

A Single Cell in an Arrangement of Convex Polyhedra in \mathbb{R}^3 *

Esther Ezra¹ and Micha Sharir^{1,2}

¹School of Computer Science, Tel Aviv University,
Tel Aviv 69978, Israel
{estere,michas}@tau.ac.il

²Courant Institute of Mathematical Sciences, New York University,
New York, NY 10012, USA

Abstract. We show that the combinatorial complexity of a single cell in an arrangement of k convex polyhedra in 3-space having n facets in total is $O(nk^{1+\varepsilon})$, for any $\varepsilon > 0$, thus settling a conjecture of Aronov et al. [5]. We then extend our analysis and show that the overall complexity of the zone of a low-degree algebraic surface, or of the boundary of an arbitrary convex set, in an arrangement of k convex polyhedra in 3-space with n facets in total, is also $O(nk^{1+\varepsilon})$, for any $\varepsilon > 0$. Finally, we present a deterministic algorithm that constructs a single cell in an arrangement of this kind, in time $O(nk^{1+\varepsilon} \log^3 n)$, for any $\varepsilon > 0$.

1. Introduction

Let $\mathcal{P} = \{P_1, \dots, P_k\}$ be a collection of k convex polyhedra in 3-space having n facets in total. Let $\mathcal{A}(\mathcal{P})$ denote the three-dimensional arrangement induced by the polyhedra in \mathcal{P} , i.e., the decomposition of 3-space into vertices, edges, faces, and three-dimensional cells, each being a maximal connected set contained in the intersection of a fixed subcollection of facets of the polyhedra of \mathcal{P} and not meeting any other facet. The *combinatorial complexity* of a cell \mathcal{C} is the number of vertices, edges, and faces appearing on its boundary. The problem studied in this paper is to obtain a sharp upper bound on the combinatorial complexity of a single cell in such an arrangement, and to construct such a cell efficiently.

* Work on this paper has been supported by NSF Grant CCR-00-98246, by a grant from the U.S.–Israeli Binational Science Foundation, 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. This paper is part of the first author’s Ph.D. dissertation, prepared under the supervision of the second author. A preliminary version of the paper has appeared in [12].

A major application of our result is to *translational motion planning*, where a rigid polyhedral robot R (*not necessarily convex*) translates in a fixed polyhedral environment, and one wishes to compute all placements of R , known as *free* placements, at which it avoids collision with any obstacle, which are reachable via a continuous collision-free motion from a given initial free placement. Assume that R can be represented as the union of k_1 convex polyhedra R_1, \dots, R_{k_1} , and that the obstacles consist of k_2 convex polyhedra A_1, \dots, A_{k_2} . (The case when R is a single convex robot admits a different treatment; see [3] and below.) For each $1 \leq i \leq k_1$ and $1 \leq j \leq k_2$, set $K_{i,j}$ to be $A_j \oplus (-R_i)$, the *Minkowski sum* of A_j and a reflected copy of R_i . Then the free portion of the configuration space is the complement of $\bigcup_{i,j} K_{i,j}$, and the subset of all placements reachable from a given initial free placement via a collision-free motion is a single cell (connected component) of this complement, which is also a single cell in the arrangement of the polyhedra $K_{i,j}$. See [3] and [20] for more details. Thus our main result gives an upper bound on the combinatorial complexity of the set of all free placements that are reachable from an initial given placement. (Note that in this application the cell under consideration lies in the complement of the union of the polyhedra.)

As we will shortly discuss, the novelty of our result is the dependence of the bound on the number of polyhedra, and not just on the total number n of their facets, which makes it (potentially) smaller than what can be derived using previously known results. In particular, we improve significantly the solution of the above motion planning problem, when the number $k_1 k_2$ of polyhedra $K_{i,j}$ is much smaller than the overall number of their facets.

Previous Results. In two dimensions, Guibas et al. [13] showed that the complexity of a single face in an arrangement of n Jordan arcs, each pair of which intersect in at most some constant number s of points, is $O(\lambda_{s+2}(n))$, where $\lambda_q(m)$ is the maximum length of *Davenport–Schinzel sequences* of order q composed of m symbols, and is nearly linear in m for any fixed q (see [23] for more details). Efficient algorithms for computing a single face in two dimensions are given in [13] and [9]. Aronov and Sharir [4] showed that the combinatorial complexity of a single face in an arrangement of k convex polygons in the plane, having n edges in total, is $O(n\alpha(k))$, where $\alpha(\cdot)$ is the inverse Ackermann function, and this bound is tight in the worst case. This (ever so slightly) improves the bound $O(\lambda_3(n)) = O(n\alpha(n))$ obtained from the general results of [13].

In higher dimensions it was shown by Aronov and Sharir [2] that the combinatorial complexity of a single cell in an arrangement of n $(d-1)$ -simplices in \mathbb{R}^d is $O(n^{d-1} \log n)$, and that this bound also holds for the complexity of the *zone* of an additional algebraic surface σ of constant degree, or any convex surface, in such an arrangement (the zone of σ is the set of all open d -dimensional cells of the arrangement that are intersected by σ ; its complexity is the sum of the complexities of its cells). They also presented a randomized nearly quadratic algorithm that constructs a single cell in an arrangement of n triangles in 3-space (see also [1]). Finally, Halperin and Sharir [15] presented a nearly quadratic bound on the combinatorial complexity of a single cell (or a zone) in an arrangement of n low-degree algebraic surface patches in 3-space. These results were later extended to any dimension $d \geq 3$ by Basu [6], where the bounds are $O(n^{d-1+\varepsilon})$, for any $\varepsilon > 0$.¹ We note that the case of a single cell is a generalization of the

¹ In bounds of this kind, the constant of proportionality depends on ε , and generally tends to ∞ when $\varepsilon \downarrow 0$.

simple problem involving the *lower envelope* of the given surfaces. This simpler problem was studied earlier [14], [22], with the same asymptotic bounds on its complexity as those mentioned above.

Applying these results to the case of convex polyhedra, we obtain bounds that depend only on the total number n of facets, and do not exploit the fact that these facets are organized in a (potentially) smaller number of convex structures. Previous results that do exploit this fact mostly study the complexity of the (boundary of the) *entire union* of \mathcal{P} . Aronov and Sharir [4] show that the complexity of the union of k convex polygons with a total of n edges in the plane is $O(k^2 + n\alpha(k))$. Aronov et al. [5] show that the complexity of the union of k convex polyhedra with a total of n facets in \mathbb{R}^3 is $O(k^3 + nk \log k)$, and it can be $\Omega(k^3 + nk\alpha(k))$ in the worst case. The bound was improved by Aronov and Sharir [3] to $O(nk \log k)$ (and $\Omega(nk\alpha(k))$) when the given polyhedra are Minkowski sums of a fixed convex polyhedron with k *pairwise-disjoint* convex polyhedra. (In the motion-planning context mentioned above, this problem arises in the case of a *convex* translating robot R .) Our new result yields a comparable bound (for only one cell, though) when the translating robot is not convex.

As noted above, the single-cell problem is a generalization of the related problem of bounding the complexity of the lower envelope of the polyhedra in \mathcal{P} . Specifically, except for vertical discontinuities of the envelope, each of its features also appears on the boundary of the unbounded cell; the number of extra features of discontinuity is also asymptotically bounded by the complexity of the unbounded cell. See, e.g., [23] for further details. The case of lower envelopes is easier to analyze, and the special case of convex polyhedra was settled by Huttenlocher et al. [17], who presented an upper bound of $O(nk\alpha(nk))$, which can easily be improved to $O(nk\alpha(k))$.

On the algorithmic front, Aronov et al. [5] presented a randomized incremental algorithm for computing the boundary of the union of k convex polyhedra in 3-space having n facets in total, whose expected running time is $O(k^3 + nk \log k \log n)$. Huttenlocher et al. [17] presented a deterministic algorithm that constructs the lower envelope in an arrangement of this kind in time $O(nk^{1+\varepsilon})$, for any $\varepsilon > 0$. Aronov and Sharir [2] show that a single cell in an arrangement of n triangles in 3-space can be computed in randomized expected time $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$. They also show that if all the triangles lie in planes having only a constant number of orientations, then a single cell in an arrangement of this kind can be constructed deterministically in $O(n^2 \log n)$ time. For the general case of algebraic surface patches of constant maximum degree, Schwarzkopf and Sharir [21] show that a single cell can be constructed in randomized expected $O(n^{2+\varepsilon})$ time, for any $\varepsilon > 0$.

