Discrete Comput Geom 12: 119–150 (1994)

Discrete & Computational
Geometry

Castles in the Air Revisited*

B. Aronov[1] and M. Sharir[2]

[1] Department of Computer Science, Polytechnic University,
Brooklyn, NY 11201, USA
aronov@ziggy.poly.edu

[2] Courant Institute of Mathematical Sciences, New York University,
New York, NY 10012, USA
and
School of Mathematical Sciences, Tel Aviv University,
Tel Aviv, Israel
sharir@math.tau.ac.il

Abstract. We show that the total number of faces bounding any one cell in an arrangement of n $(d-1)$-simplices in $\mathbb{R}^d$ is $O(n^{d-1} \log n)$, thus almost settling a conjecture of Pach and Sharir. We present several applications of this result, mainly to translational motion planning in polyhedral environments. We then extend our analysis to derive other results on complexity in arrangements of simplices. For example, we show that in such an arrangement the total number of vertices incident to the same cell on more than one "side" is $O(n^{d-1} \log n)$. We also show that the number of repetitions of a "k-flap," formed by intersecting $d-k$ given simplices, along the boundary of the same cell, summed over all cells and all k-flaps, is $O(n^{d-1} \log^2 n)$. We use this quantity, which we call the *excess* of the arrangement, to derive bounds on the complexity of m distinct cells of such an arrangement.

1. Introduction

A set T of n $(d-1)$-simplices in $\mathbb{R}^d$ decomposes $\mathbb{R}^d$ into open *cells* of dimension d (also called *d-faces*) and into relatively open *faces* of dimension k, $0 \le k < d$.

* Work on this paper by the first author has been partially supported by National Science Foundation Grant CCR-92-11541. Work on this paper by the second author has been supported by Office of Naval Research Grant N00014-90-J-1284, by National Science Foundation Grants CCR-89-01484 and CCR-91-22103, and by grants from the U.S.–Israeli Binational Science Foundation, the G.I.F.—the German–Israeli Foundation for Scientific Research and Development, and the Fund for Basic Research administered by the Israeli Academy of Sciences.

These cells and faces define the structure known as the *arrangement* $\mathscr{A}(T)$ of T; see [AS] for details concerning such arrangements in three dimensions. In this paper we prove that the maximum possible number of faces on the boundary of any single cell in $\mathscr{A}(T)$ is $O(n^{d-1} \log n)$. This is the first nontrivial upper bound for a general dimension (a trivial bound is $O(n^d)$, which is an upper bound on the total number of faces of $\mathscr{A}(T)$); it comes within a logarithmic factor of the best-known lower bound of $\Omega(n^{d-1}\alpha(n))$ [PS], where $\alpha(n)$ is the inverse Ackermann function. In addition, in the special case $d = 3$ it improves considerably the previous upper bound of $O(n^{7/3})$ established in [AS] and [AA].[1]

Our result has important applications to motion planning of mechanical systems subject to piecewise-linear "collision-constraints." An example of such a situation is a system of any number of polyhedral bodies translating independently in a polyhedral environment, avoiding the obstacles and each other. By well-known techniques (see [SS1]) the state of such a system can be mapped to a point of the "configuration space," in which each dimension corresponds to a distinct degree of freedom of the system. The set of collision-free system placements is bounded by "constraint surfaces," each representing placements where contact is being made between some specific system feature and a specific obstacle feature, or between two specific system features. In the case of translating polyhedra, the constraints on the system's behavior are Boolean combinations of linear equalities and inequalities; thus regions of the configuration space corresponding to collision-free states are separated from the remaining points by piecewise-linear surfaces. If the surfaces corresponding to all the constraints on the system, each properly triangulated, are considered, an arrangement of $(d - 1)$-simplices in the d-dimensional configuration space, d being the number of degrees of freedom of the system, are obtained; collision-free configurations thus constitute some of the cells in this arrangement, and all configurations reachable from a given initial state occupy a single cell. Thus our main theorem gives upper bounds on the descriptional complexity of the set of all configurations of a linearly constrained mechanical system (such as our translating polyhedra), which are reachable from a fixed configuration. Moreover, we provide an efficient randomized algorithm that computes such a cell, when the system has three degrees of freedom, given the constraint surfaces and the initial configuration. This leads to a near-quadratic randomized algorithm for planning translational collision-free motion of a rigid polyhedral object in a polyhedral environment. Again, this improves considerably the previous $O(n^{7/3+\delta})$ time bound of Aronov and Sharir [AS].

Motion planning was our main motivation for studying this problem. The main challenge lying further ahead in this direction is to extend our result to the general case of motion planning, in which we have a collection of n constraint surfaces or surface patches of bounded algebraic degree in d-dimensional space (d once again being the number of degrees of freedom of the moving system), and we wish to prove that the combinatorial complexity of a single cell in the arrangement of these surfaces is no more than roughly $O(n^{d-1})$. This would be a significant result,

---

[1] In this paper we regard d as a constant; thus the constants of proportionality in our bounds depend, usually exponentially, on d.

especially for systems with a small number of degrees of freedom, the first interesting case being $d = 3$; the case $d = 2$ has been satisfactorily resolved in [GSS]. Settling this conjecture for such systems would provide a significantly better upper bound on the combinatorial complexity of the portion of the configuration space that needs to be computed, and is the first step toward obtaining improved algorithms for motion planning. This paper studies the simplest case of this general problem. Some extensions of our analysis to other instances of motion planning with three degrees of freedom have recently been obtained by Halperin [H1], [H2].

In fact, the above conjecture can be formulated for an arbitrary collection of n well-behaved surfaces or surface patches in d-space (including, but not limited to, the case of bounded-degree algebraic surfaces); it states that the combinatorial complexity of one cell in the arrangement of such surfaces is no more than roughly $O(n^{d-1})$, perhaps off by a factor that involves the inverse Ackermann's function $\alpha(n)$. In the general case this appears to be a very hard problem. It extends a related conjecture that asserts that the combinatorial complexity of the *lower envelope* of such a collection of surfaces is also at most roughly $O(n^{d-1})$. The latter conjecture is also largely open, for $d \geq 3$. In the case of simplices, though, it was shown by Pach and Sharir [PS] that the maximum possible complexity of their lower envelope is $\Theta(n^{d-1}\alpha(n))$.

