The union of colorful simplices spanned by a colored point set

André Schulz* Csaba D. Tóth†

Abstract

A simplex spanned by a colored point set in Euclidean d-space is *colorful* if all vertices have distinct colors. The union of all full-dimensional colorful simplices spanned by a colored point set is called the *colorful union*. We show that for every $d \in \mathbb{N}$, the maximum combinatorial complexity of the colorful union of n colored points in $\mathbb{R}^d$ is between $\Omega(n^{(d-1)^2})$ and $O(n^{(d-1)^2}\log n)$. For $d = 2$, the upper bound is known to be $O(n)$, and for $d = 3$ we present an upper bound of $O(n^4\alpha(n))$, where $\alpha(\cdot)$ is the extremely slowly growing inverse Ackermann function. We also prove several structural properties of the colorful union. In particular, we show that the boundary of the colorful union is covered by $O(n^{d-1})$ hyperplanes, and the colorful union is the union of $d+1$ star-shaped polyhedra. These properties lead to efficient data structures for point inclusion queries in the colorful union.

1 Introduction

Given a colored set S of n points in d-dimensional Euclidean space $\mathbb{R}^d$, a simplex is *colorful* if its vertices have pairwise distinct colors. The *simplicial depth* (resp., *colorful simplicial depth*) of a point $p \in \mathbb{R}^d$ is the number of full dimensional closed simplices (resp., colorful simplices) spanned by S and containing p. By Carathéodory's theorem the set of all points with positive simplicial depth is the convex hull conv$S$. We call the set of all points with positive *colorful* simplicial depth the *colorful union* and denote it by U_S. It is the union of all colorful simplices, hence it is a polyhedron in $\mathbb{R}^d$. The study of colorful depth was pioneered by Bárány [9], who deduced a lower bound on the maximum simplicial depth by showing that a point lies in a colorful simplex for many random colorings of the point set.

The *colorful linear programming (CLP)* problem was proposed by Bárány and Onn [10]: for a colored set of n points in $\mathbb{R}^d$ and a query point $q \in \mathbb{R}^d$, find a colorful simplex that contains q or report that no such simplex exists. An important special case is that q lies in the *core* of the colored points, which is the intersection of the convex hulls of the color classes. In this special case, q is in the colorful union by the colorful Carathédory theorem [9], and so the CLP is guaranteed to be feasible. This case was thoroughly studied by Bárány and Onn [10, 11] and Deza et al. [17]. The colorful Carathéodry theorem has been strengthened.

---

*Institut für Mathematische Logik und Grundlagenforschung, Westfälische Wilhelms-Universität, Münster, Germany. `andre.schulz@uni-muenster.de`. Partially supported by the German Research Foundation (DFG) under grant SCHU 2458/1-1.

†Department of Mathematics and Statistics, University of Calgary, Calgary, AB, Canada, `cdtoth@math.ucalgary.ca`. Supported in part by NSERC grant RGPIN 35586. Research by Tóth has been conducted at Tufts University, Medford, MA.

Arocha et al. [7] and Holmsen et al. [22] have independently proved that the query point q lies in the colorful union already if q is contained in the convex hull of the union of any two color classes. Very recently, a further strengthening was obtained by Meunier and Deza [26]. However, very little is known about the general case where q is an arbitrary point in $\mathbb{R}^d$.

We design efficient data structures for a colored set S of n points in $\mathbb{R}^d$ that supports point inclusion queries for the colorful union U_S. For $d = 2$, it is easy to construct a data structure with $O(\log n)$ query time, $O(n)$ space and $O(n \log n)$ preprocessing time. Boissonnat et al. [12] proved that the union U_S of all colorful triangles for a set of n colored points in the plane is a simple polygon with at most $2n - 3$ vertices, which can be computed in $O(n \log n)$ time. Hence, a point location data structure of size $O(n)$ can support point inclusion queries for U_S. For an efficient data structure in higher dimensions one has to understand the combinatorial structure of the colorful union. We present the following results.

1. We show that for every $d \in \mathbb{N}$, the maximum combinatorial complexity of the colorful union for n colored points in $\mathbb{R}^d$ is between $\Omega(n^{(d-1)^2})$ and $O(n^{(d-1)^2} \log n)$. A tight worst case bound of $\Theta(n)$ has been known for $d = 2$, and we prove a stronger upper bound of $O(n^4 \alpha(n))$ for $d = 3$, where $\alpha(\cdot)$ is the extremely slowly growing inverse Ackermann function.

2. We show that U_S is the union of $d+1$ star-shaped polyhedra, where the star-centers are the vertices of an arbitrary colorful simplex. This reduces point inclusion queries to ray-shooting queries, and leads to efficient data structures to support point inclusion queries in arbitrary fixed dimension. In particular, in $\mathbb{R}^d$, there is a data structure of size m, $n^{d-1} \leq m \leq n^{(d-1)^2}$, that supports point inclusion queries for U_S in $O(n^{d-1+\varepsilon}/m^{1/(d-1)})$ time for any fixed $\varepsilon > 0$.

3. We show that the colorful union may have undesirable features already in $\mathbb{R}^3$. We construct colored sets S of n points in $\mathbb{R}^3$ with each of the following properties:

   (i) U_S is not star-shaped;
   (ii) U_S has a face whose edges are all adjacent to *reflex* dihedral angles;
   (iii) the boundary of U_S contains a polygonal chain of $\Omega(n)$ reflex vertices.

   On the contrary, U_S is star-shaped and has no two consecutive reflex vertices in $\mathbb{R}^2$.

Organization.

We discuss related work in the remainder of Section 1. We introduce notation and present a few preliminary lemmata in Section 2. Our lower and upper bounds for the combinatorial complexity of the colorful union in arbitrary constant dimension $d \in \mathbb{N}$ are presented in Section 3. We use our combinatorial results for designing efficient data structures for inclusion queries in Section 4. Pathological configurations for the colorful union in 3-space are presented in Section 5. An improved upper bound for the combinatorial complexity of U_S in 3-space is proved in Section 6. We conclude in Section 7 with open problems.

Related Work.

The union of geometric objects in $\mathbb{R}^d$ has applications in constructive solid modeling, motion planning, proximity problems, and conflict-free colorings. We redirect the interested reader to the excellent survey [27] on unions of various geometric objects in Euclidean space. The maximum combinatorial complexity of the union of m full-dimensional simplices in $\mathbb{R}^d$ is $O(m^d)$, this bound is attained if the simplices are flat. Research efforts focused on finding families of simplices whose union have smaller complexity. One example is the family of *fat* simplices, where the dihedral angles of all simplices are bounded from below by a constant $\delta > 0$, which is the *fatness parameter* of the family. The maximum complexity of the union of m fat triangles in $\mathbb{R}^2$ is known to be between $O(m2^{\alpha(m)} \log^* m)$ and $\Omega(m\alpha(m))$, where $\alpha(\cdot)$ is the inverse Ackermann function [19, 29, 33] (the upper bound has recently been improved from $O(m \log \log m)$ [25]). Ezra and Sharir [20] proved that the complexity of m fat tetrahedra in $\mathbb{R}^3$ is $O(m^{2+\varepsilon})$ for every $\varepsilon > 0$. Our result about the complexity of the colorful union is another example: n colored points in $\mathbb{R}^d$ determine $m = O(n^{d+1})$ colorful simplices, yet the complexity of their union is only $O(n^{(d-1)^2} \log n) = O(m^{(d-1)(1-\frac{2}{d+1})} \log m)$.

The colorful simplices in a colored points set in $\mathbb{R}^d$ can be interpreted as a complete multipartite $(d+1)$-uniform geometric hypergraph. Akiyama and Alon [5] studied the number of pairwise disjoint simplices in such a hypergraph, Dey and Pach [15] studied the intersections of hyperedges, related to the higher dimensional analogues of the crossing number. To the best of our knowledge, the combinatorial complexity of the union of colorful simplices in $\mathbb{R}^d$ has not been considered before for dimensions $d \geq 3$.

2 Preliminaries

Let S be a colored set of $n \geq d+1$ points in $\mathbb{R}^d$. We assume throughout this paper that every $d+1$ points in S are affinely independent. For $k = 0, 1, \ldots, d$, a *k-simplex* is a subset $P \subseteq S$ of size $k+1$. A d-simplex in $\mathbb{R}^d$ is also called a *simplex* for short. The convex hull of a subset $P \subseteq S$ is denoted by $\text{conv}P$. For subsets with up to three elements, we use the shorthand notation $p = \{p\}$, $pq = \{p, q\}$, and $pqr = \{p, q, r\}$; if there is no danger of confusion, we also use the same notation for the convex hulls $pq = \text{conv}\{p, q\}$ and $pqr = \text{conv}\{p, q, r\}$. We say that a simplex P *contains* a point set $Q \subset \mathbb{R}^d$ if $Q \subseteq \text{conv}P$.

The colors of the points in S are represented by positive integers. For a single point s, we denote by $\text{color}(s)$ the color of s. For $S' \subseteq S$, we denote by $\text{color}(S') \subset \mathbb{N}$ the set of colors that occur in S'. A k-simplex is *colorful* if its $k+1$ vertices have pairwise distinct colors. We assume throughout that $|\text{color}(S)| \geq d+1$, hence there is at least one full-dimensional colorful simplex in S. The *colorful union* of S is the polyhedron

$$U_S = \bigcup \{\text{conv}P : P \subseteq S, |P| = d+1, \text{ and } |\text{color}(P)| = d+1\}.$$

The *combinatorial complexity* of a polyhedron U is the number of its k-faces for all $k = 0, 1, \ldots, d$. In $\mathbb{R}^3$, in particular, if a polyhedron U is simply connected (that is, homeomorphic to a ball), then its 0-, 1-, and 2-faces form a planar graph. In this case, by Euler's formula, the combinatorial complexity of U is proportional to the number of vertices, edges, or faces.

Extremal sets and shells.

A point $p \in S$ is *extremal* if there is a halfspace H^+ bounded by a hyperplane H that contains p such that $\mathrm{color}(S \cap H^+) \subseteq \{\mathrm{color}(p)\}$. In general, a k-simplex $P \subset S$ is *extremal* if it is colorful and there is a hyperplane H containing P such that $\mathrm{color}(S \cap H^+) \subseteq \{\mathrm{color}(P)\}$, where H^+ is again a halfspace bounded by H. For an extremal $(d-2)$-simplex P, let the *wedge* $W(P)$ be the intersection of the closed halfspaces $\mathrm{cl}(H^-)$ for all hyperplanes H that witness that P is extremal. By definition, all points whose color is not in $\mathrm{color}(P)$ must lie in $W(P)$. The boundary of $W(P)$ consists of two half-hyperplanes, say H_1 and H_2, each of which contains P and one additional point of S, say $s_1, s_2 \in S$, respectively. The colors of s_1 and s_2 may be the same, but they differ from any color in $\mathrm{color}(P)$. We call the two $(d-1)$-simplices, $P \cup \{s_1\}$ and $P \cup \{s_2\}$, the *shells* of P. We also say that P is the *axis* of these two shells. It is clear that S determines $O(n^{d-1})$ extremal $(d-2)$-simplices, hence there are $O(n^{d-1})$ shells in $\mathbb{R}^d$. Figure 1 illustrates the definitions by an example.

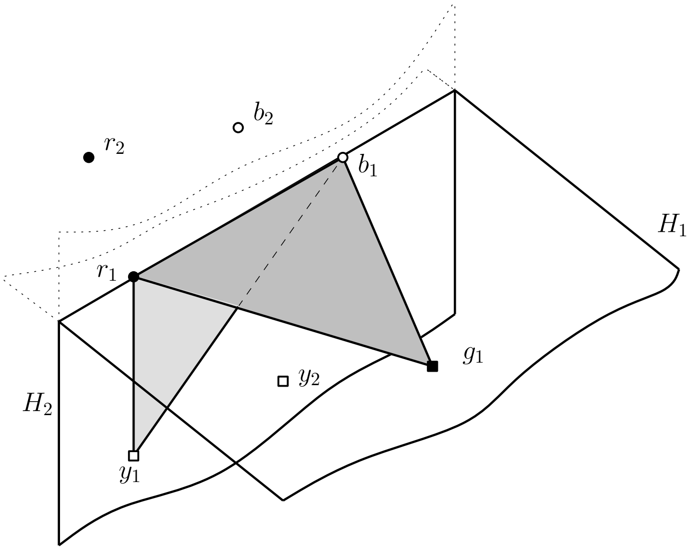

Figure 1: The 2-simplex r_1b_1 spans an extremal 2-simplex. This is witnessed by the plane H_1, which has only points of the same color as r_1 or b_1 on one side. The plane H_2 also witnesses that r_1b_1 is extremal. The the triangles $r_1b_1g_1$ and $r_1b_1y_1$ are the two shells of axis r_1b_1. The bold edges of H_1 and H_2 indicate the wedge $W(r_1b_1)$.

Parity constraints.

The following Lemma is due to Deza et al. [16, Theorem 3.5]. For completeness we include a formal proof.