Our Results. We show that the combinatorial complexity of a single cell of $\mathcal{A}(\mathcal{P})$ is $O(nk^{1+\varepsilon})$, for any $\varepsilon > 0$, thus settling a conjecture of Aronov et al. [5], who presented a lower bound of $\Omega(nk\alpha(k))$, and conjectured that the upper bound is close to $O(nk)$. We present a detailed proof for the *unbounded cell* of $\mathcal{A}(\mathcal{P})$, under the additional assumption that the polyhedra in \mathcal{P} are bounded, and then argue that this implies the same asymptotic bound on the complexity of any other cell, and also extends to the case where the polyhedra may be unbounded. We apply a variant of the *charging scheme* of Halperin and Sharir [15], which relies on the randomized technique of Clarkson and Shor [10], and is based on the proof technique of [14] and [22]. The main difference between the proof in [15] and ours is that we treat each polyhedron in \mathcal{P} as a *single surface*, and

therefore a triple of our surfaces may intersect in a large non-constant number of points, which the analysis in [15] cannot handle. We thus have to modify this analysis and adapt it to our scenario. By now, the arsenal of techniques for analyzing the complexity of substructures in higher-dimensional arrangements is quite rich: It includes, for example, the techniques of Aronov and Sharir [2] and Tagansky [24], for bounding the combinatorial complexity of a single cell in an arrangement of n $(d - 1)$ -simplices in \mathbb{R}^d , and the technique of Aronov et al. [5] for bounding the complexity of the union of k convex polyhedra in 3-space with a total of n facets. However, we were unable to extend any of these alternative techniques to our context, where we seek a bound that depends also on k .

Constructing a Single Cell. We also design an efficient deterministic algorithm that constructs a single cell of $\mathcal{A}(\mathcal{P})$, in time $O(nk^{1+\varepsilon} \log^3 n)$, for any $\varepsilon > 0$. The algorithm adapts the general recursive approach of [17], but the implementation details of the “merge step” are quite different and more involved. In contrast, the $O(n^{2+\varepsilon})$ -algorithm of [2] is randomized (ours is deterministic) and not k -sensitive.

2. The Complexity of a Single Cell

Preliminaries. Let $\mathcal{P} = \{P_1, \dots, P_k\}$ be a collection of k convex polyhedra in 3-space having n facets in total. For simplicity of the analysis, we assume that the given polyhedra are in *general position*. This excludes degenerate configurations, and allows us to assume that no four polyhedron boundaries meet at a common point, no vertex of one polyhedron lies on the boundary of another, no edges of two distinct polyhedra meet, and no edge of a polyhedron meets the polygonal curve of intersection of the surfaces of any two other polyhedra. As claimed in [5], this assumption involves no loss of generality, since we can slightly perturb the vertices of the given polyhedra, so as to move them to general position, and verify that the number of vertices, edges, and faces appearing on the boundary of a given cell \mathcal{C} of $\mathcal{A}(\mathcal{P})$ does not decrease.

We classify the vertices of $\mathcal{A}(\mathcal{P})$ as in [5]: An intersection vertex v of $\mathcal{A}(\mathcal{P})$ (i.e., not a vertex of one of the polyhedra of \mathcal{P}) is said to be an *outer* vertex if it is the intersection of an edge of one polyhedron and the relative interior of a facet of another polyhedron. Otherwise v is an *inner* vertex (that is, v is the intersection of the relative interiors of three facets of three distinct polyhedra). The number of outer vertices in the entire arrangement $\mathcal{A}(\mathcal{P})$ is $O(nk)$, since each of the $O(n)$ edges of the polyhedra of \mathcal{P} intersects at most two facets of any other polyhedron (see also [5]), so our main goal is to bound the number of inner vertices of any given (or, as in our analysis, the unbounded) cell \mathcal{C} of $\mathcal{A}(\mathcal{P})$. We note that the total number of inner vertices of the entire arrangement $\mathcal{A}(\mathcal{P})$ is $O(nk^2)$; see [5] for the easy proof.

Reducing to the Case of the Unbounded Cell in an Arrangement of Bounded Polyhedra. In this section we show that the combinatorial complexity of \mathcal{C} is $O(nk^{1+\varepsilon})$, for any $\varepsilon > 0$. In what follows we assume that the polyhedra in \mathcal{P} are *bounded* and that \mathcal{C} is the *unbounded cell* of $\mathcal{A}(\mathcal{P})$. These assumptions involve no loss of generality, since the problem of bounding the combinatorial complexity of any single cell of $\mathcal{A}(\mathcal{P})$ can be reduced to the above case. Indeed, let Δ be a cell of $\mathcal{A}(\mathcal{P})$, and assume first that Δ

lies in the complement of the union of the polyhedra in \mathcal{P} , and that the polyhedra are bounded. Let h denote a plane that intersects the interior of Δ and does not pass through any vertex of Δ ; such a plane is easily seen to exist. We cut each polyhedron P of \mathcal{P} by h , and leave an arbitrarily small gap between the two resulting pieces of P , thereby splitting P into two smaller polyhedra (P remains unchanged if it misses h). We denote the resulting collection of these polyhedra by \mathcal{P}' . It is easy to verify from the construction that Δ becomes a portion of the unbounded cell \mathcal{C}' in the new arrangement $\mathcal{A}(\mathcal{P}')$, and that the number of vertices, edges, and faces of \mathcal{C}' is not smaller than the number of vertices, edges, and faces of the original cell Δ . The new collection \mathcal{P}' contains at most $2k$ polyhedra that have at most $2(n+k)$ facets in total. Thus our asserted asymptotic upper bound on the combinatorial complexity of the unbounded cell of $\mathcal{A}(\mathcal{P})$ also holds for any cell of $\mathcal{A}(\mathcal{P})$ (in the complement of the union of \mathcal{P}).

We next consider the case where the polyhedra in \mathcal{P} may be unbounded. We note that the notion of the unbounded cell of $\mathcal{A}(\mathcal{P})$ may not be well defined in this case. However, we can transform the given polyhedra into a set of bounded polyhedra, by intersecting each unbounded polyhedron $P \in \mathcal{P}$ with a large box enclosing all vertices of $\mathcal{A}(\mathcal{P})$ in its interior, thereby obtaining a new set \mathcal{P}'' of k bounded convex polyhedra with $O(n+k)$ facets in total. It is easy to see that any bounded cell in the original arrangement $\mathcal{A}(\mathcal{P})$ appears unchanged in the resulting arrangement $\mathcal{A}(\mathcal{P}'')$, and that any unbounded cell of $\mathcal{A}(\mathcal{P})$ is contained in the unbounded cell of $\mathcal{A}(\mathcal{P}'')$. In the latter case the number of features of the new unbounded cell is not smaller than those of the original cell. (To keep the polyhedra in general position, we can intersect each polyhedron with a slightly shifted copy of the bounding box; the complexity of the resulting unbounded cell may increase in this way.)

So far, we have only considered cells in the complement of the union of \mathcal{P} . Suppose next that the cell Δ lies in the intersection K of l polyhedra of \mathcal{P} (and outside all the remaining ones). We first replace these l polyhedra by K . Clearly, this does not change the combinatorial complexity of Δ . We now observe that the vertices of Δ are either (i) vertices of ∂K , (ii) outer vertices that lie on the edges of ∂K (some of which might have been inner vertices before replacing the l original polyhedra by K), (iii) inner vertices on ∂K , or (iv) inner vertices of Δ in the interior of K . Since the complexity of K is clearly only $O(n)$, there are at most $O(n)$ vertices of type (i), at most $O(nk)$ vertices of type (ii), and all the vertices of type (iv) lie on the boundary of (one component of) the complement of the union of the $k-l$ remaining polyhedra. We now bound the number of vertices of type (iii). Each such vertex v is created by the intersection of a facet F_1 of K and the intersection edge e of a pair P_2, P_3 of polyhedra (neither of which is K). Clearly, each such edge e can intersect ∂K in at most two points, due to the convexity of K . It thus follows that the overall number of such vertices is $O(nk)$. We can thus ignore the l polyhedra containing Δ , since the cell containing Δ in the resulting subarrangement contains all but at most $O(nk)$ vertices, edges, and faces of Δ .

To summarize, we have shown that it suffices to analyze the complexity of the unbounded cell in an arrangement of k convex bounded polyhedra with a total of n facets.

In what follows with a slight abuse of notation, we denote by \mathcal{C} the closure of the unbounded cell of $\mathcal{A}(\mathcal{P})$. In the analysis we use subsets $\mathcal{P}' \subseteq \mathcal{P}$, and denote by $\mathcal{C}(\mathcal{P}')$ the (closure of the) unbounded cell of $\mathcal{A}(\mathcal{P}')$.




Fig. 1. An inner vertex v . (a), (b) The outer exterior of the four polyhedra P_1, P_2, P_3, P' , coincides with the (unbounded) cell \mathcal{C} . The polyhedra P_1, P_2, P_3 create a quadrilateral ψ , that lies on the facet F_1 of P_1 and is intersected by the fourth polyhedron P' . All four vertices of ψ lie on the boundary of \mathcal{C} . (b) A cross-section through the facet F_1 of P_1 , its intersections with P_2, P_3, P' , and the quadrilateral ψ . (c) The convex polygon ψ is formed on the cross-section through the facet F_1 (represented by the plane of the figure) and the two polyhedra P_2 and P_3 . As depicted in the figure, ψ is an octagon, and three of its edges are intersected by the two bottom polyhedra. The five remaining edges form the exposed convex chain γ , all of whose six vertices appear on $\partial\mathcal{C}$.