We now state our results in more detail and outline the structure of the paper. Let T be a collection of n $(d-1)$-simplices in $\mathbb{R}^d$, and let p be a point of $\mathbb{R}^d$ not on any simplex of T. We are interested in estimating the complexity $c(p, T)$ (i.e., total number of faces of all dimensions on the boundary) of the cell $C_p(T)$ of $\mathscr{A}(T)$ containing the point p. Let $c(T) = \max_p c(p, T)$, with the maximum taken over all points p not lying on any simplex of T (i.e., over all cells of $\mathscr{A}(T)$), and define $c_d(n) = \max\{c(T): |T| = n\}$. We aim to derive a sharp bound for $c_d(n)$. Since the complexity of the entire arrangement is $O(n^d)$, we clearly have $c_d(n) = O(n^d)$ as well. The first improvement was obtained for $c_3(n)$ by Pach and Sharir [PS], who showed a slightly subcubic bound for this quantity. Their estimate was later improved to $c_3(n) = O(n^{7/3})$ [AS], [AA]. A lower bound of $\Omega(n^{d-1}\alpha(n))$ on $c_d(n)$ is mentioned in [PS], and is conjectured in [PS] and [AS] to be tight. In this paper we almost settle this conjecture, by proving the following main theorem:

Theorem 1.1. *The number of faces of all dimensions bounding a single cell in an arrangement of n $(d-1)$-simplices in $\mathbb{R}^d$ is $O(n^{d-1}\log n)$.*

The proof proceeds in two stages. In Section 2 we derive a slightly weaker bound which is larger than the one stated above by a polylogarithmic factor. Then, in Section 3, we apply a more-refined analysis, which makes use of the results of Section 2, to prove the above theorem. An interesting feature of the argument is its close connection to the recent recurrence-based proof techniques of [ESS], [AMS], and [APS] that were introduced to analyze the complexity of zones and other structures in arrangements of hyperplanes. However, our proofs add some novel features to this approach, which we hope will find additional applications as well. For example, by applying a version of the argument used to prove our main theorem, we show that the complexity of the zone of any algebraic surface of constant degree, or of any convex surface, in an arrangement of simplices as above is also $O(n^{d-1}\log n)$.

In Sections 4 and 5 we extend our analysis to obtain sharp bounds for certain complexity measures of the entire arrangement of simplices. Some of these results require additional notation, so in this introduction we omit details concerning them and refer the reader to Section 4. One of the main results established in Section 5 is that if we sum, over all cells C of the arrangement and over all k-dimensional sets f formed by intersecting $d-k$ of the given simplices (we call these sets "k-flaps"), the number of times f appears on the boundary of C, *not counting the first such incidence*, we obtain a total "excess" that is at most $O(n^{d-1}\log^2 n)$. In the plane (for arrangements of segments) we obtain a better bound of $O(n\log n)$. Intuitively, this indicates that the large complexity ($O(n^d)$) of the entire arrangement is "caused" either by having "too many" distinct cells, or by having cells whose boundaries meet "too many" distinct k-flaps formed by the simplices. In addition to the intrinsic interest of a new complexity measure of this kind, we demonstrate its usefulness by applying it in Section 6 to derive a nontrivial bound on the complexity of m distinct cells in an arrangement of n $(d-1)$-simplices in d-space. The bound that we obtain is $O(n^{d-1}\log^2 n + m^{1/2}n^{d/2}\log^{1/2} n)$, and is probably not tight, although it compares somewhat favorably with similar bounds for arrangements of hyperplanes, obtained in [AMS]. Section 7 concludes with some applications of our main result, including an efficient randomized algorithm for computing a single cell in an arrangement of triangles in 3-space and its aforementioned application to translational motion planning, and with a discussion of possible extensions and some open problems.

2. Geometric Preliminaries and an Initial Weaker Bound

We begin our analysis with some definitions and notation. Let T be a collection of n $(d-1)$-simplices in $\mathbb{R}^d$. We view each $(d-1)$-simplex as the disjoint union of its relative interior and relatively open faces of its boundary, whose dimensions range between 0 and $d-2$; here the interior of a simplex is regarded as a $(d-1)$-face of the simplex. These faces are not necessarily faces of the arrangement $\mathcal{A}(T)$—in general they will be split into subfaces by the other simplices of T. In what follows we assume that the simplices are *in general position*, meaning that, for any $k = 2, \ldots, d$, the intersection of any collection of k relatively open simplex faces of dimension $i_1, \ldots, i_k$, respectively, is either empty or has dimension exactly $d - \sum_{j=1}^{k}(d-i_j)$. In particular, putting $i_1 = i_2 = \cdots = i_k = d-1$, any k simplices intersect, if at all, in a polytope of dimension exactly $d-k$ (as already mentioned, we refer to such an intersection polytope as a *$(d-k)$-flap*), and no $d+1$ simplices have a point in common. This assumption does not affect the quantity $c(p, T)$ that we wish to bound, since it attains its maximum value when the simplices are indeed in general position. An argument that proves this in three dimensions is given in [AS]; it can be easily generalized to arbitrary dimension.

We say that two faces of the arrangement are *incident* to each other if one is contained in the closure of the other.

We distinguish two types of faces in the arrangement—an *outer* face is contained in the relative boundary of some simplex of T, while an *inner* k-face, for $k = 0, 1, \ldots, d - 1$, lies in the intersection of exactly $d - k$ simplex interiors and avoids simplex boundaries. Note that only k-faces with $k \leq d - 2$ can be outer. It is easy to verify that the total number of outer faces is $O(n^{d-1})$. Consider, for example, the case of outer vertices. Since an outer vertex is contained in the boundary of at least one simplex, it is a vertex of the intersection of at most $d - 1$ simplices, an observation that easily yields the claimed bound. Thus it suffices to consider only inner faces.

For technical reasons, we distinguish between different *sides* of an inner face. For example, the arrangement consisting of one simplex has a single $(d - 1)$-face with two "sides." More formally, let f be an inner k-face $(0 \leq k < d)$ contained in the relative interiors of $d - k$ simplices. The hyperplanes spanning these simplices subdivide space into 2^{d-k} open regions. A *side* of f is simply a pair (f, R) where R is one of these regions. If f is a full-dimensional cell, we say that f has only one side, namely $(f, \mathbb{R}^d)$. Notice that, since we are dealing with inner faces only, a cell has exactly one side, a $(d - 1)$-face (*facet*) has two sides, a $(d - 2)$-face has four sides, and so on. A side (f, R) is called a *k-border* of a cell C if $k = d$ and $f = C$, or if f is an inner k-face on the boundary of C and some open neighborhood of f in $R \cup f$ is contained in $C \cup f$. Intuitively, this means that f is on the boundary of C and C touches f on the side of R. We define the *(inner-face) complexity* of a cell in $\mathscr{A}(T)$ to be the total number of its k-borders, for all k. As already noted, since this quantity omits outer faces, it is less than the count of all faces (or, rather, borders) bounding C by at most $O(n^{d-1})$. Thus we count inner faces f on the boundary of C with multiplicity—once for every side of f that lies in C locally near f. For example, consider two segments in the plane, crossing each other at an interior point q. Their arrangement has a single cell, and its complexity counts q four times, and each of the four subsegments incident to q is counted twice.