Lemma 1. *Let S be a colorful point set in $\mathbb{R}^d$ such that each color class has even cardinality. If a point $q \in \mathbb{R}^d \setminus S$ does not lie on any hyperplane spanned by S, then q is contained in an even number of colorful simplices.*

Proof. Consider a continuous path γ from point q to an arbitrary point r in the exterior of $\mathrm{conv}S$ avoiding all affine $(d-2)$-flats spanned by any d points in S. (In case γ runs through some $(d-2)$-flat we perturb the path slightly.) We follow γ from q to r, and keep track of the colorful simplicial depth. Clearly, the colorful depth of r is zero. The colorful depth changes only if γ crosses a colorful $(d-1)$-simplex P spanned by S. Let $S_0 \subset S$ be the set of all points whose color is missing from color(P). By our assumption, the cardinality of S_0 is even. Denote by H the hyperplane spanned by P. When γ crosses conv$P$ from halfspace H^- to H^+, the colorful simplicial depth changes by $|S_0 \cap H^+| - |S_0 \cap H^-|$, which is even, since $|S_0| = |S_0 \cap H^+| + |S_0 \cap H^-|$ is even. □

We show next that the boundary of the colorful union is covered by shells.

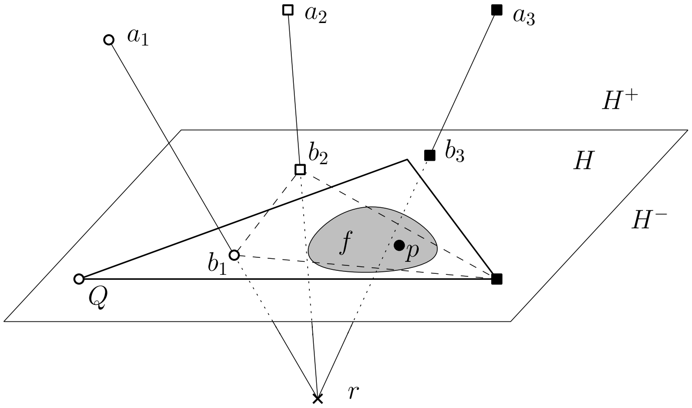

Figure 2: Construction of Lemma 2.

Lemma 2. *Every face of the colorful union U_S is contained in a shell.*

Proof. Let f be a face of the polyhedron U_S. Since U_S is the union of colorful simplices, f lies on the boundary of some colorful simplex $P \subset S$. Hence, f is contained in a colorful $(d-1)$-simplex $Q \subset P$. Assume that $P = Q \cup \{r\}$, and denote by H the hyperplane spanned by Q that induces two open halfspaces H^+ and H^-, such that $r \in H^-$. If there is a point $s \in S \cap H^+$ with color$(s) \notin$ color(Q), then f would be in the interior of the union of two colorful simplices conv$(Q \cup \{r\}) \cup$ conv$(Q \cup \{s\})$. Hence every point in $S \cap H^+$ is colored by some color from color(Q), that is, color$(S \cap H^+) \subseteq$ color(Q).

If color$(S \cap H^+) \subsetneq$ color(Q), then there is a colorful $(d-2)$-simplex $R \subset Q$ such that color$(S \cap H^+) \subseteq$ color(R). This means that R is an extremal $(d-2)$-simplex with axis R, and Q is a shell of R. Hence f is contained in shell Q, as required.

Now assume that color$(S \cap H^+) =$ color(Q). Let $A = \{a_1, \ldots, a_d\} \subset S \cap H^+$ be a set of d points with distinct colors (i.e., color$(A) =$ color(Q)). For each a_i, let $b_i = a_i r \cap H$ (refer to Fig. 2). Let $B = \{b_1, \ldots, b_d\}$, and color each b_i with the color of a_i. In the point set $Q \cup B$, each of the d color classes has cardinality 2. Pick a point p in the interior of face $f \subset H$ that does not lie on any hyperplane spanned by S. The point p is in the interior of the colorful d-simplex Q, so by Lemma 1, it is in the interior of some other colorful d-simplex Q' spanned by $Q \cup B$, which has at least one vertex in B. Let Q'' be the set obtained from Q' by replacing all points $b_i \in Q$ by their corresponding counterparts a_i. It can be easily checked that $Q'' \cup \{r\}$ is colorful and contains p in its interior. This contradicts our assumption that f (and p) are on the boundary of U_S. □

Visibility within the colorful union.

We define *visibility* with respect to the polyhedron U_S. We say that two points $p, q \in U_S$ are visible to each other if the line segment pq is contained in U_S.

Theorem 1. *Let S be a colored point set in $\mathbb{R}^d$ and let $P \subset S$ be an arbitrary colorful simplex. Then every point in U_S is contained in a colorful simplex that has a vertex in P. Consequently, every point in U_S is visible from one vertex of P.*

Proof. Let $q \in U_S$ be a point in the colorful union. If $q \in \text{conv}P$, then any vertex of P sees q. Assume that $q \notin \text{conv}P$. Since $q \in U_S$, there is a colorful simplex $Q \subseteq S$ that contains q. If P and Q have a common vertex, the statement of the theorem follows trivially. Assume that $P \cap Q = \emptyset$. Note that P and Q are each colorful, but they do not necessarily have the same $d+1$ colors. Successively pick a point in P and a point in Q whose colors are unique in $P \cup Q$, and recolor both points to a new color. The recoloring ensures that both P and Q remain colorful, and $P \cup Q$ has $d+1$ colors. Clearly, if a simplex $R \subset P \cup Q$ is colorful in the new colors, then it was colorful in the original colors, too. Now $P \cup Q$ is a colored point set where every color class has size 2. The point q is contained in the colorful simplex Q spanned by the points of $P \cup Q$. By Lemma 1, q is contained in some other colorful simplex R. Then R must have at least one common vertex with P. □

A polyhedron U in $\mathbb{R}^d$ is *star-shaped* if there is a point $p \in U$ such that every point in U is visible from p. Such a point $p \in U$ is called a *star center*. We show that in the plane, the colorful union U_S is star-shaped. In Section 5, however, we construct colored point configurations in $\mathbb{R}^3$ such that U_S is not star-shaped.

Lemma 3. *If S be a colored point set in $\mathbb{R}^2$, then U_S is star-shaped.*

Proof. Let $P \subseteq S$ be the set of all extremal points in S and let $\mathcal{W} = \{W(p) : p \in P\}$ be the set of all wedges determined by some extremal point in S. Recall that a wedge $W(p)$ contains all points in S whose color is not color(p). Since color$(S) \geq 4$, any three wedges in $\mathcal{W}$ have a non-empty intersection (in fact, their intersection contains a point of S). By Helly's theorem, all wedges in $\mathcal{W}$ have a common point, say $o \in \mathbb{R}^2$. We show that U_S is star-shaped with star center o. It is enough to show that for every point $q \in U_S$, the line segment oq lies in U_S. Consider an arbitrary point $q \in U_S$, and let e be an arbitrary edge of U_S that intersects the ray $\overrightarrow{oq}$. By Lemma 2, edge e lies on the boundary of a wedge $W(p) \in \mathcal{W}$, where o lies in the interior of $W(p)$. So $\overrightarrow{oq}$ crosses e from the interior to the exterior of U_S. It follows that $\overrightarrow{oq}$ crosses the boundary of U_S at most once, and so $\overrightarrow{oq} \cap U_S$ is a line segment. Since both o and q are in $\overrightarrow{oq} \cap U_S$, segment oq lies in U_S, as required. □

3 The combinatorial complexity of the colorful union

If the convex hull of S is a colorful simplex, then $U_S = \text{conv}S$ with $d+1$ vertices. Hence, the *minimum* combinatorial complexity (over all sets S) of U_S for $n \geq d+1$ colored points in $\mathbb{R}^d$ is $\Theta(1)$. If S is in convex position, then the minimum combinatorial complexity of U_S is $\Theta(n)$. This complexity is attained for a colored point set constructed recursively as follows. Start with the $d+1$ vertices of a colorful simplex. In each step, choose an arbitrary (colorful) face Δ of the current convex hull, place a new point near the center of Δ in the exterior of the convex hull, and color it with a color that does not occur in $\mathrm{color}(\Delta)$.

In the remainder of this section, we present lower and upper bounds for the *maximum* combinatorial complexity of the colorful union of n colored points in $\mathbb{R}^d$.

3.1 Lower bounds for the maximum combinatorial complexity

Theorem 2. *For every integer $d \geq 2$, there are $(d+1)$-colored point sets in $\mathbb{R}^d$ of size $n \geq d+1$ such that the combinatorial complexity of the colorful union is $\Omega(n^{(d-1)^2})$.*

Proof. Let $d \geq 2$ be a fixed positive integer. For every $n \geq d+1$, we construct a set S of n points of $d+1$ colors in $\mathbb{R}^d$. We have one point of color d and $d+1$ each. Let a (resp., b) be the point of color d (resp., $d+1$) on the x_d-axis at N (resp., $N+1$), for a sufficiently large N to be specified later.

The remaining $n-2$ points are evenly distributed in the first $d-1$ color classes. We construct the position of these point in three steps. Let $\mathbb{R}^{d-1}$ denote the subspace of $\mathbb{R}^d$ spanned by the first $d-1$ coordinate axes. **Step 1.** For $i = 1, \ldots, d-1$, place the points of color i in the interval $(0,1)$ of the x_i-axis. Let S_1 denote the set of these points. **Step 2.** Perturb each point in S_1 in the subspace $\mathbb{R}^{d-1}$ by a sufficiently small $\delta_1 > 0$ such that the resulting point set S_2 is in general position in $\mathbb{R}^{d-1}$. **Step 3.** Perturb the x_d coordinate of each point in S_2 by a sufficiently small $\delta_2 > 0$ such that the resulting point set S_3 is in general position in $\mathbb{R}^d$. Our point set is $S = S_3 \cup \{a, b\}$.

The points in $S_1 \subset \mathbb{R}^{d-1}$ form $\Theta(n^{d-1})$ colorful $(d-2)$-simplices in $\mathbb{R}^{d-1}$. Let $\mathcal{M}_2$ denote the $(d-1)$-tuples of pairwise disjoint colorful $(d-2)$-simplices of S_2 whose convex hulls have a non-empty intersection. Since S_2 is in general position in $\mathbb{R}^{d-1}$, the intersection of any $d-1$ pairwise disjoint $(d-2)$-simplices spanned by S_2 is either empty or a single point, and the intersection points ate distinct. We have $|\mathcal{M}_2| = \Theta((n^{d-1})^{d-1}) = \Theta(n^{(d-1)^2})$ by the second selection theorem [6, 24] (which, in turn, follows from the colorful Tverberg theorem [32]). Let $\mathcal{M}_3$ denote the corresponding $(d-1)$-tuples of $(d-2)$-simplices of S_3. Denote by v_m the intersection point of a $(d-1)$-tuple $m \in \mathcal{M}_2$, and let $V_{\mathcal{M}} = \{v_m : m \in \mathcal{M}\}$. Let $\varepsilon > 0$ be the minimum distance between the points in $V_{\mathcal{M}}$ in $\mathbb{R}^{d-1}$. For every point v_m, $m \in \mathcal{M}_2$, let $B_m \subset \mathbb{R}^d$ denote the d-dimensional ball of radius $\varepsilon/3$ centered at v_m. The balls B_m, $m \in \mathcal{M}_\in$, are pairwise disjoint.

After the second perturbation, S_3 is in general position in $\mathbb{R}^d$, and so no point is contained in $d-1$ distinct $(d-2)$-simplices. However, each colorful $(d-2)$-simplex in S_3 is extremal, with two almost vertical shells incident to points a and b, respectively. If $\delta_2 > 0$ is sufficiently small, the $d-1$ pairs of shells whose axes are the $d-1$ distinct $(d-2)$-simplices in $m \in \mathcal{M}_3$ intersects in a unique ball B_m. If N is sufficiently large, then the lower-most intersection point of these $d-1$ pairs of shells is a vertex of U_S. Since the balls B_m are pairwise disjoint, U_S has at least $\Theta(n^{(d-1)^2})$ vertices. $\square$

3.2 Upper bounds for the maximum combinatorial complexity

Boissonnat et al. [12] showed that the colorful union of a set S of n colored points in the plane is a simple polygon. They also showed that the polygon U_S has no two consecutive reflex vertices, and the convex vertices are points in S. It follows that U_S has at most $2n$

vertices in $\mathbb{R}^2$. By Lemma 2, the boundary of U_S is contained in $O(n^{d-1})$ shells for every $d \geq 2$. Aronov and Sharir [8, 31] proved that the combinatorial complexity of a single cell in the arrangement of m distinct $(d-1)$-simplices in $\mathbb{R}^d$ is $O(m^{d-1} \log m)$. The colorful union U_S has the same combinatorial complexity as the *outer face* in the arrangement of its shells, which is $O((n^{d-1})^{d-1} \log(n^{d-1})) = O(n^{(d-1)^2} \log n)$. We have shown the following.