The Number of Inner Vertices. Let v be an inner vertex of \mathcal{C} , which is incident to three facets F_1, F_2 , and F_3 of three distinct respective polyhedra P_1, P_2 , and P_3 . Then v is incident to a convex polygon ψ obtained by the intersection of one of the facets, say facet F_1 , and the two other polyhedra P_2, P_3 . In this case the polygon $\psi = F_1 \cap P_2 \cap P_3$ lies fully inside the union of \mathcal{P} , except for (some of) its vertices, one of which is v (since v lies on the boundary of the unbounded cell of $\mathcal{A}(\mathcal{P})$). The cell \mathcal{C} may contain other vertices of ψ in addition to v . See Fig. 1(a),(b) for an illustration of the case where ψ is a quadrilateral, and Fig. 1(c) for the general case.

Rather than dealing with ψ as a single entity, we break it into a collection of pairwise disjoint *exposed chains*. Such a chain is a maximal contiguous sequence of *inner vertices* v_0, v_1, \dots, v_j of ψ in counterclockwise order, such that each v_i lies on $\partial\mathcal{C}$, and each edge $v_i v_{i+1}$ does not intersect any other polyhedron of \mathcal{P} . The *length* of a chain is the number j of its edges (in particular, when $j = 0$ only one vertex of the chain lies on $\partial\mathcal{C}$). Note that we may have $v_j = v_0$ (when $j > 2$), that is, the chain may be the entire polygon ψ . Exposed chains with $v_0 \neq v_j$ are called *open*, and chains with $v_0 = v_j$ (with $j > 2$) are called *closed*. See Fig. 1(c). The length of a chain is not necessarily bounded by a constant, since two polyhedra and a facet of a third polyhedron may intersect in a convex polygon with up to $\Omega(n)$ vertices. However, we show below that the overall number of vertices of exposed chains of length at least 5, as well as of all exposed *closed* chains, is $O(nk)$. It therefore suffices to consider only exposed *open* chains of length at most 4. We define $V_0^{(j)}(\mathcal{P})$, for $j \geq 0$, to be the maximum number of inner vertices of the unbounded cell of $\mathcal{A}(\mathcal{P})$ that lie on exposed open chains of length at least j .

Our approach is to derive a recurrence relationship for the number of inner vertices, by bounding each of the functions $V_0^{(j)}$ in terms of $V_0^{(j+1)}$ (with a special handling of $V_0^{(5)}$), and the solution of the resulting system of recurrences will yield the asserted bounds. Note that we actually seek a bound on the quantity $V_0^{(0)}(\mathcal{P})$, which bounds the overall number of inner vertices of \mathcal{C} (that lie on exposed open chains).

For the analysis, we define, analogously to [15], the *level* of a vertex w of $\mathcal{A}(\mathcal{P})$ to be l if, by removing l polyhedra from \mathcal{P} , none of which is incident to w , we make w a vertex of the unbounded cell in the resulting subarrangement, and if l is the smallest number with that property. Clearly, all vertices of \mathcal{C} are at level 0. (In general, the set of l polyhedra whose removal “exposes” w need not be unique.)

Lemma 2.1. *For each $j = 0, \dots, 4$, and for any parameter $\xi \leq k$, we have*

$$V_0^{(j)}(\mathcal{P}) = O(\xi^2 \mathbf{E}[V_0^0(\mathcal{R})] + \xi^3 \mathbf{E}[V_0^{(j+1)}(\mathcal{R})] + nk), \quad (1)$$

where \mathcal{R} is a random sample of k/ξ polyhedra of \mathcal{P} , and $\mathbf{E}[\cdot]$ denotes expectation with respect to the choice of \mathcal{R} .

Proof. We fix $0 \leq j \leq 4$, and let γ be an exposed open chain of length $j' \geq j$ but less than 5, that is contained in an intersection polygon ψ (if $j' \geq 5$ then, as we will show in Lemma 2.2, the overall size of these chains is $O(nk)$). Let $v \in \partial\mathcal{C}$ be the *last* (most counterclockwise) vertex of γ , let e be the edge of ψ that emanates from v in the counterclockwise direction, and let q denote the other endpoint of e .

We traverse the line containing e from v towards q , and stop as soon as one of the critical events listed below is encountered. In each case we charge v to certain features encountered along the traced portion β . (We may stop either before or after reaching q , or at q itself.)

Case (a): β ends at an outer vertex z (in this case $z = q$), before reaching any vertex of \mathcal{C} . We charge v to z . Note that z is charged in this manner only a constant number of times, since along each edge emanating from z it is charged at most twice. Thus the overall number of vertices v of this kind is $O(nk)$; see Fig. 2(a).

Case (b): β ends at a vertex u of \mathcal{C} . Clearly, the relative interior of vq is disjoint from \mathcal{C} . If $u = q$, then, since v is the last vertex on the chain, the edge vq must intersect a fourth polyhedron (other than the three that create the chain). Otherwise, if q is not a




Fig. 2. The charging scheme. (a) We trace the segment $\beta = vz$ from a vertex v of a polygon ψ , and charge v to the outer vertex z . (b) We trace the segment $\beta = vu$ from a vertex v of a polygon ψ , and encounter the polyhedra R_1, R_2 before reaching $u \in \partial\mathcal{C}$. We charge the first ξ vertices w_i , assuming there are at least ξ of them. (c) When we remove these polyhedra, q becomes a vertex of \mathcal{C} , and the edge vq becomes a (portion of an) exposed chain.

vertex of \mathcal{C} , β must contain q as an interior point; see Fig. 2(b). Let t denote the number of distinct polyhedra of \mathcal{P} that intersect β (excluding the three defining ψ). We consider the following two subcases:

Case (b.1): $t \geq \xi$. In this case we charge v to a block of ξ points of intersection between β and the first ξ polyhedra that we encounter during the traversal. Note that each polyhedron intersects β in at most two points, and we choose the point of intersection that lies nearer to v along β (at which β enters the corresponding polyhedron). Any such point w is an inner vertex of $\mathcal{A}(\mathcal{P})$, and it can be charged up to three times (we omit the easy details). By construction, each of the charged vertices w lies at level at most ξ , as is easily verified.

Next we obtain an upper bound for the number of inner vertices of $\mathcal{A}(\mathcal{P})$ that lie at level at most ξ , by applying the probabilistic analysis technique of Clarkson and Shor [10], in a manner similar to that in [15]. That is, we choose a random sample \mathcal{R} of $r = k/\xi$ polyhedra of \mathcal{P} , and construct the arrangement $\mathcal{A}(\mathcal{R})$. Let w be an inner vertex of $\mathcal{A}(\mathcal{P})$ at level $l \leq \xi$, and let \mathcal{L} be a collection of l polyhedra whose removal makes w a vertex of $\mathcal{C}(\mathcal{R})$. The probability that w shows up as a vertex of $\mathcal{C}(\mathcal{R})$ is at least $\binom{k-l-3}{r-3} / \binom{k}{r}$, since among all samples of r polyhedra, those that contain the three polyhedra that form w and do not contain any of the polyhedra of \mathcal{L} make w a vertex of $\mathcal{C}(\mathcal{R})$. Hence, we have $\sum_{l=0}^{\xi} \binom{k-l-3}{r-3} / \binom{k}{r} V_l \leq \mathbf{E}[V_0^0(\mathcal{R})]$, where $V_l = V_l(\mathcal{P})$ is the number of vertices w of $\mathcal{A}(\mathcal{P})$ at level l . As in [10] and [15], this implies that, for $r = k/\xi$, we have $\sum_{l=0}^{\xi} V_l = O(\xi^3 \mathbf{E}[V_0^0(\mathcal{R})])$. Since we charge v to a block of ξ of these points, and none of these points is charged more than three times, this implies that the number of inner vertices of \mathcal{C} that fall into this subcase is $O(\xi^2 \mathbf{E}[V_0^0(\mathcal{R})])$.

Case (b.2): $t < \xi$. In this case if we remove the t encountered polyhedra from the arrangement (while retaining the three polyhedra that define ψ), the next vertex q of ψ that we meet during the traversal must appear on the boundary of the unbounded cell \mathcal{C}' in the reduced arrangement, and the edge $e = vq$ of ψ does not intersect in its interior any of the remaining polyhedra. Thus $\gamma \cup e$ becomes (possibly a prefix of) an exposed chain of length at least $j + 1$, and we charge v to q (note that q is charged up to six times in this manner); see Fig. 2(c) for an illustration.

Applying standard arguments, as in [10], [15], and above, one can easily show that the number of vertices v that fall into this case is $O(\xi^3 \mathbf{E}[V_0^{(j+1)}(\mathcal{R})])$, where \mathcal{R} is, as above, a random sample of k/ξ polyhedra of \mathcal{P} .

Since we deal with chains of bounded length, considering only the *last* vertex in a chain affects the bound by only a constant factor. Combining cases (a), (b.1), and (b.2), the lemma follows. \square

Lemma 2.2.