For $0 \leq k \leq i < d$, we define a (k, i)-border of a cell C to be a pair $((f, R), (g, Q))$ of borders of C of dimension k and i, respectively, with $f \subset \bar{g}$ and $R \subset Q$. Note that once f, g, and R are specified, the side Q is uniquely defined. A (k, d)-border is a pair $((f, R, (C, \mathbb{R}^d))$, where (f, R) is a k-border of C. Intuitively, a (k, i)-border is a pair of oriented inner faces of C, with the first face incident to the second, so that their orientations agree. For instance, in the arrangement formed in the plane by the two coordinate axes, the northeast side of the origin together with the upper side of the positive x-axis forms a $(0, 1)$-border of the northeast quadrant.

We now extend the notion of *popularity*, introduced in [APS], to the context studied here. Fixing the cell C, we call an inner k-face f of C *popular* if all 2^{d-k} sides of f are k-borders of C. For instance, a popular (in fact, the only popular) cell is C, and a popular facet is one that touches C on both sides. A (k, i)-border $((f, R), (g, Q))$ of C is *popular* if g is a popular i-face. Let $\tau_k^{(i)}(p; T)$ denote the number of popular (k, i)-borders of the cell $C_p = C_p(T)$ containing a specified point p, not lying on any simplex. Notice that the problem of bounding the complexity of C_p reduces to bounding the quantities $\tau_k^{(d)}(p; T)$, for all $0 \le k \le d$, as they refer to the number of borders of various dimensions bounding C_p. We put $\tau_k^{(i)}(n) = \max \tau_k^{(i)}(p; T)$, with the maximum taken over all collections of n $(d-1)$-simplices in $\mathbb{R}^d$ and over all choices of point p not on any simplex. The main result of this section, which is slightly weaker than Theorem 1.1, is:

Theorem 2.1. $\tau_k^{(i)}(n) = O(n^{d-1} \log^{i-1} n)$, *for* $0 \le k \le i \le d$.

Proof. Let T be a collection of simplices as above, and let p be a point not on any simplex. We use the following grand scheme for obtaining the desired bounds for $\tau_k^{(i)}(p; T)$, for all $0 \le k \le i \le d$. We first show, in Lemma 2.4, that $\tau_i^{(i)}(p; T) = O(n^{d-1})$, for all $0 \le i < d$, and that $\tau_d^{(d)}(p; T) = 1$. We then observe that, for any i, we trivially have $\tau_0^{(i)}(p; T) \le 2\tau_1^{(i)}(p; T)$, since any face has at most twice as many vertices as edges. Next we derive a recurrence for $\tau_1^{(i)}(p; T)$, which involves $\tau_0^{(i-1)}(p; T)$ and thus, by the above observation, $\tau_1^{(i-1)}(p; T)$. Solving this recurrence inductively on i, we obtain the desired bounds for $k = 0, 1$ and for all i. For $k > 1$, we "charge" a popular (k, i)-border $((f, R), (g, Q))$ to one of the edges e of f, regarded as the $(1, i)$-border $((e, R'), (g, Q))$, for an appropriate side R'; clearly, this is a popular $(1, i)$-border. By the assumption of general position, such a $(1, i)$-border cannot be charged by more than $\binom{d-1}{d-k}$ (k, i)-borders, which implies that $\tau_k^{(i)}(p; T) = O(\tau_1^{(i)}(p; T))$. This gives us the desired bounds for the quantities $\tau_k^{(i)}(p; T)$, for all $0 \le k \le i \le d$, thus implying the assertion of the theorem.

We therefore begin the proof by providing a bound on $\tau_i^{(i)}(p; T)$, for all $0 \le i \le d$. We first need a technical result that extends a theorem of Aronov and Sharir [AS] to arbitrary dimensions.

Theorem 2.2 (Chopping Theorem). *Let T be a collection of n $(d-1)$-simplices in $\mathbb{R}^d$. Let $\mathscr{C}$ be any collection of m cells in the arrangement of these simplices. Then a decomposition of the cells of $\mathscr{C}$ into $m + O(n^{d-1})$ convex polyhedra exists, i.e., a collection of $m + O(n^{d-1})$ open pairwise disjoint convex polyhedra, each fully contained in some cell of $\mathscr{C}$, so that their closures cover (the union of) $\mathscr{C}$.*

Proof. Clearly, we can ignore the convex cells of $\mathscr{C}$. We subdivide *all* nonconvex cells in $\mathscr{A} = \mathscr{A}(T)$ into convex subcells, in a manner similar to that in the proof of the Slicing Theorem of [AS], and show that this process does not increase the number of cells by more than $O(n^{d-1})$.

The construction proceeds as follows. We assume, without loss of generality, that no simplex is parallel to the x_d axis. Intuitively, the only nonconvexities present in cells of $\mathscr{A}$ are caused by relative boundaries of the simplices protruding into a cell. We erect vertical "walls," each fully contained in the interior of a cell, extending from each such boundary, thereby eliminating all nonconvex features. Let σ be one of the given simplices, and let σ_1 be one of the d $(d-2)$-simplices bounding σ; throughout the proof we refer to such a $(d-2)$-simplex as a "simplex-facet" of σ. Construct the hyperplane h_1 passing through σ_1 and parallel to the x_d axis (we refer to h_1 as the *vertical hyperplane spanned by* σ_1) and consider the $(d-1)$-dimensional arrangement $\mathscr{A}_1$ obtained by intersecting $\mathscr{A}$ with h_1; note that σ_1 itself is a simplex in $\mathscr{A}_1$. We take all the cells of $\mathscr{A}_1$ whose boundary meets σ_1 in a $(d-2)$-face (these cells constitute the so-called *zone* of σ_1 in $\mathscr{A}_1$), and add them to the arrangement $\mathscr{A}$. We refer to these cells ($(d-1)$-faces) as the *walls erected from* σ_1. Intuitively, we have drawn a vertical hyperplane h_1 through σ_1, and have added to $\mathscr{A}$ all the cells of the cross-sectional arrangement $\mathscr{A} \cap h_1$ that touch σ_1.

We now repeat this process in an incremental fashion, for each of the $d-1$ remaining simplex-facets σ_j bounding σ and for each simplex-facet bounding the remaining simplices in T. Whenever we process such a simplex-facet, the walls erected from it become part of $\mathscr{A}$. Thus when we process a simplex-facet σ', the vertical hyperplane that it spans has to be intersected with the original arrangement $\mathscr{A}$ as well as with all the vertical walls erected from previously processed simplex-facets. Note that the resulting decomposition will depend on the order in which simplex-facets are being processed, so it is not unique. It is easy to check that the introduction of the walls erected from a simplex-facet σ' eliminates σ' as a source of local nonconvexity and does not create any further points of local nonconvexity. It follows that the end result is a decomposition of the cells of $\mathscr{A}$ into open pairwise disjoint convex polyhedra. (A simpler decomposition scheme could be achieved by adding the vertical hyperplanes passing through each of the $(d-2)$-dimensional facets of the given simplices, and by using them to decompose the nonconvex cells of $\mathscr{A}$; however, this might result in the formation of too many subcells.)

