then $C_{3,4}$ (which is clearly contained in this wedge, since it meets the wedge, and, being a cell of $\mathcal{A}(\mathcal{H})$, cannot cross h_1 or h_2) cannot meet all three facets of W, contrary to assumption; see Figure 3(b). This contradiction implies that G is planar, and this, as argued above, implies the assertion of the lemma. $\Box$

Remark: As already noted, Lemma 2.3 is fairly general, and makes no assumption about fatness of T. In fact, we believe that in certain circumstances it might also be generalized to situations where T is the boundary of a non-polyhedral convex shape. In this case, the assertion would be that the number of cells of $\mathcal{A}(R)$ that touch at least three pairwise disjoint connected sub-regions on T is $O(r \log r)$ (perhaps with some additional restrictions on these sub-regions, or with a larger number of sub-regions). We consider the lemma to be of independent interest, and believe that it (and/or extensions of it of the kind just suggested) will find additional applications in related problems.

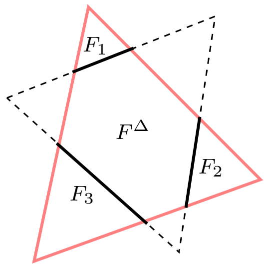

Figure 4: One facet F_Δ of Δ meets all three facets of T, in the depicted manner.

Lemma 2.4 *Let $T \in \mathcal{T}$ be a fixed tetrahedron, and let C be a cell of $\mathcal{A}(R)$ that meets at least three facets of T, but not any vertex of T. Then the number of simplicial subcells Δ of C that meet at least three facets of T, each as a long facet in Δ, is $O(\log r)$.*

Proof: As in Lemma 2.3, it is sufficient to assume that T is a trihedral wedge, and to show that the number of simplicial subcells Δ of C that meet all three facets of T, each as a long facet in Δ, is $O(\log r)$.

We first claim that if all the three facets F_1, F_2, F_3 of T are long in Δ, there must be one (triangular) facet F^Δ of Δ that meets all these facets. This easily follows from the fact that each of these facets intersects Δ in either a triangle or a quadrilateral, which yields at least 9 intersections between facets of T and facets of Δ. Since Δ has four facets, at least one of them must meet all three facets of T, as claimed. In addition, each of F_1, F_2, F_3 intersects F^Δ in a distinct pair of edges, as is easily verified; see Figure 4.

Let C_i denote the convex polytope obtained from C after $i-1$ steps of the Dobkin-Kirkpatrick hierarchical decomposition, for $i \geq 1$ (see [11] and earlier in this section). Recall that a simplicial subcell Δ is said to be generated at step i if it has a vertex v that is removed from C_i; that is, v belongs to the independent set of vertices of C_i collected at the i-th step. Recall also that v and its adjacent edges and facets are removed, they leave a hole in C_i. The convex hull of the other vertices of that hole is constructed, and its (triangular) facets are connected to v to form $O(1)$ simplices that fill up the corresponding gap between C_i and C_{i+1}, and Δ is one of these simplices. Note that the three other vertices of Δ belong to C_{i+1}, and that all three edges of Δ incident to v lie on the boundary of C_i. See Figure 1.

In what follows, we fix a decomposition step i, and show that there are only $O(1)$ simplices Δ of C that are generated at step i and have the properties in the lemma. The discussion above implies that for each such simplex Δ, the corresponding facet F^Δ appears either on the boundary of C_i, or on the boundary of C_{i+1}, or as an "internal" facet of a hole of C_i that is connected to the peeled-off vertex v of Δ, as described above.

Let u denote the apex of T. By assumption, $u \notin C$. Let $\mathcal{F}^{(i)}$ denote the collection of all facets F^Δ of simplicial subcells Δ of C that are generated at step i, such that F^Δ meets all three facets of T and such that these facets are all long in Δ. To simplify the analysis, we first prune away facets from $\mathcal{F}^{(i)}$, until $\mathcal{F}^{(i)}$ has the property that, for each peeled-off vertex v of C_i there is at most one simplex Δ incident to v, generated at step i, and contributing a facet

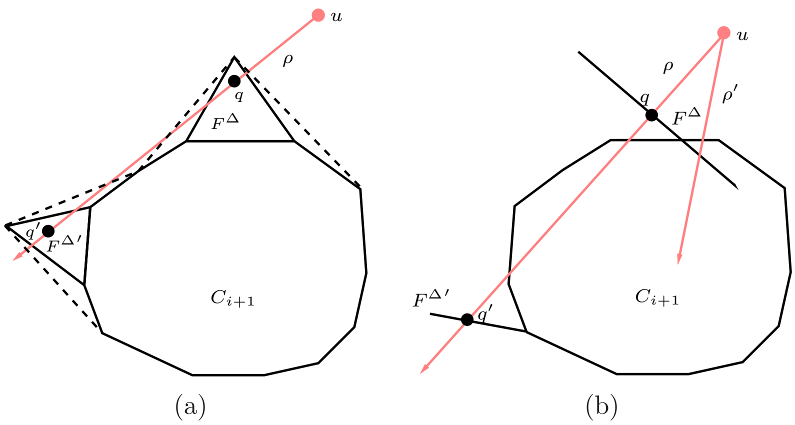

Figure 5: (a) The facets F^{Δ}, $F^{\Delta\prime}$ are internal to two distinct holes generated at step i, and are fully visible from u (with C_{i+1} opaque). The ray ρ that hits F^{Δ} at q and then $F^{\Delta\prime}$ at q' must cross ∂C_i at least twice between q and q', which is impossible. (b) Illustrating the proof that no ray from u can cross two distinct facets in $\mathcal{F}_7^{(i)}$.

to $\mathcal{F}^{(i)}$. By construction, this reduces the size of $\mathcal{F}^{(i)}$ by at most a constant factor.

We partition $\mathcal{F}^{(i)}$ into the following seven subcollections:

- $\mathcal{F}_1^{(i)}$, which consists of all facets of C_i in $\mathcal{F}^{(i)}$ that are visible from u (regarding C_i itself as opaque); that is, the relative interiors of all the segments connecting u to points on any $F^{\Delta} \in \mathcal{F}_1^{(i)}$ do not meet ∂C_i;
- $\mathcal{F}_2^{(i)}$, which consists of all facets of C_i in $\mathcal{F}^{(i)}$ that are invisible from u; that is, all the segments connecting u to any $F^{\Delta} \in \mathcal{F}_2^{(i)}$ cross ∂C_i (once) before reaching F^{Δ};
- $\mathcal{F}_3^{(i)}$, which consists of all facets of C_{i+1} in $\mathcal{F}^{(i)}$ that are not facets of C_i and are visible from u (regarding C_{i+1} itself as opaque); that is, the relative interiors of all the segments connecting u to points on any $F^{\Delta} \in \mathcal{F}_3^{(i)}$ do not meet ∂C_{i+1}; any such segment crosses ∂C_i (once) before reaching F^{Δ};
- $\mathcal{F}_4^{(i)}$, which consists of all facets of C_{i+1} in $\mathcal{F}^{(i)}$ that are not facets of C_i and are invisible from u; that is, all the segments connecting u to any $F^{\Delta} \in \mathcal{F}_4^{(i)}$ cross ∂C_{i+1} (once) before reaching F^{Δ}; as in the previous case, any such segment also crosses ∂C_i (once) before reaching F^{Δ};
- $\mathcal{F}_5^{(i)}$, which consists of all facets in $\mathcal{F}^{(i)}$ that are internal to the holes (components of $C_i \setminus C_{i+1}$) generated at step i, and are fully visible from u (in the presence of C_{i+1} as an opaque object);
- $\mathcal{F}_6^{(i)}$, same as $\mathcal{F}_5^{(i)}$, but consisting of facets that are fully invisible from u (fully occluded by C_{i+1}); and
- $\mathcal{F}_7^{(i)}$, same as $\mathcal{F}_5^{(i)}$, $\mathcal{F}_6^{\Delta}$, but consisting of facets that are partially visible from u (partially occluded by C_{i+1}).