- (a) $V_0^{(j)}(\mathcal{P}) = O(nk)$, for $j \geq 5$.
- (b) *The number of vertices on exposed closed chains is $O(nk)$.*

Proof. (a) Let γ be an exposed open chain of at least five edges, lying on an intersection polygon $\psi = F_1 \cap P_2 \cap P_3$, where F_1 is a facet of P_1 , and P_1, P_2, P_3 are three distinct polyhedra in \mathcal{P} . By definition, γ consists of only inner vertices, and thus its edges must alternate between edges incident to ∂P_2 and edges incident to ∂P_3 ; without loss of




Fig. 3. Proof of Lemma 2.2(a). The convex polygon $M = F_1 \cap P_2$ is drawn in the plane of the facet F_1 . The two chords c_1, c_3 lying on the boundaries of the two respective polyhedra P_3, P' , connect the two edges e, e' of M along the two respective intersection polygons ψ, ψ' . The additional chord c_2 , connecting e and e' , lies on an intersection quadrilateral, and the chord c_2' lies on a pentagon having an outer vertex (in the interior of M).

generality, assume that at least three of the edges of γ lie on ∂P_2 . That is, the convex polygon $M = F_1 \cap P_2$ has at least three edges that lie on ∂P_2 , and γ contains at least two chords (edges of γ lying in the interior of M whose endpoints lie on ∂C) that connect pairs of these edges. We now argue that, once we fix a pair e, e' of edges of M , they can be connected by at most two chords, each of which belongs to a distinct exposed chain γ of length at least 5, with at least three edges that lie on ∂P_2 (two of which are e, e'). Indeed, suppose that there are three chords c_1, c_2, c_3 of this kind that connect e and e' . Since these are edges of exposed chains, they do not cross each other, so one of them, say c_2 , lies in between the other two. Moreover, the intersection polygon ϕ that contains c_2 as an edge must also be fully contained in the region bounded by e, e', c_1, c_3 . See Fig. 3. This implies that either ϕ is a quadrilateral, or ϕ is a closed polygon with at least five edges, all of whose vertices that lie in the interior of M are outer vertices of $\mathcal{A}(\mathcal{P})$. In either case, c_2 does not lie on an exposed chain with at least 5 edges (as assumed above).

We now fix a facet F and a polyhedron $P \in \mathcal{P}$, different from the one containing F in its boundary, and define the graph $G_M = (V_M, E_M)$, associated with the polygon $M = F \cap P$, as follows. The set V_M of vertices of G_M consists of all edges of the polygon M that lie on ∂P . For each exposed chain γ that visits at least three edges of M as above, we connect two vertices $v_1, v_2 \in V_M$ by an edge in E_M , if they represent two distinct respective edges of M that are connected by a chord along γ . In addition, if there are two chords connecting v_1, v_2 , we arbitrarily choose one of them to represent the edge (v_1, v_2) , ensuring that G_M is simple. Our construction implies that the number of edges $|E_M|$ is proportional to the overall size of all exposed chains γ that involve at least three edges of M . The “natural” drawing of G_M shows that it is planar. Hence, by Euler’s formula, $|E_M| = O(|V_M|)$. That is, the overall size of the chains under consideration is proportional to $|V_M|$. Since $\sum_{F,P} |V_M|$, summed over all facets F and polyhedra $P \in \mathcal{P}$, is $O(nk)$, we obtain that the overall size of all the exposed chains of length at least 5 is $O(nk)$, as asserted.




Fig. 4. From special quadrilaterals to special vertices. (a) The outer exterior of the three polyhedra P_1 , P_2 and P_3 coincides with the unbounded cell \mathcal{C} . These polyhedra create a special quadrilateral ψ of \mathcal{C} . (b) The cross-section through the facet F_1 of P_1 , P_2 , P_3 , and ψ . (c) The only surviving facets defined by ψ are the base F_1 of P_1 , and the two adjacent walls F_2 of P_2 and F_3 of P_3 , and thus v is a special vertex of the unbounded cell \mathcal{C}' . All four “upper” sides of v are accessible in \mathcal{C}' .

(b) In view of (a), it suffices to consider exposed closed chains of length 3, 4, and 5 (for longer chains, we can pretend that they are open, by discarding one edge, and then apply (a)). We first note that such a chain, i.e., the polygon ψ containing it, cannot be a triangle or a pentagon: as is easily verified, one of the vertices of ψ must then be the intersection of two adjacent edges that lie on the boundary of the same polyhedron, implying that this vertex is an outer vertex of $\mathcal{A}(\mathcal{P})$, contrary to the definition of chains.

Let us assume then that ψ is a quadrilateral. Thus either ψ is empty (i.e., does not meet any polyhedron other than P_1 , P_2 , P_3) or else it must contain an outer vertex of $\mathcal{A}(\mathcal{P})$ in its interior (that is obtained by the intersection of the interior of ψ and another polyhedron P'). In the latter case we simply charge ψ to the outer vertex that it contains (note that this vertex is charged only once in this manner), and thus the overall number of charges in this case is $O(nk)$. In the former case, adapting a term from [5], which was also used in [3] and [18], we call ψ a *special quadrilateral* of \mathcal{C} ; see Fig. 4(a),(b).

We claim that the number of special quadrilaterals of \mathcal{C} is $O(nk)$. In the analysis we “open up” the polyhedra of \mathcal{P} by sampling only some of their facets. As a consequence, we now have to regard \mathcal{C} as an *open* cell. We first define the following notions by modifying similar definitions in [15]: (i) An *x-extreme* vertex v of \mathcal{C} is a vertex whose x -coordinate is the smallest or largest in the closure of some connected component of $N \cap \mathcal{C}$, where N is a sufficiently small ball centered at v ; see Fig. 5(a) for an example (in two dimensions). (ii) A point $p \in \partial\mathcal{C}$, with x -coordinate x_0 , is said to be *critical* (with respect to the x -direction) if for every sufficiently small ball N centered at p there exists a connected component K of $N \cap \mathcal{C}$, such that $K \cap \pi_{x_0}$ is disconnected, where π_ξ denotes the plane $x = \xi$, but $K \cap \pi_x$ is connected either for all $x < x_0$, or for all $x > x_0$ sufficiently close to x_0 ; see Fig. 5(b) for an example.

We distinguish between different *sides* of an inner vertex v : Suppose that v is created by the intersection of three facets F_1 , F_2 , and F_3 . The planes spanning these facets subdivide space into eight open octants. A *side* of v is a pair (v, O) , where O is one of these octants. Let \mathcal{F} denote the set of all the n facets of the polyhedra in \mathcal{P} . Consider any subset $\mathcal{F}' \subset \mathcal{F}$ and its arrangement $\mathcal{A}(\mathcal{F}')$. We say that (v, O) is *incident* to (an open) cell \mathcal{C}' of $\mathcal{A}(\mathcal{F}')$, if v is a vertex of the closure of \mathcal{C}' , and the intersection of O with a sufficiently small neighborhood of v is contained in \mathcal{C}' .




Fig. 5. (a) The outer exterior of the two rectangles R_1, R_2 coincides with the unbounded cell C of their arrangement. The intersection vertex v is a locally x -extreme vertex of C . (b) The outer exterior of the two tetrahedra Δ_1, Δ_2 coincides with the unbounded cell C of their arrangement. The intersection point q of an edge of Δ_2 and a facet of Δ_1 is a critical point.

Now let $\psi = F_1 \cap P_2 \cap P_3$ be a special quadrilateral (with four vertices on the boundary of C), where F_1 is a facet of some polyhedron P_1 , and P_2, P_3 are two other polyhedra in \mathcal{P} . Suppose that $\partial\psi$ is formed by the intersection of F_1 with the facets F_2 and F_2' of P_2 , and F_3 and F_3' of P_3 . We refer to the facet F_1 as the *base* of ψ , and to the facets F_2, F_2', F_3, F_3' as the *walls* bounding ψ . See Fig. 4.

Following the above notation, and continuing to assume that $\psi = F_1 \cap P_2 \cap P_3$ is a special quadrilateral, let \mathcal{F}' be a subset of \mathcal{F} that contains F_1, F_2 , and F_3 , and does not contain F_2', F_3' . Let C' denote the unbounded (open) cell of $\mathcal{A}(\mathcal{F}')$. We then say that the vertex $v = F_1 \cap F_2 \cap F_3$ is a *special vertex* of C' . Clearly, v is a vertex of C' . The halfspace bounded by the plane through F_1 and not containing P_1 contains four of the eight sides of v , and all these four sides are incident to C' . Indeed, removal of F_2', F_3' exposes all these sides of v to the unbounded cell. See Fig. 4.

Following the approach of Clarkson and Shor [10], we choose a random subset $\mathcal{F}' \subset \mathcal{F}$ of expected size $n/2$, by selecting each facet of \mathcal{F} independently with probability $\frac{1}{2}$. Let $\varphi'(\mathcal{F}')$ be the overall number of special vertices of C' . We then bound the expected number $\mathbf{E}[\varphi'(\mathcal{F}')] of special vertices of the unbounded cell of $\mathcal{A}(\mathcal{F}')$, as defined above, and show that the actual number $\varphi(\mathcal{P})$ of special quadrilaterals of C (before the sampling) is at most proportional to $\mathbf{E}[\varphi'(\mathcal{F}')]$.$

We claim that each special vertex v of C' is a locally x -extreme point of C' . Indeed, let F_1, F_2, F_3 be the facets incident to v , and let O_1, O_2, O_3, O_4 be the four octants of v that lie in the exterior halfspace of F_1 (the one which does not contain P_1). All four sides (v, O_i) of v , for $i = 1, \dots, 4$, are incident to C' . The plane π that passes through v and is orthogonal to the x -axis misses exactly one of these octants, showing that v is a locally x -extreme vertex of C' .