Theorem 3. *For every $d \geq 2$, the combinatorial complexity of the union of colorful tetrahedra spanned by a set of n colored points in $\mathbb{R}^d$ is $O(n^{(d-1)^2} \log n)$.*

In Section 6, we improve this general bound for $d = 3$ from $O(n^4 \log n)$ to $O(n^4 \alpha(n))$, where $\alpha(\cdot)$ is the inverse of the Ackermann function.

4 Efficient data structures for point inclusion queries

Using Theorem 1, we can build a data structure for point inclusion queries in the colorful union U_S. Let S be a colored set of n points in $\mathbb{R}^d$, and let $G = \{g_1, \ldots, g_{d+1}\} \subset S$ be an arbitrary colorful simplex. For $i = 1, \ldots, d+1$, let $S_i = \{s \in S : \text{color}(s) \neq i\} \cup \{g_i\}$, and let U_i be the colorful union of S_i. Note that g_i is the only point of color i in S_i. By Theorem 1, $U_S = \bigcup_{i=1}^{d+1} U_i$. That is, for a query point $q \in \mathbb{R}^d$, we have $q \in U_S$ if and only if $q \in U_i$ for some $i = 1, \ldots, d+1$. It is easy to test whether $q \in U_i$ with a ray shooting query.

Lemma 4. *For every $i = 1, 2 \ldots, d+1$, we have $q \in U_i$ if and only if $q = g_i$ or the ray emitted from g_i in the direction of q passes through q* before *reaching a shell of $S_i \cup \{g_i\}$.*

Proof. Assume that $q \neq g_i$. Recall that U_i is star-shaped with star-center g_i. If $q \in U_i$, then the ray $\overrightarrow{g_i q}$ passes through q before reaching the boundary of U_i, and the boundary of U_i is a shell of S_i by Lemma 2. Conversely, suppose that the ray $\overrightarrow{g_i q}$ hits a shell convΔ of S_i. Since the ray starts from g_i, Δ is spanned by $S_i \setminus \{g_i\}$. If the ray passes through q before hitting convΔ, then q is contained in the colorful simplex $\Delta \cup \{g_i\}$. □

Let T_i denote the set of shells of S_i. Since $|S_i| = O(n)$, we have $|T_i| = O(n^{d-1})$. A ray shooting query for T_i would report the *first* shell hit by a ray, not the *last* one. Nevertheless, the problem can be reduced to vertical ray shooting. If g_i is on the convex hull of S, then a projective transformation can map g_i to infinity such that rays emitted from g_i become vertical rays directed downwards. The *last* shell hit by a vertical downward ray passing through q is the *first* shell hit by a vertical upward ray starting from infinity. If g_i is not on the convex hull, we can partition $\mathbb{R}^d$ into two halfspaces by a hyperplane containing g_i, and build a ray shooting data structure for the set of shells in T_i clipped in each halfspace.

The currently available data structures for ray shooting queries among a set of $(d-1)$-simplices in $\mathbb{R}^d$ are based on range spaces of finite VC-dimension, multi-level partition trees, and Megiddo's parametric search technique [2, 3, 14, 30]. Given s not necessarily disjoint $(d-1)$-dimensional simplices in $\mathbb{R}^d$, one can construct a data structure for vertical ray shooting with $O(s^{n-1}\alpha(s))$ space and $O(\log s)$ query time; if the space is reduced to m, $\Omega(s) \leq m \leq O(s^{d+1})$, then the query time increases to $O(s^{1+\varepsilon}/m^{1/(d-1)})$ for any constant $\varepsilon > 0$. Since $O(d)$ vertical ray shooting data structures, each for $s = O(n^{d-1})$ shells in $\mathbb{R}^d$, can jointly answer a containment query for U_S, we have the following result.

Theorem 4. *For a set of n colored points in $\mathbb{R}^d$, there are data structures for answering point inclusion queries for the colorful union U_S. There is a data structure with $O(n^{(d-1)^2}\alpha(n))$ space and $O(\log n)$ query time. If the available space is reduced to m, $\Omega(n^{d-1}) \leq m \leq O(n^{(d-1)^2})$, then the query time increases to $O(n^{d-1+\varepsilon}/m^{1/(d-1)})$ for any constant $\varepsilon > 0$.*

Remark. Edelsbrunner [18] (see also [28, 31]) proved that the combinatorial complexity of the *upper envelope* of m possibly intersecting $(d-1)$-simplices in $\mathbb{R}^d$ is $\Theta(m^{d-1}\alpha(m))$. Since $U_S = \cup_{i=1}^{d+1} U_i$, then U_S is the union of $d+1$ star-shaped polyhedra, each of which has $O(n^{(d-1)^2}\alpha(n))$ combinatorial complexity. This, however, does not imply the same upper bound for the complexity of U_S.

5 Colored point configurations with undesirable features

A point configuration whose colorful union is not star-shaped.

In the plane, colorful union U_S is always star shaped by Lemma 3. The star center is typically not a point in S, and U_S is not necessarily the union of two star shaped polygons with star centers at S. However, U_S is the union of three star shaped polygons with star centers in S by Theorem 1.

We show that in 3-space there are colored point sets whose colorful union is not star-shaped. One example is shown in Table 1 and Fig. 3(a).

vertex	x	y	z	color
p_1	-0.9	-1	1	1
p_2	-0.9	1	1	1
p_3	0.9	-1	1	2
p_4	0.9	1	1	2

vertex	x	y	z	color
p_5	-1	-0.9	-1	3
p_6	1	-0.9	-1	3
p_7	-1	0.9	-1	4
p_8	1	0.9	-1	4

Table 1: The coordinates of a point set whose colorful union is not star shaped.

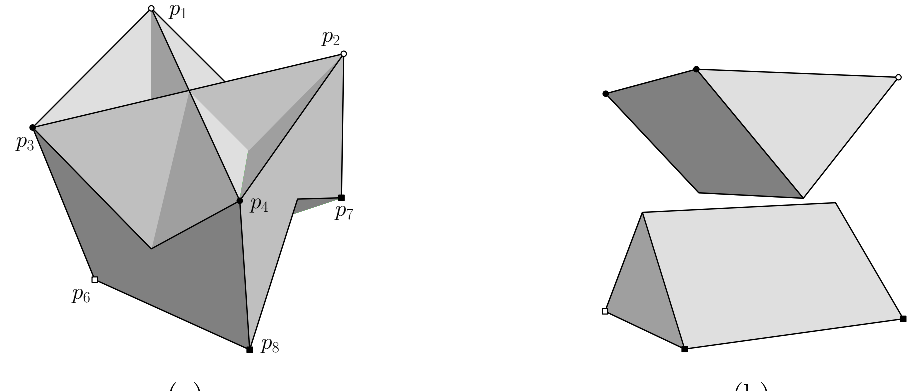

(a) (b)

Figure 3: (a) The colorful union of the point set S. (b) A superset of the visibility regions.

In the point set in Fig. 3(a), axes p_1p_3 and p_2p_4 have color $\{1,2\}$. Every point in U_S that sees the relative interior of p_1p_3 must lie in wedge $W(p_1,p_3)$. Similarly, the relative interior of p_2p_4 can only be seen from wedge $W(p_2,p_4)$. Thus any point in U_S that sees both axes must lie in $W(p_2,p_4) \cap W(p_1,p_3)$. The intersection of these two wedges, however, is strictly below the xy-plane. A superset of the intersection of the wedge with U_S is depicted in Fig. 3(b). Similarly, the relative interior of p_5p_7 and p_6p_8 is visible from the intersection of two wedges, which is a region strictly above the xy-plane (Fig. 3(b)). It follows that no point can see all these four edges of U_S.

A face bounded by edges with reflex dihedral angles.

In the plane, every face of U_S is incident to a point in S (c.f. [12]). Since U_S is simply connected, this property implies that the combinatorial complexity of U_S is $O(n)$ in $\mathbb{R}^2$. In 3-space, however, there are colored point sets S such that a face of U_S is not incident to any colorful edge. Fig. 4 shows an explicit example with 8 colored points in $\mathbb{R}^3$. The white face in Fig. 4(right) is bounded by four edges with reflex dihedral angles.

vertex	x	y	z	color
p_1	7	-10	-3	1
p_2	8	-9	0	1
p_3	0	-14	10	2
p_4	9	-15	6	2
p_5	4	-13	0	2
p_6	2	-10	-1	3
p_7	4	-6	0	3
p_8	5	-6	-4	4

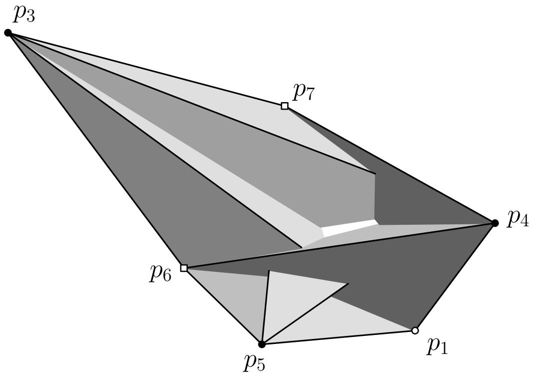

Figure 4: An example where a face is not incident to any axis.

A chain of reflex vertices.

In the plane, U_S has no two adjacent reflex vertices [12]. This property no longer holds in 3-space. Fig. 5 depicts a family of point sets in $\mathbb{R}^3$ where the boundary of a face may contain an arbitrary long chain of reflex vertices. Points of colors 1 and 2 are arranged along two parallel lines in the xy-plane, such that the complete bipartite graph between the first two color classes forms a convex chain of length $\Omega(n)$ (Fig. 5, left). Two points of color 3 and 4 are placed below the xy-plane on opposite sides of all vertical planes through edges of color $\{1,2\}$. A small perturbation can make any 4 points affinely independent while preserving the reflex chain of length $\Omega(n)$.

6 Colorful union in 3-space

In this section we prove an upper bound for the combinatorial complexity of the colorful union in 3-space. Recall that Theorem 5 gives an upper bound of $O(n^4 \log n)$ in 3-space. We can

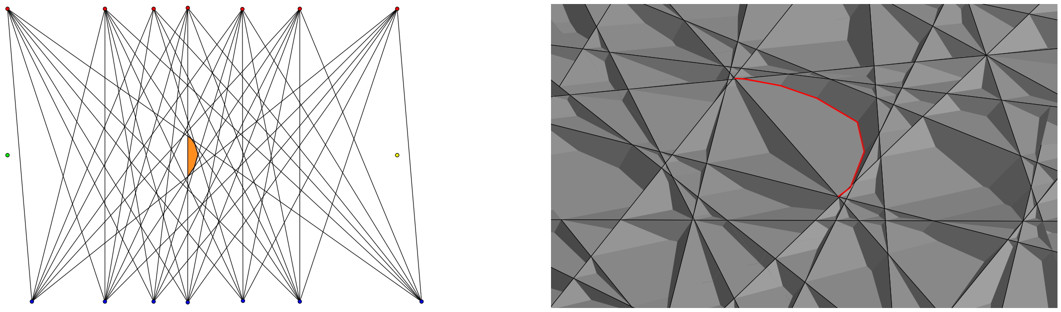

Figure 5: A chain of reflex vertices of U_S along the boundary of a single face.

replace the logarithmic factor with the extremely slowly growing inverse of the Ackermann function.

Theorem 5. *The combinatorial complexity of the union of colorful tetrahedra spanned by a set of n colored points in $\mathbb{R}^3$ is $O(n^4\alpha(n))$, where $\alpha(\cdot)$ is the inverse of the Ackermann function.*

The proof builds on the following lemma, whose proof is postponed to Subsection 6.2.

Lemma 5. *If S is a set of n colored points in $\mathbb{R}^3$, then the relative interior of every shell of S contains $O(n^2\alpha(n))$ vertices of U_S.*

Proof of Theorem 5. By Lemma 2 every face of the colorful union U_S is contained in a shell. A set of n colored points spans $O(n^2)$ colorful edges. So there are $O(n^2)$ extremal colorful edges, hence $O(n^2)$ shells. It is enough to prove that each shell contains $O(n^2\alpha(n))$ vertices, edges, and faces. By Euler's formula it is enough to count the number of vertices lying on each shell. First, we count the vertices of U_S lying on the boundaries of shells. The boundaries of shell triangles consist of $O(n^2)$ colorful segments. If a vertex of U_S is contained in a colorful segment pq, then it is an endpoint $p, q \in S$ or an intersection point of segment pq with the relative interior of one of the $O(n^2)$ shells. Therefore, each colorful segment contains $O(n^2)$ vertices of U_S, and U_S has $O(n^4)$ vertices on the boundaries of shells. There are $O(n^2\alpha(n))$ vertices of U_S in the relative interior of each shell triangle by Lemma 5. It follows that each shell contains $O(n^2\alpha(n))$ vertices of U_S, as required. □

6.1 Auxiliary results in the plane

Before the proof of Lemma 5, we present auxiliary results in the plane. We assume throughout this section that S is a finite set of colored points in the plane such that no three points are collinear and $|\text{color}(S)| \geq 3$. For a closed polygonal curve γ in the plane, the *faces* of γ are the connected components of $\mathbb{R}^2 \setminus \gamma$. Clearly, one face is unbounded and all other faces are bounded.

