The Union of Congruent Cubes in Three Dimensions*

János Pach[†] Ido Safruti[‡] Micha Sharir[§]

Abstract

A *dihedral (trihedral) wedge* is the intersection of two (resp. three) half-spaces in $\mathbf{R}^3$. It is called α*-fat* if the angle (resp., solid angle) determined by these half-spaces is at least $\alpha > 0$. If, in addition, the sum of the three face angles of a trihedral wedge is at least $\gamma > 4\pi/3$, then it is called (γ, α)*-substantially fat*. We prove that, for any fixed $\gamma > 4\pi/3, \alpha > 0$, the combinatorial complexity of the union of n (a) α-fat dihedral wedges, (b) (γ, α)-substantially fat trihedral wedges is at most $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, where the constants of proportionality depend on ε, α (and γ). We obtain as a corollary that the same upper bound holds for the combinatorial complexity of the union of n (nearly) congruent cubes in $\mathbf{R}^3$. These bounds are not far from being optimal.

1 Introduction

The *combinatorial complexity* (or, simply, *complexity*) of a polyhedral set is the total number of its faces of all dimensions. To obtain an upper bound on the complexity of a polyhedral set in $\mathbf{R}^3$, by Euler's Polyhedral Formula, it is sufficient to bound the number of its *vertices*. The problem of bounding the combinatorial complexity of the *union* of various geometric objects has a long history. It is partly motivated by questions in robot motion planning and manufacturing. Specifically, let $A_1, \ldots, A_n$ be n pairwise disjoint convex objects ('obstacles'), and let B be another convex body ('robot'), free to translate amid the obstacles and constrained not to intersect any of them. The space of all *collision-free* translations of B (at which it does not intersect any obstacle) is the complement of $\bigcup_{i=1}^{n} A_i \oplus (-B)$, where $K_i = A_i \oplus (-B) = \{x - y \mid x \in A_i,\ y \in B\}$ is the *Minkowski sum* of the two objects A_i and $-B$. Hence, the problem of computing the space of all free positions of B reduces to that of computing the union of these Minkowski sums. The first task towards the design of an efficient algorithm for this problem is to obtain a sharp bound on the combinatorial complexity of the union.

*Work by János Pach and Micha Sharir has been supported by NSF Grant CCR-97-32101 and by a grant from the U.S.-Israeli Binational Science Foundation. Work by János Pach was also supported by a grant of the Hungarian Science Foundation and by a CUNY Research Award. Work by Micha Sharir was also supported by a grant from the Israeli Academy of Sciences for a Center of Excellence in Geometric Computing at Tel Aviv University, and by the Hermann Minkowski–MINERVA Center for Geometry at Tel Aviv University.

[†]Department of Computer Science, City College, CUNY, New York, NY, USA, Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA, and Hungarian Academy of Sciences, Budapest, Hungary; `pach@cims.nyu.edu`.

[‡]School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel; `idos@math.tau.ac.il`. The research by this author was done while the author was an M.Sc. student under the supervision of Micha Sharir.

[§]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; `sharir@math.tau.ac.il`.

In the plane, the complexity of the union of Minkowski sums was shown to be linear by Kedem et al. [11]. In $\mathbf{R}^3$, we know the following. (i) If the A_i's and B are convex polyhedra, then the complexity of the union is $O(Nn \log n)$, where N denotes the overall complexity of the Minkowski sums K_i [3]; see also [4]. (ii) If the A_i's are convex polyhedra consisting of a total of n faces, and B is a ball, then the complexity of the union is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$ [1]. In other words, all known results concerning unions of *Minkowski sums* yield linear upper bounds in the plane and near-quadratic upper bounds in 3-space.

The above results cannot be extended to the union of general convex objects without imposing any further restriction on their shapes or relative position. Indeed, it is easy to see that the union of n triangles in the plane (tetrahedra in 3-space) can have quadratic (resp., cubic) complexity. Since all constructions realizing these bounds use very 'thin' objects, it is a natural question to ask what happens if we restrict our attention to unions of 'fat' convex polytopes. For bounded objects, fatness means that the ratio between the circumradius and the inradius of any input object is bounded by a fixed constant. For unbounded objects, another definition is needed—see Definition 1.1 below.

The case of *planar* fat objects has been studied extensively in [2, 7, 8, 9, 10, 13, 14]. It was shown that the complexity of the union of n fat triangles is $O(n \log \log n)$ [13, 14] and that of n fat wedges is $O(n)$ [2, 9]. For general convex fat objects of 'constant description complexity', the combinatorial complexity of the union is $O(n^{1+\varepsilon})$, for any $\varepsilon > 0$ [10] (see also [7, 8] for slight improvements and extensions).

In contrast, in three and higher dimensions, very few non-trivial bounds are known. It is an easy consequence of the Upper Bound Theorem for convex polytopes that the combinatorial complexity of the union of n balls in $\mathbf{R}^d$ is $O(n^{\lceil d/2 \rceil})$. Asymptotically the same upper bound is known for the complexity of the union of n *axis-parallel* hypercubes [5], which can be improved to $O(n^{\lfloor d/2 \rfloor})$ when all cubes have the same size.

In spite of many efforts, even in three dimensions no non-trivial (i.e., subcubic) upper bound was known for the complexity of the union of n congruent cubes, not necessarily in parallel position. The aim of this paper is to establish a nearly quadratic upper bound on this quantity. Actually, we will prove a more general result. For this we need some preparation.

Definition 1.1 *The intersection of two (three) half-spaces is called a* dihedral *(resp.* trihedral*)* wedge. *The boundary of a dihedral wedge consists of a straight line* edge *and two half-planes. The boundary of a trihedral wedge consists of a vertex (apex), three* edges *and three* faces *that are half-lines and 2-dimensional wedges, respectively.*

For any $\alpha > 0$, a dihedral (trihedral) wedge is called α-fat *if its dihedral angle (resp., solid angle) is at least α. For any $\gamma > 4\pi/3$, an α-fat trihedral wedge is said to be* (γ, α)-substantially fat, *if the sum of the angles of its three faces is at least $\gamma > 4\pi/3$.*

Note that a right-angle octant, obtained by taking the intersection of three half-spaces bounded by mutually orthogonal planes, is $(3\pi/2, \pi/2)$-substantially fat. However, a trihedral wedge defined by three planes supporting different faces of a regular tetrahedron is *not* (γ, α)-fat for any $\gamma > 4\pi/3$, because the angles of its faces are too small. The requirement that $\gamma > 4\pi/3$ is technical, made in order to facilitate our proof.

All families studied in this paper consist of n convex polyhedral objects in $\mathbf{R}^3$, each having a constant number of vertices, edges, and faces. As we pointed out earlier, to give an upper bound for the combinatorial complexity of the union of such families, it is sufficient to bound the number of *vertices* of the union. Such a vertex is either a vertex of an input polyhedron, or it can be obtained as the intersection of an edge of a polyhedron with a face of another, or it is the intersection point of three faces belonging to three distinct polyhedra. Clearly, the number of vertices of the first two types is $O(n^2)$, so the main task is to estimate the number of vertices of the third type.

We prove the following three results. In all of them, the constants of proportionality hidden in the O-notation depend on the relevant fixed parameters (ε, α, γ, λ).

Theorem 1.2 *For any $\alpha, \varepsilon > 0$, the combinatorial complexity of the union of n α-fat dihedral wedges in 3-space is $O(n^{2+\varepsilon})$.*

Theorem 1.3 *For any $\gamma > 4\pi/3$, $\alpha, \varepsilon > 0$, the combinatorial complexity of the union of n (γ, α)-substantially fat trihedral wedges in 3-space is $O(n^{2+\varepsilon})$.*

In Section 5, we apply Theorem 1.3 to deduce

Theorem 1.4 *Let $\lambda > 1, \varepsilon > 0$. The combinatorial complexity of the union of any family of n cubes in 3-space, whose edge lengths differ only by a factor of at most λ, is $O(n^{2+\varepsilon})$.*

All of these results are nearly tight in the worst case. That is, an $\Omega(n^2)$ lower bound can easily be established in each of these cases.

An important new tool in our analysis is the concept of *special cubes.*

Definition 1.5 *Given a family $\mathcal{P}$ of convex polyhedra in 3-space, a* special cube *C (with respect to $\mathcal{P}$) is the intersection of three members of $\mathcal{P}$ such that (i) C is disjoint from every other member of the family, and (ii) C has the combinatorial structure of a cube, with each of the three polyhedra contributing two* opposite *faces to C.*

Cubes that satisfy only condition (ii) are referred to as quasi-special cubes*. The* level *of a quasi-special cube C is the number of members of $\mathcal{P}$ that intersect C, other than the three members whose intersection equals to C.*

This notion is related to the concept of *special quadrilaterals* used in [3, 4]. The significance of special cubes, which extends beyond the applications given in this paper, is shown by the following theorem, whose somewhat technical proof is postponed to Section 7.

Theorem 1.6 *Let $\mathcal{P}$ be a family of n convex polyhedra in 3-space, each having at most some constant number of faces. Suppose that the number of special cubes determined by any m members of $\mathcal{P}$ is $O(m^\gamma)$, for some $\gamma > 2$. Then the number of vertices on the boundary of the union of $\mathcal{P}$ is $O(n^\gamma)$.*

Here is a brief overview of the approach we follow. Consider a family of fat dihedral wedges. First, we 'deform' the wedges to new 'canonical' wedges, without losing more than quadratically many special cubes in the process. We reduce the problem to the case when there exists a plane P intersecting every (3-dimensional) wedge in a fat 2-dimensional wedge, whose bounding rays belong to a fixed set of constantly many 'canonical' directions. In this way, we obtain a constant number of families, each consisting of wedges with *isothetic* cross sections (i.e., whose cross sections are translates of each other, lying in planes parallel to P), and it suffices to bound the complexity of the union of at most *three* such families. This is done in Section 3, by first handling the (trivial) case of a single family, then passing to the case of two families, and finally tackling the general case.

For trihedral wedges, the analysis is more elaborate, since our current machinery works only when, for any vertex v of the union, there exists a (canonical) plane P, so that all three wedges incident to v intersect P in *unbounded* regions. The reason for this is quite technical, and it originates in the method developed in [13, 14] for studying the case of fat triangles in the plane. This is why we can handle only trihedral wedges that are substantially fat (with the sum of their face angles being greater than $4\pi/3$). Even with this assumption, the canonization process is more involved than for dihedral wedges. We eventually manage to transform each trihedral wedge to a new canonical wedge, all of whose cross sections by planes of some canonical direction are either empty or isothetic to some canonical 2-dimensional wedge. This allows us to apply the arguments used for dihedral wedges, with only minor modifications.

The case of nearly equal cubes is an easy consequence of the result for substantially fat trihedral wedges, specialized to right-angle octants. More specifically, we lay a grid whose size is slightly smaller than that of the cubes, consider the union within each cell of the grid separately, replace each cube whose boundary crosses such a cell by an octant, and apply the bound on the complexity of the union of such octants.