We now apply the analysis of [15], which shows that the number of x -extreme vertices in a single cell of $\mathcal{A}(\mathcal{F}')$ is proportional to one plus the number of critical points of that cell. Thus it is sufficient to bound the number of critical points of C' . Each critical point

must be either a vertex of a polyhedron in \mathcal{P} , or an outer vertex of $\mathcal{A}(\mathcal{P})$, because no inner vertex can be critical, as is easily verified. Hence, the overall number of critical points of \mathcal{C}' , and thus also the number of special vertices of \mathcal{C}' , is only $O(nk)$.

We now bound the number $\varphi(\mathcal{P})$ of special quadrilaterals in terms of $\mathbf{E}[\varphi'(\mathcal{F}')]$. A special quadrilateral ψ becomes a special vertex if its base and an adjacent pair of its walls are chosen in \mathcal{F}' and neither of the two other walls defining ψ is chosen in \mathcal{F}' . Thus the probability that ψ becomes a special vertex is at least $4\left(\frac{1}{2}\right)^5 = \frac{1}{8}$, since there are four pairs of adjacent walls defining ψ , and the four corresponding events are pairwise disjoint. Hence $\mathbf{E}[\varphi'(\mathcal{F}')] \geq \frac{1}{8}\varphi(\mathcal{P})$, which implies that $\varphi(\mathcal{P}) = O(nk)$, as asserted. This completes the proof of part (b). \square

Remarks. 1. The proof of (a) is similar to the one given in [18], although the context in which we apply it, and some of the details, are quite different.

2. The number of special quadrilaterals is tight in the worst case, as follows from a construction of Aronov et al. [5], where the number of special quadrilaterals in the complement of the union of the polyhedra in the resulting collection is $\Omega(nk)$. In their construction the complement of the union is connected, and thus coincides with the unbounded cell in the arrangement of these polyhedra, so the lower bound immediately follows; see [5] for further details.

3. The analysis of [15], which shows that the number of x -extreme vertices in a single cell is proportional to one plus the number of critical points of that cell, can be extended to any number m of cells (see [15] for the technical details). Thus the analysis in (b) applies to any number m of cells in the arrangement, and yields a bound of $O(m + nk)$ on the overall number of special quadrilaterals of these cells (where all four vertices of each special quadrilateral appear on the boundary of a common cell).

The solution of the recurrences derived in Lemmas 2.1 and 2.2 is $V_0^{(j)}(\mathcal{P}) \leq A_j(\varepsilon)nk^{1+\varepsilon}$, for any $\varepsilon > 0$ and $0 \leq j \leq 4$, where the constants $A_j(\varepsilon)$ depend only on ε and j . In particular, the number of inner vertices of the unbounded cell of $\mathcal{A}(\mathcal{P})$ is at most $A_0(\varepsilon)nk^{1+\varepsilon}$, for any $\varepsilon > 0$. The proof of this claim is routine but technical, and is a variant of earlier proofs that solve similar recurrences (see, e.g., the analysis of [15]).

Since the number of outer vertices is only $O(nk)$, routine arguments imply that the overall complexity (number of vertices, edges, and faces) of \mathcal{C} is $O(nk^{1+\varepsilon})$, for any $\varepsilon > 0$, and that, as asserted above, this carries over to any cell of $\mathcal{A}(\mathcal{P})$. The lower bound follows from a construction of Aronov et al. [5]. In summary, we have obtained our main result:

Theorem 2.3. *Let \mathcal{P} be a collection of k convex polyhedra in \mathbb{R}^3 with n facets in total. The combinatorial complexity of a single cell of $\mathcal{A}(\mathcal{P})$ is $O(nk^{1+\varepsilon})$, for any $\varepsilon > 0$, where the constant of proportionality depends on ε . This bound is almost tight in the worst case, since there are constructions where the complexity of a single cell is $\Omega(nk\alpha(k))$.*

Zone Complexity. The preceding analysis can be extended to derive an upper bound on the overall complexity of the *zone* of a surface σ in $\mathcal{A}(\mathcal{P})$. (Recall that the zone of σ in $\mathcal{A}(\mathcal{P})$ is the set of all open cells of $\mathcal{A}(\mathcal{P})$ that are intersected by σ , and the complexity of the zone is the sum of the complexities of these cells.) We obtain:

Theorem 2.4. *Let \mathcal{P} be a collection of k convex polyhedra in \mathbb{R}^3 with n facets in total, and let σ be an algebraic surface of constant degree, or the boundary of an arbitrary convex set in 3-space. Then the combinatorial complexity of the zone of σ in $\mathcal{A}(\mathcal{P})$ is $O(nk^{1+\varepsilon})$, for any $\varepsilon > 0$, where the constant of proportionality depends on ε , and on the degree of σ in the algebraic case.*

Proof. Extending the technique presented in [11] and in [15], we cut the boundary of each polyhedron $P \in \mathcal{P}$ along its curve of intersection with σ , thus obtaining a new collection \mathcal{P}' of $O(n)$ subpatches whose overall complexity is $O(n+k)$. We leave an arbitrarily small gap between the resulting subpatches of P , so that all cells of the zone of σ form a single cell \mathcal{C}' in the arrangement $\mathcal{A}(\mathcal{P}')$. However, each of these patches does not necessarily have the shape of a convex polyhedron, and thus the result of Section 2 is not directly applicable in this case. Nevertheless, we can extend the preceding analysis and show that the combinatorial complexity of \mathcal{C}' is $O(nk^{1+\varepsilon})$, for any $\varepsilon > 0$.

We first claim that the overall number of intersection vertices on σ itself is $O(nk)$. Any such vertex is the intersection of σ with some edge $e = F_1 \cap F_2$, where F_1, F_2 are two facets of two distinct polyhedra $P_1, P_2 \in \mathcal{P}$. Since either σ has constant description complexity or is the boundary of a three-dimensional convex set, it intersects e in a constant number of points (two points in the latter case). Since the overall number of edges e of this kind is $O(nk)$, the asserted bound follows. We now extend the charging scheme described in Section 2 as follows. We trace a segment β starting at a vertex v on $\partial\mathcal{C}'$ (that does not lie on σ), and stop when one of the critical events listed in Lemma 2.1 is encountered along β . Note that in case (b) β ends at a vertex u of \mathcal{C}' (see Lemma 2.1), which may lie on σ . In this case we charge v to u , and, as is easily verified, u is charged at most twice in this manner, and thus the overall number of charges of this kind is $O(nk)$. Combining this case with the cases listed in Lemma 2.1, we obtain the same recurrences as those given in Lemma 2.1. In these recurrences we also use the fact that the number of locally x -extreme vertices of \mathcal{C}' is $O(nk)$. This follows by bounding the number of locally x -extreme vertices by one plus the number of critical points of \mathcal{C}' , using the general relationship of [15]. In this case each critical point is either an outer vertex, as above, or a vertex on σ , so the number of these points is still $O(nk)$. We thus conclude that the combinatorial complexity of \mathcal{C}' is $O(nk^{1+\varepsilon})$, for any $\varepsilon > 0$. \square

3. Constructing a Single Cell

In this section we present a divide-and-conquer algorithm that constructs the unbounded cell in an arrangement of *bounded* convex polyhedra. Specifically, let $\mathcal{P} = \{P_1, \dots, P_k\}$ be a collection of k bounded convex polyhedra in 3-space having n facets in total, and let \mathcal{C} be the unbounded cell of $\mathcal{A}(\mathcal{P})$. We present below an efficient algorithm that constructs \mathcal{C} , whose overall running time is $O(nk^{1+\varepsilon} \log^3 n)$, for any $\varepsilon > 0$. At the end of this section we present a slight modification of the algorithm so that it constructs any specified cell in $\mathcal{A}(\mathcal{P})$, e.g., the cell containing a given marking point p not lying on any polyhedron boundary.

We apply a variant of the divide-and-conquer scheme presented in [17]. We partition the given polyhedra into r roughly equal subsets $\mathcal{P}_1, \dots, \mathcal{P}_r$, each of which consists of

$\lceil k/r \rceil$ polyhedra, where r is a sufficiently large constant that we will fix shortly. Let n_i be the overall number of facets of the polyhedra of \mathcal{P}_i , for $i = 1, \dots, r$, and let $\mathcal{P}_{i,j} = \mathcal{P}_i \cup \mathcal{P}_j$, for $1 \leq i < j \leq r$. We recursively compute the unbounded cells $\mathcal{C}_{i,j}$ of the arrangements $\mathcal{A}(\mathcal{P}_{i,j})$, and “merge” them (that is, superimpose them) to extract from them the unbounded cell \mathcal{C} in the entire arrangement $\mathcal{A}(\mathcal{P})$. As we will show, the merge step can be performed in overall time $O(rnk^{1+\varepsilon} \log^3 n) = O(nk^{1+\varepsilon} \log^3 n)$, for any $\varepsilon > 0$, which implies the following recurrence for the maximum time $T(k, n)$ for constructing the unbounded cell in an arrangement of k convex polyhedra with n facets in total:

$$T(k, n) \leq \sum_{1 \leq i < j \leq r} T\left(\frac{2k}{r}, n_i + n_j\right) + cnk^{1+\varepsilon} \log^3 n, \quad (2)$$

where c is a sufficiently large constant that depends on r , and where $\sum_{i=1}^r n_i = n$. Using induction on k , it is easy to see that $T(k, n) = O(nk^{1+\varepsilon} \log^3 n)$, for any $\varepsilon > 0$, where the constant of proportionality depends on our choice of r and ε .