Lemma 6. *Let S be a colored point set in $\mathbb{R}^2$.*

(i) *If γ is a simple closed curve contained in U_S, then each bounded face of γ is contained in U_S.*

(ii) *If γ is a closed polygonal chain whose vertices are points in S and every edge is colorful, then every bounded face of γ lies in the interior of U_S.*

(iii) *If $s_1, s_2 \in S$ and $s_2 \in W(s_1)$, then segment s_1s_2 is contained in U_S.*

Proof. Statement (i) follows since U_S is simply connected [12]. For part (ii), a closed polygonal chain γ can be decomposed into simple (possibly overlapping) curves. Every simple curve lies in U_S and hence all the induced bounded faces are contained in U_S, as required. For part (iii), consider $s_1, s_2 \in S$ with $s_2 \in W(s_1)$. If s_1 and s_2 have different colors, then s_1s_2 is colorful and is clearly contained in U_S. Assume that $\mathrm{color}(s_1) = \mathrm{color}(s_2)$, and let $s', s'' \in S$ be the two points lying on the shells with axis s_1. The colors of s' and s'' each are different from $\mathrm{color}(s_1)$ and hence from $\mathrm{color}(s_2)$. It follows that the closed polygonal chain (s_1, s', s_2, s'') has colorful edges and $s_1s_2 \subset U_S$ by part (ii) of the lemma. □

We say that a line segment e is *fully visible* from a point $p \in U_S$ in $\mathbb{R}^2$, if $\mathrm{conv}(e \cup \{p\}) \subseteq U_S$. We show that if S is a colored point set in the plane, then almost all edges of U_S are fully visible from at least one of two given points in S.

Lemma 7. *Let S be a colored point set in $\mathbb{R}^2$. Let $s_1, s_2 \in S$ with* $\mathrm{color}(s_1) = 1$ *and* $\mathrm{color}(s_2) = 2$.

(1) *If an edge e of U_S is fully visible from neither s_1 nor s_2, then e is an edge of* $\mathrm{conv}S$, *hence the endpoints of e are some extremal points $p_1, p_2 \in S$. Furthermore,* $\mathrm{color}(p_1, p_2) = \{1, 2\}$, *the point set $\{p_1, p_2, s_1, s_2\}$ is in convex position, and the wedges $W(p_1)$ and $W(p_2)$ each contain exactly one of s_1 and s_2.*

(2) *All but at most two edges of U_S are fully visible from s_1 or s_2.*

(3) *If an edge e of U_S is fully visible from s_1 or s_2, then e is incident to some extremal vertex $p_1 \in S$ such that the convex hull of e and the vertices $\{s_1, s_2\} \cap W(p_1)$ is contained in U_S.*

Proof. (1): Let e be an edge of U_S that is fully visible from neither s_1 nor s_2 (see Fig. 6(a)). By Lemma 2, e lies on a shell segment p_1p_2, for some $p_1, p_2 \in S$. We shall prove that $e = p_1p_2$. At this point, however, we only assume $e \subseteq p_1p_2$. It is clear that $\{p_1, p_1\} \cap \{s_1, s_2\} = \emptyset$, otherwise the entire segment p_1p_2 would be fully visible from s_1 or s_2. We have $\mathrm{color}(\{p_1, p_2\}) = \{1, 2\}$, otherwise the triangle $p_1p_2s_1$ or $p_1p_2s_2$ would be colorful, and segment p_1p_2 would be fully visible from s_1 or s_2.

If $\mathrm{conv}\{p_1, p_2, s_1, s_2\} = (p_1, p_2, s_1, s_2)$ (see Fig. 6(b)), then $\mathrm{conv}\{p_1, p_2, s_1, s_2\} \subset U_S$ by Lemma 6(ii), and so p_1p_2 is fully visible from both s_1 and s_2. Similarly, if $\mathrm{conv}\{p_1, p_2, s_1, s_2\}$ is a triangle (Fig. 6(c)), then the non-convex quadrilateral (p_1, p_2, s_1, s_2) or (p_2, p_1, s_2, s_1) is a simple polygon, which contains triangle $p_1p_2s_1$ or $p_1p_2s_2$. In both cases, the entire quadrilateral is contained in U_S by Lemma 6(ii), and so p_1p_2 is fully visible from s_1 or s_2. Therefore, $\mathrm{conv}\{p_1, p_2, s_1, s_2\}$ has to be a quadrilateral with some monochromatic edges. Assume, without loss of generality, that $\mathrm{conv}\{p_1, p_2, s_1, s_2\} = (p_1, s_1, s_2, p_2)$ in counterclockwise order.

Let H be the supporting line of p_1p_2, and assume that $s_1, s_2 \in H^-$. We claim that p_1p_2 is an edge of $\mathrm{conv}S$. Suppose, to the contrary, that p_1p_2 is not an edge of $\mathrm{conv}S$. Then there is

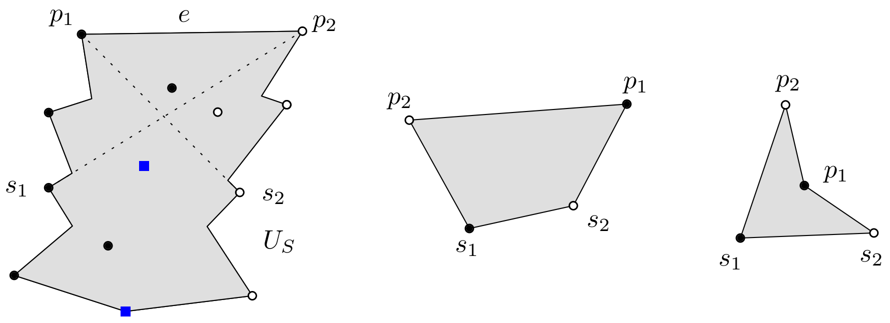

(a) (b) (c)

Figure 6: The colorful union of the planar point configurations in the proof of Lemma 7(1).

a point $r \in S \cap H^+$. We have $\text{color}(r) \in \{1, 2\}$, otherwise the relative interior of p_1p_2 would lie in the interior of U_S, and e would not be an edge of U_S. Assume that $\text{color}(r) = 1$ (the case that $\text{color}(r) = 2$ is analogous). Let r' be the intersection point of segment rs_2 and line H. We distinguish three cases based on the order of point p_1, p_2, r' along H. Case 1: the order of these points is (r', p_1, p_2) (see Fig. 7(a)). Then p_1p_2 is contained in a face of the non-simple polygon (r, p_2, s_1, s_2), and by Lemma 6(ii) p_1p_2 lies in the interior of U_S. Case 2: the order is (p_1, r', p_2) (see Fig. 7(b)). Then $p_1r' \subset p_1p_2$ is fully visible from s_2 by applying Lemma 6(ii) for (p_1, p_2, s_1, s_2); and the relative interior of $r'p_2 \subset p_1p_2$ lies in the interior of U_S by applying Lemma 6(ii) for (r, p_2, s_1, s_2). Case 3: the order is (p_1, p_2, r'). Then (p_1, s_2, r, p_2) is a simple polygon, whose interior lies in U_S by Lemma 6(ii), and so p_1p_2 is fully visible from s_2. All three cases lead to a contradiction, so we conclude that $S \cap H^+ = \emptyset$ and p_1p_2 is an edge of $\text{conv}S$. It follows p_1 and p_2 are extremal points in S, and $e = p_1p_2$.

Since the segments p_1s_2 and p_2s_1 are colorful, we have $s_2 \in W(p_2)$ and $s_1 \in W(p_2)$. However, $s_1 \notin W(p_2)$, otherwise segment p_2s_1 would be contained in U_S (by Lemma 6(iii)), and so triangle $p_1p_2s_1 \subset U_S$ (by Lemma 6(i)), hence $e = p_1p_2$ would be fully visible from s_1. One can show analogously that $s_2 \notin W(p_1)$, as required.

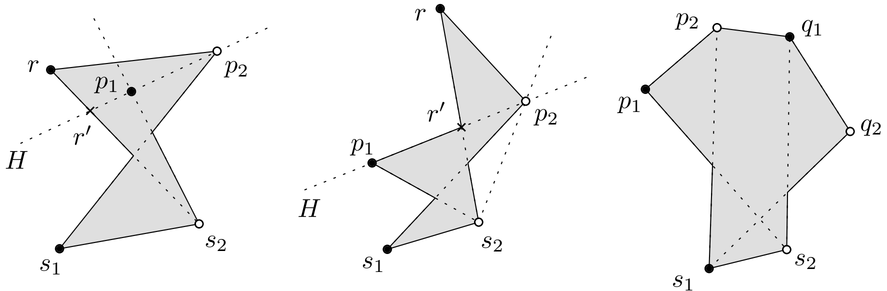

(a) (b) (c)

Figure 7: The colorful union of the planar point configurations in the proof of Lemma 7(1) and (2).

(2): Suppose that U_S has two edges, p_1p_2 and q_1q_2, that are fully visible from neither s_1 nor s_2. We have shown in Lemma 7(1) that $p_1, p_2, q_1, q_2 \in S$ with colors in $\{1, 2\}$, and both (p_1, p_2, s_2, s_1) and (q_1, q_2, s_1, s_1) are convex quadrilaterals. We may assume without loss of generality that color(p_1) = color$(q_1) = 1$ and color(p_2) = color$(q_2) = 2$. Hence p_1p_2 and q_1q_2 each lie in a halfplane bounded by the supporting line of s_1s_2. We now show that they must lie on opposite sides of s_1s_2. Suppose, to the contrary, that p_1p_2 and q_1q_2 are on the same side of s_1s_2. From Lemma 7(1), both p_1p_2 and q_1q_2 are edges of conv$S$. We may assume that conv$\{s_1, s_2, p_1, p_2, q_1, q_2\} = (s_1, s_2, p_2, p_1, q_2, q_1)$ in counterclockwise order (Fig. 7(c)). The vertices of the convex quadrilaterals (s_1, p_2, q_1, s_2) and (s_2, p_1, p_2, q_1) are colored alternately 1 and 2, and so they are contained in U_S. Therefore p_1p_2 is fully visible from s_1, which is a contradiction. We conclude that p_1p_2 and q_1q_2 lie on opposite sides of s_1s_2. Hence U_S has at most two edges (at most one each side of s_1s_2) that are visible from neither s_1 nor s_2.

(3): Assume without loss of generality that edge e is fully visible from s_1. By Lemma 2, e lies on a shell segment p_1p_2, for some $p_1, p_2 \in S$. Recall that U_S has no two consecutive reflex vertices and every convex vertex is a point in S [12]. Hence we may assume that one endpoint of e is $p_1 \in S$. Since p_1p_2 is a shell, then color$(p_1) \neq$ color(p_2). We distinguish two cases.

Case 1: $e = p_1p_2$**.** By symmetry, we may assume that color$(p_1) \neq 2$ and color$(p_2) \neq 1$. This implies that $s_1 \in W(p_2)$ and $s_2 \in W(p_1)$. If $W(p_1)$ or $W(p_2)$ contains exactly one of s_1 and s_2, then statement (3) follows immediately. So assume that both $W(p_1)$ and $W(p_2)$ contain $\{s_1, s_2\}$. By Lemma 6(iii) and by construction every edge of the complete graph spanned by p_1, p_2, s_1, s_2 is contained in U_S, which in turn shows (by Lemma 6(ii)) that conv$(e \cup \{s_1, s_2\})$ lies in U_S.

Case 2: p_2 **is not in** e**.** Then $e = p_1r$, where $r \in$ conv p_1p_2 is a reflex vertex of U_S. If $s_1 \notin W(p_1)$ the lemma follows trivially, so assume that $s_1 \in W(p_1)$. The reflex endpoint r is followed by an extremal point, say $q \in S$, on the boundary of U_S. Since $q \notin W(p_1)$ we have color(q) = color(p_1). The segments s_1p_1, p_1r, and rs_1 are contained in U_S since e is fully visible from s_1. Segment p_1s_1 is also contained in U_S by Lemma 6(iii), and segment s_1s_2 is contained in U_S because it is colorful. Furthermore, p_2s_2 is colorful, or else qs_2 is colorful (since color(q) = color$(p_1) \neq$ color(p_2)). Due to Lemma 6(ii), in all situations the convex hull of e and $\{s_1, s_2\}$ lies in a polygon whose edges are contained in U_S, and hence conv$\{s_1, s_2, e\} \subseteq U_S$ (see Figure 8). Notice that a colorful edge rs_2 cannot cut through e. □

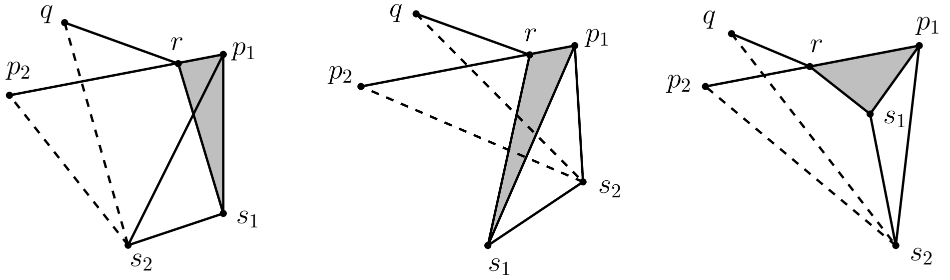

Figure 8: Proof of Lemma 7(3). The solid edges are contained in U_S as well as one of the dashed edges of each scenario. The figure depicts the three interesting configurations of s_1, s_2 with respect to e.