We also consider the algorithmic problem of efficient construction of the union of a family of, say, n nearly equal cubes. Using the algorithm of Aronov et al. [4], together with our new combinatorial bounds, we obtain a randomized algorithm that computes the union in expected time $O(n^{2+\varepsilon})$.

Three interesting problems remain unsolved: In the first two, we wish to obtain near-quadratic upper bounds for the combinatorial complexity of the union of (1) *any* collection of n cubes (of wildly different sizes), and (2) any collection of n α*-fat* (rather than substantially fat) trihedral wedges in 3-space. (3) Is there a superquadratic lower bound for any of the functions discussed above? As noted above, quadratic lower bounds are known for each of them. We expect that positive answers to (1) and (2) will lead to a near-quadratic bound on the complexity of the union of any family of fat convex polytopes in 3-space.

2 Canonization of Dihedral Wedges

Let $\mathcal{W}$ be a family of n α-fat dihedral wedges in 3-space. Let $U_{\mathcal{W}}$ denote the union of $\mathcal{W}$, and let $\mathcal{A}(\mathcal{W})$ denote the arrangement of the (faces bounding the) wedges in $\mathcal{W}$.

For a constant parameter σ, let $\mathcal{D}(\sigma)$ be a set of $O(1)$ directions (points on the unit sphere) such that any spherical cap of radius larger than σ contains a direction $d \in \mathcal{D}(\sigma)$.

The following lemma holds for all triples of (not necessarily fat) wedges.

Lemma 2.1 *There exists an absolute constant $\sigma > 0$ such that for any three dihedral wedges w_1, w_2, w_3, there exists a direction $d \in \mathcal{D}(\sigma)$ such that $|\langle d, \mathbf{e}_{w_i} \rangle| \geq \frac{1}{4}$, for $i = 1, 2, 3$, where $\mathbf{e}_{w_i}$ denotes the unit vector in the direction of the edge e_{w_i} of w_i.*

Proof: Consider the set of directions d such that $|\langle d, u \rangle| \leq \frac{1}{4}$ for a fixed direction u. This set is a band of width $\pi - 2\arccos(1/4)$ centered at the great circle orthogonal to u. The area of such a band is $4\pi \cdot (1/4) = \pi$. Hence, the area of the set of directions d where the asserted condition on d is not satisfied is at most 3π. Thus, the complement set of 'good' directions is of area at least π. Since the union of three such bands (each around a great circle) has at most eight holes, there exists at least one hole of area larger than $\frac{\pi}{8}$. Since any such hole is bounded by at most a constant number of circular arcs, the claim readily follows. □

Lemma 2.2 *Let P_d be a plane orthogonal to a direction d satisfying the conditions in the previous lemma. If w_1, w_2, w_3 are all α-fat then the three planar wedges $w_i \cap P_d$, for $i = 1, 2, 3$, are all $(\alpha/4)$-fat.*

Proof: Let w be one of these wedges. It is a routine exercise in stereometry to show that the angle of the cross-sectional wedge $w \cap P_d$ is minimized when the bisector plane of w is orthogonal to P_d. Assume that w does indeed attain this minimum. Let α' denote the angle of $w \cap P_d$. Let e' denote the orthogonal projection of e_w onto P_d. Let γ denote the angle between e_w and e'; note that $\sin\gamma = \langle \mathbf{e}_w, d\rangle$. Denote by A the point $e_w \cap P_d$; let B be a point on e' at distance 1 from A, let C be the foot of the perpendicular from B onto e_w, and let D, E be the points of intersection between ∂w and the line within P_d passing through B and orthogonal to e'; see Figure 1.

We have $BC = \sin\gamma$, $BD = BC\tan(\alpha/2) = \sin\gamma \cdot \tan(\alpha/2)$, and thus $\tan(\alpha'/2) = \sin\gamma \cdot \tan(\alpha/2)$, or $\alpha' = 2\arctan(\sin\gamma \cdot \tan(\alpha/2))$. We claim that $\arctan(\lambda x) \geq \lambda \arctan x$ for any $x \geq 0$ and $\lambda \leq 1$. Indeed, the function $f(x) = \arctan(\lambda x) - \lambda \arctan x$ vanishes at 0, and its derivative is

$$f'(x) = \frac{\lambda}{1 + \lambda^2 x^2} - \frac{\lambda}{1 + x^2} \geq 0,$$

from which the preceding inequality follows. We thus obtain

$$\alpha' \geq 2\sin\gamma \cdot \alpha/2 = \alpha \sin\gamma = \alpha\langle \mathbf{e}_w, d\rangle \geq \alpha/4.$$

This completes the proof of the lemma. □

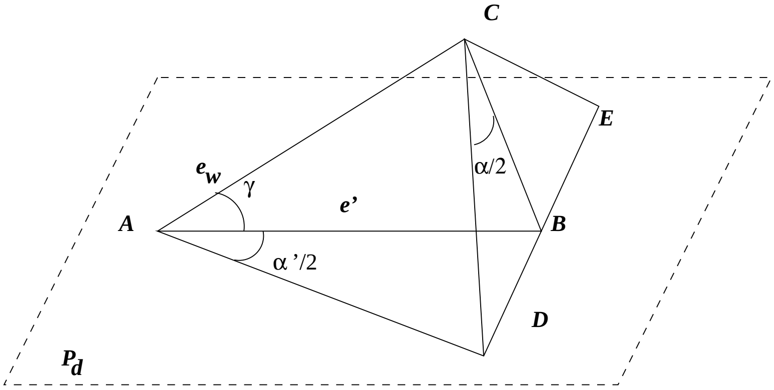

Figure 1: The setup in the proof of Lemma 2.2.

A direction satisfying the properties of Lemma 2.1 (and of Lemma 2.2) is called a *good direction* for the triple $w_1, w_2, w_3 \in \mathcal{W}$.

For each $d \in \mathcal{D}$, let $\mathcal{W}_d \subseteq \mathcal{W}$ denote the subfamily consisting of all members of $\mathcal{W}$ that cross the planes orthogonal to d in $(\alpha/4)$-fat 2-dimensional wedges, i.e., in angles of size at least $\alpha/4$. For simplicity, we refer to these planes as *horizontal*. Construct on a horizontal unit circle $O(1/\alpha)$ pairwise disjoint 'canonical' arcs, each of length $\Omega(\alpha)$ (say, $\alpha/16$), so that (i) each horizontal line through the origin meets at most one of these arcs, and (ii) each arc of length at least $\alpha/4$ on the unit circle fully contains at least one of these arcs.

For each wedge $w \in \mathcal{W}_d$, rotate its faces inwards about its edge (which remains fixed) until the directions of their horizontal cross sections coincide with the endpoints of one of these arcs. Let $\mathcal{W}'_d$ denote the resulting collection of wedges. By this process, we decompose $\mathcal{W}'_d$ into $O(1/\alpha)$ *canonical* subfamilies so that the horizontal cross sections of any two wedges belonging to the same subfamily are isothetic.

The proof of the following lemma, which justifies the canonization process, hinges on the fact that the edges of the given wedges *are not moved* during the deformation.

Lemma 2.3 *The number of special cubes for $\mathcal{W}_d$ is smaller than or equal to the number of special cubes for $\mathcal{W}'_d$.*

Proof: Let C be a special cube for $\mathcal{W}_d$, formed by the intersection of three wedges $w_1, w_2, w_3 \in \mathcal{W}_d$. Let w'_i denote the canonical image of w_i, for $i = 1, 2, 3$. Put $C' = w'_i \cap w'_2 \cap w'_3$. We claim that after the canonization, C' remains a (nonempty) special cube (for $\mathcal{W}'_d$). This is shown as follows.

(i) The canonization process ensures that $w' \subset w$ for any $w \in \mathcal{W}$, from which it follows that $C' \subset w_1 \cap w_2 \cap w_3 = C$. Moreover, since C, being a special cube, is disjoint from all other wedges, it follows that the same holds for C'.

(ii) Let h_1, h_2, h_3 be three half-planes such that h_i is bounded by e_{w_i} and lies fully in w_i, for $i = 1, 2, 3$. The intersection point $v = h_1 \cap h_2 \cap h_3$ lies in C, by construction. Since the new faces of each wedge are both half-planes of this kind, this implies that during the canonization process, each special cube shrinks but does not disappear ($C' \neq \emptyset$).

(iii) The boundary $\partial C'$ does not meet any of the edges of the w_i's, because C does not meet them and they do not move during the canonization process. Thus, C' must have the combinatorial structure of a cube and the two faces bounding each wedge w_i contribute opposite faces to C'. $\square$

Theorem 2.4 *The number of special cubes for $\mathcal{W}'_d$ is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$.*

Combined with Lemma 2.3, this theorem implies that the number of special cubes for $\mathcal{W}_d$ is also $O(n^{2+\varepsilon})$. This, combined with Theorem 1.6, implies that the complexity of the union of $\mathcal{W}_d$ is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$. Finally, since each vertex $v \in U_{\mathcal{W}}$ is also a vertex of $U_{\mathcal{W}_d}$, for some canonical d, it follows that the complexity of the union of $\mathcal{W}$ is also $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, thus establishing Theorem 1.2.

3 The Complexity of the Union of Fat Dihedral Wedges

The aim of this section is to prove Theorem 2.4. As we have argued above, it is sufficient to bound the number of special cubes determined by at most *three canonical* subfamilies of $\mathcal{W}'_d$.

The union of one canonical family. In this case, it is easy to see that a single canonical family admits *no* special cubes. Indeed, let C be the intersection of three wedges from the same subfamily of $\mathcal{W}'_d$. Then any intersection of C with a plane orthogonal to d is unbounded, which is impossible for a special cube.

The union of two canonical families. Let $\mathcal{R}$ and $\mathcal{B}$ be two canonical subfamilies of $\mathcal{W}'_d$. We refer to their wedges as red and blue, respectively.

Theorem 3.1 *The number of special cubes in the union of two canonical subfamilies of n α-fat dihedral wedges in $\mathcal{W}'_d$ is $O(n^2\alpha(n))$, where $\alpha(\cdot)$ is the inverse Ackermann function.*

Proof: We refer to the actual wedges as *3-wedges* (not to be confused with trihedral wedges), and to their horizontal cross sections as *2-wedges* (as above, d is assumed to be the direction of the z-axis). We also assume that the red 2-wedges point to the right (in the horizontal planes containing them, so that the apex of a red 2-wedge is its rightmost point and its symmetry axis is parallel to the x-axis) and that the blue 2-wedges point upwards (the apex of a blue 2-wedge has maximum y-coordinate and its symmetry axis is parallel to the y-axis). We will thus refer to the two edges of a red 2-wedge (and to the corresponding faces of the 3-wedge) as the 'top' and 'bottom' edges (and faces) and similarly use 'left' and 'right' for the blue wedges; cf. Figure 2 for an illustration.