The merge step of the algorithm uses the following ray shooting technique (similar ideas were used in [1]). Suppose we are at a vertex $v \in \partial\mathcal{C}$ that lies on three facets F, F', F'' of three respective polyhedra $P \in \mathcal{P}_i, P' \in \mathcal{P}_j, P'' \in \mathcal{P}_l$, for some triple of distinct indices i, j, l . (Vertices that lie on facets of polyhedra in only two subcollections $\mathcal{P}_i, \mathcal{P}_j$, are easier to handle.) We shoot from v along the intersection segment $e = F \cap F'$ in the direction that proceeds from v along $\partial\mathcal{C}$ (that is, away from P''), until we hit a new polyhedron, thereby reaching a new vertex v' of $\partial\mathcal{C}$. We then shoot from v' along its two other incident edges along $\partial\mathcal{C}$, and keep exploring $\partial\mathcal{C}$ in this manner until no new vertices are found. See Fig. 6(a). At this point we have traced a complete connected component of 1-skeleton E of $\partial\mathcal{C}$, i.e., the union of the edges of $\partial\mathcal{C}$. We now look for a new vertex of $\partial\mathcal{C}$ that lies on a different connected component of E , and repeat the above procedure until the collection entire 1-skeleton E is constructed. Completing the representation of \mathcal{C} is then quite routine (and thus omitted).

To complete the presentation of the algorithm, we need to describe two steps: (i) How to perform the ray shootings. (ii) How to find a point on each connected component of E .

We note that a variant of both of these steps was previously introduced by Aronov and Sharir [1], where they present an algorithm that constructs m cells in an arrangement of




Fig. 6. The merge step. (a) Shooting along an edge of $\partial\mathcal{C}$ (the white portion) within F . (b) The short blue segment crosses a contiguous subsequence of long red segments within $\sigma(\mu)$. (c) Vertical extensions connect different components of E on F .

n convex plates that lie in planes with only a constant number of orientations. However, we improve and generalize both steps of the algorithm: (i) Our ray-shooting machinery supports $O(\log^3 n)$ query time in more general setups, whereas the query time in the ray-shooting machinery of [1] is $O(\log n)$ times the number of distinct orientations (which can be very expensive in our case). (ii) We simplify the second step of the algorithm of [1] by exploiting the topological properties stated in Lemma 3.1—see below.

Consider task (i). Fix a facet F of some polyhedron $P \in \mathcal{P}_i$, and consider the problem of shooting within F from a vertex v along an edge e of some $\partial\mathcal{C}_{i,j}$. Fix a third subcollection \mathcal{P}_l , and consider the subtask of finding the first point (if any) where the portion of e along which we shoot hits some polyhedron of \mathcal{P}_l . This is equivalent to asking for the first segment of $\partial\mathcal{C}_{i,l}$ lying on F that this portion of e hits. We need to repeat this step for each $l \neq i, j$, and take among all the (at most) $r - 2$ output vertices the one that is nearest to v , or the other endpoint of e if it does not hit any other segment.

We thus face the following “trichromatic” subproblem. We have three fixed subcollections, $\mathcal{P}_i, \mathcal{P}_j, \mathcal{P}_l$, to which we refer as the *green collection*, the *red collection*, and the *blue collection*, respectively. We fix a green facet F , and denote by $\mathcal{R} = \mathcal{R}_F$ the set of all edges of $\partial\mathcal{C}_{i,j}$ on F (green–red edges) and by $\mathcal{B} = \mathcal{B}_F$ the set of all edges of $\partial\mathcal{C}_{i,l}$ on F (green–blue edges). To simplify the notation, we refer to the edges of \mathcal{R} (resp., \mathcal{B}) simply as *red* (resp., *blue*). We thus need to preprocess these two sets to support efficient ray-shooting queries that involve the red and the blue edges—see below.

Efficient Ray-Shooting Machinery. Our ray-shooting machinery uses a *hereditary segment tree* data structure \mathcal{T} (see [8] and [19]) that stores the set $\mathcal{S} = \mathcal{R} \cup \mathcal{B}$ of the red and the blue segments on F , using their projections on the x -axis (in an appropriate generic two-dimensional coordinate frame attached to F). Put $N = |\mathcal{S}|$. Each interior node μ of \mathcal{T} represents an interval along the x -axis, which is the union of the “atomic intervals” associated with the leaves of the subtree rooted at μ . Let $\sigma(\mu)$ be the vertical slab containing all points whose x -coordinates lie in the interval that μ represents. A segment $s \in \mathcal{S}$ is said to be *long* in a slab $\sigma(\mu)$ if s crosses $\sigma(\mu)$ from side to side, otherwise, if at least one endpoint of s lies in $\sigma(\mu)$, s is *short* in $\sigma(\mu)$. We store in each node μ of \mathcal{T} two respective *long lists* of red and of blue long segments, and two additional respective *short lists* of red and of blue short segments, as follows. If a segment s is long in $\sigma(\mu)$ but not in the parent slab of $\sigma(\mu)$, then we store s in the appropriate red or blue long list of μ (the standard segment tree rule). The short red (blue) list of μ contains all red (blue) segments stored in all long red (blue) lists of the *proper descendants* of μ . As shown in [8] and [19], the overall size of the lists stored in all the nodes of \mathcal{T} , and the time to construct them, is $O(N \log N)$. As observed in [8] and [19], each red–blue intersection point arises as the intersection of a long segment and another (short or long) segment in exactly one slab $\sigma(\mu)$. We sort each of the long red lists within each slab $\sigma(\mu)$ in their (well-defined total) ascending vertical order (i.e., the order in the y -direction, in the F -coordinate frame), and store them, in this order, at the leaves of a balanced binary tree \mathcal{T}_μ . Since the red long segments are pairwise disjoint, the set of the red long segments that a blue (short or long) segment b crosses (within $\sigma(\mu)$) is a *contiguous* subsequence of the red long list, and can be represented as the disjoint union of $O(\log N)$ subtrees of \mathcal{T}_μ ; see Fig. 6(b) and [8] and [19]. We thus query \mathcal{T}_μ with each blue segment b that is

stored at μ , and store b at each of the $O(\log N)$ subtrees of \mathcal{T}_μ that constitute the output to the query. We thus obtain in each node v of \mathcal{T}_μ a *complete bipartite graph* $\mathcal{R}_v \times \mathcal{B}_v$, where \mathcal{R}_v is the set of red long segments stored at the leaves of the subtree rooted at v , and \mathcal{B}_v is the set of all blue segments (long and short) that reach v in their query. Every pair of segments $e \in \mathcal{R}_v, e' \in \mathcal{B}_v$ intersect each other (within $\sigma(\mu)$). We construct an analogous structure for the blue long list and the red segments in $\sigma(\mu)$. The overall running time needed to construct the trees \mathcal{T}_μ and the corresponding bipartite graphs, over all nodes μ of T , is $O(N \log^2 N)$.

Suppose now, that we query with a ray that starts at some point u , and proceeds from u along a portion ρ' of a red segment ρ . For each node μ of \mathcal{T} that stores ρ in one of its lists we do the following. If ρ is stored in the short list of μ , we locate the two endpoints of ρ' in the blue long list of μ , and report the first blue long segment (nearest to u) lying between these two endpoints. If ρ is stored in the long list of μ , then we visit each of the $O(\log N)$ nodes v of \mathcal{T}_μ on the search path to ρ , and, searching in the corresponding \mathcal{B}_v , find the first blue segment in that list hit by ρ' . The overall query time in \mathcal{T}_μ is $O(\log^2 N)$, thus the query takes a total of $O(\log^3 N)$ time. The final output of the query is the first blue segment that the ray hits, among all reported blue segments.

In summary, we maintain for each facet F of a polyhedron $P \in \mathcal{P}_i$, for $1 \leq i \leq r$, $r - 1$ trees \mathcal{T} as above, which correspond to all $r - 1$ sets of edges of $\partial\mathcal{C}_{i,j}$, for $j \neq i$, that lie on F , and each ray-shooting query accesses $r - 2$ such trees. Since r is a constant, a ray-shooting query can be performed in overall time $O(\log^3 n)$.

Tracing all the Components of E . Consider next task (ii), of finding a point on each connected component of 1-skeleton E of $\partial\mathcal{C}$. We first show

Lemma 3.1. *Let \mathcal{K} be a connected component of $\bigcup \mathcal{P}$. Then the intersection of $\partial\mathcal{K}$ and $\partial\mathcal{C}$ (i.e., the common boundary of \mathcal{K} and \mathcal{C}) is connected.*

Proof. In the proof we replace \mathcal{C} by its closure, and thus regard \mathcal{C} as a closed set. Recall also that we assume that the polyhedra in \mathcal{P} are in general position. We use the three-dimensional version of the Jordan curve theorem, which asserts that a closed surface in \mathbb{R}^3 (a connected embedded 2-manifold without boundary) separates the three-dimensional space into two connected (disjoint) components (see, e.g., Proposition 2B.1 of [16]). Specifically, we proceed as follows.