6.2 Proof of Lemma 5

We are given a set S of n colored points in $\mathbb{R}^3$. Let $\Delta = v_1v_2v_3$ be a shell triangle with axis v_1v_2. We show that the relative interior of convΔ contains $O(n^2\alpha(n))$ vertices of U_S. Assume without loss of generality that color$(v_i) = i$ for $i = 1, 2, 3$. Let H be the plane spanned by Δ. Let $U_\Delta = \text{cl}(\text{conv}\Delta \cap \text{int}(U_S))$ be the restriction of U_S to the triangle convΔ. It is enough to show that U_Δ has $O(n^2\alpha(n))$ vertices in the interior of convΔ.

By Lemma 2, every edge of U_Δ lies in the intersection of convΔ and another shell triangle. Let T denote the set of all shell triangles t such that $t \neq \Delta$ and conv$\Delta \cap$ conv t contains at least one edge of U_Δ lying in the relative interior of convΔ. Since there are $O(n^2)$ shells, we have $|T| = O(n^2)$. For every $t \in T$, the line segment conv$\Delta \cap$ conv t lies in U_Δ, and it may contain several collinear edges of U_Δ. Let trace(t) be the convex hull of all edges of U_Δ along conv$\Delta \cap$ conv t. We say that a trace(t) is *k-visible*, for $k = 0, 1, 2$, if U_Δ contains the convex hull of trace(t) and k vertices of Δ, and k is the maximum such integer. In particular, a k-visible trace is fully visible from k vertices of Δ, where visibility is understood with respect to the polygon U_Δ. If trace(t), $t \in T$ is 0-visible but it can be decomposed into two line segments which are each fully visible from a vertex of Δ, then fix one such decomposition, and call the two segments the *half-traces* of trace(t). Every half-trace is 1-visible.

Outline of the proof of Lemma 5.

Every vertex of U_Δ lying in the interior of convΔ is at the intersection of two traces. We proceed with a systematic classification of different types of traces, and devise upper bounds for the number of intersection points that lie on the boundary of U_Δ. Our key tool will be a reduction to the lower envelope of line segments in the plane. It is well known that the lower envelope of k segments in the plane has $O(k\alpha(k))$ vertices, and this bound is the best possible [4]. The reduction will be fairly easy for the intersections of traces that are fully visible from the same vertex of Δ (Lemma 8), and for intersections where at least one of two intersecting traces is 2-visible (Lemma 14). The greatest challenge is to count the vertices of U_Δ at intersections of 0- or 1-visible traces which are not visible from the same vertex of Δ. To estimate the crossings among 0- and 1-visible traces, we define for every such trace a wedge $w(t)$ such that trace(t) lies on the boundary of $w(t)$. We prove a serious restriction on the intersection pattern of these wedges (Lemma 16): if a vertex of U_Δ lies at the intersection of two such wedges but no vertex of Δ is contained in both wedges, then one of the wedges contains the apex of the other. This restriction would immediately imply a linear bound on the intersection points in terms of the number of wedges if $w(t) \cap \text{conv}\Delta \subseteq U_\Delta$ for every such trace(t). However, $w(t) \cap \text{conv}\Delta$ is not always contained in U_Δ, and we need proceed with care. We continue with the detail.

Traces fully visible from a single vertex of Δ.

The intersection of two traces is either the endpoint of one of the traces or the crossing point of the two traces (i.e., it lies in the relative interior of both traces). There are $O(n^2)$ traces, and so U_Δ has $O(n^2)$ vertices at the endpoints of traces. It remains to show that U_Δ has $O(n^2\alpha(n))$ vertices at crossings of traces. To begin with, we show that U_Δ has $O(n^2\alpha(n))$ vertices at the crossings of traces which are fully visible from the *same* vertex of Δ. For $i = 1, 2, 3$, let L_i be the set of all traces and half-traces fully visible from v_i. The convex hull of each segment in L_i and v_i is a triangle lying in U_Δ; let D_i denote the union of all of these triangles. It is clear that $D_i \subseteq U_\Delta$.

Lemma 8. *For $i = 1, 2, 3$, the set D_i has $O(n^2\alpha(n))$ vertices.*

Proof. By definition, D_i is the union of $O(n^2)$ triangles that lie in convΔ and share vertex v_i. Apply a projective transformation that maps v_i to infinite, and maps the incident edges of Δ to vertical rays pointing up. Every triangle incident to v_i and lying in convΔ is mapped to a region vertically above a line segment. The number of vertices of D_i is the combinatorial complexity of the lower envelope of these segments. It is known that the lower envelope of $O(n^2)$ line segments has $O(n^2\alpha(n^2)) = O(n^2\alpha(n))$ vertices, which is the maximum length of a Davenport-Schinzel sequence of order 3 over $O(n^2)$ symbols [4]. $\square$

To prove Lemma 5, it remains to bound the number of crossings of pairs of traces which are (1) either fully visible from different vertices of Δ, or (2) one of them is not fully visible from any vertex of Δ. In order to better understand the crossings between traces, we distinguish three types of traces and study their properties.

Three types of traces.

Recall that H denotes the plane spanned by Δ. As usual, H defines two open halfspaces, H^+ and H^-. Let H^- be the open halfspace containing all points of colors other than $\{1, 2, 3\}$. Since v_1v_2 is an axis of Δ, we have color$(S \cap H^+) \subseteq \{1, 2\}$. A triangle $t \in T$ is of

- type A if exactly one vertex of t is in $S \cap H^+$,
- type B if exactly two vertices of t are in $S \cap H^+$ and an axis of t crosses H,
- type C if exactly two vertices t are in $S \cap H^+$ and the axis of t is in H^+.

Denote by T_A, T_B, and T_C, respectively, the set of shell triangles of type A, B, and C. We have $T = T_A \cup T_B \cup T_C$. We say that trace$(t)$ is of *type A* (resp., *type B* or *C*) if $t \in T_A$ (resp., T_B or T_C).

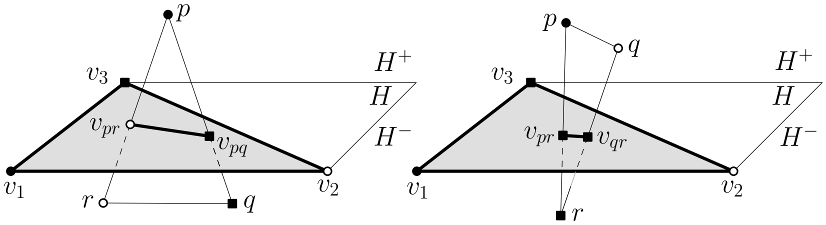

Figure 9: Triangle $\Delta = v_1v_2v_3$ with a triangle pqr of type A (left), and of type B or C (right).

Traces of type A.

For every $p \in S \cap H^+$ and $q \in S \cap \mathrm{cl}(H^-)$, let $v_{pq} = pq \cap H$. That is, v_{pq} is the intersection point of segment pq and plane H. For every $p \in S \cap H^+$, let U_p denote the union of all colorful tetrahedra spanned by p and three points in $S \cap \mathrm{cl}(H^-)$. It is clear that $U_p \subseteq U_S$. We also define the planar point set $S(p) = \{v_{pq} : q \in S \cap \mathrm{cl}(H^-), \mathrm{color}(q) \neq \mathrm{color}(p)\}$ and color each point v_{pq} with color(q). In particular, the three vertices of Δ are in $S(p)$ with their original colors. Observe that $H \cap U_p = U_{S(p)}$.

Lemma 9. *If $t \in T_A$ and p is the vertex of t in $S \cap H^+$, then* trace(t) *is contained in a single edge of $U_{S(p)}$.*

Proof. Let $t = prq$. Since $t \in T_A$, we have $q, r \in S \cap \mathrm{cl}(H^-)$ and $t \cap H = v_{pq}v_{pr}$. Segment $v_{pq}v_{pr}$ is colorful in $S(p)$, and so it is contained in $U_{S(p)}$. The edges of U_Δ along trace(t) must be on the boundary of $U_{S(p)}$ since $U_{S(p)} \cap \mathrm{conv}\Delta \subseteq U_\Delta$. By the result of Boissonnat et al. [12], no two edges of $U_{S(p)}$ are collinear. Therefore all of trace(t) is contained in a single edge of $U_{S(p)}$. □

Consider a colorful segment pq with $p \in S \cap H^+$ and $q \in S \cap \mathrm{cl}(H^-)$. Recall that if pq is extremal (with respect to S), then $W(pq)$ denotes the minimal 3-dimensional wedge with axis pq that contains all points in S whose colors are not in color(pq). Let $w_{pq} = W(pq) \cap H$, and observe that it is a planar wedge with apex v_{pq}. If pq is extremal with respect to $\{p\} \cup (S \cap \mathrm{cl}(H^-)) \subseteq S$, then let W^-_{pq} denote the minimal 3-dimensional wedge with axis pq that contains all points of $S \cap \mathrm{cl}(H^-)$ whose colors are not in color(pq). Let $w^-_{pq} = W^-_{pq}$, and observe that w^-_{pq} is also a wedge with apex v_{pq}. It is clear that $w^-_{pq} \subseteq w_{pq}$ if pq is extremal, however the two wedges are not necessary the same.

We are now ready to classify the traces of type A.

Lemma 10. *Every* trace(t) *of type A satisfies one of the following three conditions:*

- *it is 2-visible;*
- *it is 1-visible and it lies on the boundary of a wedge w^-_{pq}, where pq is an edge of t, and w^-_{pq} contains exactly one vertex of Δ;*
- *it is 0-visible and lies on the boundary of two wedges, w^-_{pq} and w^-_{pr}, where both pq and pr are edges of t, furthermore, the wedges w^-_{pq} and w^-_{pr} each contain exactly one vertex of Δ, which are two distinct vertices of Δ.*

There are $O(n)$ 0-visible traces of type A.

Proof. Let $t = prq$ be a shell triangle of type A with $p \in \cap H^+$ and $q, r \in S \cap \mathrm{cl}(H^-)$. Assume without loss of generality that color$(p) = 1$. By Lemma 9, trace(t) lies on a single edge of $U_{S(p)}$, which we denote by e. We have $v_2, v_3 \in S(p)$, since $v_2, v_3 \in S \cap \mathrm{cl}(H^-)$. Invoke Lemma 7 for the point set $S(p)$ with $s_1 = v_2$ and $s_2 = v_3$.

First assume that e is fully visible (with respect to the polygon $U_{S(p)}$) to neither v_4 nor v_3. By Lemma 7(1), both endpoints of e are extremal points in $S(p)$. Since trace$(t) \subseteq v_{pq}v_{pr}$, the endpoints of e are v_{pq} and v_{pq}. Therefore, trace(t) lies on the boundary of the wedges w^-_{pq} and w^-_{pr}. If trace$(t) \subseteq e$ is 2-visible, then there is nothing left to prove; and if trace(t) is 1- or 0-visible, then the claim follows from Lemma 7(1).

Now assume that e is fully visible from v_2 or v_3 (with respect to the polygon $U_{S(p)}$). This implies that trace$(t) \subset e$ is not 0-visible (with respect to the polygon U_Δ). If trace(t) is 2-visible then there is nothing left to prove. So we may assume that e is 1-visible from v_2 or v_3. By Lemma 7(3), e is incident to some extremal vertex $v \in S(p)$ such that the convex hull of e and vertices $\{v_2, v_3\} \cap W(v)$ is contained in $U_{S(p)}$. Since trace$(t) \subseteq v_{pq}v_{pr}$, we may assume without loss of generality that $v = v_{pq}$, hence $W(v) = w^-_{pq}$. On one hand, w^-_{pq} cannot contain both v_2 and v_3, otherwise trace(t) would be 2-visible. On the other hand w^-_{pq} contains at least one of v_2 and v_3, since the colors of p, v_2, and v_3 are distinct. Finally w^-_{pq} does not contain v_1, otherwise it would be 2-visible by Lemma 6. Hence w^-_{pq} contains exactly one vertex of Δ.