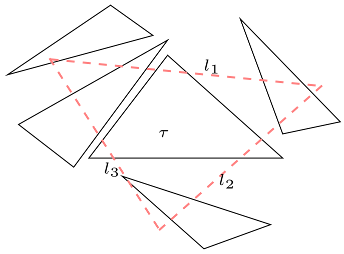

Figure 6: The centrally projected facets of some $\mathcal{F}_k^{(i)}$, and the central projections l_1, l_2, l_3 of the three respective facets F_1, F_2, F_3 of T. The triangle τ is the unique triangle that meets all three edges l_1, l_2, l_3.

We next claim that each subset $\mathcal{F}_k^{(i)}$ consists of at most one facet. This implies that $\mathcal{F}^{(i)}$ has constant size, which, since the decomposition has only $O(\log r)$ steps, implies the bound stated in the lemma. We first need the following easy technical claim.

Claim: If we project the triangles of $\mathcal{F}_k^{(i)}$, for any fixed $1 \le k \le 7$, centrally from u, the projected triangles are pairwise disjoint.

Proof: The claim easily follows for $\mathcal{F}_1^{(i)}$, $\mathcal{F}_2^{(i)}$, $\mathcal{F}_3^{(i)}$, $\mathcal{F}_4^{(i)}$ by definition and by the convexity of C_i, C_{i+1}. Consider $\mathcal{F}_5^{(i)}$, and assume to the contrary that it contains two facets F^Δ, $F^{\Delta\prime}$, such that a ray ρ emanating from u meets both of them, hitting, say, first F^Δ and then $F^{\Delta\prime}$ at two respective points q, q'. By the initial pruning process, F^Δ and $F^{\Delta\prime}$ lie in different holes of $C_i \setminus C_{i+1}$. By definition of $\mathcal{F}_5^{(i)}$, qq' is disjoint from C_{i+1}, and is fully contained in C_i, by convexity. This, however, is impossible, because qq' has to cross from some hole of $C_i \setminus C_{i+1}$ to a different one, and the boundary of such a hole is contained in $\partial C_i \cup \partial C_{i+1}$, and thus qq' must cross ∂C_i (at least twice), a contradiction; see Figure 5(a) for an illustration.

The case of $\mathcal{F}_6^{(i)}$ is argued similarly. Here again qq' is disjoint from C_{i+1}, because uq must have already crossed ∂C_{i+1} twice. Finally, for $\mathcal{F}_7^{(i)}$, we argue as follows. As above, the segment qq' is fully contained in C_i and crosses from one hole of $C_i \setminus C_{i+1}$ to another hole, so it must cross ∂C_{i+1} twice. Since F^Δ is partially occluded by C_{i+1}, there exists another ray ρ' from u that first crosses ∂C_{i+1} and then hits F^Δ. This, however, is impossible, since it would have implied that F^Δ and ∂C_{i+1} cross each other; this is proved by continuity, moving ρ towards ρ', within the plane that they span, and is illustrated in Figure 5(b). This completes the proof of the claim. $\square$

Let us now fix one of the subsets $\mathcal{F}_k^{(i)}$. The central projection of ∂T from u is a triangle whose three edges l_1, l_2, l_3 are the "head-on" projections of the respective facets F_1, F_2, F_3. Each facet $F^\Delta \in \mathcal{F}_k^{(i)}$ projects to a triangle that meets all three edges l_1, l_2, l_3. However, since the projections of the facets F^Δ of $\mathcal{F}_k^{(i)}$ are pairwise disjoint, at most one of them can touch all three edges l_1, l_2, l_3, as is easily checked; see Figure 6.

As argued above, this completes the proof of the lemma. $\square$

We have thus established the following theorem:

Theorem 2.5 *For any tetrahedron T, the overall number of simplicial cells Δ of Ξ that meet at least three facets of T is $O(r \log^2 r)$.*

Proof: Lemma 2.3 shows that only $O(r \log r)$ cells C of $\mathcal{A}(R)$ meet three facets of T. Of those, at most four contain an apex of T, and they have a total of $O(r \log r)$ sub-simplices. For any other cell, only $O(\log r)$ of its simplices have this property, as shown in Lemma 2.4. $\Box$

Remark: With some additional care, the proof of Lemma 2.4 can be extended to the case where C contains an apex of T. However, as just argued, the validity of Theorem 2.5 does not require this stronger property.

2.3 The overall recursive analysis

We now apply the following recursive scheme (a similar scheme has been presented in [16]). Each step of the analysis involves a simplex Δ_0, which, in the initial step, is the entire 3-space, (or, rather, a sufficiently large simplex that contains all the vertices in the arrangement of the tetrahedra), and in further recursive steps is a cell of a cutting of some larger simplex, from the preceding recursive level.

We construct a $(1/r)$-cutting Ξ of the arrangement of the planes that support facets of the tetrahedra that cross Δ_0, using the Dobkin-Kirkpatrick hierarchical decomposition of each cell of the corresponding arrangement, as described above. Let Δ be a simplicial cell of Ξ and let $\mathcal{D}^\Delta$ (resp., $\mathcal{T}^\Delta$) denote the set of D-tetrahedra (resp., T-tetrahedra) within Δ. Put $N_D = N_D^\Delta := |\mathcal{D}^\Delta|$, and $N_T = N_T^\Delta := |\mathcal{T}^\Delta|$. (In the actual construction of the cutting we draw two respective random samples of $O(r \log r)$ planes, where the first sample is taken from the facets of the D-tetrahedra in Δ_0, and second one is taken from the facets of the T-tetrahedra in Δ_0. This guarantees, with high probability, the property that the number of D-tetrahedra (resp., T-tetrahedra) in each subcell of Δ_0 is at most $\frac{|\mathcal{D}^{\Delta_0}|}{r}$ (resp., $\frac{|\mathcal{T}^{\Delta_0}|}{r}$)).