Let H be a connected component of the common boundary of \mathcal{K} and \mathcal{C} . We claim that H is a closed 2-manifold (without boundary). Indeed, being a manifold is a local condition. Let p be a point on H , and let N be a sufficiently small neighborhood of p . Then, due to the general position assumptions, $N \cap \mathcal{K}$ is either (locally) (i) a halfspace, in case p lies in the interior of a two-dimensional face of H , (ii) a convex dihedral wedge (or its complement), if p appears on an edge of H , or (iii) a trihedral cone, if p coincides with a vertex of H (the cone is the complement of a convex cone in case p coincides with an inner vertex, and non-convex (with a non-convex complement) if p coincides with an outer vertex). Therefore in all the above cases $N \cap H$ is locally a disk, and thus \mathcal{K} is a connected polyhedral set whose boundary is a (possibly not connected) polyhedral 2-manifold without boundary, and so is \mathcal{C} .

Assume now, for the sake of contradiction, that the common boundary of \mathcal{K} and \mathcal{C} has at least two connected components. Let H_1, H_2 be two of these boundary components. Then, applying the three-dimensional variant of the Jordan curve theorem, we obtain that H_1 separates \mathbb{R}^3 into two disjoint open connected regions R_1, R_2 , where R_1 contains \mathcal{K} and R_2 contains \mathcal{C} . Since \mathcal{C} is connected, the second boundary component must be contained in R_2 . However, then it is clear that \mathcal{K} cannot be connected, which results in a contradiction. \square

We first describe how we construct, for each connected component \mathcal{K} of $\bigcup \mathcal{P}$, the entire portion $E_{\mathcal{K}}$ of E that lies on $\partial\mathcal{K}$. Let F be a fixed facet of some $P \in \mathcal{P}_i$. F contains $r - 1$ planar subdivisions S_j , each formed by the edges of $\partial\mathcal{C}_{i,j}$ on F , for some $j \neq i$. We first remove those vertices of each union $\mathcal{C}_{i,j}$ on F that lie inside other unions $\mathcal{C}_{i,l}$ (this can be accomplished by $r - 2$ point locations of each vertex in the other subdivisions). We then draw one or two vertical extensions (upwards and/or downwards in the y -direction of the F -frame) from each surviving vertex of each subdivision S_j , into the interior of the corresponding regions $F \cap \partial\mathcal{C}_{i,j}$, and stop as soon as these extensions meet another segment (that belongs to one of the $r - 1$ subdivisions), or reach the boundary of F ; see Fig. 6(c) for an example. This can be easily done in time $O(r|F| \log n) = O(|F| \log n)$, where $|F|$ is the total number of edges of the unions $\partial\mathcal{C}_{i,j}$ on F , for $j \neq i$.

Consider the graph \mathcal{G} whose vertices are the vertices of $E_{\mathcal{K}}$ and the endpoints of the vertical extensions that lie on edges of $E_{\mathcal{K}}$, and whose edge set is the union of the edges of $E_{\mathcal{K}}$ and the above vertical extensions (where edges that contain vertical extension endpoints are partitioned into subedges at these points). We claim that \mathcal{G} is connected. Indeed, we observe, using Lemma 3.1, that each pair u, v of vertices of \mathcal{G} is connected along a path π on the common boundary of \mathcal{C} and \mathcal{K} . Consider any connected portion π' of π that intersects some facet F . It is easily checked that the endpoints of π' lie on edges of $\partial\mathcal{C}$ (and of $\partial\mathcal{K}$). Since the vertical extensions on F decompose $F \cap \partial\mathcal{C}$ into simply connected regions, we can replace π' by the concatenation of appropriate boundary portions of the regions intersected by π' , and thus obtain a new path π^* that connects u to v and is composed only of edges of \mathcal{G} . (Technically, some portions of π^* may not traverse complete edges of \mathcal{G} , in which case we simply discard these portions.) Hence \mathcal{G} is connected.

Let \mathcal{K} be a connected component of $\bigcup \mathcal{P}$, and let $E_{\mathcal{K}}$ denote $E \cap \partial\mathcal{K} \subseteq \partial\mathcal{C}$, as above. Suppose we are given a starting point vertex of $E_{\mathcal{K}}$. Having found such a vertex, we start the repeated ray-shooting procedure along $\partial\mathcal{C}$, in the manner described above, but with the following modification. When the shooting reveals a new edge e of $E_{\mathcal{K}}$, we check whether e contains any endpoint of some vertical extension. For any such extension, we add its other (one or two) contacts with $E_{\mathcal{K}}$ to the set of points from which further ray shootings should be attempted. See Fig. 6(c). In other words, the modified ray-shooting mechanism traces the graph \mathcal{G} defined above. Since \mathcal{G} is connected and the union of its edges contains $E_{\mathcal{K}}$, we construct the entire $E_{\mathcal{K}}$ in this manner.

Finding a starting point on each $E_{\mathcal{K}}$ is done as follows. The highest vertex v of a polytope that has not yet been traversed is a candidate for a new starting point, provided that the upward vertical ray from v either does not hit any other polyhedra or hits a polyhedron at a point that belongs to the already traced boundary of \mathcal{C} . This shooting can be accomplished in overall time $O(k^2 \log n)$, using standard

techniques involving vertical ray-shooting in 3-space (see, e.g., [1] for a variant of this technique).

In conclusion, the overall running time of the merge step is easily seen to be $O(nk^{1+\varepsilon} \log^3 n)$, for any $\varepsilon > 0$. As argued above, this also bounds the running time of the entire algorithm.

Constructing Any Specified Cell in $\mathcal{A}(\mathcal{P})$. We now show how to modify the previous algorithm slightly so that it constructs the cell of $\mathcal{A}(\mathcal{P})$ containing a given marking point p (not lying on any polyhedron boundary). With a slight abuse of notation, let \mathcal{C} denote this cell. We denote by \mathcal{P}' the subset of the polyhedra of \mathcal{P} that contain p in their interior (as in the case of the unbounded cell, we may have $\mathcal{P}' = \emptyset$); clearly, \mathcal{C} is contained in $\bigcap \mathcal{P}'$.

We now construct $\bigcap \mathcal{P}'$ in $O(n \log n)$ time, as the intersection of $O(n)$ halfspaces. Next, we pass some plane π_0 through p , and split any polyhedron $P \in \mathcal{P} \setminus \mathcal{P}'$ that π_0 crosses into two subpolyhedra P^-, P^+ , both bounded by π_0 , which the algorithm regards symbolically as slightly separated from each other. Let \mathcal{P}^* denote the new set of polyhedra (excluding those in \mathcal{P}'). We next construct the unbounded cell of $\mathcal{A}(\mathcal{P}^* \setminus \mathcal{P}')$ in $O(nk^{1+\varepsilon} \log^3 n)$ time, for any $\varepsilon > 0$. Finally, we merge these two cells using our merge procedure (in this case we only merge two cells rather than $O(r^2)$ cells as in the original description of this procedure).

We summarize below our analysis:

Theorem 3.2. *Let \mathcal{P} be a collection of k bounded convex polyhedra in \mathbb{R}^3 with n facets in total, and let p be a given marking point not lying on any polyhedron boundary. Then the cell \mathcal{C} of the arrangement $\mathcal{A}(\mathcal{P})$ that contains p can be constructed in time $O(nk^{1+\varepsilon} \log^3 n)$, for any $\varepsilon > 0$, where the constant of proportionality depends on ε .*

Discussion. We note that our algorithm, though efficient, is not output-sensitive, since the number of vertices that appear on the boundary of some of the recursively computed unbounded cells $\mathcal{C}_{i,j}$ may be significantly larger than those that appear on the boundary of the (final) unbounded cell \mathcal{C} . In this case the merge step eliminates most of these vertices when \mathcal{C} is constructed.

4. Concluding Remarks

We have presented a nearly tight bound on the combinatorial complexity of a single cell in an arrangement of k convex polyhedra in \mathbb{R}^3 with n facets in total, thus settling a conjecture of Aronov et al. [5], which asserts that the complexity of a single component of the complement of the union of these polyhedra is close to $O(nk)$. We have also extended our result to derive an upper bound on the overall complexity of the zone of a low-degree algebraic surface patch, or the boundary of an arbitrary convex set, in an arrangement of k convex polyhedra in 3-space with n facets in total. Finally, we presented a deterministic algorithm that constructs a single cell of an arrangement of this kind in overall running time $O(nk^{1+\varepsilon} \log^3 n)$, for any $\varepsilon > 0$.