How many subcells have we created? More precisely, how many more subcells are there, compared with the original cell count in $\mathscr{C}$? Recall that a wall is a $(d-1)$-face, so its addition to the arrangement increases the cell count by at most one—either it splits the cell in which it is contained into two subcells, or it does not affect the number of cells at all, if it cuts the cell without splitting it in two; in the latter case the wall alters the topology of the cell. In any case, it suffices to show that the number of walls created during our construction is $O(n^{d-1})$. This will be established by arguing that no simplex-facet σ' has more than $O(n^{d-2})$ walls erected from it.

The claim is immediate for the first simplex-facet, σ_1. A wall erected from σ_1 is a full-dimensional cell in the arrangement $\mathscr{A}_1 = \mathscr{A} \cap h_1$ whose boundary meets σ_1 in a $(d-2)$-face. Thus the number of walls is at most twice the number of full-dimensional cells in the $(d-2)$-dimensional arrangement $\mathscr{A}_1 \cap \sigma_1 = \mathscr{A} \cap \sigma_1$. Since this arrangement is formed by a collection of at most n polytopes, $\{\sigma \cap \sigma_1 : \sigma \in T\}$, each of small constant complexity, it has $O(n^{d-2})$ faces altogether, which implies our assertion for σ_1.

The situation is more complicated for walls constructed later in the process, as the $(d-1)$-dimensional cross-sectional arrangements, in which new walls are constructed, consist of not only (cross sections of) original simplices but also of previously constructed walls. We treat this case in the following indirect manner.

Suppose we are currently processing a simplex-facet σ' of some simplex $\sigma \in T$.

Let h' be the vertical hyperplane passing through σ', and consider the arrangement $\mathscr{A}'$ formed in h' by its intersections with all the simplices of T and with the vertical hyperplanes spanned by all the simplex-facets bounding the simplices of T. Clearly, $\mathscr{A}'$ is a refinement of the intersection of h' and the current version $\mathscr{A}_c$ of $\mathscr{A}$. It is easy to see that the number of faces of $\mathscr{A}'$ touching σ' in a $(d-2)$-dimensional set is at least as large as the number of such faces in the coarser arrangement $\mathscr{A}_c \cap h'$. Moreover, as argued above, the former quantity is bounded by twice the number of full-dimensional cells in the $(d-2)$-dimensional arrangement $\mathscr{A}' \cap \sigma'$. Since this arrangement is formed by at most n polytopes of complexity $O(1)$ and by dn additional hyperplanes, the number of its cells is clearly $O(n^{d-2})$.

This completes the proof of the Chopping Theorem. □

Corollary 2.3. *In an arrangement of n $(d-1)$-simplices in $\mathbb{R}^d$, any single cell can be decomposed into $O(n^{d-1})$ convex polyhedra. The same claim also holds for the collection of all nonconvex cells, all cells met by an arbitrary hyperplane, all cells met by an arbitrary convex surface, or all cells met by an arbitrary algebraic surface of bounded degree.*[2]

Proof. Immediate once it is noted that the number of cells in all cell collections listed above is $O(n^{d-1})$. In the case of cells met by an algebraic surface, this can be deduced by first considering the arrangement of the n hyperplanes spanning simplices of T, applying there the simple cell-counting argument given in [APS], and observing that each cell of $\mathscr{A}$ is a union of some cells of this arrangement.

Remark. The Chopping Theorem is not a generalization of the three-dimensional Slicing Theorem of Aronov and Sharir [AS]; it is rather an extension of one of its corollaries. The Slicing Theorem proper bounds the increase in the *complexity* (i.e., total number of faces of all dimensions) of the collection of subcells resulting from the decomposition of the cells into convex polyhedra, constructed as above. This increase, in the three-dimensional case, is $O(n^2\alpha(n))$. It seems reasonable to conjecture that, in d dimensions, this increase, as effected by our construction, is $O(n^{d-1}\alpha(n))$. Proving this, however, seems to require considerably more careful analysis.

Lemma 2.4. $\tau_i^{(i)}(p; T) = O(n^{d-1})$, *for* $i = 0, 1, \ldots, d-1$, *and* $\tau_d^{(d)}(p; T) = 1$.

Proof. First, observe that, since C_p is the only popular cell and has only one side, $\tau_d^{(d)}(p; T) = 1$. Now let $i = 0, 1, \ldots, d-1$. Recall that $\tau_i^{(i)}(p; T)$ is simply the number of popular i-borders of C_p, i.e., 2^{d-i} times the number of inner i-faces all of whose 2^{d-i} sides occur on the boundary of C_p. To prove our claim, we associate each such face with a vertex of C_p and argue that:

(1) No vertex is charged more than a constant number of times.
(2) The number of charged vertices is $O(n^{d-1})$.

---

[2] Clearly, in the last case the constants will also depend on the degree of the surface.

We set up the correspondence as follows: Rotate the arrangement in such a fashion that every inner i-face has a unique lowest vertex, where the height of a vertex is its x_d coordinate. Let f be an inner popular i-face and let v_f be its lowest vertex. We claim that v_f is either a locally lowest vertex of C_p (meaning that there is a side R of v_f so that R lies fully above v_f and (v_f, R) is a border of C_p) or an outer vertex of C_p. Indeed, if v_f lies in the interior of d simplices, the hyperplanes spanning them partition space into 2^d orthants, and exactly one of these orthants, call it Q, has v_f as its lowest point. The intersection of f with a sufficiently small neighborhood of v_f is equal to the intersection of that neighborhood with $d - i$ of these hyperplanes and with the upper half-spaces bounded by the remaining i hyperplanes. Since f is popular, all 2^{d-i} of its sides touch C_p, and it is clear that Q must be contained in one of these sides. Hence Q, locally near f, lies in C_p, and v_f is indeed a locally lowest vertex of C_p.

We have assigned each popular i-face to its lowest vertex. The above argument implies that no vertex is charged by more than $\binom{d}{i}$ popular i-faces, since we assume general position.

As already noted, the number of outer vertices in the entire arrangement is $O(n^{d-1})$. Thus it remains to show that the number of locally lowest inner vertices of C_p is also $O(n^{d-1})$. The Chopping Theorem implies that C_p can be decomposed into a collection of $O(n^{d-1})$ disjoint open convex polyhedra, the union of whose closures covers C_p. Then a locally lowest vertex of C_p is necessarily a lowest vertex of one of these polyhedra. Applying an appropriate rotation as necessary, we can assume that each of these convex polyhedra has at most one lowest vertex. Hence the number of locally lowest vertices of C_p is also $O(n^{d-1})$, which completes the proof of the lemma. □

Remarks. (1) As will be seen shortly, the proof of Theorem 2.1 uses the preceding lemma only in the special case $i = 1$.