During each step of the recursion, after partitioning Δ_0 into smaller subcells Δ, we immediately dispose of any *new* DDD and DDT vertices within each subcell Δ, and show that the overall number of these vertices is $O\left((N_D + N_T){N_T}^{1+\varepsilon}\right)$, for any $\varepsilon > 0$. These vertices are not considered during any further recursive substep, but only the remaining DTT and TTT vertices.

The recursion bottoms out when $N_T \le c$, for some absolute constant $c \ge 3$. In this case we bound the number of the remaining inner DTT and TTT vertices of the union in a brute-force manner, and thus obtain an overall bound of $O({N_T}^2 N_D + {N_T}^3) = O(N_D + 1)$ on the number of these vertices.

To bound the number of DDD vertices in Δ, we replace each tetrahedron in $\mathcal{D}^\Delta$ by the equivalent halfspace or dihedral wedge, and face the problem of bounding the overall number of vertices appearing on the boundary of the union of N_D halfspaces and α'-fat dihedral wedges[4] (where, as argued earlier, $\alpha' > 0$ is a constant that depends only on α). As shown in [24], the number of such vertices is $O(N_D^{2+\varepsilon})$, for any $\varepsilon > 0$, where the constant of proportionality

[4]Clearly, any DDD-vertex of the full union in Δ is also a vertex of the union of $\mathcal{D}^\Delta$.

depends on ε and α. In an additional major step of the analysis, we derive, in Section 2.4 below, a similar bound on the number of DDT-vertices.

The recursive scheme. With all this machinery at hand, we can now proceed to the proof of Theorem 2.1. The analysis begins with the initial cell Δ_0, with $N_T^{\Delta_0} = n$, $N_D^{\Delta_0} = 0$, and recursively subdivides Δ_0, using a $(1/r)$-cutting of the preceding kind, for some sufficiently large constant parameter r.

For simplicity, write in what follows $n_T = N_T^{\Delta_0}$, and $n_D = N_D^{\Delta_0}$. Theorem 2.5 implies that there are only $O(rn_T \log^2 r)$ crossings between the cells of Ξ and their T-tetrahedra. It thus follows that, for any $r^2 \log^2 r \le s \le M = O(r^3 \log^3 r)$, where M is the size of Ξ, the number of cells in Ξ that are crossed by at least $\dfrac{rn_T \log^2 r}{s}$ T-tetrahedra is at most $O(s)$. (The case $s < r^2 \log^2 r$ cannot arise, since each cell of Ξ is intersected by at most $\frac{n_T}{r}$ tetrahedra of $\mathcal{T}^{\Delta_0}$.) We partition the set of cells of Ξ into at most $\log\left(\frac{M}{r^2 \log^2 r}\right) = \Theta(\log r)$ subsets, so that the i-th subset Ξ_i contains $O(2^i r^2 \log^2 r)$ cells Δ of Ξ, each of which satisfies

$$\frac{n_T}{2^i r} \le N_T^{\Delta} = |\mathcal{T}^{\Delta}| \le \frac{2n_T}{2^i r},$$

for $i = 1, \ldots, \log\left(\frac{M}{r^2 \log^2 r}\right)$. Note that $|\mathcal{T}^{\Delta}| = O\left(\frac{n_T}{r^2}\right)$ for each of the $O(r^3 \log^3 r)$ cells Δ in the last subset. For the number of D-tetrahedra in any cell Δ, we use the bound $N_D^{\Delta} = |\mathcal{D}^{\Delta}| \le \frac{n_D + n_T}{r}$, for each $\Delta \in \Xi$.

As in [16], we recurse in each cell Δ of Ξ, where the goal of the recursive step at Δ is to obtain an upper bound for the number of DTT and TTT vertices in Δ (including vertices of these kinds that appear on $\partial\Delta$). Thus, before entering the recursion, we need to bound the number of *new* DDD and DDT vertices within Δ (or on its boundary). These are vertices v that were DTT or TTT vertices at the parent cell Δ_0 of Δ, but have become DDT or DDD vertices at Δ. Partition $\mathcal{D}^{\Delta_0}$ into $k = \lceil n_D/n_T \rceil$ subsets $\mathcal{D}_1, \ldots, \mathcal{D}_k$, each consisting of at most n_T D-tetrahedra. The preceding observations imply that any new DDD or DDT vertex in Δ must be a DDD or a DDT vertex of the union of the wedges and tetrahedra of $\mathcal{D}_j \cup \mathcal{T}^{\Delta_0}$, clipped to Δ, for some $j = 1, \ldots, k$ (because it was a DTT or a TTT-vertex in Δ_0, and so involves at most one wedge of $\mathcal{D}^{\Delta_0}$). By the results of [24] and of Section 2.4, the number of such vertices, for any fixed j, is thus $O({n_T}^{2+\varepsilon})$, for any $\varepsilon > 0$. Hence, summing over $j = 1, \ldots, k$, the overall number of new DDD and DDT vertices in Δ is $O(k{n_T}^{2+\varepsilon}) = O((n_D + n_T){n_T}^{1+\varepsilon})$, for any $\varepsilon > 0$. Repeating the analysis to each subcell Δ of Δ_0, and recalling that r is a constant, the overall number of new DDD and DDT vertices within the children of Δ_0 is

$$O(r^3 \log^3 r \cdot (n_D + n_T){n_T}^{1+\varepsilon}) = O((n_D + n_T){n_T}^{1+\varepsilon}),$$

for any $\varepsilon > 0$.

Let $U(N_T, N_D)$ denote the maximum number of DTT- and TTT-vertices that appear on the boundary of the union at a recursive step involving up to N_T T-tetrahedra and N_D D-tetrahedra. Then U satisfies the following recurrence:

$$U(N_T, N_D) \leq \begin{cases} O\left((N_D + N_T)N_T^{1+\varepsilon}\right) \\ + \sum_{i=0}^{\log\left(\frac{M}{r^2 \log^2 r}\right)} O(2^i r^2 \log^2 r) U\left(\frac{2N_T}{2^i r}, \frac{N_D+N_T}{r}\right), & \text{if } N_T > c, \\ \\ O(N_D + 1), & \text{if } N_T \leq c, \end{cases}$$

where $\varepsilon > 0$ is arbitrary, $c \geq 3$ is an appropriate constant, and where the constant of proportionality in the first expression depends on (ε, α and on) r. It is straightforward to verify (see also [16]), that the solution of this recurrence is $U(N_T, N_D) = O(N_T(N_T + N_D)^{1+\varepsilon})$, for any $\varepsilon > 0$, with a constant of proportionality that depends on ε and on α.

Substituting the initial values $N_T = n$, $N_D = 0$, we conclude that the overall combinatorial complexity of the union is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, as asserted. This completes the proof of Theorem 2.1. modulo the still missing analysis of the number of DDT-vertices.