It suffices to estimate the number of *red-red-blue* special cubes of the union (i.e., cubes formed by two red wedges and one blue wedge).

We will regard the z-axis as the 'time-axis' and regard the 2-wedges as translating in the xy-plane at constant (though possibly different) velocities. The vertices (of the third kind; see the remark before Theorem 1.2) of the union then become critical events, at which three edges bounding the moving 2-wedges become concurrent.

Let C be a special cube formed by the intersection of two red wedges r, r', and of one blue wedge b. Thus, C has six faces: each of r, r', b contributes one pair of opposite faces to its boundary. Four faces of C are red, and are arranged in a cycle. As is easily verified, up to the possible permutation of r and r', the cycle has the form (top face of r, top face of r', bottom face of r, bottom face of r'). In other words, the cube has an edge where the two top faces of r and r' meet. The two endpoints of this edge are vertices of C that are also vertices of the union. We refer to them as *special* vertices.

Let v_1, v_2 be two special red-red-blue vertices, so that v_1 is top-top-right, v_2 is top-top-left, and both are vertices of the same special cube (and are the endpoints of a common edge of that cube). Then v_1, v_2 lie on the top boundaries e, e' of two red 2-wedges r, r' (which are overlapping) and on the right and left boundaries f_1, f_2 of a blue 2-wedge b. We classify the blue 2-wedges as being 'short' or 'long', where b is *short* if the portion of f_1 between its apex u and v_1 does not meet any red 2-wedge, and is *long* otherwise.

We first bound the number of special red-red-blue cubes, for which the corresponding blue 2-wedge b is short. We use a 2-dimensional coordinate frame F_{f_1} to represent points on f_1 by (t, ξ), where t is the time and ξ is the distance along f_1 from its apex u. Each red 2-wedge r shows up in F_{f_1} as a (portion of a) wedge — its top and bottom edges trace straight-line segments in F_{f_1}, where the trace of the top (resp., bottom) edge of r is the bottom (resp., top) segment of the traced wedge in F_{f_1}. In general, vertices of the union of the 3-wedges along f_1 appear as vertices of the boundary of the union of these representing wedges in F_{f_1}. However, when b is short, the vertex v_1 under consideration is a vertex of the *lower envelope* of the traced wedges in F_{f_1}. Thus, the number of such vertices is $O(n\alpha(n))$ (see, e.g., [15]). Summing over all blue 2-wedges, we conclude that the number of (not necessarily special) top-top-right vertices for which the blue 2-wedge is short is $O(n^2\alpha(n))$.

Suppose next that v_1, v_2 are two special vertices, as above, for which the corresponding blue 2-wedge b is long. Suppose, with no loss of generality, that at the time when these vertices appear,

r is contained in r'. Put $e_1 = f_1 \cap r$ and $e_2 = f_2 \cap r$. Since v_1 and v_2 lie on a special cube, e_1 and e_2 are segments that lie on faces of that cube, and hence, by definition, they do not meet any other red or blue 2-wedge. Thus, the respective lower endpoints w_1, w_2 of e_1, e_2 lie only in the interior of r' (and on the boundaries of r and b) but are outside all other red and blue wedges. See Figure 2.

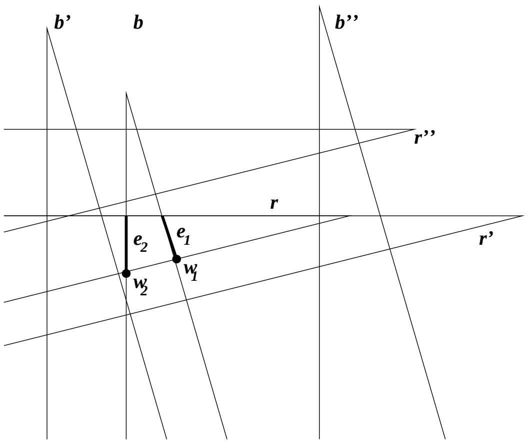

Figure 2: The cross section of a two-colored special cube, at the time when the top edges of r, r' overlap. The bottom vertices w_1, w_2 lie inside r' only and on the boundaries of r and b.

We now apply the analysis of Matoušek et al. [13], developed for studying the complexity of the union of fat planar triangles (see also [14]). Let p denote the apex of r and let q be the leftmost point on the top edge e of r that does not lie in the interior of any other red 2-wedge (since the red 2-wedges are homothetic, q is uniquely defined). Let ρ denote the ray emanating from q to the right in the direction of the bottom edges of the red 2-wedges. It is easily verified that the apex of any blue 2-wedge b which is long with respect to r must lie above ρ; see Figure 3. Consider the collection of long blue 2-wedges that form special vertices along e. Then, by the preceding observation, the segments of intersection of these blue 2-wedges with the bottom edge of r are pairwise disjoint. It now follows from the analysis of [13, Lemma 3.5] and from its improvement in [14, Lemma 2.5] that the number of such 2-wedges is $O(\frac{1}{\alpha}\log\frac{1}{\alpha})$. Since the number of overlaps between top edges of red 2-wedges is $O(n^2)$, we conclude that the number of red-red-blue special cubes for which the corresponding blue 2-wedge is long is $O(\frac{n^2}{\alpha}\log\frac{1}{\alpha})$. This completes the proof of Theorem 3.1. □

Note that Theorem 3.1, in combination with Theorem 1.6, implies that the combinatorial complexity of the union of two canonical families of n α-fat dihedral wedges in $\mathcal{W}'_d$ is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$. This fact is used in the next stage of the analysis.

The union of three canonical families. Let $\mathcal{R}, \mathcal{G}, \mathcal{B} \subset \mathcal{W}'_d$ be three canonical subfamilies of α-fat dihedral wedges, and refer to their members as *red, green* and *blue,* respectively. Our goal is to bound the number of special cubes of $\mathcal{R} \cup \mathcal{B} \cup \mathcal{G}$, formed by the intersection of a red wedge, a blue wedge and a green wedge.

Definition 3.2 *(a) For any $r \in \mathcal{R}, b \in \mathcal{B}$, we say that r is* semi-free within b *(with respect to $\mathcal{R} \cup \mathcal{B}$), if each of the two sides of r contain a point (a 'semi-free point') that lies inside b and*

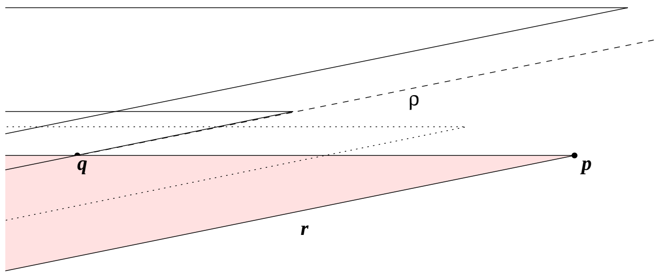

Figure 3: No red 2-wedge above r can have a point below ρ, the dashed ray amanating from q to the right. Thus, the apex of any long blue wedge with respect to r lies above ρ.

outside all other red and blue 2-wedges. (b) We say that r is initial *for b (with respect to $\mathcal{R} \cup \mathcal{B}$), if there is a side of b that no other red 2-wedge intersects between the apex of b and r. (See Figure 4.)*

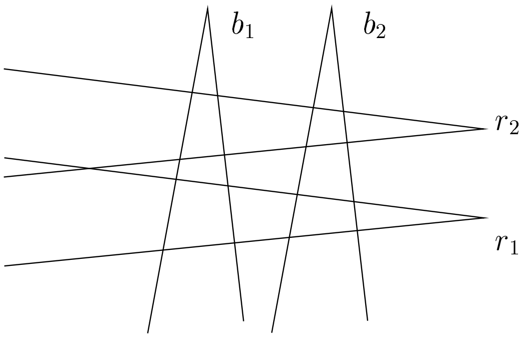

Figure 4: r_2 is initial for b_2, whereas r_1 is not. Each red wedge is semi-free within every blue wedge, and vice versa.

Arguing as in the 2-family case and as in [13], we obtain:

Lemma 3.3 *There exists a constant $c = O(\frac{1}{\alpha}\log\frac{1}{\alpha})$ with the property that for every red 2-wedge r and any fixed time t, the number of blue 2-wedges b for which r is semi-free within b and r is not initial for b at time t, is at most c. Similar properties hold for all other kinds of (ordered) bichromatic pairs of wedges.*

Proof: For a red 2-wedge r, let B denote the collection of all blue 2-wedges b such that r is semi-free within b and r is not initial for b. Order the elements of B in the order of their containment of semi-free points along the upper or lower ray of r (this order is clearly well defined). Let B' be the subsequence of B consisting of every other element. For each 2-wedge b in B', r is not initial for b and $b \cap \partial r$ consists of two segments that do not meet any other 2-wedge of B'. Indeed, if one of these segments intersects another blue 2-wedge b' in B' then, as is easily verified, no intermediate element b'' of B can have any free point on the same ray of r, contrary to assumption (see Figure 5).

Let e_r^+ denote the edge of r that is closer to the apex of b, and let e_r^- be the other edge. Let r' be the first red 2-wedge encountered when traversing e_r^+ from the apex of r. Consider the three

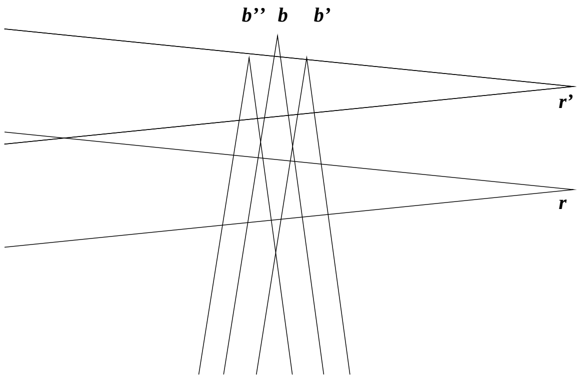

Figure 5: Since r is semi-free within b, the two blue neighbors b', b'' cannot overlap along ∂r.

lines that contain $e_{r'}^-$, e_r^+, and e_r^-, respectively. Each $b \in B'$ crosses the second line between the apex of r and the intersection vertex $e_r^+ \cap e_{r'}^-$, and no two elements of B' intersect each other within the slab formed between the first and third lines. Hence, arguing as in the case of two canonical families, we have $|B'| = O(\frac{1}{\alpha} \log \frac{1}{\alpha})$. Since $|B| \le 2|B'| + 1$, we have that $|B| = O(\frac{1}{\alpha} \log \frac{1}{\alpha})$ as well. $\square$

Definition 3.4 *An ordered pair (r, b) of 2-wedges is called* exposed *at time t, if at least one of the following two conditions is satisfied (with respect to the union of the two respective canonical subfamilies): (i) b is initial for r, or (ii) r is semi-free within b and r is not initial for b.*

Clearly, Lemma 3.3 implies that for any 2-wedge $r \in \mathcal{R}$, the number of exposed pairs (r, b), for $b \in \mathcal{B}$, at any fixed t, is at most $c + 1$.