(2) The proof of Lemma 2.4 also applies, if, instead of a single cell C_p, we consider the collection of all cells crossed by some simplex, hyperplane, convex surface, or bounded-degree algebraic surface, provided we modify the notion of popularity by defining a k-face to be popular if all of its 2^{d-k} sides lie in cells of the given collection. See [APS] for details.

We next proceed by induction on i and derive a recurrence for $\tau_1^{(i)}(p; T)$, for $i = 2, \ldots, d$, using an approach similar to that used in [ESS], [AMS], and [APS]. Fix a simplex $\sigma \in T$ and consider a popular $(1, i)$-border $((f_0, R), (g_0, Q))$ of $C_p(T)$ with $f_0 \not\subset \sigma$. When we remove σ, the face g_0 becomes part of a possibly larger inner i-face g, which is clearly also popular. Moreover, f_0 (resp. (f_0, R)) is a part of some inner edge f (resp. 1-border) of g. Thus $((f, R), (g, Q))$ is necessarily a popular $(1, i)$-border of $C_p(T\backslash\{\sigma\})$.

So let $((f, R), (g, Q))$ be a popular $(1, i)$-border of $C_p(T\backslash\{\sigma\})$, and consider what happens to it when σ is reinserted into the arrangement. Let (g, Q_l), $l = 1, \ldots, 2^{d-i}$, be the sides of g in $\mathscr{A}(T\backslash\{\sigma\})$ ((g, Q) is one of these sides). The following cases may occur:

$\sigma \cap g = \varnothing$. In this case g may or may not occur on the boundary of $C_p(T)$, but $((f, R), (g, Q))$ contributes at most one popular $(1, i)$-border to $\tau_1^{(i)}(p; T)$, namely itself.

$\sigma \cap g \neq \varnothing$ and $\sigma \cap f = \varnothing$. Introduction of σ splits g into one or more pieces, more than one of which may contain f on its boundary. It follows that only one component g^+ of $g \backslash \sigma$ has the property that $((f, R), (g^+, Q))$ is a $(1, i)$-border of some cell in $\mathscr{A}$. Thus $((f, R), (g, Q))$ can contribute at most one popular $(1, i)$-border to $C_p(T)$, namely $((f, R), (g^+, Q))$.

$\sigma \cap g \neq \varnothing$ and $\sigma \cap f \neq \varnothing$. Let h^+, h^- denote the two open half-spaces bounded by the hyperplane spanned by σ. Since f is an edge not contained in σ, σ splits f into two subedges, $f^+ = f \cap h^+$, $f^- = f \cap h^-$. As above let g^+ (resp. g^-) be the unique component of $g \backslash \sigma$ (i.e., subface of g in $\mathscr{A}$) which is incident to f^+ (resp. f^-) on the correct side of f; it may be that $g^+ = g^-$. Let g^* be the unique subface of $g \cap \sigma$ which is incident to $f \cap \sigma$ and is contained in the closure of R near it. Consider the two $(1, i)$-borders $((f^+, R), (g^+, Q))$ and $((f^-, R), (g^-, Q))$. We are only interested in cases where both of them become popular borders in $C_p(T)$, for only then will our count go up. Let $Q_l^+ = Q_l \cap h^+$, $Q_l^- = Q_l \cap h^-$, for $l = 1, \ldots, 2^{d-i}$. Thus we are interested in situations where C_p meets all 2^{d+1-i} orthants Q_l^+, Q_l^- locally near g. Notice that all these orthants are incident to g^*, an $(i-1)$-face in $\mathscr{A}$. Hence g^* is a popular $(i-1)$-face of C_p and $((f \cap \sigma, R \cap h^+), (g^*, Q \cap h^+))$ and $((f \cap \sigma, R \cap h^-), (g^*, Q \cap h^-))$ are popular $(0, i-1)$-borders of $C_p(T)$.

To sum up, the number of popular $(1, i)$-borders in $C_p(T)$ which are not contained in σ is bounded by

$$\tau_1^{(i)}(p; T \backslash \{\sigma\}) + \tfrac{1}{2}\rho_\sigma,$$

where ρ_σ is the number of popular $(0, i-1)$-borders $((f', R'), (g', Q'))$ with $g' \subset \sigma$. A factor of $\frac{1}{2}$ appears because the increase of 1 in $\tau_1^{(i)}$ is charged to two popular $(0, i-1)$-borders. If we sum these bounds over all simplices $\sigma \in T$ and observe that every popular $(1, i)$-border in $C_p(T)$ is counted exactly $n - d + 1$ times (it is not counted if and only if σ is one of the $d - 1$ simplices containing the 1-face of the border), we obtain

$$(n - d + 1)\tau_1^{(i)}(p; T) \leq \sum_{\sigma \in T} \tau_1^{(i)}(p; T \backslash \{\sigma\}) + \frac{d + 1 - i}{2} \tau_0^{(i-1)}(p; T),$$

where the factor $(d + 1 - i)$ comes from the fact that a popular $(0, i-1)$-border is charged at most $d + 1 - i$ times, once for each simplex σ containing its $(i-1)$-face.

Passing to the maximum quantities $\tau_k^{(i)}(n)$, we thus obtain

$$\tau_1^{(1)}(n) = O(n^{d-1})$$

and

$$\tau_1^{(i)}(n) \le \frac{n}{n-d+1}\,\tau_1^{(i)}(n-1) + \frac{d+1-i}{n-d+1}\,\tau_1^{(i-1)}(n), \qquad i = 2, \ldots, d,$$

where we have used an earlier observation that $\tau_0^{(i)}(n) \le 2\tau_1^{(i)}(n)$, for $1 \le i \le d$.

We first transform these equations into simpler ones, by assuming that $n \ge d$ and substituting

$$\tau_1^{(i)}(n) = \binom{n}{d-1}\psi_1^{(i)}(n).$$

This yields the following relations, as is easily verified:

$$\psi_1^{(1)}(n) = O(1) \tag{1}$$

and

$$\psi_1^{(i)}(n) \le \psi_1^{(i)}(n-1) + \frac{d+1-i}{n-d+1}\,\psi_1^{(i-1)}(n), \qquad i = 2, \ldots, d. \tag{2}$$

We claim that, for $i = 1, 2, \ldots, d$, $\psi_1^{(i)}(n) = O(\log^{i-1} n)$, with the constant of proportionality depending on i and d. This easily follows from (1) and (2) by induction on i. By definition of ψ, this yields $\tau_1^{(i)}(n) = O(n^{d-1} \log^{i-1} n)$. The argument given in the beginning of the proof of the theorem implies that the same asymptotic bound applies to $\tau_k^{(i)}(n)$ for all $0 \le k \le i$, thereby completing the proof of Theorem 2.1. $\square$

Theorem 2.1 already implies that the overall complexity of a single cell is $O(n^{d-1} \log^{d-1} n)$. This bound is further improved in the next section.