Since a cube in 3-space can be partitioned into a constant number of α-fat tetrahedra, for some appropriate constant parameter $\alpha > 0$, we obtain the following extension of the bound in [24]:

Corollary 2.6 *The complexity of the union of n arbitrarily oriented cubes in $\mathbb{R}^3$, of arbitrary side lengths, is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, where the constant of proportionality depends on ε.*

Similar results can be obtained for any collection of polyhedral objects that can be decomposed into, or covered by, a total of n α-fat tetrahedra.

A similar, almost verbatim, analysis yields the bound $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, for the complexity of the union of n α-fat trihedral wedges. For the sake of completeness, we state this result explicitly:

Corollary 2.7 *The complexity of the union of n α-fat trihedral wedges is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, where the constant of proportionality depends on ε and α.*

Our analysis can easily be extended to the problem of bounding the complexity of the union of n α-fat arbitrarily oriented triangular prisms, with cross-sections of arbitrary sizes. In this case, we apply a similar decomposition scheme as in the case of α-fat tetrahedra, exploiting similar properties to those stated in Lemmas 2.3 and 2.4 for the case where T is a triangular prism (rather than a trihedral wedge) — simply think of a prism as a wedge with apex at infinity. A lower bound construction that yields $\Omega(n^2\alpha(n))$ vertices on the boundary of the union can be supplied using similar techniques as in the original problem. We thus obtain:

Corollary 2.8 *The complexity of the union of n arbitrarily oriented α-fat triangular prisms, with cross-sections of arbitrary sizes, in $\mathbb{R}^3$, is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, where the constant of proportionality depends on ε and α. The bound is almost tight in the worst case.*

Similar results can be obtained for any collection of polyhedral prisms that can be decomposed into, or covered by a total of n α-fat triangular prisms.

2.4 The number of DDT-vertices

In this section we provide the missing ingredient of the preceding analysis, showing that the number of DDT-vertices is nearly-quadratic.

We thus have, at each step of the analysis, a simplex Δ, a set $\mathcal{D} = \mathcal{D}^{\Delta}$ of N_D α'-fat dihedral wedges, and a set $\mathcal{T} = \mathcal{T}^{\Delta}$ of N_T α-fat tetrahedra[5]. Our goal is to obtain a nearly-quadratic bound on the number of DDT vertices on the boundary of the union of $\mathcal{D} \cup \mathcal{T}$ within Δ. We may assume that $N_T \leq N_D$. Otherwise, we apply (in "reverse") the partitioning trick used in the preceding analysis. That is, we partition $\mathcal{T}$ arbitrarily into $k = \lceil N_T/N_D \rceil$ subsets $\mathcal{T}_1, \ldots, \mathcal{T}_k$, of size at most N_D each, establish a bound $O(N_D^{2+\varepsilon})$ on the number of DDT-vertices on each union $\mathcal{D} \cup \mathcal{T}_i$, $i = 1, \ldots, k$, and add up the bounds, to obtain an overall bound of $O(N_T N_D^{1+\varepsilon})$. In other words, the goal is to establish the upper bound $O(N_D^{2+\varepsilon})$ on the number of DDT-vertices, assuming that $N_T \leq N_D$.

Let L denote the set of the xy-projections of the edges (lines) of the wedges in $\mathcal{D}$. (We assume that the coordinate system is generic, so none of these lines projects to a single point.) We construct a $(1/r)$-cutting Ξ of the planar arrangement $\mathcal{A}(L)$, by taking a random sample R of $O(r \log r)$ lines of L, for some sufficiently large constant parameter r, constructing the planar arrangement $\mathcal{A}(R)$, and triangulating each of its cells using the two-dimensional version of the Dobkin-Kirkpatrick hierarchical decomposition of a convex polygon [11]. We obtain $O(r^2 \log^2 r)$ triangles, and we may assume that the sample R is such that each triangle is crossed by at most N_D/r lines of L (this indeed happens with high probability). We lift each cell of $\mathcal{A}(R)$, and each of its sub-triangles, into a vertical prism (or rather its portion within the current simplex Δ). Each triangular prism σ is crossed by at most N_D/r edges of the wedges of $\mathcal{D}$. Any other wedge either misses σ altogether, or each of its bounding halfplanes that meets σ crosses σ completely, cutting it into two disconnected pieces (as if it were a plane).

We now claim that, given a fixed tetrahedron $T \in \mathcal{T}$, the overall number of vertical triangular prisms σ, erected over the cells of Ξ, such that σ meets at least three facets of T, is $O(r \log^2 r)$. Indeed, a facet F of T whose bounding edges do not meet σ must intersect σ in a triangle whose boundary is contained in $\partial\sigma$. Since $T \cap \sigma$ is convex, there can be at most two such facets. Hence σ meets at least one of the edges e of T. That is, the projection e^* of e crosses the triangular cell of Ξ which is the base of σ. However, applying a simplified version of the argument used above for 3-dimensional arrangements, e^* can cross only $O(r \log r)$ cells of $\mathcal{A}(R)$, and only $O(\log r)$ triangles within each cell, from which the claim follows.

Summing up, we have so far $O(r^2 \log^2 r)$ subproblems, each defined within a triangular prism σ, and involves the following sets of objects: (i) The set $\mathcal{D}^\sigma$ of dihedral wedges of $\mathcal{D}$ whose edges cross σ; (ii) the set $\mathcal{P}^\sigma$ of dihedral wedges of $\mathcal{D}$ that cross σ but whose edges do not cross σ (we can think of each member of $\mathcal{P}^\sigma$ as either a halfspace or a region enclosed between a pair of planes crossing σ); (iii) the set $\mathcal{T}^\sigma$ of tetrahedra such that at least three of their facets cross σ; and (iv) the set $\mathcal{W}^\sigma$ of tetrahedra that cross σ, and at most two of their facets cross σ. Put $N_{\mathcal{D}^\sigma} := |\mathcal{D}^\sigma|$, $N_{\mathcal{P}^\sigma} := |\mathcal{P}^\sigma|$, $N_{\mathcal{T}^\sigma} := |\mathcal{T}^\sigma|$, and $N_{\mathcal{W}^\sigma} := |\mathcal{W}^\sigma|$. As just argued, $\sum_\sigma N_{\mathcal{T}^\sigma} = O(N_T \cdot r \log^2 r)$.

[5] Some of these tetrahedra may be trihedral wedges, in case only three of the facets of a tetrahedron T appear in Δ. However, to simplify the presentation, we will refer in what follows to all the elements of $\mathcal{T}$ as tetrahedra.