As t varies, a pair (r, b) may start or stop being exposed. It follows from the analysis of the 2-family case that the number of events at which, say, a blue 2-wedge, b, starts or stops being initial for some red 2-wedge, r, is $O(n^2\alpha(n))$. Indeed, this event corresponds to a vertex of the lower envelope of blue wedges in the appropriate 2-dimensional frame attached to one of the red edges.

Consider an event at which a red 2-wedge r starts or stops being semi-free within a blue 2-wedge b. It is easily checked that at such an event three sides of red and blue 2-wedges become concurrent at a vertex that is contained only in (the interior of) b. Clarkson-Shor's technique [6], combined with the statement at the end of the analysis of the 2-family subcase, implies that the number of such events is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$. Let T denote the sorted list of all critical times at which the overall set of exposed pairs, with respect to all possible pairs of families, changes. As just argued, we have $|T| = O(n^{2+\varepsilon})$, for any $\varepsilon > 0$.

We now return to the analysis of 'tricolored' special cubes. Let $r \in \mathcal{R}$, $b \in \mathcal{B}$, and $g \in \mathcal{G}$ be three 3-wedges that form such a cube $C = r \cap b \cap g$. Denote the cross section of C at time t by $C(t)$. Excluding times at which vertices of C occur, $C(t)$ is a convex polygon with at most 6 sides, so that each side is a portion of an edge of one of these 2-wedges, and no two successive sides of $C(t)$ belong to the boundary of the same 2-wedge. It is easily verified that if we sweep a plane through any convex polytope with the combinatorial structure of a cube, so that the plane is not parallel to any of its facets, then there always exists a cross section which is either a pentagon or a hexagon. (Such a cross section arises when the plane has three vertices of the polytope on one side and five on the other.)

Let $C(t_0)$ be a cross section of our cube which is either a pentagon or a hexagon. Then at least two of the 2-wedges, say r and b, contribute two sides to $\partial C(t_0)$.

Lemma 3.5 *For each unordered pair w, w' in $\{r, b, g\}$, either (w, w') or (w', w) is exposed at time t_0.*

Proof: If r contributes two sides to $\partial C(t_0)$ then r is semi-free both within b and within g (with respect to corresponding bichromatic collection of 2-wedges). If r is not initial for b then, by definition, (r, b) is exposed, and if r is initial for b then (b, r) is exposed. The claim thus holds for $\{r, b\}$ and, arguing similarly, for $\{r, g\}$. The case of $\{b, g\}$ follows from the fact that either b or g also contributes two sides to $\partial C(t_0)$. $\Box$

We now apply a technique similar to that used in [12]. The list T of critical times partitions the time-axis into $O(n^{2+\varepsilon})$ atomic intervals. Let $C = r \cap b \cap g$ be a tricolored special cube, as above, and let I be some atomic interval containing a time t_0 where the property of Lemma 3.5 holds for C. That is, for each of the unordered pairs $\{r, b\}$, $\{r, g\}$, $\{b, g\}$, (at least) one of its ordered pairs is exposed over I. Hence, up to a permutation of the 2-wedges r, b, g, either

(a) (r, b) and (r, g) are exposed, or

(b) (r, b), (b, g) and (g, r) are exposed.

We first estimate the number of special cubes of type (a). Fix a red 2-wedge r_0, and let T_{r_0} denote the sublist of critical times at which some bichromatic pair (r_0, w) starts or stops being exposed.

The following procedure computes a superset of all special cubes of the form $r_0 \cap b \cap g$ that satisfy the condition in (a). Iterate over the list T_{r_0}. For each time t in that list, at which a pair (r_0, b_0), for some $b_0 \in \mathcal{B}$, becomes exposed, output all triples (r_0, b_0, g), for $g \in \mathcal{G}$, for which (r_0, g) is currently exposed (there are at most $c + 1$ such triples). Apply a symmetric step when a pair (r_0, g_0), for $g_0 \in \mathcal{G}$, becomes exposed. It is easy to see that every special cube $r_0 \cap b \cap g$ that satisfies the condition in (a) will be output by this procedure, and that the total output size is at most $(c+1)|T_{r_0}|$. This is easily seen to imply that the number of tricolored special cubes that satisfy (a) is $O(n^{2+\varepsilon})$.

Consider next special cubes of type (b). The following procedure computes a superset of those cubes. Iterate over the list T. For each time t in that list, at which a pair (r, b), for $r \in \mathcal{R}$, $b \in \mathcal{B}$, becomes exposed, output all triples (r, b, g), for $g \in \mathcal{G}$, for which (b, g) is currently exposed (there are at most $c + 1$ such triples). Apply an appropriately symmetric step when any other type of bichromatic ordered pair becomes exposed at t. It is easy to see that every special cube $r \cap b \cap g$ that satisfies the condition in (b) will be output by this procedure, and that the total output size is at most $(c + 1)|T| = O(n^{2+\varepsilon})$.

This shows that the overall number of tricolored special cubes is $O(n^{2+\varepsilon})$, from which Theorems 2.4 and 1.2 follow.

4 The Union of Substantially Fat Trihedral Wedges

We next extend the analysis given in the preceding section to the case of (γ, α)-*substantially fat trihedral wedges.* Substantial fatness is required to ensure the following property:

Lemma 4.1 *There exists a canonical set $\mathcal{D}$ of $O(1)$ directions on the unit sphere with the following property. Let w_1, w_2, w_3 be three (γ, α)-substantially fat trihedral wedges, for $\gamma > 4\pi/3$. Then there exists $d \in \mathcal{D}$ such that, for any plane π orthogonal to d and for each $i = 1, 2, 3$, the cross section $\pi \cap w_i$ is unbounded (any such section is either a 2-wedge or a truncated 2-wedge), and the angle between its bounding rays is at least α_0, for some constant α_0 that depends on α and γ.*

Proof: Let w be a trihedral wedge whose edges emanate from its apex in directions a, b, c. A direction d has the property that any plane orthogonal to d crosses w in an unbounded region if and only if

$$\text{the signs of the scalar products } \langle a, d\rangle,\ \langle b, d\rangle,\ \langle c, d\rangle, \text{ are not all equal;} \tag{1}$$

see Figure 7(a).

Moreover, arguing as in the proof of Lemma 2.1, the angle between the two rays, bounding any intersection of w with a plane orthogonal to d, will be fat if the following holds:

$$\min\{|\langle a, d\rangle|, |\langle b, d\rangle|, |\langle c, d\rangle|\} \geq \delta, \tag{2}$$

for some fixed $\delta > 0$. A *direction* d that satisfies (1) and (2) will be called *good* for w.

We next estimate the probability that a randomly selected direction is good for the 3-wedge w. We first calculate the probability that d satisfies (1).

For directions x, y, represented as points on the unit sphere, let $P_{x:y}$ denote the probability that a plane through the origin 0 separates x from y. Let $P_{x:yz}$ be the probability that a plane through 0 separates x from y and z.

Claim 4.2 *(i) $P_{x:y} = \theta_{x,y}/\pi$, where $\theta_{x,y}$ is the angle between the vectors x and y.*

(ii) $P_{x:y} = P_{x:yz} + P_{y:xz}$

Indeed, to see (i), consider the plane P spanned by $x, y, 0$. For any plane π passing through the origin, π separates x from y if and only if the intersection line $\ell = \pi \cap P$ separates x from y in P, and the probability for this to happen is $\theta_{x,y}/\pi$, as asserted.

To see (ii), we note that the event that a plane π through the origin separates x from y is the disjoint union of the events that π separates x from y and z and that π separates y from x and z.

The Claim implies that the probability that d satisfies (1) is

$$P_{a:bc} + P_{b:ac} + P_{c:ab} = \frac{P_{a:b} + P_{a:c} + P_{b:c}}{2} = \frac{\theta_{a,b} + \theta_{a,c} + \theta_{b,c}}{2\pi}.$$

Thus, for a substantially fat trihedral wedge, we have that this probability is at least $\gamma/(2\pi)$, thus a direction violates (1) with probability at most $1 - \gamma/(2\pi)$.

As mentioned in the proof of Lemma 2.1, for a given direction x, the measure of the set of directions d for which $|\langle d, x\rangle| \leq \delta$ is $4\pi\delta$. Since the total area of the sphere is 4π, a random direction violates the inequality $|\langle d, x\rangle| \leq \delta$ with probability δ. Repeating this argument for each of the three edges of w, the probability for violating the inequality (2) is thus at most 3δ.

It follows that a direction is bad for a substantially fat wedge w with probability at most $1 - \gamma/(2\pi) + 3\delta$.

Let w_1, w_2, w_3 be three substantially fat wedges. The preceding argument implies that a direction will be good for all three wedges w_1, w_2, w_3 with probability larger than $1 - 3(1 - \gamma/2\pi + 3\delta) = 3\gamma/2\pi - 9\delta - 2$. Hence, assuming that $\gamma > 4\pi/3$ and that $\delta < \gamma/(12\pi) - 1/9$, the above probability is at least $3\gamma/(4\pi) - 1 > 0$.

This implies, as above, that there exists a set $\mathcal{D}$ of size $O(1)$ (which depends on γ and increases as γ approaches $4\pi/3$) such that, for any three substantially fat wedges, $\mathcal{D}$ contains a direction that is good for all of them. The cross sections of the three wedges by any plane orthogonal to d satisfy the properties asserted in the lemma, with $\alpha_0 > \delta\alpha$, as shown in Section 2. $\square$

4.1 Canonization of trihedral wedges

A significant step in the proof of Theorem 1.3 is the canonization process, which is considerably more intricate than in the case of dihedral wedges. Here is a brief overview of the canonical process. We fix a direction d and focus on the subset $\mathcal{W}_d$ of wedges that satisfy the properties in Lemma 4.1 with respect to d. For each $w \in \mathcal{W}_d$, we fold inwards each of the three faces of w, and then replace the shrunk wedge by the union of $O(1)$ new wedges, so that (i) we lose at most quadratically many special cubes, and (ii) the intersection of a plane orthogonal to d with any new wedge is either a canonical 2-wedge (as in the case of dihedral wedges) or empty. This will allow us to apply a variant of the arguments used for the case of dihedral wedges, from which the asserted near-quadratic bound will follow.

In more details, we proceed as follows. Fix a direction $d \in \mathcal{D}$, and consider the family $\mathcal{W}_d$ of wedges for which d is a good direction. The canonization process of $\mathcal{W}_d$ has two stages.

First canonization stage: Folding backward faces. Let w be a trihedral wedge in $\mathcal{W}_d$. Define the *forward face* of w to be the face spanned by those two edges of w, call them e_0 and e_1, for which the scalar product with d is of the same sign (say positive), and call the two other faces of w the *backward faces*. In what follows we refer to a ray ρ emerging from the apex o of w as *positive* (resp., *negative*) if the scalar product of a vector along ρ with d is positive (resp., negative). Thus, the forward face of w is spanned by the two positive edges of w, and each of the two backward faces is spanned by one positive edge and by the unique negative edge, e_2, of w.