3. The Improved Bound

Using a more refined analysis, we next improve the bound derived in the preceding section to $O(n^{d-1} \log n)$. The first step toward this goal is to show that the quantities $\tau_k^{(i)}(n)$, for $k \le i$, are proportional to $\tau_2^{(i)}(n)$. More precisely, we have

Proposition 3.1. *For $i > 2$ and $k = 0, \ldots, i$, we have $\tau_k^{(i)}(p; T) = O(\tau_2^{(i)}(p; T) + n^{d-1})$.*

Proof. By the analysis of Section 2, it suffices to show that $\tau_1^{(i)}(p; T) = O(\tau_2^{(i)}(p; T) + n^{d-1})$. For this, let $((e, R), (f, Q))$ be a popular $(1, i)$-border in $C_p(T)$. The edge e is contained in the intersection of $d - 1$ simplices of T. By definition of popularity, a 3-border (b, W) exists such that $((e, R), (b, W))$ is a $(1, 3)$-border of

$C_p(T)$, and $((b, W), (f, Q))$ is a *popular* $(3, i)$-border of that cell. The number of ways of choosing b is $\binom{d-1}{2}$; we assign $((e, R), (f, Q))$ to $((b, W), (f, Q))$, for an arbitrary choice of b.

Now fix the popular $(3, i)$-border $((b, W), (f, Q))$. Observe that b is a connected polyhedron in 3-space, constituting a single cell in an arrangement of $O(n)$ 2-simplices (i.e., triangles), obtained by intersecting each simplex of T with the 3-flap containing b, and then decomposing each intersection into $O(1)$ triangles. The analysis of Aronov and Sharir [AS, Appendix A] implies that the number $E(b)$ of 1-borders and the number $F(b)$ of 2-borders of b satisfy the relationship

$$E(b) \leq c(F(b) + g(b) + 1),$$

where c is an appropriate absolute constant and $g(b) \geq 0$ is the *genus* of b. Moreover, Aronov and Sharir show that $\sum_b g(b) = O(n^2)$, where the sum is taken over all 3-faces b contained in the same 3-flap.[3]

Now the total number of popular $(1, i)$-borders is clearly bounded by

$$O\left[\sum_{((b, W), (f, Q))} E(b)\right],$$

where the sum is taken over all popular $(3, i)$-borders $((b, W), (f, Q))$ of $C_p(T)$. We thus obtain

$$\tau_1^{(i)}(p; T) = O\left[\sum_{((b, W), (f, Q))} (F(b) + g(b) + 1)\right].$$

However, $\sum F(b) = O(\tau_2^{(i)}(p; T))$, because the left-hand side is proportional to the number of triples of the form (φ, b, f), where φ is a 2-border bounding a 3-face b which bounds a popular i-face f of $C_p(T)$. We can charge each such triple to the appropriate $(2, i)$-border involving φ and f, which is being counted in the right-hand side, and each such border gets charged only a constant number of times, through a constant number of intermediary 3-faces b. Also,

$$\sum g(b) = O(n^2) \cdot \binom{n}{d-3} = O(n^{d-1}).$$

Hence

$$\begin{aligned} \tau_1^{(i)}(p; T) &= O(\tau_2^{(i)}(p; T) + n^{d-1} + \tau_3^{(i)}(p; T)) \\ &= O(\tau_2^{(i)}(p; T) + n^{d-1}), \end{aligned}$$

[3] The bound on $E(b)$ is essentially a consequence of Euler's formula. The bound on $\sum_b g(b)$ is obtained by incrementally adding the simplices of T one by one, and by keeping track of the changes in the overall genus of their union; see [AS] for more details.

because $\tau_3^{(i)}(p; T) = O(\tau_2^{(i)}(p; T))$, by a straightforward charging argument similar to that given in the preceding section. This completes the proof of the proposition. □

Remark. The preceding proof is similar to (albeit somewhat more involved than) a step in the new proof of the zone theorem in arrangements of hyperplanes, as given in [ESS].

Our next step is to produce a recurrence relating $\tau_2^{(i)}$ to $\tau_2^{(i-1)}$ in a manner analogous to that of Section 2. We require the following technical result. It extends similar observations proved in [EGS1] and [HKS], and we believe it to be of independent interest.

Lemma 3.2. *Consider a face f in an arrangement $\mathscr{A}$ of segments in $\mathbb{R}^2$. Add a new segment s, to form a new arrangement $\mathscr{A}'$. In $\mathscr{A}'$, f is decomposed into one or more faces. Let $\mathscr{F}$ be a* subset *of them. Then $|\mathscr{F}|$, the number of faces in $\mathscr{F}$, is at most $1 + p + r$, where p is the number of edges of $\mathscr{A}'$ contained in s and having faces of $\mathscr{F}$ on both sides of them, and r is the number of reflex corners of f that do* not *occur on the boundary of any face of $\mathscr{F}$.*

Proof. The splitting of f into subfaces is caused by the presence of the edges of $s \cap \operatorname{int}(f)$. The case where s avoids f altogether is trivial, so we can assume that each subface touches s. These edges can be classified into four categories:

(i) Extreme edges of s that meet ∂f at just one point.
(ii) Edges that are incident on both sides to subfaces in $\mathscr{F}$.
(iii) Edges that are incident on both sides to subfaces not in $\mathscr{F}$.
(iv) Edges that are incident on one side to a subface in $\mathscr{F}$ and on the other side to a subface not in $\mathscr{F}$.

We ignore edges of type (i), since they do not affect the number of subfaces. The number of edges of type (ii) is p, by definition. If we erase all these edges, the number of subfaces in $\mathscr{F}$ will decrease by at most p, so it suffices to show that this number, after the erasures, is at most $1 + r$. We also erase all edges of type (iii) and merge, for each such edge, the two non-$\mathscr{F}$ subfaces incident to it; note that it is possible for a type-(iii) edge to have the same original subface on both sides.

We are now left with a partition of f into a collection of subfaces, some of which are "surviving" (i.e., made of subfaces in $\mathscr{F}$) while the others are "nonsurviving" (made of subfaces not in $\mathscr{F}$); moreover, no two surviving subfaces (resp. no two nonsurviving subfaces) are adjacent along a remaining edge of $s \cap \operatorname{int}(f)$.

Let G be the bipartite *connectivity graph*, whose nodes are the surviving and nonsurviving subfaces of f, and each of whose arcs connects a surviving subface and a nonsurviving subface adjacent along a type-(iv) edge of $s \cap \operatorname{int}(f)$. Notice that this definition allows for multiple arcs between two nodes. Since f is a connected face, it is clear that G is also connected.