The goal is to bound the number of inner vertices v, within σ, of the union $\mathcal{D}^\sigma \cup \mathcal{P}^\sigma \cup \mathcal{T}^\sigma \cup \mathcal{W}^\sigma$, such that v is incident to the boundaries of two objects in $\mathcal{D}^\sigma \cup \mathcal{P}^\sigma$, and of one object in $\mathcal{T}^\sigma \cup \mathcal{W}^\sigma$. We classify each inner vertex v in σ of this kind as either DDW, if v is incident to the boundaries of two objects in $\mathcal{D}^\sigma$ and one object in $\mathcal{W}^\sigma$, DPW, if the three objects whose boundaries are incident to v are in $\mathcal{D}^\sigma$, $\mathcal{P}^\sigma$, and $\mathcal{W}^\sigma$, respectively, PPW, if two of these objects are in $\mathcal{P}^\sigma$ and one in $\mathcal{W}^\sigma$, DDT, if two of these objects are in $\mathcal{D}^\sigma$ and one in $\mathcal{T}^\sigma$, DPT, if the objects are in $\mathcal{D}^\sigma$, $\mathcal{P}^\sigma$, and $\mathcal{T}^\sigma$, respectively, or PPT, if two of the objects are in $\mathcal{P}^\sigma$ and one in $\mathcal{T}^\sigma$.

Since each vertex of type DDW, DPW, or PPW lies on the boundary of the union of α'-fat dihedral wedges (or halfspaces), it follows, by applying the results of [24], that the number of such vertices is $O((N_{\mathcal{D}^\sigma}+N_{\mathcal{P}^\sigma}+N_{\mathcal{W}^\sigma})^{2+\varepsilon})$, for any $\varepsilon > 0$. Summing over all prisms σ, and using the facts that r is constant and that $N_T \le N_D$, we obtain the overall bound $O(N_D^{2+\varepsilon})$, for any $\varepsilon > 0$. The general bound, using the partitioning trick described above, is $O((N_D + N_T)N_D^{1+\varepsilon})$, for any $\varepsilon > 0$. We next show how to bound the number of the remaining types of vertices.

Assume for the moment that we have managed to establish a nearly-quadratic bound, of the form $O((N_D + N_T)N_D^{1+\varepsilon})$, for any $\varepsilon > 0$, on the number of PPT-vertices and DPT-vertices (which will be accomplished in the next two steps of the analysis). Then we are only left with the task of bounding the number of DDT-vertices within σ, which we do recursively. To recap, there are $O(r^2 \log^2 r)$ such recursive subproblems, over all triangular prisms σ, in each of which we apply the nearly-quadratic bound on the number of vertices of all the remaining types, and continue to bound the number of DDT-vertices in a recursive manner.

The recursion bottoms out when either $N_T \le c$ or $N_D \le c$, for some absolute constant $c \ge 3$. We then bound the number of the remaining vertices of the union under consideration (that is, inner vertices whose type is still DDT) in a brute-force manner, and thus obtain an overall bound of $O({N_D}^2 N_T) = O(N_D^2 + N_T)$ on the number of these vertices.

Let $U_1(N_T, N_D)$ denote the maximum number of DDT-vertices that appear on the boundary of the union at a recursive step involving N_D dihedral wedges and N_T tetrahedra. Then U_1 satisfies the recurrence:

$$U_1(N_T, N_D) \le \begin{cases} O\left((N_D + N_T)N_D^{1+\varepsilon}\right) + \sum_{i=0}^{\log\left(\frac{M}{r \log^2 r}\right)} O(2^i r \log^2 r) U_1\left(\frac{2N_T}{2^i}, \frac{N_D}{r}\right), & \text{if } \min\{N_T, N_D\} > c, \\ O(N_D^2 + N_T), & \text{if } \min\{N_T, N_D\} \le c, \end{cases} \tag{1}$$

where $\varepsilon > 0$ is arbitrary, $c \ge 3$ is an appropriate constant, $M = O(r^2 \log^2 r)$ is the overall number of prisms in the decomposition, and where the constants of proportionality in the non-recursive terms depend on r (and on ε, α). This follows from the fact that the number of prisms σ with $\frac{N_T}{2^i} \le N_{\mathcal{T}^\sigma} < \frac{2N_T}{2^i}$ is at most $O(2^i r \log^2 r)$. As above, it is easy to verify that the solution of this recurrence is $U_1(N_T, N_D) = O((N_D + N_T)N_D^{1+\varepsilon})$, for any $\varepsilon > 0$, with a constant of proportionality that depends on ε and on α.

The number of PPT-vertices. To bound the number of PPT vertices, we launch a new recursive analysis, which, as the analysis in Section 2.3, is based on cuttings in arrangements of planes in 3-space. Recycling for the moment some of the previous notations, we have, at each step, a subproblem within some simplex Δ_0, involving a set $\mathcal{P} = \mathcal{P}^{\Delta_0}$ of pairs of planes, at least one of which crosses Δ_0 (but their intersection line does not cross Δ_0), and a set $\mathcal{T} = \mathcal{T}^{\Delta_0}$ of tetrahedra, so that, for each $T \in \mathcal{T}$, at least three of its facets cross Δ_0. Put $N_P = |\mathcal{P}^{\Delta_0}|$, $N_T = |\mathcal{T}^{\Delta_0}|$. Initially, Δ_0 is the (clipped) vertical triangular prism σ of some specific recursive instance of the above recursion that handles DDT vertices.

We first draw a random sample $R \subset \mathcal{P}^{\Delta_0}$ of $O(r \log r)$ pairs of planes, for some sufficiently large constant parameter r, and construct the sampled arrangement $\mathcal{A}(R)$ within Δ_0. We then collect only the cells in the complement of the union of the wedges enclosed between each sampled pair of planes (within Δ_0). Since the wedges are all α'-fat, the analysis of [24] implies that the overall number of these cells is $O(r^{2+\varepsilon})$, for any $\varepsilon > 0$. Furthermore, since R is a collection of planes within Δ_0, each cell that we consider is a convex (possibly unbounded) polyhedron. We next triangulate each of these cells C using the Dobkin-Kirkpatrick hierarchical decomposition of convex polytopes (see the beginning of this section and [11]), and obtain an overall number of $O(r^{2+\varepsilon})$ simplicial subcells. Using similar considerations as in the original problem, we may assume that each simplicial cell Δ of the resulting decomposition is crossed by at most N_P/r wedge boundaries (pairs of planes) in $\mathcal{P}^{\Delta_0}$, and each edge of any tetrahedron in $\mathcal{T}^{\Delta_0}$ crosses at most $O(r \log^2 r)$ cells. In addition, Theorem 2.5 implies that the overall number of simplicial cells, each of which meets at least three facets of any fixed tetrahedron $T \in \mathcal{T}^{\Delta_0}$, is $O(r \log^2 r)$.

Note that, in this problem, the decomposition generates only two types of vertices, PPW and PPT. We thus apply the above decomposition recursively, where we dispose immediately (i.e., derive a nearly-quadratic bound on the number) of all PPW-vertices within each cell Δ of the decomposition, and bound the number of the (remaining) PPT-vertices recursively. At the bottom of the recurrence we bound the number of PPT-vertices by brute force, as above. An appropriate variant of the preceding analysis leads to a recurrence relationship similar to (1), with the difference that (i) N_P replaces N_D, and (ii) the upper bound on M is $O(r^{2+\varepsilon})$, rather than $O(r^2 \log^2 r)$; this, however, has no effect on the asymptotic solution of the recurrence. That is, we obtain that the maximum number of PPT-vertices that appear on the boundary of the union at a recursive step, involving N_P α'-fat dihedral wedges (which behave like pairs of planes) and N_T α-fat tetrahedra, is $O((N_P + N_T)N_P^{1+\varepsilon})$, for any $\varepsilon > 0$.