By Lemma 7(2), $U_{S(p)}$ has at most two edges that are not fully visible from any vertex of Δ, and at most one such edge intersects the interior of Δ. Therefore, there is at most one 0-visible trace of type A on the boundary of each $U_{S(p)}$, $p \in S \cap H^+$. This proves that there are $O(n)$ 0-visible traces of type A. □

Traces of types B and C.

We show in Lemma 12 below that every trace of type B or C is 1- or 2-visible. Therefore, all 0-visible traces are of type A.

Lemma 11. *Let pqr be a colorful triangle such that $p, q \in S \cap H^+$ and $r \in S \cap H^-$, with* color$(p) = 1$ *and* color$(q) = 2$*. If* color$(r) \geq 4$*, then the interior of $v_1v_3v_{qr} \cup v_3v_{qr}v_{pr} \cup v_2v_3v_{pr}$ is contained in the interior of U_S. If* color$(r) = 3$*, then every bounded face of the closed polygonal chain $(v_1, v_3, v_2, v_{pr}, v_{qr})$ in Δ is contained in the interior of U_S.*

Proof. If color$(r) \geq 4$, then the colorful simplices conv$\{v_1, q, v_3, r\}$, conv$\{p, q, r, v_3\}$, and conv$\{p, v_2, v_3, r\}$ contain triangles $v_1v_3v_{qr}$, $v_3v_{pr}v_{qr}$, and $v_2v_3v_{pr}$, respectively.

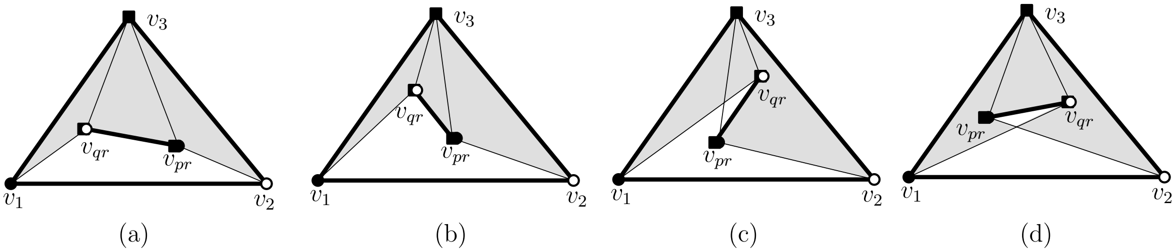

Figure 10: Triangle $\Delta = v_1v_2v_3$ and the closed polygonal chain $(v_2, v_3, v_1, v_{qr}, v_{pr})$ where pqr it a triangle of type B or C.

Now suppose that color$(r) = 3$. Refer to Fig. 10. The closed polygonal chain (v_1, v_2, p, q) is alternately colored 1 and 2. Every edge of this chain forms a colorful triangle with each of v_3 and r. These eight triangles of color $\{1, 2, 3\}$ form a combinatorial octahedron. By adding a point of a fourth color to each triangle, we obtain colorful tetrahedra that jointly contain every bounded region determined by the immersion of this octahedron into $\mathbb{R}^3$. Every bounded face of the planar polygonal chain $(v_1, v_3, v_2, v_{pr}, v_{qr})$ is on the boundary of one such region lying in H^+. Since Δ forms a colorful tetrahedron with a point in H^-, the bounded faces of $(v_1, v_3, v_2, v_{pr}, v_{qr})$ in Δ are in the interior of U_S. □

Lemma 12. *Every* trace(t) *of type B is either*

- *2-visible; or*
- *1-visible and lies on the boundary of a wedge w_{pr} such that pr is an axis of t, $w_{pr} \cap$* conv$\Delta \subset U_\Delta$*, and w_{pr} contains exactly one vertex of Δ.*

Proof. Let $t = pqr$ such that $p, q \in H^+$ and $r \in H^-$, with $\mathrm{color}(p) = 1$ and $\mathrm{color}(q) = 2$. Assume without loss of generality that pr is the axis of t. If $\mathrm{color}(r) \geq 4$, then by Lemma 11, we have $v_3 v_{pr} v_{qr} \cup v_2 v_3 v_{pr} \subset U_\Delta$. Since wedge $W(pr)$ contains v_2, v_3 and q, the union of these two triangles is a convex quadrilateral, and so $\mathrm{trace}(t)$ is 2-visible, from v_2 and v_3.

Assume now that $\mathrm{color}(r) = 3$. By Lemma 11, U_Δ contains the bounded faces of the closed polygonal chain $(v_1, v_3, v_2, v_{pr}, v_{qr})$. Since wedge $W(pr)$ contains both v_2 and q, this polygon has a convex interior angle at v_{pr}. If $v_{qr} \notin \angle(\overrightarrow{v_{pr}v_2}, \overrightarrow{v_{pr}v_3})$ (Fig. 10(b)), then $\mathrm{conv}(\mathrm{trace}(t) \cup v_2 v_3) \subset U_\Delta$ and $\mathrm{trace}(t)$ is 2-visible. If $v_{qr} \in \angle(\overrightarrow{v_{pr}v_2}, \overrightarrow{v_{pr}v_3})$ (Fig. 10(c)), then $w_{pr} \cap \mathrm{conv}\Delta \subset U_\Delta$. Since $\mathrm{trace}(t)$ lies on the boundary of wedge $w_{pr} = W(pr) \cap H$, it is 1-visible from v_2. Furthermore, v_2 is the only vertex of Δ in the wedge w_{pr}. □

Lemma 13. *Every trace of type C is either*

- *2-visible; or*
- *1-visible and lies on the boundary of a halfplane* h^- *such that* $h^- \cap \mathrm{conv}\Delta \subset U_\Delta$.

Proof. Let $t = pqr$ such that $p, q \in H^+$ and $r \in H^-$, with $\mathrm{color}(p) = 1$ and $\mathrm{color}(q) = 2$, where pq is the axis of t. Note that v_3 lies in the wedge $W(pq)$. Hence, for every edge of U_Δ along $\mathrm{trace}(t)$, the interior of U_Δ lies on the same side as v_3. Let h be the supporting line of $v_{pr}v_{qr}$, and let h^- be the halfplane containing v_3. If h intersects the segments v_1v_3 and v_2v_3 (Fig. 10(a) and (d)), then by Lemma 11 we have $h^- \cap \Delta \subset U_\Delta$.

Assume now that h intersects segments v_1v_2 and, without loss of generality, v_1v_3 (Fig. 10(b)). By Lemma 11, every bounded face of the closed polygonal chain $(v_1, v_3, v_2, v_{pr}, v_{qr})$ lies in U_Δ. Since $\mathrm{trace}(t) \subseteq v_{pr}v_{qr}$ contains some edges of the boundary of U_Δ, segment $v_{pr}v_{qr}$ is adjacent to an unbounded face of this polygonal chain, and so the segments v_2v_{pr} and v_3v_{qr} cannot cross. Therefore the quadrilateral $(v_2, v_{pr}, v_{qr}, v_3)$ is convex, and it is contained in U_Δ by Lemma 11. This shows that $\mathrm{trace}(t)$ is 2-visible, from v_2 and v_3. □

2-visible traces.

We show that 2-visible traces are incident to at most $O(n^2\alpha(n))$ vertices of U_Δ.

Lemma 14. *Let* $t_1, t_2 \in T$ *such that their traces intersect at point* x *and let* $\mathrm{trace}(t_1)$ *be 2-visible.*

(a) *If* $\mathrm{trace}(t_2)$ *is 2-visible, then both traces are fully visible from a vertex of* Δ.

(b) *If* $\mathrm{trace}(t_2)$ *is 1-visible, then both traces are fully visible from a vertex of* Δ*; or* x *is an endpoint of* $\mathrm{trace}(t_1)$ *or* $\mathrm{trace}(t_2)$.

(c) *If* $\mathrm{trace}(t_2)$ *is 0-visible, then* x *is an endpoint of* $\mathrm{trace}(t_2)$*; or* $\mathrm{trace}(t_2)$ *can be decomposed into two half-traces, each of which is fully visible from some vertices of* Δ.

Proof. We may assume without loss of generality that U_Δ contains a convex quadrilateral $Q = \mathrm{conv}(\mathrm{trace}(t_1) \cup v_2v_3)$.

(a) Since Δ has three vertices, any two 2-visible traces are fully visible from a vertex of Δ.

(b) The statement clearly holds if $\mathrm{trace}(t_2)$ is visible from v_2 or v_3. Assume that $\mathrm{trace}(t_2)$ is fully visible from vertex v_1 only. Recall that a trace is the union of collinear edges of

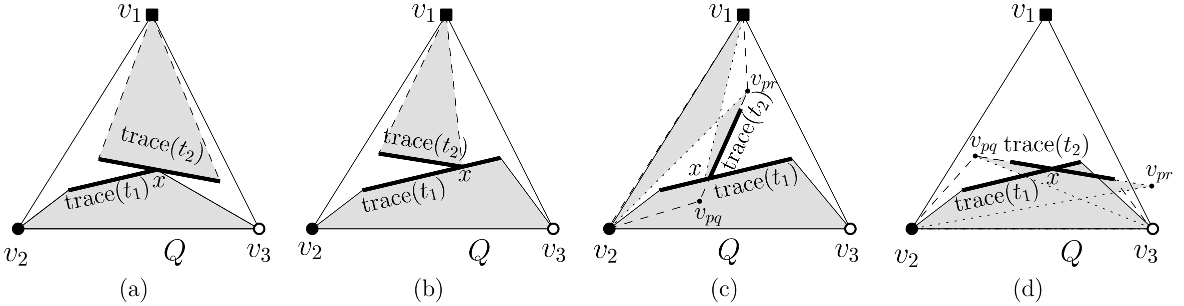

Figure 11: Triangle $\Delta = v_1v_2v_3$ where trace(t_1) is fully visible from both v_2 and v_3.

U_Δ. Since an endpoint of a trace cannot be in the interior of $Q \subseteq U_\Delta$ or the triangle conv(trace(t_2) $\cup\, v_3$), one of trace(t_1) and trace(t_2) has to end at x. See Fig. 11(a)-(b).

(c) By Lemmata 12 and 13, trace(t_2) is of type A. By Lemma 10, trace(t_2) lies on the boundary of two wedges w^-_{pq} and w^-_{pr} that each contain a unique vertex of Δ, say $v_i \in w^-_{pq}$ and $v_j \in w^-_{pr}$, $i \neq j$. Hence the segments v_iv_{pq} and v_jv_{pr} cross, and they are both contained in $U_{S(p)} \subseteq U_\Delta$ by Lemma 6(iii). By Lemma 6(ii), the bounded faces of the nonsimple quadrilateral $v_iv_jv_{pr}v_{pq}$ are also contained in $U_{S(p)} \subseteq U_\Delta$. If $v_iv_j \neq v_1v_2$, then x is the endpoint of trace(t_2), since trace(t_2) cannot enter the interior of Q (Fig. 11(c)). If $v_iv_j = v_1v_2$, then x partitions segment $v_{pq}v_{pr}$ (hence trace(t_2)) into two segments that are each fully visible from v_1 and v_2, respectively (Fig. 11(d)). □

Lemma 14 implies that every vertex of U_Δ that lies on a 2-visible trace is either an endpoint of some trace or it is a crossing of two traces fully visible from the same vertex of Δ. The total number of these vertices of U_Δ is $O(n^2\alpha(n)$ from Lemma 8. It remains to count the number of crossings between 0- or 1-visible traces that are not visible from a common vertex of Δ. By Lemma 13, every 1-visible trace(t) of type C lies on the boundary of a halfplane h^- with $h^- \cap \mathrm{conv}\Delta \subset U_\Delta$. Clearly, such a trace($t$) cannot cross any other trace (a crossing is in the relative interior of both traces). We conclude that all 2-visible traces and all traces of type C jointly contain $O(n^2\alpha(n))$ vertices of U_Δ. Let X denote the set of all vertices of U_Δ located at crossings between 0- or 1-visible traces of type A or B that are not visible from a common vertex of Δ. It remains to show that $|X| = O(n^2\alpha(n))$.

Directed traces.

To every 1-visible trace(t) of type A or B, we associate a wedge $w(t)$ such that

- t lies on the boundary of $w(t)$;
- the apex of $w(t)$ is on the an edge pq of t, with $p \in S \cap H^+$ and $q \in S \cap \mathrm{cl}(H^-)$;
- the edge pq of t is extremal with respect to the point set $\{p\} \cup (S \cap \mathrm{cl}(H^-)) \subseteq S$;
- $w(t)$ contains exactly one vertex of Δ, whose color is not in color(pq).

If trace(t) is of type A, then let $w(t) = w^-_{pq}$ from Lemma 10. If trace(t) is of type B, then let $w(t) = w_{pr}$ from Lemma 12. To every 0-visible trace of type A, can associate two wedges, w^-_{pq} and w^-_{pr} from Lemma 10. Direct the boundary of each wedge $w(t)$ towards their apices (see Fig. 12). This induces a unique direction on each edge along 1-visible traces, and two possible directions on the edges along 0-visible traces. We define the *in-degree* of a crossing $x \in X$ as the number of incoming edges along 1-visible traces. Note that 0-visible edges are not counted in the in-degree, but for each crossing $x \in X$ on a 0-visible trace, one of the associated wedges is directed away from 0. Let $Y \subseteq X$ be the set of crossings that have in-degree 1 or 2. Lemma 15 below implies that it is enough to prove that $|Y| = O(n^2\alpha(n))$.