Claim. *Let g be a nonsurviving subface of f with degree* $\deg(g)$ *in* G. *Then its boundary ∂g contains at least* $\deg(g) - 1$ *reflex vertices.*

Assume the claim to be true. If there are v surviving subfaces and t nonsurviving subfaces, then the total number r of reflex vertices in nonsurviving subfaces satisfies the inequality

$$r \geq \sum_{j=1}^{t} \deg(g_j) - t = E - t,$$

where $g_1, \ldots, g_t$ are the nonsurviving subfaces, and E is the number of arcs of G. Since G is connected, we have $E \geq v + t - 1$, which implies that $r \geq v - 1$, which, as noted above, completes the proof of the lemma.

It therefore suffices to prove the above claim. Let g be a nonsurviving subface. Let γ be a connected component of ∂g that touches s, and assume first that γ is the external boundary of g. Orient γ so that as we traverse it g lies to our right. Let $e_1, \ldots, e_k$ be all the edges of type (iv) contained in $\gamma \cap s \cap \operatorname{int}(f)$ and separating g from a surviving face, numbered in the order of their occurrence along γ. The total angle through which we turn as γ is fully traversed is -2π. Pick a point z_i in the interior of e_i, for $i = 1, \ldots, k$, and consider the k arcs into which these points partition γ. We denote by γ_j the portion of γ from z_j to z_{j+1}, for $j = 1, \ldots, k$, where we take $z_{k+1} = z_1$. Let $\pi\varepsilon_j$ be the total turning angle of γ_j as we traverse it from z_j to z_{j+1}. Observe that the quantities ε_j are all integers, and $\sum_j \varepsilon_j = -2$.

It is easy to establish the following properties:

1. If $\varepsilon_j \geq 0$, then γ_j contains at least $\varepsilon_j + 1$ reflex vertices. Indeed, since we cannot turn by more than $+\pi$ at any single vertex, γ_j must have at least ε_j reflex vertices to reach a total turning angle of $\pi\varepsilon_j$. Moreover, since the first turn along γ_j must be a negative (i.e., clockwise) turn, it must be compensated by at least one additional positive turn. Thus γ_j must contain at least $\varepsilon_j + 1$ reflex vertices.
2. If $\varepsilon_j = -1$, then γ_j also contains at least one reflex vertex. Indeed, suppose to the contrary that there is such an arc γ_j with no reflex vertex. Then γ_j starts at z_j along e_j and ends at z_{j+1} along e_{j+1}, where e_j and e_{j+1} are collinear and oppositely oriented. Since γ_j has only right turns, it has the shape of a spiral, and it is easily checked that the turning angle of such a spiral must be negative and with absolute value at least 3π, a contradiction. (See Fig. 1 for an illustration.)

We now have all the tools needed to complete the proof of the claim. Denote by Z (resp. P, N, N_1, and N^*) the set of arcs γ_j with $\varepsilon_j = 0$ (resp. $\varepsilon_j > 0$, $\varepsilon_j < 0$, $\varepsilon_j = -1$, and $\varepsilon_j \leq -2$).

By the properties noted above, the number of reflex corners along γ is at least

$$|Z| + |N_1| + \sum_{\gamma_j \in P} (\varepsilon_j + 1) = |Z| + |N_1| + |P| + \sum_{\gamma_j \in P} \varepsilon_j.$$

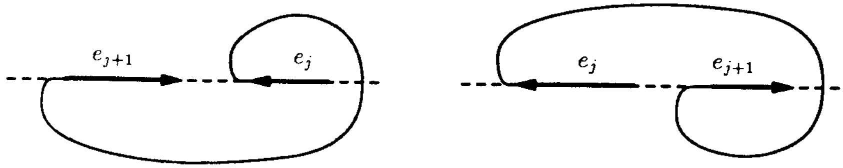

Fig. 1. The proof of property 2 in Lemma 3.2

Since $\sum_j \varepsilon_j = -2$, we have

$$\sum_{\gamma_j \in P} \varepsilon_j = |N_1| + \sum_{\gamma_j \in N^*} |\varepsilon_j| - 2 \geq 2|N^*| + |N_1'| - 2.$$

Hence the number of reflex corners along γ is at least

$$\begin{aligned} |Z| + |N_1| + |P| + \sum_{\gamma_j \in P} \varepsilon_j &\geq |Z| + |N_1| + |P| + 2|N^*| + |N_1| - 2 \\ &= |Z| + |P| + 2|N| - 2 \\ &= k + |N| - 2 \geq k - 1, \end{aligned}$$

since there must be at least one arc with $\varepsilon_j < 0$.

If γ is an internal component of ∂g, the analysis proceeds in much the same way, except that the total turning angle along γ is 2π, instead of -2π. This is even more favorable, because the preceding calculations then imply that the number of reflex corners along γ is at least $k_\gamma + 2$, where k_γ is the number of type-(iv) edges e_j that lie on γ. Summing these inequalities over all components of ∂g, the claim follows, and this completes the proof of the lemma. $\square$

We can finally turn to the derivation of the desired recurrence. Let $d \geq 3$, and fix $i > 2$. In the preceding section we obtained a recurrence for $\tau_1^{(i)}$ in terms of $\tau_1^{(i-1)}$ by removing a simplex σ from T, adding it back, estimating the increase in the number of popular $(1, i)$-borders not contained in σ, in the cell under consideration, and finally averaging this increase over all choices of σ. For technical reasons, we use a modified strategy here: first, we consider $(2, i)$-borders instead of $(1, i)$-borders; next, instead of analyzing the insertion of a single simplex, we start with very few simplices, just enough to define a $(2, i)$-border. We then fix the order in which all the other simplices are to be inserted, one by one, estimate the increase in the number of popular $(2, i)$-borders *contained in the initial set of simplices*, as caused by the insertion of all other simplices, and finally average this increase over all choices of initial sets and insertion sequences. To this end, we define

$$\varphi_k^{(i)}(p; T) = \frac{\tau_k^{(i)}(p; T)}{\binom{n}{d-k}}.$$

This function generalizes the function ψ defined in the proof of Theorem 2.1. As in the preceding section, we also denote by $\varphi_k^{(i)}(n)$ the maximum value of $\varphi_k^{(i)}(p; T)$ over all choices of points p and sets T of n $(d-1)$-simplices in d-space.

The number of k-flaps in $\mathscr{A}(T)$, including empty ones, is exactly $\binom{n}{d-k}$, so that $\varphi_k^{(i)}(p; T)$ can be interpreted as the average contribution to $\tau_k^{(i)}$ by a k-flap, counting empty flaps as well. We now construct a recurrence for $\varphi_2^{(i)}$. For each permutation π of the simplices in T let $\chi(\pi)$ denote the number of popular $(2, i)$-borders of $C_p(T)$ whose 2-face is contained in the 2-flap formed by the first $d-2$ simplices of π. The average value of $\chi(\pi)$ over all choices of π is $\varphi_2^{(i)}(p; T)$, by definition. Thus our approach is to estimate this average value. When we consider only the first $d-2$ simplices of π, which define the relevant 2-flap F, the number of relevant popular $(2, i)$-borders is constant, so it suffices to analyze only the increase in this number as the remaining simplices are added.