The number of DPT-vertices. Here too we bound the number of DPT-vertices using a separate recursive analysis, where, at each step, we have a subproblem within some simplex Δ_0, involving a set $\mathcal{D} = \mathcal{D}^{\Delta_0}$ of dihedral wedges whose boundary edges cross Δ_0, a set $\mathcal{P} = \mathcal{P}^{\Delta_0}$ of pairs of planes, at least one of which crosses Δ_0, and a set $\mathcal{T} = \mathcal{T}^{\Delta_0}$ of tetrahedra, so that, for each $T \in \mathcal{T}$, at least three of its facets cross Δ_0. Put $N_D = |\mathcal{D}^{\Delta_0}|$, $N_P = |\mathcal{P}^{\Delta_0}|$, and $N_T = |\mathcal{T}^{\Delta_0}|$. Initially, Δ_0 is a (clipped) vertical triangular prism, as above.

We choose some sufficiently large constant parameter r, and draw three random samples, each of which consists of $O(r \log r)$ planes, which contain the facets of the wedges of $\mathcal{D}^{\Delta_0}$, the facets of the wedges of $\mathcal{P}^{\Delta_0}$, and the facets of the tetrahedra of $\mathcal{T}^{\Delta_0}$, respectively. Let R denote the union of the three samples. We form the arrangement $\mathcal{A}(R)$, and triangulate each of its cells, using, as usual, the Dobkin-Kirkpatrick hierarchical decomposition.

We obtain $O(r^3 \log^3 r)$ simplicial cells in the decomposition. Assuming that the drawn samples are good, we may assume consequently that each of these cells Δ is crossed by at most N_D/r dihedral wedges of $\mathcal{D}^{\Delta_0}$, at most N_P/r dihedral wedges (bounded by the pairs of planes) of $\mathcal{P}^{\Delta_0}$, and at most N_T/r tetrahedra of $\mathcal{T}^{\Delta_0}$. Each edge of any tetrahedron in $\mathcal{T}^{\Delta_0}$ crosses at most $O(r\log^2 r)$ cells, and the overall number of simplicial cells, each of which meets at least three facets of a fixed tetrahedron $T \in \mathcal{T}^{\Delta_0}$, is $O(r\log^2 r)$. Each edge ℓ of a dihedral wedge of $\mathcal{D}^{\Delta_0}$ crosses only $O(r\log^2 r)$ cells, so the overall number of wedge-cell crossings, for which the edge of the wedge appears in the cell, is $O(N_D \cdot r\log^2 r)$. In any other crossing of a cell Δ by a dihedral wedge, the wedge behaves like a pair of planes, at least one of which crosses Δ.

The decomposition therefore generates, within each simplicial cell Δ, vertices of type PPW, DPW, PPT, and DPT. The number of vertices of the first three types is nearly-quadratic, by the bound in [24] and by the preceding analysis, and the number of DPT-vertices, within each cell Δ, is bounded recursively. As in the original recursive scheme, presented in Section 2.3, at each step of the recursion, we only need to bound the number of new PPW, DPW, and PPT vertices within Δ. Each such vertex is incident to the boundary of a dihedral wedge in $\mathcal{D}^{\Delta_0}$, a plane in (a pair in) $\mathcal{P}^{\Delta_0}$, and a tetrahedron in $\mathcal{T}^{\Delta_0}$. We can refine the quadratic bound on the number of these vertices using the following variant of the partition trick. That is, suppose, without loss of generality, that $N_D \le N_P \le N_T$. Partition $\mathcal{T}^{\Delta_0}$ into $\lceil N_T/N_P \rceil$ sets, each of size at most N_P. Bound separately the number of DPT vertices of the above kind for $\mathcal{D}^{\Delta_0}$, $\mathcal{P}^{\Delta_0}$, and each subset of $\mathcal{T}^{\Delta_0}$. The bound is $O(N_P^{2+\varepsilon})$, for any $\varepsilon > 0$, within each of these subproblems, for a total of

$$O\left(N_P^{2+\varepsilon} \cdot \frac{N_T}{N_P}\right) = O(N_T N_P^{1+\varepsilon}),$$

for any $\varepsilon > 0$. Considering also the other symmetric cases, we have thus shown that the number of new PPW, DPW, and PPT vertices is $O\left((N_D N_P + N_D N_T + N_P N_T) \cdot \max\{N_D, N_P, N_T\}^{\varepsilon}\right)$, for any $\varepsilon > 0$.

We do not have to process in further recursive steps dihedral wedges of $\mathcal{D}^{\Delta_0}$ whose edge does not meet the current cell Δ, as well as tetrahedra of $\mathcal{T}^{\Delta_0}$ with at most two facets meeting Δ, since the DPT-vertices that they induce have already been counted.

At the bottom of the recursion (when $\min\{N_D, N_P, N_T\} \le c$, for some absolute constant $c \ge 3$), we bound the number of the remaining inner DPT-vertices of the union in a brute-force manner, and thus obtain an overall bound of $O(N_D N_P N_T) = O(N_D N_P + N_D N_T + N_P N_T)$ on this number.

Let $U_2(N_D, N_P, N_T)$ denote the maximum number of DPT-vertices that appear on the boundary of the union at a recursive step within some simplex Δ_0, involving a set $\mathcal{D}^{\Delta_0}$ of N_D dihedral wedges, a set $\mathcal{P}^{\Delta_0}$ of N_P pairs of planes (dihedral wedges whose edge does not cross Δ_0), and a set $\mathcal{T}^{\Delta_0}$ of N_T tetrahedra. For each $i \ge 1$, consider those cells Δ for which

$$\frac{1}{2^i r} < \max\left\{\frac{N_{\mathcal{D}^\Delta}}{N_D}, \frac{N_{\mathcal{T}^\Delta}}{N_T}\right\} \le \frac{1}{2^{i-1} r}.$$