We fold inwards each of the backward faces of w, so that any intersection of the modified w by a plane orthogonal to d will be bounded by rays whose orientations belong to some fixed set of constant size (this property is not enforced on any *bounded* segment on the boundary of such a cross-section). This deformation creates two new concave edges, e', e''. This folding requires some care, and is done as follows.

We first fold each of the backward faces of w inwards along its positive edge (e_0 or e_1). Let f_1 (resp., f_4) denote the folded face incident to e_1 (resp., e_0). Let π_1, π_4 be the two planes containing f_1, f_4, respectively, and put $e^* = \pi_1 \cap \pi_4 \cap w$. Note that the orientation of e^* varies continuously as a function of the folding angles of f_1 and f_4, and that the angle between e_2 and planes orthogonal to d is at least some fixed constant (that is, $\arcsin \delta$, where δ is as defined in (2)). This implies that we can perform the folding so that the intersections of f_1 and f_4 with planes orthogonal to d have canonical orientations, and e^* is negative. Let w' denote the wedge spanned by the edges e_0, e_1, e^*. Clearly, the respective portions f_1', f_4' of f_1 and f_4 between e^* and the respective edges e_1, e_0 are backward faces of w' and f remains its forward face.

Next, fold the two backward faces of w again, but this time about their common edge e_2. Again, using continuity and the fact that the angles that e_0 and e_1 make with planes orthogonal to d are bounded away from 0, we can perform this folding so that the following property holds: Denote the folded faces as f_2 and f_3, where f_2 (resp., f_3) is folded from the face of w between e_1 and e_2 (resp., between e_0 and e_2). Let e'_1 be the edge of intersection of f_1 and f_2, and let e'_0 be the edge of intersection of f_3 and f_4. Then we require that the intersections of f_2 and f_3 with planes orthogonal to d be at canonical orientations, and that the edges e'_1 and e'_0 be both positive. Let w'' be the wedge spanned by e'_0, e'_1 and e_2. It is easily verified that e^* must be contained in (the interior of) w'', and that $w^* = w' \cup w''$ is a (nonconvex) pentahedral wedge bounded by the faces f, f'_1, f_2, f_3, f'_4. Moreover, f_2 and f_3 are the two backward faces of w'' and its third face, which we denote as f_5, is a forward face. See Figure 6.

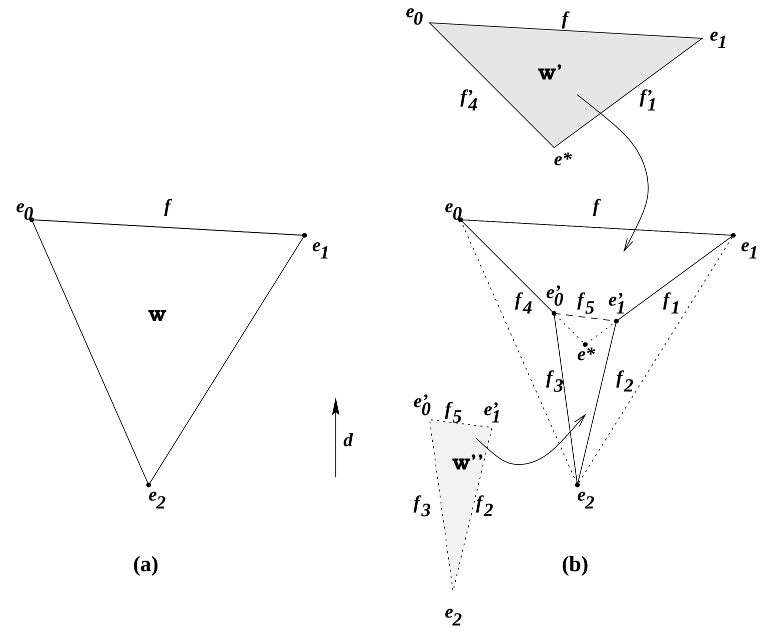

Figure 6: First canonization stage of trihedral wedges. The good direction d is upwards (within the page). (a) A cross section of w by a plane parallel to d. (b) The same cross section after the folding.

We repeat this construction to each wedge $w \in \mathcal{W}_d$, and consider the collection $\mathcal{W}'_d$ consisting of all new wedges like w', w'', constructed above. The size of $\mathcal{W}'_d$ is at most twice that of $\mathcal{W}_d$.

We now relate the number of special cubes in $\mathcal{W}_d$ to those in $\mathcal{W}'_d$. First, arguing as in the proof of Lemma 2.3, it follows that the total number of special cubes in the collection of the deformed pentahedral wedges w^* decreases by at most $O(n^2)$. Indeed, a special cube $C = w_1 \cap w_2 \cap w_3$ in $\mathcal{W}_d$ can stop being a special cube only if a concave edge of one of the deformed w_1^*, w_2^*, w_3^* appears on the boundary of the cube. In particular, this edge crosses the original C.

Let us then bound the number of special cubes in $\mathcal{W}_d$ crossed by a concave edge t, say of w_1. Since t is contained in w_1, by definition, any special cube crossed by t is formed by w_1 and by two other wedges w_2, w_3. For each other wedge w, let I_w denote the interval $w \cap t$. The endpoints of these intervals partition t into at most $2n$ *atomic intervals*, and the intersection of t with a special cube C consists of one or of several consecutive atomic intervals. Actually, such an intersection must be a single atomic interval, for otherwise a fourth wedge would have intersected C, contrary to the properties of special cubes. For the same reason, no atomic interval can lie in two distinct special cubes. It follows that the number of special cubes crossed by t is at most $2n$, so the total number of such special cubes, over all possible choices of w_1, is $O(n^2)$. This establishes the claim.

Next we relate the number of special cubes in the collection of deformed pentahedral wedges to the number of special cubes in $\mathcal{W}'_d$, through the following lemma.

Lemma 4.3 *Let C be a special cube formed by the intersection of w^* with two other deformed wedges. Then C is, or contains, a quasi-special cube in $\mathcal{W}'_d$ formed by the intersection of these two other wedges and one of the (undeformed) wedges w', w''.*

Proof: w^* has five faces. Since the boundary of any special cube formed by the intersection of w^* with two other wedges meets exactly two faces of w^*, there are $\binom{5}{2} = 10$ cases to consider.

Since faces f'_1, f_2 form a concave angle between them, no special cube as above can meet both f'_1 and f_2, so this case is impossible. The same argument rules out the pair f_3, f'_4.

The faces f, f'_1, f'_4 are all contained in w'. Thus, any special cube that meets a pair of these faces is a quasi-special cube in $\mathcal{W}'_d$ that involves w'.

The pair of faces f_2, f_3 are faces of w'', so a special cube C that meets these faces will be a quasi-special cube in $\mathcal{W}'_d$ defined by w'' and the only new wedge that intersects C (but does not define it) is w'.

The remaining four cases involve pairs of faces, one from $\{f_2, f_3\}$ and one from $\{f, f'_1, f'_4\}$ (excluding the pairs (f'_1, f_2) and (f_3, f'_4)). Any such special cube C must cross the forward face f_5 of w'' and cannot meet any edge of f_5, since those are also edges of w^*. Thus, $C \cap w''$ has the structure of a cube, with opposite faces belonging to the same subwedge, so it is a quasi-special cube in $\mathcal{W}'_d$, and the only new wedge that intersects C (but does not define it) is w'. This completes the proof of the lemma. $\Box$

We note that the weight of any quasi-special cube constructed in the preceding proof is at most three: it can be intersected (and not defined) by at most one new wedge that replaces each of the three pentahedral wedges that formed the original cube.

We have thus transformed the trihedral wedges of $\mathcal{W}_d$ into a family $\mathcal{W}'_d$ of new trihedral wedges, whose size is at most $2|\mathcal{W}_d|$ and such that (a) the two backward faces of any new wedge have cross sections orthogonal to d at canonical orientations; and (b) the number of special cubes in the original family is at most the number of quasi-special cubes of level at most three in the new collection $\mathcal{W}'_d$, plus $O(n^2)$. Following Clarkson-Shor's technique [6], the number of quasi-special cubes at level at most three is at most proportional to the number of special cubes in an appropriate random subfamily of $\mathcal{W}'_d$. Hence, it suffices to bound the number of special cubes in $\mathcal{W}'_d$.

Second canonization stage: Folding forward faces. Let $w \in \mathcal{W}'_d$ and let π be a plane orthogonal to d. The cross section $w \cap \pi$ is either a (canonical) 2-wedge (of angle $\geq \alpha_0$) or a truncated 2-wedge (whose rays have canonical orientations). There is a unique plane π_0 orthogonal to d (passing through the apex of w) so that, as we sweep π parallel to itself from infinity to π_0, the cross section $w \cap \pi$ is a 2-wedge that translates at constant velocity. After reaching π_0, the apex of this 2-wedge 'flattens out' and is replaced by a new edge that keeps widening as we sweep; see Figure 7(a). We note though that the above description fits wedges whose forward face points at the direction of *increasing* time. Handling wedges whose forward face points at the direction of decreasing time can be accomplished in a fully symmetric manner.

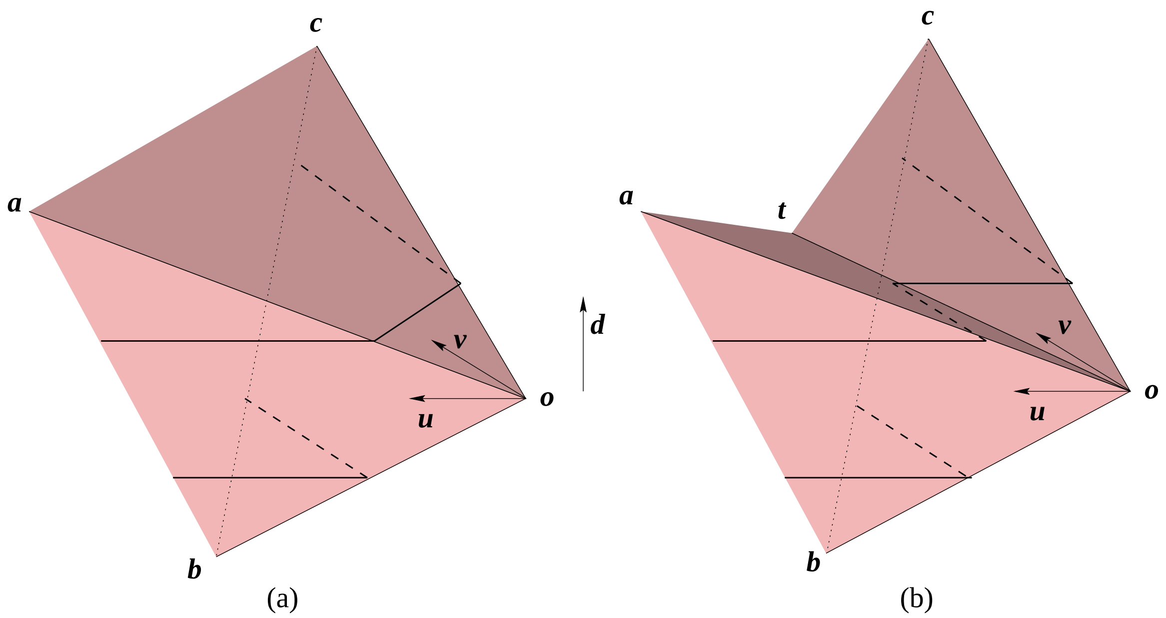

Figure 7: The cross sections of a trihedral wedge w by planes in a good direction: (a) for the original w; (b) after folding inwards a face of w.