We define two *neighbors* for each vertex $x \in X$. Assume that x is the crossing of trace(t_1) and trace(t_2). Let $x_1 \in X$ be the vertex of U_Δ along trace(t_1) such that the relative interior is xx_1 lies in the interior of U_Δ. Similarly, let $x_2 \in X$ be the vertex along trace(t_2) such that the relative interior is xx_2 lies in the interior of U_Δ. The vertices $x_1 \in$ trace(t_1) and $x_2 \in$ trace(t_2) are the neighbors of x.

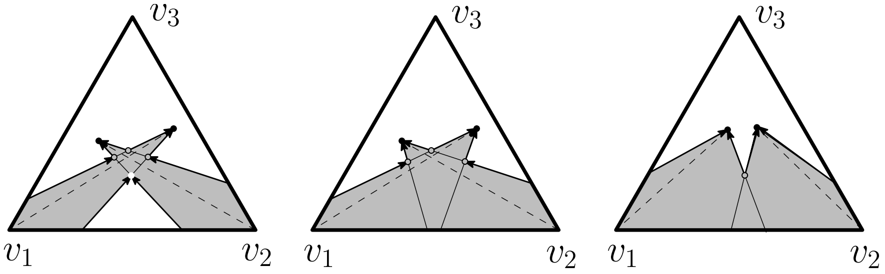

Figure 12: The boundaries of wedges are directed towards their apices. The boundaries of two wedges have up to four crossings, but at most three crossings have in-degree 0 or 1.

Lemma 15. *We have* $|X| \leq O(|Y| + n^2\alpha(n))$.

Proof. Recall that every trace is the union of collinear edges of U_Δ, and trace(t) is contained in a segment conv$\Delta \cap$ conv t for some $t \in T$. There are $O(n)$ 0-visible traces by Lemma 10, and so U_Δ has at most $O(n^2)$ vertices at crossings between pairs of 0-visible traces.

If $x \in X$ lies on a 1-visible trace and its in-degree is 0, then one of its neighbors has in-degree at least 1. Charge x to one such neighbor x' arbitrarily. The neighbor x' is either in X or it is one of the $O(n^2\alpha(n))$ vertices of U_Δ that are not in X. Since every vertex of U_Δ has at most two neighbors in X, each vertex of U_Δ is charged at most twice. We have $|X| \leq 3|Y| + O(n^2\alpha(n))$. □

Crossing wedges.

Consider a vertex $x \in Y$ of U_Δ. It is at the crossing of some 0- or 1-visible traces trace(t_1) and trace(t_2). Since the in-degree of x is 1 or 2, we may assume that trace(t_1) is 1-visible and directed towards x, and trace(t_2) is either 0-visible undirected or 1-visible directed arbitrarily. Recall that trace(t_1) and trace(t_2) lie on the boundaries of some wedges $w(t_1)$ and $w(t_2)$, respectively, each of which contains exactly one vertex of Δ. If trace(t_2) is 1-visible, then $w(t_1)$ and $w(t_2)$ contain distinct vertices of Δ, otherwise the two traces would be visible from the same vertex of Δ. If trace(t_2) is 0-visible, then it lies on the boundary of two wedges, each of which contains a distinct vertex of Δ, and we can choose either of them to be $w(t_2)$.

Let the *core* of wedge $w(t_1)$ (resp., $w(t_2)$) be the line segment between the apex and the vertex of Δ it contains. The following lemma restricts how the wedges $w(t_1)$ and $w(t_1)$ can intersect if the cores of $w(t_1)$ and $w(t_2)$ cross each other.

Lemma 16. *Let $x \in Y$ be the crossing of two traces,* $\mathrm{trace}(t_1)$ *and* $\mathrm{trace}(t_2)$*. Let v_{pr} and v_{sq} be the apices of $w_1 = w(t_1)$ and $w_2 = w(t_2)$, respectively, with $p, r \in S \cap H^+$ and $q, s \in S \cap \mathrm{cl}(H^-)$. Assume that the cores of w_1 and w_2 cross. Then* $\mathrm{color}(p) = \mathrm{color}(q)$*,* $\mathrm{color}(r) \neq \mathrm{color}(s)$*, and we have $v_{pr} \in w_2$ or $v_{qs} \in w_1$.*

Proof. Let $v_i \in w_1$ and $v_j \in w_2$ denote the unique vertices of Δ contained in the two wedges. We have $i \neq j$ as $v_i v_{pr}$ and v_{qs} cross. Since $\{p, q, r, s\}$ has two vertices on each side of H, then $H \cap \mathrm{conv}\{p, q, r, s\}$ is a convex quadrilateral (Fig. 13, left), in which v_{pr} and v_{qs} are opposite vertices. Assume without loss of generality that $H \cap \mathrm{conv}\{p, q, r, s\} = (v_{pr}, v_{ps}, v_{qs}, v_{qr})$ in counterclockwise order.

Recall that $\mathrm{color}(S \cap H^+) \subseteq \{1, 2\}$. Since w_1 and w_2 each contain exactly one vertex of Δ, we have $\mathrm{color}\{p, q, r, s\} \subset \{1, 2, 3\}$. Since they contain distinct vertices of Δ, we have $\mathrm{color}(pr) \neq \mathrm{color}(qs)$. Under these constraints, there are still several possible distributions of the colors among the points p, q, r, and s. By symmetry, it is enough to consider three cases, depending on whether the two points on one side of H have different colors or not.

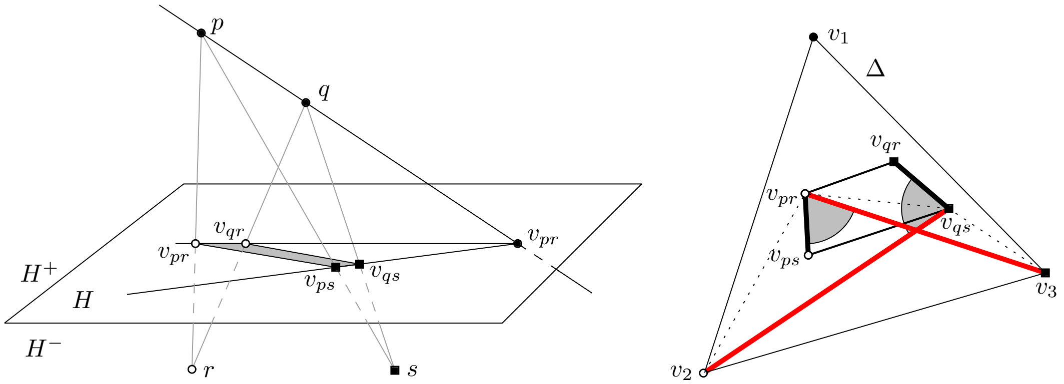

Figure 13: Configuration of the points p, q, r and s in Lemma 16.

Case 1: $\mathrm{color}(p) = \mathrm{color}(q)$**.** Without loss of generality, assume that $\mathrm{color}(p) = \mathrm{color}(q) = 1$, $\mathrm{color}(r) = 2$ and $\mathrm{color}(s) = 3$. Since the core segments $v_{pr}v_3$ and $v_{qs}v_2$ cross (Fig. 13, right), then $Q = (v_2, v_3, v_{qs}, v_{pr})$ is a convex quadrilateral. Since $(v_{pr}, v_{ps}, v_{qs}, v_{qr})$ is also a convex quadrilateral, one of v_{qr} and v_{ps} must lie outside of Q. Recall that wedge w_1 contains all points in $S(p)$ whose color is not in $\mathrm{color}(pr) = \{1, 2\}$. Therefore $v_3, v_{ps} \in w_1$ and similarly $v_2, v_{qr} \in w_2$. If v_{qr} is outside of Q, then $v_{pr} \in \angle(\overrightarrow{v_{qs}v_{qr}}, \overrightarrow{v_{qs}, v_2}) \subseteq w_2$. Otherwise $v_{qs} \in \angle(\overrightarrow{v_{pr}v_{ps}}, \overrightarrow{v_{pr}, v_3}) \subseteq w_1$.

Case 2: $\mathrm{color}(r) = \mathrm{color}(s)$**.** Without loss of generality, assume that $\mathrm{color}(p) = 1$, $\mathrm{color}(q) = 2$ and $\mathrm{color}(r) = \mathrm{color}(s) = 3$. Then $v_{pr}v_2 \subset w_{pr}$ and $v_{qs}v_1 \subset w_{qs}$. Since the core segments $v_{pr}v_2$ and $v_{qs}v_1$ cross, then x is in a bounded face of the polygon $(v_2, v_3, v_1, v_{qs}, v_{pr})$. (See Fig. 14.) By Lemma 11, the bounded faces of the polygons $(v_2, v_3, v_1, v_{qr}, v_{pr})$ and $(v_2, v_3, v_1, v_{qs}, v_{ps})$ are contained in U_Δ. They jointly contain all bounded faces of $(v_2, v_3, v_1, v_{qs}, v_{pr})$, including x. That is, x is not on the boundary of U_Δ, contradicting $x \in Y$.

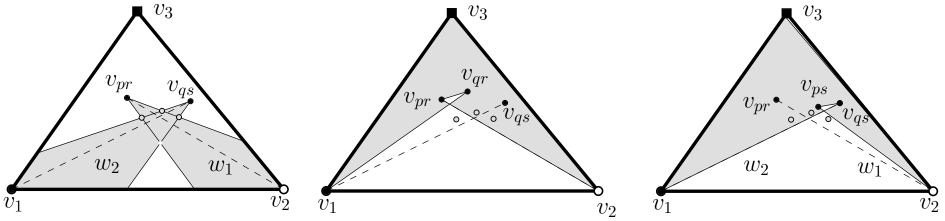

Figure 14: In Case 2, the crossing $x \in Y$ of trace(t_1) and trace(t_2) lies in a bounded face of the pentagon $(v_1, v_3, v_2, v_{pr}, v_{qr})$ or $(v_1, v_2, v_3, v_{qs}, v_{qr})$.

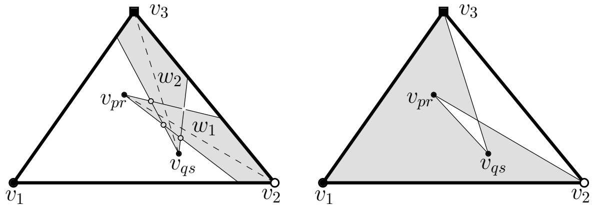

Figure 15: In Case 3, the crossing $x \in Y$ of trace(t_1) and trace(t_2) lies in a bounded face of the pentagon $(v_1, v_2, v_3, v_{qs}, v_{pr})$, which is contained in the interior of U_S.

Case 3: color$(p) \neq$ color(q) **and** color$(r) \neq$ color(s)**.** Without loss of generality, we may assume that color$(p) = 1$, color$(q) = 2$, color$(r) = 3$ and color$(s) = 1$. Then $v_{pr}v_2 \subset w_1$ and $v_{qs}v_3 \subset w_2$. Since the core segments $v_{pr}v_2$ and $v_{qs}v_3$ cross, x is in a bounded face of the pentagon $(v_3, v_1, v_2, v_{pr}, v_{qs})$. (See Fig. 15.) Similarly to the proof of Lemma 11, the closed polygonal chain (v_2, v_3, q, r) is alternately colored 2 and 3. Every edge of this chain forms a colorful triangle with each of v_1, p, and s. By adding a point of a fourth color to each triangle, we obtain colorful tetrahedra that jointly contain every bounded region in $\mathbb{R}^3$ determined by these triangles. Every bounded face of the planar polygonal chain $(v_3, v_1, v_2, v_{pr}, v_{qs})$ is on the boundary of one such region lying in H^+. Since Δ forms a colorful tetrahedron with a point in H^-, all bounded faces of $(v_3, v_1, v_2, v_{pr}, v_{qs})$ in Δ are in the interior of U_S, contradicting $x \in Y$. $\square$

We are now ready to bound the number of crossings in Y. We devise a charging scheme where we charge each crossing $x \in Y$ to a vertex of one of the sets $\widehat{D}_k$, $k = 1, 2, 3$, defined below. For $k = 1, 2, 3$, let $\widehat{L}_k$ denote the set of all traces, half-traces and all edges of the polygons $U_{S(p)}$, $p \in S \cap H^+$ that are fully visible from v_k. Note that $|\widehat{L}_k| = O(n^2)$. The convex hull of each segment in $\widehat{L}_k$ and v_k is a triangle contained in U_Δ; let $\widehat{D}_k$ denote the union of all these triangles. Analogously to Lemma 15, $\widehat{D}_k$ has $O(n^2\alpha(n))$ vertices for $k = 1, 2, 3$.