Since, as argued above, $\sum_\Delta N_{\mathcal{D}^\Delta} = O(N_D \cdot r\log^2 r)$, and $\sum_\Delta N_{\mathcal{T}^\Delta} = O(N_T \cdot r\log^2 r)$, it follows that the number of cells Δ, satisfying the above inequalities, is $O(2^i r^2 \log^2 r)$. Moreover, all cells appear in these counts, because, by construction, we have $N_{\mathcal{D}^\Delta} \le N_D/r$ and $N_{\mathcal{T}^\Delta} \le N_T/r$, for each cell Δ. Hence, U_2 satisfies the following recurrence:

$$U_2(N_D, N_P, N_T) \le \begin{cases} O\left((N_D N_P + N_D N_T + N_P N_T) \cdot \max\{N_D, N_P, N_T\}^{\varepsilon}\right) + \\ \sum_{i=0}^{\log\left(\frac{M}{r^2 \log^2 r}\right)} O(2^i r^2 \log^2 r) U_2\left(\frac{2N_D}{2^i r}, \frac{N_P}{r}, \frac{2N_T}{2^i r}\right), & \text{if } \min\{N_D, N_P, N_T\} > c, \\ O(N_D N_P + N_D N_T + N_P N_T), & \text{if } \min\{N_D, N_P, N_T\} \le c, \end{cases}$$

where $\varepsilon > 0$ is arbitrary, $c \ge 3$ is an appropriate constant, $M = O(r^3 \log^3 r)$ is the overall number of cells in the decomposition (at the current recursive step), and the non-recursive terms also depend on r (and on ε, α). Again, it is easy to verify that the solution of this recurrence is $U_2(N_D, N_P, N_T) = O\left((N_D N_P + N_D N_T + N_P N_T) \cdot \max\{N_D, N_P, N_T\}^{\varepsilon}\right)$, for any $\varepsilon > 0$, with a constant of proportionality that depends on ε and on α (see also [16] for similar considerations).

This finally completes the analysis, and establishes Theorem 2.1. $\square$

3 The Union of α-Fat Triangles in the Plane

Let $\mathcal{T}$ be a collection of n α-fat triangles in the plane (i.e., each angle of any triangle is at least α). In this section we follow a simple variant of our approach to the three-dimensional problem, and derive a nearly-linear bound on the combinatorial complexity of the union of the triangles in $\mathcal{T}$.

We first draw a random sample R of $O(r \log r)$ of the lines containing the edges of the triangles in $\mathcal{T}$, for some sufficiently large constant parameter r, and form the arrangement $\mathcal{A}(R)$. We then triangulate each cell of the arrangement, using, e.g., bottom-vertex triangulation. The number of the resulting triangles (we will call them simplices, to avoid confusion with the triangles of $\mathcal{T}$) is $M = O(r^2 \log^2 r)$, and, with high probability, each of these simplices is crossed by at most n/r edges of the triangles in $\mathcal{T}$. We can therefore assume that our sample has this property. Thus the resulting decomposition is a $(1/r)$-cutting for the edges of $\mathcal{T}$, and the simplices are the cells of this cutting.

Similarly to the original problem, we fix a triangle $T \in \mathcal{T}$, and a cell Δ of the cutting that T meets, and classify T as being either a *W-triangle* in Δ, if Δ meets only one or two edges of T, or a *T-triangle* in Δ, if Δ meets all the three edges of T. As a consequence, each intersection vertex v of the union boundary that appears in Δ is classified as being either WW, if the two edges that are incident to v belong to two respective W-triangles in Δ, WT, if one of these edges belongs to a W-triangle and the other belongs to a T-triangle, or TT, if both of these edges belong to two respective T-triangles. (In all three cases, the relevant triangles are distinct.)

We next observe the easy fact that, for any triangle $T \in \mathcal{T}$, there is only a *single* cell of the cutting that meets all three edges of T; see Lemma 2.4 and Figure 6.

We now apply a recursive scheme, similar to that used in the three-dimensional setup. Let Δ be a simplex of the cutting and let $\mathcal{W}^\Delta$ (resp., $\mathcal{T}^\Delta$) denote the set of W-triangles (resp., T-triangles) within Δ. Put $N_W^\Delta := |\mathcal{W}^\Delta|$, and $N_T^\Delta := |\mathcal{T}^\Delta|$. The preceding observation implies that $\sum_\Delta N_T^\Delta \le N_T$, where N_T is the overall number of triangles.

During each step of the recursion, we immediately dispose of any new WW- and WT-vertices within each subcell Δ, and continue to bound the number of TT-vertices recursively. The recursion bottoms out when $N_T \le c$, for some absolute constant $c \ge 2$. In this case the number of the remaining intersection vertices of the union (within the current Δ) is $O(1)$.

To bound the number of WW-vertices in Δ, we replace each triangle in $\mathcal{W}^\Delta$ by the equivalent halfplane or wedge, and face the problem of bounding the overall number of vertices appearing on the boundary of the union of N_W^Δ halfplanes and *α-fat* wedges (that is, wedges whose angle is at least α). As shown in [14], the number of such vertices is $O(N_W^\Delta)$. As we will shortly show, the number of WT-vertices in Δ is $O\left((N_T^\Delta + N_W^\Delta)\cdot(N_W^\Delta)^\varepsilon\right)$, for any $\varepsilon > 0$. Summing over all Δ, and using the fact that r is a constant, we get the overall bound $O(N_T^{1+\varepsilon})$, for any $\varepsilon > 0$.

Let $U_1(N_T)$ denote the maximum number of intersection vertices that appear on the boundary of the union at a recursive step involving N_T triangles. For each $1 \le i \le \log(M/r)$, where $M = O(r^2 \log^2 r)$ is the overall number of cells in the cutting, the number of cells Δ with $\frac{N_T}{r2^i} < N_T^\Delta \le \frac{2N_T}{r2^i}$ is at most $2^i r$, and this also holds for the set of all the remaining cells (with $N_T^\Delta \le \frac{2N_T}{r2^i}$, where $i = \log(M/r) + 1$). Recall also that $N_T^\Delta, N_W^\Delta \le \frac{N_T}{r}$ always holds, by construction. Hence U_1 satisfies the recurrence:

$$U_1(N_T) \le \begin{cases} O\left(N_T^{1+\varepsilon}\right) + \sum_{i=1}^{1+\log\left(\frac{M}{r}\right)} 2^i r \cdot U_1\left(\frac{N_T}{r2^{i-1}}\right), & \text{if } N_T > c, \\ \\ O(1), & \text{if } N_T \le c, \end{cases}$$

where $\varepsilon > 0$ is arbitrary, $c \ge 2$ is an appropriate constant, and the constants of proportionality in the non-recursive terms depend on r (and on ε, α). Note that we process recursively only the T-triangles, since vertices incident to W-triangles are estimated and discarded immediately before processing the new recursive step.

It is easy to verify that the solution of this recurrence is $U_1(N_T) = O(N_T^{1+\varepsilon})$, for any $\varepsilon > 0$ (slightly larger than the ε in the non-recursive term, but still arbitrarily close to 0), with a constant of proportionality that depends on ε and on α.

The number of WT-vertices. To complete the analysis, we next establish a near-linear bound on the number of WT-vertices. The analysis somewhat resembles the derivation of the upper bound on the number of PPT-vertices in Section 2.4. We have, at each step, a subproblem within some simplex Δ_0, involving a set $\mathcal{W} = \mathcal{W}^{\Delta_0}$ of wedges, whose boundaries cross Δ_0, and a set $\mathcal{T} = \mathcal{T}^{\Delta_0}$ of triangles, so that, for each $T \in \mathcal{T}$, all three of its edges cross Δ_0. Put $N_W = |\mathcal{W}^{\Delta_0}|$, $N_T = |\mathcal{T}^{\Delta_0}|$.