We replace the family of the cross sections of w that are truncated wedges by another family, as follows. For a plane π preceding π_0 we leave the cross section $\pi \cap w$ unchanged. For a plane π succeeding π_0, let w' denote the truncated 2-wedge $w \cap \pi$, with bounding rays ρ_1, ρ_2 and bounding segment e. We replace w' by two 2-wedges w'_1, w'_2, whose apices are at the endpoints of e and whose bounding rays are parallel to ρ_1 and ρ_2; see Figure 7(b).

Lemma 4.4 *There exist three trihedral wedges w_1, w_2, w_3, all contained in w, so that:*

(i) $w_1 \cup w_2 \cup w_3$ has one vertex—the apex o of w.

(ii) $w_1 \cup w_2 \cup w_3$ has four (unbounded) edges, where three of these edges are convex and coincide with the edges of w and the fourth is concave.

(iii) $w_1 \cup w_2 \cup w_3$ has four faces, two of which coincide with two faces of w and the other two are obtained by folding the third face of w inwards about each of its edges.

(iv) The cross sections $\pi \cap (w_1 \cup w_2 \cup w_3)$, over all planes π orthogonal to d, coincide with the modified cross sections of w.

(v) Each of the intersections $\pi \cap w_1$, $\pi \cap w_2$, $\pi \cap w_3$, for planes π orthogonal to d, is empty on one side of the apex o, and is a 2-wedge translating at constant velocity on the other side of o. All three kinds of 2-wedges are homothetic (and canonical).

Proof: Let a, b, c be unit vectors along the edges of w, as above. We may assume, without loss of generality, that $\langle a, d \rangle$ and $\langle c, d \rangle$ have the same sign, and that the sign of $\langle b, d \rangle$ is opposite; that is, the face between a and c is the forward face of w.

Let $F_{x,y}$ denote the planar wedge bounded by the two rays that emanate from the apex o of w in the directions x, y, respectively. Take the backward face $F_{a,b}$ (resp., $F_{b,c}$) of w and draw in it the ray u (resp., v) orthogonal to d. Let w_1 be the convex hull of the edges a, u, v, let w_2 be the convex hull of the edges c, u, v, and let w_3 be the convex hull of the edges b, u, v. Note that any intersection of $F_{a,v}$ with a plane orthogonal to d is empty if the plane lies in the negative side of o, and is a ray parallel to v, otherwise. This, and a symmetric statement concerning $F_{c,u}$, imply that all cross sections orthogonal to d of each of w_1, w_2, w_3, if nonempty, are all homothetic to each other. This is easily seen to imply all five properties asserted in the lemma. □

Let $\mathcal{W}''_d$ denote the collection of the transformed, canonical wedges, obtained, as in Lemma 4.4, from the wedges of $\mathcal{W}'_d$. Arguing as in the proof of Lemma 2.3, and in the preceding analysis of the first canonization stage, we obtain:

Lemma 4.5 *The number of special cubes for $\mathcal{W}'_d$ is smaller than or equal to the number of special cubes for $\mathcal{W}''_d$ plus $O(n^2)$.*

Proof: Let w_1, w_2, w_3 be three trihedral wedges in $\mathcal{W}'_d$, forming a special cube C, and let w'_1, w'_2, w'_3 be their canonical images, that is, the union of the three partial wedges that replace each original wedge, as in Lemma 4.4. Denote by t_{w_i} the direction of the new concave edge of w'_i, for $i = 1, 2, 3$. By definition, $C = w_1 \cap w_2 \cap w_3$. We may assume that C does not meet any plane orthogonal to d and passing through the apex of one of the w_i's. Indeed, since the overall number of such planes is $O(n)$, the number of special cubes that violate this assumption is $O(n^2)$, so we may ignore them. Define $C' = w'_1 \cap w'_2 \cap w'_3$. Since each $w'_i \subset w_i$, it follows that $C' \subset C$. The cube C has two opposite faces from the boundary of each w_i. Since only one face of w_i has been folded in the canonization process, it follows that C has three faces, one on the boundary of each w_i, such that these faces also lie on the respective boundaries $\partial w'_i$. Let a be the point of intersection of these faces. Then a is a vertex of C and also a vertex of C'.

Consider the connected component C'' of C' that contains a. If C'' has the combinatorial structure of a cube with pairs of opposite faces lying on the boundary of the same w'_i then, arguing as in the proof of Lemma 4.3, C'' is easily seen to be a quasi-special cube in $\mathcal{W}''_d$. Hence, each special cube for $\mathcal{W}'_d$ of this kind is mapped (in a 1-1 manner) to a quasi-special cube for $\mathcal{W}''_d$, whose level is, as above, at most three.

If C'' does not have the combinatorial structure of a cube, as above, then $\partial C''$ must meet one of the new concave rays t_{w_i}. Suppose, without loss of generality, that t_{w_1} meets $\partial C''$. In particular, t_{w_1} crosses C. Arguing as in the analysis of the first canonization stage, the overall number of such special cubes is $O(n^2)$. This completes the proof of the lemma. □

To recap, we have taken the original family $\mathcal{W}_d$, for a fixed direction d, and have deformed each wedge $w \in \mathcal{W}_d$ in two steps. We first have folded inwards its two backward faces and replaced the 'pinched' wedge w^* by the union of two new wedges, so that the backward faces of each new wedge have planar cross sections orthogonal to d with canonical orientations. Then we have taken each new wedge w' and folded inwards its forward face, so that this pinched wedge can be replaced by the union of three other wedges, so that any cross section orthogonal to d of any new wedge is either empty or is a canonical 2-wedge, and all these canonical 2-wedges (from the same w') are homothetic to each other. The resulting set $\mathcal{W}''_d$ is partitioned to $O(1)$ subfamilies, each consisting of 3-wedges with homothetic cross sections, as above.

4.2 Special cubes and complexity of the union for canonical trihedral wedges

The estimation of the number of special cubes for $\mathcal{W}''_d$ is similar to that for dihedral wedges, with the following significant difference. Let w be a trihedral wedge in $\mathcal{W}''_d$. From the point of view of the sweeping plane π, the modified cross section $\pi \cap w$ is a 2-wedge translating at some fixed velocity until encountering some critical plane π_0, after which it disappears altogether, or, symmetrically, the intersection is empty until the 2-wedge suddenly appears in π and then translates at some constant velocity. We thus need to modify the preceding analysis so that it also handles these appearances and disappearances.

The case of a single canonical family is trivial, and is handled as in the case of dihedral wedges.

The union of two canonical families. For two families, denoted red and blue, we estimate the number of special red-red-blue cubes.

We first bound the number of special red-red-blue cubes, for which the corresponding blue 2-wedge is short (see the analysis of dihedral wedges for the definition). Using the same 2-dimensional frames introduced there, it is easily verified that the number of such cubes is $O(n^2\alpha(n))$. Indeed, as shown there, this number is proportional to the overall complexity of the lower envelopes, over time, of the red 2-wedges as seen along some ray bounding a blue 2-edge, summed over all blue 2-wedges. Since a red 2-wedge r can appear or disappear at some critical time, it means that r is represented in such a 2-D frame by a segment that may start or stop at that critical time. Hence, each envelope is still an envelope of $O(n)$ segments and rays, and the claim follows as above.

Consider next special red-red-blue cubes for which the corresponding blue 2-wedge is long. Any such cube has, as in the dihedral case, a top-top-right vertex. This vertex lies on a 'top-top' red edge (in a cross-section where the boundaries of two red 2-wedges overlap). Since appearances and disappearances of 2-wedges do not affect such overlaps, the number of overlaps is $O(n^2)$, as before.

Consider the planar cross section at the time of a red top-top overlap. This is a planar arrangement of two families of $\alpha/4$-fat wedges. Thus, there are at most $c = O(1/\alpha \log(1/\alpha))$ blue 2-wedges that are long and form a red-red-blue special cube in this arrangement.

Thus, the total number of special cubes for two canonical families is $O(n^2(\alpha(n)+1/\alpha \log(1/\alpha)))$. This, combined with Theorem 1.6, implies that the complexity of the union of two canonical families of trihedral wedges is $O(n^{2+\varepsilon})$.

The union of three canonical families. For three canonical families, denoted red, green and blue, we estimate the number of special red-green-blue cubes, following and adapting the analysis of the dihedral case.

Let $C = r \cap g \cap b$, for $r \in \mathcal{R}, g \in \mathcal{G}, b \in \mathcal{B}$, be such a 'tricolored' special cube. The following properties, established for the case of dihedral wedges, continue to hold for the trihedral wedges of $\mathcal{W}''_d$, as is easily verified.

(i) There exists a plane π such that $\pi \cap C$ has five or six edges.

(ii) At the time (i) occurs, for any $w, w' \in \{r, g, b\}$, either (w, w') is exposed or (w', w) is exposed.

(iii) For each wedge w, the number of bichromatic exposed pairs (w, w'), at any fixed time, is at most some constant c (equal to $O(\frac{1}{\alpha} \log \frac{1}{\alpha})$).

The number of events at which some wedge b starts or stops being initial for some other wedge r is $O(n^2\alpha(n))$, as follows from the arguments used in the case of two families of trihedral wedges.

Consider next an event at which some 2-wedge r starts or stops being semi-free within another 2-wedge b. It is easily checked that at such an event either (a) three sides of red and blue 2-wedges become concurrent at a vertex that is contained only in b (as in the case of dihedral wedges), or (b) some 2-wedge appears or disappears. Clarkson-Shor's technique, combined with the preceding result for two families of trihedral wedges, implies that the number of events of type (a) is also $O(n^{2+\varepsilon})$. Concerning events of type (b), we note that the 2-wedge(s) that newly appear coincide at that critical time with the 2-wedge(s) that disappear. It follows that the status of being semi-free can change at this time only for pairs that involve one of the 2-wedges that appear or disappear. Since the number of such pairs is $O(n)$ (at the time of appearance/disappearance), and there are only $O(n)$ events of appearance/disappearance, it follows that the number of changes of type (b) is only $O(n^2)$.