We note that the combinatorial upper bound that we obtained in this paper can be used to improve other combinatorial bounds on substructures in arrangements of convex

polyhedra in higher dimensions. Consider the problem of bounding the combinatorial complexity of the lower envelope of k convex polyhedra in \mathbb{R}^4 with n facets in total. The three-dimensional version of this problem was solved by Huttenlocher et al. [17], who presented an upper bound of $O(nk\alpha(nk))$. In their technique they divide the polyhedra into two subcollections of approximately equal size, referred to as the “red” collection R and the “blue” collection B . Then they bound the *increase* in the number of red faces caused by the addition of the facets of the blue polyhedra to the already existing red envelope, and vice versa, and, using an appropriate recursion scheme, show that the overall combinatorial complexity of the lower envelope is proportional to the increase (when each of R, B contains roughly $k/2$ polyhedra). As shown in [17], the increase is bounded by the number of reflex vertices in the arrangement and the overall complexity of the *zone* of the boundary of each blue facet F in the three-dimensional arrangement \mathcal{A}_F induced by intersecting F (currently being added to the envelope) with all original red facets and previously added blue facets. An inspection of the proof given in [17] shows that it can be extended to four dimensions. The number of reflex vertices in the four-dimensional arrangement is easily seen to be $O(n^2k)$. The overall complexity of all these zones is bounded as follows. Let F be a blue facet. The intersection of F with each previously added polyhedron (or with an added portion of a blue polyhedron) is a three-dimensional convex polyhedron (possibly bounded by a portion of the boundary of F). Let \mathcal{P}_F be the collection of these polyhedra, and let t_F be the overall number of their facets. Applying the bounds stated in this paper, the zone complexity of the boundary of F in the arrangement $\mathcal{A}(\mathcal{P}_F)$ is $O(t_F k^{1+\varepsilon})$, for any $\varepsilon > 0$, and since $\sum_{P \in B} \sum_{F \in P} t_F = O(n^2)$ (as can be easily verified), it follows that the overall complexity of these zones, over all blue facets F , is $O(n^2 k^{1+\varepsilon})$, for any $\varepsilon > 0$. Thus the overall increase is bounded by $O(n^2 k^{1+\varepsilon})$, which implies that the combinatorial complexity of the lower envelope is $O(n^2 k^{1+\varepsilon})$, for any $\varepsilon > 0$. That is, we have

Theorem 4.1. *The combinatorial complexity of the lower envelope of k convex polyhedra in \mathbb{R}^4 with n facets in total is $O(n^2 k^{1+\varepsilon})$, for any $\varepsilon > 0$, where the constant of proportionality depends on ε .*

An open problem that this paper leaves is to tighten the small remaining gap between the upper and the lower bounds on the complexity of a single cell, or at least to improve the bound to $O(nk \cdot \text{polylog}(k))$. We note that in an attempt to derive a tighter bound, we tried to apply the techniques of Aronov and Sharir [2] and Tagansky [24] for upper bounding the complexity of a single cell in an arrangement of n $(d-1)$ -simplices in \mathbb{R}^d , as well as the technique of Aronov et al. [5] for bounding the complexity of the union of k convex polyhedra in \mathbb{R}^3 with n facets in total. Nevertheless, we did not manage to adapt these techniques so that they yield bounds that depend also on k . However, for small values of k , we can use the bound $O(k^3 + nk \log k)$, presented by Aronov et al. [5], on the complexity of the union of the polyhedra. In this case one can easily conclude that the complexity of the unbounded cell in the arrangement of the polyhedra is $O(nk \log k)$, for any $k \leq \sqrt{n}$.

It would be interesting to extend our results to higher dimensions. An easy lower bound in d -dimensions, where $d > 2$ is even, is obtained by the construction presented by Aronov et al. [5], in which we consider a planar collection of $k/(d/2)$ convex

polygons with a total of $n/(d/2)$ edges, so that the unbounded face in the two-dimensional arrangement induced by these polygons has $\Omega(n\alpha(k))$ vertices on its boundary (see [4]). Next, we consider a family of $d/2$ mutually orthogonal 2-flats in d -space, and place one copy of the two-dimensional configuration into each 2-flat. We then extend each polygon into a prism in the remaining $d - 2$ coordinates. As shown in [5], the number of vertices on the boundary of the unbounded cell in the arrangement induced by the k prisms, having a total of n facets, is at least $\Omega(n^{d/2}\alpha^{d/2}(k))$. The construction for odd $d > 3$ is similar and yields $\Omega(kn^{\lfloor d/2 \rfloor}\alpha^{\lfloor d/2 \rfloor}(k))$ vertices on the boundary of the unbounded cell in the arrangement (see [5] for further details). The best known upper bound in $d > 3$ dimensions is $O(n^{d-1+\varepsilon})$, for any $\varepsilon > 0$, presented by Basu [6]. Nevertheless, obtaining a sharp upper bound seems at the moment elusive, and requires developing new topological arguments that help to bound the number of extreme points and critical points that appear on the boundary of a single cell in the arrangement.

Finally, another technical problem that this paper raises is to improve the dependence of the running time on n by a logarithmic factor. It would seem that this could be done using fractional cascading [7] in the ray-shooting machinery, but so far we have not been able to apply this enhancement.

Acknowledgments

The authors thank Evgenii Shustin, Zur Izhakian, and Boris Aronov for useful discussions concerning Lemma 3.1. The authors also thank two anonymous referees for their useful comments, and (once again) to Boris Aronov for his thorough reading of the paper and for his numerous helpful comments and suggestions.

References

1. B. Aronov and M. Sharir. Triangles in space or building (and analyzing) castles in the air. *Combinatorica*, 10(2):137–173, 1990.
2. B. Aronov and M. Sharir. Castles in the air revisited. *Discrete Comput. Geom.*, 12:119–150, 1994.
3. B. Aronov and M. Sharir. On translational motion planning of a convex polyhedron in 3-space. *SIAM J. Comput.*, 26(6):1785–1803, 1997.
4. B. Aronov and M. Sharir. The common exterior of convex polygons in the plane. *Comput. Geom. Theory Appl.*, 6(3):139–149, 1997.
5. B. Aronov, M. Sharir, and B. Tagansky. The union of convex polyhedra in three dimensions. *SIAM J. Comput.*, 26(6):1670–1688, 1997.
6. S. Basu. On the combinatorial and topological complexity of a single cell. *Discrete Comput. Geom.*, 29(1):41–59, 2003.
7. B. Chazelle and H. Edelsbrunner. Fractional cascading: I. A data structuring technique. *Algorithmica*, 1(2):133–162, 1986.
8. B. Chazelle, H. Edelsbrunner, L. J. Guibas, and M. Sharir. Algorithms for bichromatic line segment problems and polyhedral terrains. *Algorithmica*, 11:116–132, 1994.
9. B. Chazelle, H. Edelsbrunner, L. J. Guibas, M. Sharir, and J. Snoeyink. Computing a face in an arrangement of line segments and related problems. *SIAM J. Comput.*, 22(6):1286–1302, 1993.
10. K. L. Clarkson and P. W. Shor. Applications of random sampling in computational geometry, II. *Discrete Comput. Geom.*, 4(5):387–421, 1989.
11. H. Edelsbrunner, L. J. Guibas, J. Pach, R. Pollack, R. Seidel, and M. Sharir. Arrangements of curves in the plane: topology, combinatorics, and algorithms. *Theoret. Comput. Sci.*, 92:319–336, 1992.

12. E. Ezra. Almost tight bound for a single cell in an arrangement of convex polyhedra in \mathbb{R}^3 . In *Proc. 21st Annu. ACM Sympos. Comput. Geom.*, pages 22–31, ACM Press, New York, 2005.
13. L. J. Guibas, M. Sharir, and S. Sifrony. On the general motion planning problem with two degrees of freedom. *Discrete Comput. Geom.*, 4:491–521, 1989.
14. D. Halperin and M. Sharir. New bounds for lower envelopes in three dimensions with applications to visibility of terrains. *Discrete Comput. Geom.*, 12:313–326, 1994.
15. D. Halperin and M. Sharir. Almost tight upper bounds for the single cell and zone problems in three dimensions. *Discrete Comput. Geom.*, 14(4):385–410, 1995.
16. A. Hatcher. *Algebraic Topology*. Cambridge University Press, Cambridge, 2002.
17. D. P. Huttenlocher, K. Kedem, and M. Sharir. The upper envelope of Voronoi surfaces and its applications. *Discrete Comput. Geom.*, 9:267–291, 1993.
18. J. Pach, I. Safruti, and M. Sharir. The union of congruent cubes in three dimensions. *Discrete Comput. Geom.*, 30:133–160, 2003.
19. L. Palazzi and J. Snoeyink. Counting and reporting red/blue segment intersections. *CVGIP: Graph. Models Image Process.*, 56(4):304–310, 1994.
20. J. T. Schwartz, and M. Sharir. On the “Piano Movers” problem: II. General techniques for computing topological properties of real algebraic manifolds. *Adv. Appl. Math.*, 4:298–351, 1983.
21. O. Schwarzkopf and M. Sharir. Vertical decomposition of a single cell in a three-dimensional arrangement of surfaces. *Discrete Comput. Geom.*, 18:269–288, 1997.
22. M. Sharir. Almost tight upper bounds for lower envelopes in higher dimensions. *Discrete Comput. Geom.*, 12:327–345, 1994.
23. M. Sharir and P. Agarwal. *Davenport–Schinzel Sequences and Their Geometric Applications*. Cambridge University Press, New York, 1995.
24. B. Tagansky. A new technique for analyzing substructures in arrangements of piecewise linear surfaces. *Discrete Comput. Geom.*, 16(4):455–479, 1996.

Received August 17, 2005, and in revised form January 12, 2006, and March 23, 2006.

Online publication December 5, 2006.