We first draw a random sample $R \subset \mathcal{W}^{\Delta_0}$ of $O(r \log r)$ wedges, for some sufficiently large constant parameter r, form and triangulate the arrangement $\mathcal{A}(R)$, and collect only the cells in the complement of the union of these wedges. Since the wedges are all α-fat, it follows by the analysis of [14] that the overall number of these cells is $O(r \log r)$ (where the constant of proportionality depends on α). Arguing as above, we may assume that each subcell Δ of the resulting decomposition is crossed by at most N_W/r wedge boundaries in $\mathcal{W}^{\Delta_0}$, and there is at most one cell that meets all three edges of any fixed triangle $T \in \mathcal{T}^{\Delta_0}$.

In this problem, the decomposition generates only two types of vertices that we need to bound: *new* WW-vertices (those that are incident to the boundary of at least one triangle that had been a T-triangle in the parent cell), which we dispose of immediately, and WT-vertices, which we process recursively. Using [14] and arguing as above, the number of new WW-vertices is $O(N_W + N_T)$. The recurrence bottoms out when $\min\{N_T, N_W\} \leq c$, for some constant $c \geq 2$; we then bound the number of the (remaining) WT-vertices, using brute force, by $O(N_W + N_T)$.

Let $U_2(N_T, N_W)$ denote the maximum number of WT-vertices that appear on the boundary of the union at a recursive step within some simplex Δ_0, involving a set $\mathcal{W}^{\Delta_0}$ of N_W wedges and a set $\mathcal{T}^{\Delta_0}$ of N_T triangles. Arguing as above, one can show that U_2 satisfies the following recurrence:

$$U_2(N_T, N_W) \leq \begin{cases} O(N_W + N_T) + \sum_{i=1}^{1+\log M} 2^i \cdot U_2\left(\frac{2N_T}{2^i}, \frac{N_W}{r}\right), & \text{if } \min\{N_T, N_W\} > c, \\ \\ O(N_W + N_T), & \text{if } \min\{N_T, N_W\} \leq c, \end{cases}$$

where $M = O(r \log r)$ is the overall number of cells in the decomposition (at the current recursive step), $c \geq 2$ is an appropriate constant, and the non-recursive terms depend also on r (and on α). Again, it is easy to verify that the solution of this recurrence is $U_2(N_T, N_W) = O\left((N_T + N_W) \cdot {N_W}^{\varepsilon}\right)$, for any $\varepsilon > 0$, with a constant of proportionality that depends on ε and on α. This completes the analysis, and establishes the following result:

Theorem 3.1 *The complexity of the union of n α-fat triangles in the plane is $O(n^{1+\varepsilon})$, for any $\varepsilon > 0$, where the constant of proportionality depends on ε and α. The bound is almost tight in the worst case.*

Remark: The bound is not as sharp as the one obtained in [22, 25]. We note that the dependence on α is the same as that for the union of α-fat wedges. Since the latter dependence is only proportional to $\frac{1}{\alpha} \log \frac{1}{\alpha}$ [25], the same holds for the union of fat triangles.

4 Conclusions

In this paper we have solved a major open problem in the study of the union of objects in three dimensions, which has resisted a solution for over a decade. Yet, there is still a small remaining gap between our upper bound on the complexity of the union of n α-fat tetrahedra and the corresponding lower bound $\Omega(n^2\alpha(n))$ (which we conjecture to be tight). Closing this gap remains a challenging open problem.

A natural open problem is to extend the new machinery presented in this paper to the problem of bounding the union of other families of geometric objects in 3-space. One such problem concerns the union of cylinders (with arbitrary radii); a nearly-quadratic bound is known only when all the cylinders have equal radii [2]. Another related problem is to obtain a nearly-quadratic bound on the complexity of the union of n arbitrary *α-fat* convex objects of constant description complexity (that is, convex objects c, for which there exist two concentric balls, $B \subseteq c \subseteq B'$, such that the ratio between the radii of B' and B is at most α, for some fixed $\alpha > 1$). Weaker goals would be to prove this only for nearly equal objects of this kind, or obtaining a subcubic bound for the union of such objects of arbitrary sizes.

Acknowledgments. The authors wish to thank Boris Aronov for useful discussions concerning the problem, which have led to an earlier simple proof of Lemma 2.3. The authors also wish to thank Mark de Berg for useful discussions on the problem at the 2005 Carleton-Eindhoven Workshop on Computational Geometry (where work on the current paper has begun).

References

[1] P. K. Agarwal, J. Pach and M. Sharir. State of the union, of geometric objects: A review, to appear in *Proc. Joint Summer Research Conference on Discrete and Computational Geometry—Twenty Years Later*, Contemp. Math, AMS, Providence, RI.

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

[3] B. Aronov, A. Efrat, V. Koltun, and M. Sharir. On the union of κ-round objects in three and four dimensions. *Discrete Comput. Geom.*, 36:511–526, 2006.

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

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

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

[7] B. Chazelle and J. Friedman. A deterministic view of random sampling and its use in geometry. *Combinatorica*, 10:229–249, 1990.

[8] K. L. Clarkson. New applications of random sampling in computational geometry. *Discrete Comput. Geom.*, 2:195–222, 1987.

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

[10] M. de Berg, M. van Kreveld, M. Overmars, and O. Schwarzkopf. *Computational Geometry: Algorithms and Applications*. Springer-Verlag, Berlin, 2000.

[11] D. P. Dobkin and D. G. Kirkpatrick. Determining the separation of preprocessed polyhedra: a unified approach. *Proc. 17th Internat. Colloq. Automat. Lang. Prog.*, pages 400–413, 1990.

[12] A. Efrat. The complexity of the union of (α, β)-covered objects. *SIAM J. Comput.*, 34:775–787, 2005.

[13] A. Efrat and M. Katz. On the union of α-curved objects. *Comput. Geom. Theory Appl.*, 14:241–254, 1999.

[14] A. Efrat G. Rote. and M. Sharir. On the union of fat wedges and separating a collection of segments by a line. *Computat Geom: Theory Appl.*, 3:277–288, 1994.

[15] A. Efrat and M. Sharir. On the complexity of the union of fat objects in the plane. *Discrete Comput. Geom.*, 23:171–189, 2000.

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

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

[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.*, 1:59–71, 1986.

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

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

[21] J. Matoušek, N. Miller, J. Pach, M. Sharir, S. Sifrony, and E. Welzl. Fat triangles determine linearly many holes. In *Proc. 32nd Annu. IEEE Sympos. Found. Comput. Sci.*, pages 49–58, 1991.

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

[23] J. Pach. On the complexity of the union of geometric objects. *Lecture Notes in Computer Science*, 2098:292–307, 2001.

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

[25] J. Pach and G. Tardos. On the boundary complexity of the union of fat triangles. *SIAM J. Comput.*, 31:1745–1760, 2002.

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