Let T denote the sorted list of all critical events at which the set of exposed pairs changes. Following the procedure presented for the case of dihedral wedges, it is easily verified that the number of special cubes in this case is at most $(c+1)|T|$, where c is the constant given in (iii) above. Since $|T| = O(n^{2+\varepsilon})$, we conclude that the number of special cubes in the case of three canonical families is $O(n^{2+\varepsilon})$.

Hence, the number of special cubes in $\mathcal{W}_d''$, and thus also in $\mathcal{W}_d$, is $O(n^{2+\varepsilon})$ for any $\varepsilon > 0$. Summing this over all directions in $\mathcal{D}$, the same asymptotic bound also holds for the overall number of special cubes in $\mathcal{W}$. This, combined with Theorem 1.6, implies that the complexity of the union of the original $\mathcal{W}$ is also $O(n^{2+\varepsilon})$ for any $\varepsilon > 0$.

This completes the proof of Theorem 1.3.

Remark: Trying to extend the proof of Theorem 1.3 to the case of wedges that are not substantially fat faces the difficulty that we might have planar cross sections that are *bounded triangles*. In this case it is not necessarily true that the number of exposed pairs (r, b) involving a fixed 2-wedge r is constant at any given time. This is the main reason why substantial fatness is needed in our analysis, and an obvious open problem is to extend the present technique, so that it can also handle bounded cross sections of wedges.

5 The Union of Nearly Equal Cubes

In this section, we apply Theorem 1.3 to derive Theorem 1.4. Without loss of generality, we may assume that the side length of any cube in the given collection $\mathcal{C}$ is between 1 and λ. Fix some constant parameter $t < 1/\sqrt{3}$. Construct a grid G of cubes with side length t. Clearly, any cube $c \in \mathcal{C}$ intersects only a constant number of grid cubes. Consider the collection G' of grid cubes Q, for which the collection $\mathcal{C}_Q$ of cubes of $\mathcal{C}$ that intersect Q is nonempty. Then $|G'| = O(n)$ and $\sum_{Q\in G'} |\mathcal{C}_Q| = O(n)$.

Let Q be a grid cube in G', and put $n_Q = |\mathcal{C}_Q|$. Let c be a cube in $\mathcal{C}_Q$. By the choice of t, Q cannot contain a pair of points that lie on opposite faces of c. It follows that there exists a vertex v of c so that the intersection of ∂c with Q is contained in the union of the three faces of c incident to v. Moreover, any edge of c that meets Q must be incident to v. Replace c by the trihedral wedge w that has v as an apex and is spanned by c (formally, $w = \{v + \rho(x - v) \mid x \in c,\ \rho \geq 0\}$). Then $c \cap Q = w \cap Q$. Let $\mathcal{W}_Q$ denote the resulting collection of trihedral wedges, for all $c \in \mathcal{C}_Q$. Then any vertex of the union of $\mathcal{C}_Q$ within Q is also a vertex of the union of $\mathcal{W}_Q$.

By Theorem 1.3, the complexity of the union of $\mathcal{W}_Q$ is $O(n_Q^{2+\varepsilon})$, for any $\varepsilon > 0$. Summing over all grid cubes Q, Theorem 1.4 follows.

Theorem 1.4 can be extended in several ways, using essentially the same proof.

Theorem 5.1 *Let $\mathcal{B}$ be a family of n boxes so that the ratio between the side lengths of any pair of edges belonging to distinct members or to the same member of $\mathcal{B}$ is at most λ, for some constant parameter $\lambda > 1$. Then the complexity of the union of $\mathcal{B}$ is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$.*

Theorem 5.2 *Let $\mathcal{P}$ be a family of n convex polytopes, each bounded by a constant number of faces, so that the solid angles at the vertices of these polytopes are all (γ, α)-substantially fat, for some constants $\gamma > 4\pi/3$, $\alpha > 0$, and the ratio between the distance from a vertex of some polytope to a non-incident edge of the same polytope, and any other similar distance (within the same or another polytope) is at most λ, for some fixed constant parameter $\lambda \geq 1$. Then the complexity of the union of $\mathcal{P}$ is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$.*

6 Efficient Construction of the Union

In this section we consider the problem of constructing efficiently the (boundary of the) union of n nearly congruent cubes (or of any of the other kinds of objects studied in this paper). For this we adapt the randomized algorithm of Aronov et al. [3, 4], which constructs the boundary of the union along each face of each cube separately, and then 'stitches' together these boundary portions.

Let F be a face of one of the cubes. The algorithm intersects all other cubes with F, thereby obtaining a collection of convex polygons, and then computes the union of these polygons by a straightforward randomized incremental construction that inserts these polygons one by one in a random order. By adapting the analysis in [4] to the case at hand, it is easily seen that the expected running time of the algorithm is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$. The reader is referred to [4] for further details. In other words, we have shown:

Theorem 6.1 *The union of n nearly equal cubes can be computed in randomized expected time $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$. Similar near quadratic bounds hold for the computation of the union of fat dihedral wedges or of substantially fat trihedral wedges.*

7 The Complexity of the Union and Special Cubes

In this section we conclude the paper by proving Theorem 1.6. This provides a general-purpose analysis that obtains a bound on the complexity of the union of an arbitrary family of convex polyhedra in three dimensions, which depends on bounds on the number of special cubes in any subfamily.

We first recall the technique of Aronov et al. [4] for analyzing the complexity of the union of arbitrary convex polyhedra in 3-space. We then extend it and show that the complexity depends on bounding the number of special cubes.

Let $\mathcal{W} = \{w_1, \ldots, w_n\}$ be a collection of n fat wedges in 3-space. Let w_i, w_j, w_k be three distinct wedges in $\mathcal{W}$. Let F_i^1, F_i^2 denote the two faces of w_i. The triple (F_i^a, w_j, w_k), for $a \in \{1, 2\}$, defines a *special quadrilateral*, denoted as $Q_{i,a:j,k}$, if the following conditions hold:

(i) $Q_{i,a:j,k} = F_i^a \cap w_j \cap w_k$ is a quadrilateral.

(ii) Each of the intersections $F_i^a \cap \partial w_j \cap w_k$ and $F_i^a \cap \partial w_k \cap w_j$ consists of two opposite edges of $Q_{i,a:j,k}$.

(iii) $Q_{i,a:j,k} \cap w_l = \emptyset$ for any $w_l \in \mathcal{W} \setminus \{w_i, w_j, w_k\}$.

Let $Q(\mathcal{W})$ denote the number of special quadrilaterals for $\mathcal{W}$, and let $Q(n)$ denote the maximum value for $Q(\mathcal{W})$, taken over all collections of n α-fat wedges (with α fixed).

The level of a vertex v of the arrangement $\mathcal{A}(\mathcal{W})$ is the number of polyhedra in $\mathcal{W}$ that contain v in their interiors. Let us denote by $C_0(\mathcal{W})$ the number of vertices on $\partial U_{\mathcal{W}}$ (which is equal to the number of level-0 vertices of $\mathcal{A}(\mathcal{W})$), and by $C_0(n)$ the maximum value of $C_0(\mathcal{W})$, taken over all collections $\mathcal{W}$ of n α-fat wedges (with α fixed).

Following Aronov et al. [4], we call a triple (f, e, e') *special* if f is a level-1 2-dimentional face of $\mathcal{A}(\mathcal{W})$, and e, e' are 0-level edges of f, and we can trace the boundary of f from e to e' without passing through any other level-0 edge. We denote by $C^{(1)}(\mathcal{W})$ the number of special triples in $\mathcal{A}(\mathcal{W})$, and by $C^{(1)}(n)$ the maximum value of $C^{(1)}(\mathcal{W})$, taken over all collections $\mathcal{W}$ of n α-fat wedges (with α fixed).

The following relation is established in [4] between $C_0(n)$ and $Q(n)$. First we have

$$\frac{n-5/3}{n} C_0(n) \le C_0(n-1) + O(n^2) + \frac{4}{n} C^{(1)}(n) \ . \tag{3}$$

The number of special triples is bounded in turn by the recurrence:

$$\frac{n-2}{n} C^{(1)}(n) \le C^{(1)}(n-1) + \frac{1}{n} O(n^2 + Q(n)) \ . \tag{4}$$

Notice that Definition 1.5 implies that every face of a special cube is a special quadrilateral. The proof proceeds by charging special quadrilaterals to special cubes or to vertices at shallow levels in the arrangement $\mathcal{A}(\mathcal{W})$. Let $Q = Q_{1,a:2,3}$ be a special quadrilateral. Consider the corresponding intersection $C = w_1 \cap w_2 \cap w_3$. We distinguish between two cases:

Case (a): C has the combinatorial structure of a cube. Let $p_{ij} = F_1^1 \cap F_2^i \cap F_3^j$, for $i, j = 1, 2$, denote the four vertices of Q, and let $s_{ij} = w_1 \cap F_2^i \cap F_3^j$ denote the edge of C emanating from p_{ij} 'away' from Q. Let t denote the total number of intersections of the edges s_{ij} with faces of other wedges. Fix some threshold parameter k, to be specified later, and consider the following two subcases:

(i) $t > k$: We charge Q to the first k intersections encountered along the incoming edges, and note that each of the charged vertices is a vertex of the arrangement at level at most k. Moreover, any such vertex v can be charged by at most six special quadrilaterals. Indeed, any such quadrilateral contains a vertex (at level 0) that lies on one of the intersection edges incident to v, and the portion of that edge between the quadrilateral and v does not contain any other vertex at level 0.

Denote by $C_\xi(n)$ (and $C_{\le\xi}(n)$) the maximum number of vertices at level ξ (resp., at most ξ) in an arrangement of n α-fat wedges. Applying Clarkson-Shor's probabilistic analysis technique [6], we have

$$C_{\le k}(n) = \sum_{\xi=0}^{k} C_\xi(n) = O(k^3 C_0(n/k)) \ .$$

This, and the argument in the preceding paragraph, imply that the number of special quadrilaterals of this type is $O(k^2 C_0(n/k))$.

(ii) $t \le k$: Suppose first that C is crossed by an edge of some polyhedron in $\mathcal{W}$ (that is, the edge intersects ∂C at two points). We claim that there exists an intersection point q between such an edge and ∂C that lies at level $\le k$ in $\mathcal{A}(\mathcal{W})$. Indeed, if an edge e crosses ∂C then, since it does not cross Q, it must cross one of the four 'side faces' of C adjacent to Q. Let F be such a face. There exist at most k polyhedra that cross the two side edges of F (those adjacent to Q). The cross section of any other polyhedron with F must be a convex polygon which is either fully contained in F or 'exits' it only through its bottom edge. It is easily verified that the boundary of the union of the cross sections $F \cap w$, over all such polyhedra w, must contain a vertex v, which is thus an intersection of a polyhedron edge with F that lies at level at most k in $\mathcal{A}(\mathcal{W})$. We then charge Q to v, and note that v can be charged by at most $O(k^2)$ special quadrilaterals Q. Indeed, v determines the face F and thus one of the three polyhedra that induce Q. The other two must be two of the at most k other polyhedra that contain v in their interior. The number of intersections between edges and faces of polyhedra is $O(n^2)$, which implies that the number of special quadrilaterals Q under consideration is $O(k^2n^2)$.