Lemma 17. *We have* $|Y| = O(n^2\alpha(n))$.

Proof. We will charge every crossing $x \in Y$ to a vertex of $\widehat{D}_k$, for some $k \in \{1, 2, 3\}$ such that each vertex is charged at most twice.

Consider a crossing $x \in Y$ of some trace(t_1) and trace(t_2). We may assume that trace(t_1) is 1-visible and directed towards x, and trace(t_2) is either 0-visible undirected or 1-visible directed arbitrarily. The traces lie on the boundary of wedges $w(t_1)$ and $w(t_2)$ such that they each contain exactly one vertex of Δ, say $v_i \in w(t_1)$ and $v_j \in w(t_2)$, $i \neq j$. This means that $\mathrm{trace}(t_1)$ is 1-visible from v_i, and $\mathrm{trace}(t_1)$ enters the interior of $w(t_2)$ at x. Note, however, that $x \in Y$ is an interior point of $\mathrm{trace}(t_1)$, and so it has to reach the boundary of U_Δ between x and the apex of $w(t_1)$. Specifically, the neighbor of x along $\mathrm{trace}(t_1)$, say x', is on the boundary of U_Δ. Denote the apices of $w(t_1)$ and $w(t_2)$ by v_{pr} and v_{qs}, respectively, such that $p, q \in S \cap H^+$ and $r, s \in S \cap \mathrm{cl}(H^-)$.

We show that the cores of $w(t_1)$ and $w(t_2)$ cross. Suppose, to the contrary, that the core of $\mathrm{trace}(t_1)$ does not cross the core of $\mathrm{trace}(t_2)$. Then $\mathrm{trace}(t_1)$ enters $w(t_2)$ at x. It cannot exit from the wedge $w(t_2)$, otherwise both $\mathrm{trace}(t_1)$ and the core of $w(t_1)$ would cross the core of $w(t_2)$. If $\mathrm{trace}(t_2)$ is of type B, then $w(t_1) \cap \mathrm{conv}\Delta \subset U_\Delta$ by Lemma 12, and so x is an endpoint of $\mathrm{trace}(t_1)$, contradicting $x \in X$. If $\mathrm{trace}(t_2)$ is of type A, then it lies on the boundary of $U_{S(q)}$ by Lemma 10. Recall that $U_{S(q)}$ denotes the colorful union of the point set $S(q)$, and we have $U_{S(q)} \cap \mathrm{conv}\Delta \subseteq U_\Delta$. Now $\mathrm{trace}(t_1)$ enters the interior of $U_{S(q)}$ at x, and it cannot exit from $U_{S(q)}$ by Lemma 7. Again, x is an endpoint of $\mathrm{trace}(t_1)$, contradicting $x \in X$. We conclude that the cores of $w(t_1)$ and $w(t_2)$ cross.

By Lemma 16, we have $\mathrm{color}(p) = \mathrm{color}(q)$ and $\mathrm{color}(r) \neq \mathrm{color}(s)$; and we also have $v_{pr} \in w(t_2)$ or $v_{qs} \in w(t_1)$. Without loss of generality, assume that $\mathrm{color}(p) = \mathrm{color}(q) = 1$, $\mathrm{color}(r) = 2$, and $\mathrm{color}(s) = 3$. Then the wedge $w(t_1)$ with apex v_{pr} contains $v_3 \in S$, and $w(t_2)$ with apex v_{qs} contains vertex $v_2 \in S$. In particular, $\mathrm{trace}(t_1)$ is fully visible from v_3. We may assume without loss of generality that the apex of $\mathrm{trace}(t_1)$ lies in $w(t_2)$. If $\mathrm{trace}(t_2)$ is of type B, then $w(t_1) \cap \mathrm{conv}\Delta \subset U_\Delta$, and so x is an endpoint of $\mathrm{trace}(t_1)$, contradicting $x \in X$. Therefore, $\mathrm{trace}(t_2)$ is of type A.

Now $\mathrm{trace}(t_1)$ enters the interior of $U_{S(q)}$ at x, it reaches the boundary of $U_{S(q)}$ on or before arriving to x', but it remains in wedge $w(t_2)$. We claim that the directed segment xx' reaches the boundary of $U_{S(q)}$ at an edge e of $U_{S(q)}$ fully visible from v_3. Invoke Lemma 7 for the planar point set $S(q)$ with $s_1 = v_2$ and $s_2 = v_3$. Every edge of $U_{S(q)}$ is either fully visible from v_3 or lies on the boundary of some wedge that contains v_2. If $\mathrm{trace}(t_1)$ enters the interior of any such wedge w that contains v_2, then it cannot exit from w by Lemma 16. Hence, $\mathrm{trace}(t_1)$ must exit from $U_{S(q)}$ at a point x'' located on an edge of $U_{S(q)}$ fully visible from v_3.

Note that x'' is the intersection of $\mathrm{trace}(t_1)$ and an edge e of $U_{S(q)}$, both of which are fully visible from v_3. Both $\mathrm{trace}(t_1)$ and e are in the set of segments $\widehat{L}_3$, and so part of $\mathrm{trace}(t_1)$ lies in the interior of $\widehat{D}_3$ in some small neighborhood of x''. However, $\widehat{D}_3 \subseteq U_\Delta$, and so the directed segment xx' reaches the boundary of $\widehat{D}_3$. Let $\hat{x}$ be the first intersection point of $\mathrm{trace}(t_1)$ with the boundary of $\widehat{D}_3$ along $x''x'$. Since $\mathrm{trace}(t_1) \in \widehat{L}_3$, $\hat{x}$ is a vertex of $\widehat{D}_3$. We charge the crossing x to vertex $\hat{x}$ of $\widehat{D}_3$. Similarly, charge every crossing $x \in Y$ to some vertex of $\widehat{D}_k$ for some $k = 1, 2, 3$. Each point $\hat{x}$ is charged at most twice. As noted above, $\widehat{D}_k$, $k = 1, 2, 3$, have a total of $O(n^2\alpha(n))$ vertices. This proves that $|Y| = O(n^2\alpha(n))$. □

We have shown that U_Δ has a total of $O(n^2\alpha(n))$ vertices in the interior of $\mathrm{conv}\Delta$. This completes the proof of Lemma 5. □

7 Open problems

We have tightened the gap between the lower and upper bound for the combinatorial complexity of the colorful union in $\mathbb{R}^3$. We do not know whether the term corresponding to the inverse of the Ackermann function is necessary. It is an obvious open problem to simplify and extend our results to higher dimensions.

We have transformed inclusion queries for U_S into vertical ray shooting queries, which lead to the data structure proposed in Theorem 4. In particular, this data structure is a ray shooting data structure for shell $(d-1)$-simplices spanned by the point set. As a consequence, the data structure ignores the fact the simplices are spanned by a *ground set* of only n points, and it also ignores the colors. It remains an interesting open question whether these two structural properties can be exploited to design a more efficient data structure.

Depth queries are a more general form of inclusion queries, because a point is contained in the colorful union if it has colorful simplicial depth of at least one. Little is known about general colorful simplicial depth queries and it is desirable to come up with efficient data structures for this problem. For the monochromatic case, the simplicial depth can be computed in $\mathbb{R}^2$ in $O(n \log n)$ time [21, 23], and in $\mathbb{R}^3$ in $O(n^2)$ time [13, 23]. In higher dimensions no better strategy than the trivial $O(n^{1+d})$ brute force test seems to be known. Afshani and Chang [1] proposed a data structure for approximate simplicial depth queries. There are no similar approximate results for the colorful simplicial depth.

Acknowledgements.

We thank Kushan Ahmadian, Boris Aronov, Gill Barequet, Marina Gavrilova, Andreas Holmsen, Mashhood Ishaque, Natan Rubin, Diane Souvaine, and Andrew Winslow for useful conversations about these topics.

References

[1] P. Afshani and T. M. Chan, On approximate range counting and depth, *Discrete Comput. Geom.* **42** (1) (2009), 3–21.

[2] P. K. Agarwal, Range searching, in: *Handbook of Discrete and Computational Geometry*, J. Goodman and J. O'Rourke, Eds. CRC Press, 2004, chapter 36, pp. 809–838.

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

[4] P. K. Agarwal and M. Sharir, Davenport-Schinzel sequences and their geometric applications, in: *Handbook of Computational Geometry*, J.-R. Sack and J. Urrutia, Eds. Elsevier, 2000, chapter 1, pp. 1–47.

[5] J. Akiyama and N. Alon, Disjoint simplices and geometric hypergraphs, in: *Proceedings of the third international conference on Combinatorial mathematics.* New York Academy of Sciences, 1989, pp. 1–3.

[6] N. Alon, I. Bárány, Z. Füredi, and D. J. Kleitman, Point selections and weak ε-nets for convex hulls, *Combin. Probab. Comput.* **1** (3) (1992), 189–200.

[7] J. L. Arocha, I. Bárány, J. Bracho, R. Fabila, and L. Montejano, Very colorful theorems, *Discrete Comput. Geom.* **42** (2) (2009), 142–154.

[8] B. Aronov and M. Sharir, Castles in the air revisited, *Discrete Comput. Geom.* **12** (1) (1994), 119–150.

[9] I. Bárány, A generalization of Carathéodory's theorem, *Discrete Math.* **40** (2-3) (1982), 141–152.

[10] I. Bárány and S. Onn, Carathéodory's theorem, colourful and applicable, *Bolyai Soc. Math. Stud., János Bolyai Math. Soc., Budapest* **6** (1997), 11–21.

[11] I. Bárány and S. Onn, Colourful linear programming and its relatives, *Math. Oper. Res.* **22** (3) (1997), 550–567.

[12] J.-D. Boissonnat, O. Devillers, and F. P. Preparata, Computing the union of 3-colored triangles, *Intern. J. Comput. Geom. Appl.* **1** (1991), 187–196.

[13] A. Y. Cheng and M. Ouyang, On algorithms for simplicial depth, in: *Canadian Conf. Comput. Geom.*, 2001, pp. 53–56.

[14] M. de Berg, *Ray Shooting, Depth Orders and Hidden Surface Removal*, vol. 703 of *LNCS*. Springer, 1993, chapter 5 "Ray shooting from a fixed point", pp. 53–65.

[15] T. K. Dey and J. Pach, Extremal problems for geometric hypergraphs, *Discrete Comput. Geom.* **19** (1998), 473–484.

[16] A. Deza, S. Huang, T. Stephen, and T. Terlaky, Colourful simplicial depth, *Discrete Comput. Geom.* **35** (4) (2006), 597–615.

[17] A. Deza, S. Huang, T. Stephen, and T. Terlaky, The colourful feasibility problem, *Discrete Appl. Math.* **156** (11) (2008), 2166–2177.

[18] H. Edelsbrunner, The upper envelope of piecewise linear functions: Tight bounds on the number of faces, *Discrete Comput. Geom.* **4** (1) (1989), 337–343.

[19] E. Ezra, B. Aronov, and M. Sharir, Improved bound for the union of fat triangles, in: *SODA*, 2011, D. Randall, Ed., SIAM, pp. 1778–1785.

[20] E. Ezra and M. Sharir, Almost tight bound for the union of fat tetrahedra in three dimensions, in: *Proc. 48th Sympos. on Foundations of Comp. Sci. (FOCS)*, 2007, IEEE, pp. 525–535.

[21] J. Gil, W. Steiger, and A. Wigderson, Geometric medians, *Discrete Math.* **108** (1-3) (1992), 37–51.

[22] A. F. Holmsen, J. Pach, and H. Tverberg, Points surrounding the origin, *Combinatorica* **28** (6) (2008), 633–644.

[23] S. Khuller and J. S. B. Mitchell, On a triangle counting problem, *Inf. Process. Lett.* **33** (6) (1990), 319–321.

[24] J. Matoušek, *Lectures on Discrete Geometry*, vol. 212 of *Graduate Texts in Math.* Springer, 2002.

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

[26] F. Meunier and A. Deza, A further generalization of the colourful Carathéodory theorem, technical report AdvOL-Report 2011/4, McMaster University, 2011.

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

[28] J. Pach and M. Sharir, The upper envelope of piecewise linear functions and the boundary of a region enclosed by convex plates: Combinatorial analysis, *Discrete Comput. Geom.* **4** (1989), 291–310.

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

[30] M. Pellegrini, Ray shooting and lines in space, in: *Handbook of Discrete and Computational Geometry*, J. Goodman and J. O'Rourke (eds.), CRC Press, 2004, chapter 36, pp. 809–838.

[31] B. Tagansky, A new technique for analyzing substructures in arrangements of piecewise linear surfaces, *Discrete Comput. Geom.* **16** (4) (1996), 455–479.

[32] R. T. Živaljević and S. T. Vrećica, The colored Tverberg's problem and complexes of injective functions, *J. Combin. Theory Ser. A* **61** (2) (1992), 309–318.

[33] A. Wiernik and M. Sharir, Planar realization of nonlinear Davenport Schinzel sequences by segments, *Discrete Comput. Geom.* **3** (1988), 15–47.