Suppose next that C is intersected by an edge e of a polyhedron in $\mathcal{W}$ which has an endpoint inside v. By assumption, the polyhedron bounded by e is either fully contained within C or 'exits' C only through its bottom face (the one opposite to Q). As above, the boundary of the union of the portions within C of all such polyhedra must contain a vertex of one of them, which is thus a vertex at level at most k in $\mathcal{A}(\mathcal{W})$. Arguing as above, the number of special quadrilaterals Q for which this subcase applies is only $O(k^2n)$.

We may thus assume that C is not crossed by any edge of a polyhedron in $\mathcal{W}$, so the only polyhedra that intersect C are those $t \le k$ polyhedra that intersect some of the four side edges of C. Define the *level* of a cube C that satisfies conditions (i) and (ii) of Definition 1.5 to be the number of polyhedra of $\mathcal{W}$, that intersect C. other than those three that define C. Hence, in the case at hand, C is a cube at level at most k. We charge Q to C (which can be charged in this manner at most six times). Denote the maximum number of cubes at level ξ (resp., at most ξ) in a collection of n α-fat wedges by $\psi_\xi(n)$ (resp., by $\psi_{\le\xi}(n)$). In particular, ψ_0 bounds the number of special cubes in the given collection. Applying again Clarkson-Shor's technique, we obtain

$$\psi_{\le k}(n) = \sum_{\xi=0}^{k} \psi_\xi(n) = O(k^3\psi_0(n/k))$$

Hence, the number of special quadrilaterals under consideration is $O(k^3\psi_0(n/k))$.

Case (b): C does not have the combinatorial structure of a cube. (This case is easy to analyze in the case of wedges, and most of the foregoing analysis is not required for that special case.)

For convenience, assume that the face F_1^a of w_1 that contains Q is F_1^1. Denote by F_2^1, F_2^2 the two faces of w_2 that contain two opposite edges of Q and by F_3^1, F_3^2 the two faces of w_3 that contain the other two opposite edges of Q. If C is not a cube then one of the following cases has to arise:

(i) One of the four intersection edges $w_1 \cap F_2^i \cap F_3^j$, for $i, j = 1, 2$, is unbounded.

(ii) One of those four intersection edges ends within the interior of w_1.

(iii) Not all four of those intersection edges leave w_1 from the same face.

(Note that case (iii) cannot occur for dihedral wedges, since any such wedge has only two faces.)

In subcases (i) and (ii) we can charge Q to the corresponding intersection edge $F_2^i \cap F_3^j$. It is clear that any such edge can be charged in this manner at most twice. Indeed, if it is charged by a special quadrilateral Q as above, then the portion of the edge between its intersection point v with Q and its endpoint, or from v to infinity, is fully contained in the third polyhedron w_1. Since v is a vertex at level 0, the claim is immediate. It follows that the number of special quadrilaterals in subcases (i) and (ii) is $O(n^2)$.

In subcase (iii) at least one of the 'side faces' F_2^i, F_3^j (say, F_2^1), has the property that $\varphi = C \cap F_2^1$ is not a quadrilateral. Moreover, if e is the edge of φ that is also an edge of Q, then the edges of φ adjacent to e both lie on ∂w_3 and the two (necessarily distinct!) edges adjacent to these edges lie on ∂w_1; φ may have additional edges that lie on either boundary.

We first assume that neither of the two edges $w_1 \cap F_2^1 \cap F_3^j$ is crossed by more than k other polyhedra, for the threshold parameter k that we have chosen. If this does occur, we use the same charging scheme employed in Case (a) above.

Suppose that the remaining portion of $\partial\varphi$ contains two successive edges that lie on ∂w_1. Then φ has a vertex that is an intersection of an edge of w_1 with F_2^1. We can then charge Q to such an intersection v, and note that v cannot be charged more than $2k$ times. Indeed, consider the face $K = F_2^1 \cap w_1$. This is a convex polygon with $O(1)$ edges and with v as a vertex, and φ is obtained by intersecting K with w_3. Let w be another polyhedron in $\mathcal{W}$ that contains v and induces a face $\varphi' = K \cap w$ with the same structure as above, so that φ' is adjacent to a special quadrilateral Q' along some edge of K. Then, as is easily verified, at least one of the two edges of φ lying on ∂w_3 and one of the two edges of φ' lying on ∂w must cross each other (see Figure 8(a)), which implies the asserted property. Hence, the number of special quadrilaterals Q in this subcase is $O(n^2 k)$.

Consider next the case where φ has two successive edges that lie on ∂w_3. In this case φ has a vertex that is an intersection of an edge of w_3 with F_2^1. We charge Q to such an intersection v. Given v, we know w_3 and F_2^1. Their intersection is a convex polygon K' with $O(1)$ sides, and φ touches at least three of its sides, so that one of the 'chords' of φ, i.e., an edge of φ lying in the interior of K', is disjoint from any other polyhedron—this is the edge incident to Q (see Figure 8(b)). It is easily checked that, once the two edges of K' connected by this chord are fixed (there are $O(1)$ choices for such a pair of edges) the chord is unique, from which the claim follows. Hence, the number of special quadrilaterals Q in this subcase is only $O(n^2)$.

Otherwise, the edges of φ alternate between edges incident to ∂w_1 and edges incident to ∂w_3, and their total number is at least six. We claim that, when F_2^1 and w_3 are fixed, there can be only $O(1)$ polyhedra w_1 that generate a special quadrilateral Q with F_2^1 and w_3, as above. Indeed, put $K = F_2^1 \cap w_3$. K is a convex polygon with $O(1)$ edges but with at least three edges that lie on ∂w_3, and $\partial\varphi$ has at least three chords of ∂w_1 that connect pairs of these edges, with one of the chords (the one incident to Q) being disjoint from any other polyhedron. Arguing as above, it is easy to see that, once the two edges of K connected by this chord are fixed (there are $O(1)$ choices for such a pair of edges), the chord is unique (see Figure 8(c)). Indeed, if two chords connect the same pair of edges of K' then one of them bounds a quadrilateral within K', contrary to assumption. This implies that the number of special quadrilaterals Q in this subcase is only $O(n^2)$.

Thus, if we add all the bounds obtained so far, we obtain the following recurrence for the maximum number $Q(n)$ of special quadrilaterals:

$$Q(n) = O(k^2 n^2 + k^2 C_0(n/k) + k^3 \psi_0(n/k)) \ .$$

By assumption, $\psi_0(n/k) = O((n/k)^\gamma)$, so we have

$$Q(n) = O(k^2 n^2 + k^{3-\gamma} n^\gamma + k^2 C_0(n/k)). \tag{5}$$

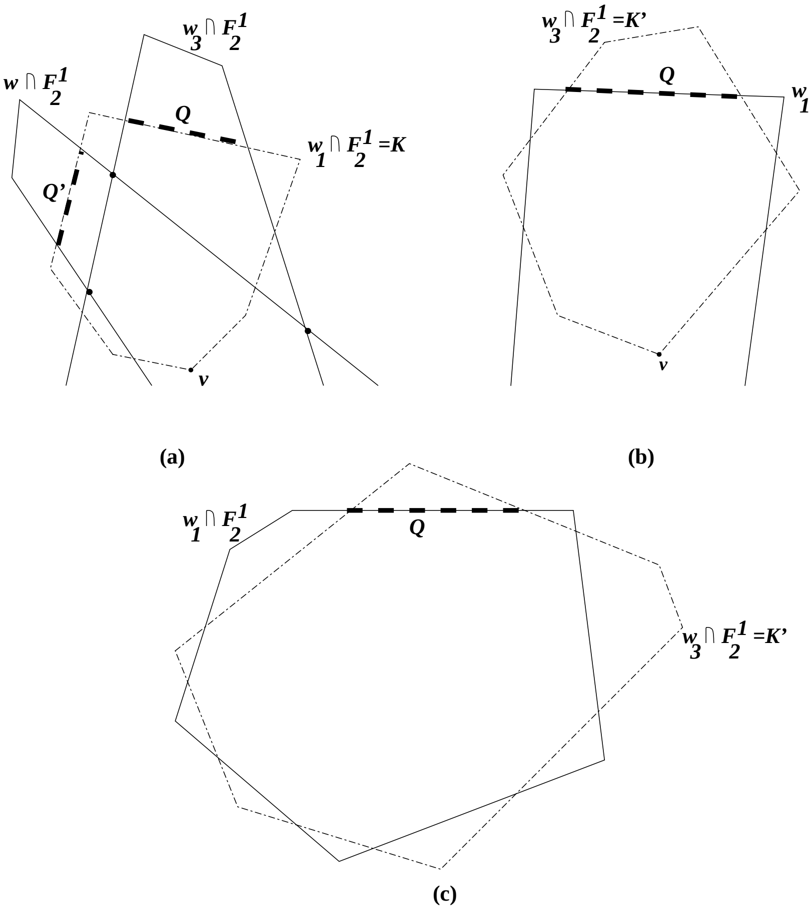

Figure 8: The cases where φ is not a quadrilateral: (a) φ has a vertex lying on an edge of w_1; (b) φ has a vertex lying on an edge of w_3; (c) $\partial\varphi$ alternates between edges lying on ∂w_1 and edges lying on ∂w_3.

Arguing as in [15], for example, and using the fact that $\gamma > 2$, the solution of the combined recurrences (3), (4) and (5) can be shown to be $O(n^\gamma)$. This completes the proof of Theorem 1.6. $\square$

References

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

[2] H. Alt, R. Fleischer, M. Kaufmann, K. Mehlhorn, S. Näher, S. Schirra and C. Uhrig, Approximate motion planning and the complexity of the boundary of the union of simple geometric figures, *Algorithmica* 8 (1992), 391–406.

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

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

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

[6] K. Clarkson and P. Shor, Applications of random sampling in computational geometry II, *Discrete Comput. Geom.* 4 (1989), 387–421.

[7] A. Efrat, The complexity of the union of (α, β)-covered objects, *Proc. 15th ACM Symp. on Computational Geometry*, 1999, 134–142.

[8] A. Efrat and M. Katz, On the union of α-curved objects, *Proc. 14th ACM Symp. on Computational Geometry*, 1998, 206–213.

[9] A. Efrat, G. Rote and M. Sharir, On the union of fat wedges and separating a collection of segments by a line, *Comput. Geom. Theory Appls.* 3 (1993), 277-288.

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

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

[12] D. Leven and M. Sharir, On the number of critical free contacts of a convex polygonal object moving in two-dimensional polygonal space, *Discrete Comput. Geom.* 2 (1987), 255–270.

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

[14] J. Pach and G. Tardos, On the boundary complexity of the union of fat triangles, *Proc. 41st Annual Symp. Foundations of Comp. Science*, IEEE Comp. Soc., Los Alamitos, 2000, pp. 423–431.

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