Let us first consider a single step, say the last one, in this incremental process. Thus let $f \subseteq F$ be the 2-face of some appropriate popular $(2, i)$-border $((f, R), (g, Q))$ in $\mathscr{A}(T \setminus \{\sigma\})$, and consider what happens to this border when σ is inserted into the arrangement. (Note that, once f and R are fixed, there is only a constant number of choices for g and Q.) Let F^* denote the plane containing F, and let $\mathscr{A}' = \mathscr{A}(T \setminus \{\sigma\}) \cap F^*$ be the arrangement restricted to F^*; note that f is a face of $\mathscr{A}'$. We are interested in the number of popular $(2, i)$-borders of the form $((f', R), (g', Q))$ in the cell under consideration in $\mathscr{A}(T)$, with $f' \subset f$. Consider the collection $\mathscr{F}$ of all such faces f', for a given choice of f, g, R, and Q. By the previous lemma, the increase in the contribution of f to $\tau_2^{(i)}$, namely $|\mathscr{F}| - 1$, is bounded by the number of times two faces of $\mathscr{F}$ are adjacent across a segment of $\sigma \cap F^*$, plus the number of reflex corners of f that do not appear on the boundary of any $f' \in \mathscr{F}$. Repeating this argument for all 2-faces $f \subseteq F$, we conclude that the increase of F's contribution to $\tau_2^{(i)}$, after inserting a single simplex σ, is bounded by $q + r$, where:

- r is the *reflex-corner count*: it is the number of triples (v, R, g), where R is a side of F, v is a reflex corner of a face $f \subseteq F$ in $\mathscr{A}(T)$, and g is an i-face of $\mathscr{A}(T)$ incident to f, such that g is not popular in $\mathscr{A}(T)$ but, when σ is removed, g becomes (a portion of) a popular face.
- q is the *face-pair count*: it is the number of triples (e, R, g), where R is a side of F, g is a popular i-face of $\mathscr{A}(T \setminus \{\sigma\})$, e is an edge contained in $F \cap \sigma$ and incident on both sides to 2-faces $f', f'' \subseteq F$, so that i-faces $g', g'' \subseteq g$ in $\mathscr{A}(T)$ and a side Q of g exist such that $((f', R), (g', Q))$ and $((f'', R), (g'', Q))$ are both popular $(2, i)$-borders in $\mathscr{A}(T)$.

Suppose we start the insertion process with the $d-2$ simplices $\sigma_1, \ldots, \sigma_{d-2}$ defining F and insert all other simplices one by one, say in the order σ_{d-1} through σ_n. Then the total sum of the r-counts is bounded by 2^{d-2} (the total number of sides of F) times $\binom{d-2}{d-i}$ (the number of ways to choose the i-face g) times the total number of reflex corners on all faces of $F \cap \mathscr{A}$, which is clearly $O(n)$. This follows from the fact that, for a fixed side R and i-face g, a reflex corner can be charged only at the step where g is first created (g has to be nonpopular but a portion of a face that was popular just before the step). Notice that this bound is independent of the choice of F and of the order of insertion of the simplices. Thus the same argument holds for the *average* change to reflex corners made in accounting for the increase in $\chi(\pi)$, where the average is taken over all permutations π of $\sigma_1, \ldots, \sigma_n$.

Fix $d - 2 < j \le n$, and consider the step of inserting σ_j into $\mathscr{A}(\{\sigma_1, \ldots, \sigma_{j-1}\})$. Let (e, R, g) be one of the triples that contribute to the q-count portion of the increase in $\tau_2^{(i)}$. Note that e is contained in $\sigma_j \cap F^*$ and bounds on $(i - 1)$-border (an appropriate portion of $g \cap \sigma_j$) that lies between two popular i-borders in $\mathscr{A}(T)$; hence e and this $(i - 1)$-border form a popular $(1, i - 1)$-border, contained in σ_j. We average the number of such $(1, i - 1)$-borders over all permutations with the same set J of the first j simplices. If we further fix the set D of first $d - 2$ simplices and the jth simplex, then $F = \bigcap D$ and $\sigma = \sigma_j$ are fixed, so exactly the same $(1, i - 1)$-borders arise in the jth step of all these permutations. If we now vary σ_j, the last of the j simplices, stepping through every simplex of $J \backslash D$ in turn, each popular $(1, i - 1)$-border of $\mathscr{A}(J)$ whose 1-border lies in F will be charged for at most a constant number of simplices. If we repeat the same process over all choices of a $(d - 2)$-element subset D of J, every popular $(1, i - 1)$-border of $\mathscr{A}(J)$ arises in this process. However, there is a multiplicative factor here, as the charging scheme only depends on the choice of $D \subset J$ and of $\sigma_j \in J \backslash D$ and not on the order in which the members of D and of $J \backslash (D \cup \{\sigma_j\})$ appear in π. Hence the total charge, summed over all possible permutations of the simplices of J, is at most $(d - 2)!$ $(j - d + 1)! \cdot ((d + 1 - i)/2)\tau_1^{(i-1)}(p; J)$ (the factor $(d + 1 - i)/2$ arises for the same reason as in the preceding section). The average charge over all permutations fixing the set J would then be

$$\frac{(d-2)!\,(j-d+1)!}{j!} \cdot \frac{d+1-i}{2}\, \tau_1^{(i-1)}(p; J) \le \frac{(d-2)!\,(j-d+1)!}{j!} \cdot \frac{d+1-i}{2}\, \tau_1^{(i-1)}(j).$$

As the latter quantity is independent of J, the average charge over all permutations of the n, rather than just the first j, simplices is also at most

$$\frac{(d-2)!\,(j-d+1)!}{j!} \cdot \frac{d+1-i}{2}\, \tau_1^{(i-1)}(j).$$

To summarize, the average "q-increase" at the jth step is bounded by

$$O\left(\frac{(d-2)!\,(j-d+1)!}{j!}\, \tau_1^{(i-1)}(j)\right) = O\left(\frac{\tau_1^{(i-1)}(j)}{\binom{j}{d-1}}\right).$$

If we now sum these increases over all j, add the average reflex-corner count, and recall that the contribution of the first $d-2$ simplices to $\chi(\pi)$ is a constant, we conclude that the average number of popular $(2, i)$-borders whose 2-faces are contained in the 2-flap defined by the first $d-2$ simplices in the permutation is

$$\varphi_2^{(i)}(p; T) = O\left(1 + n + \sum_{j=d-1}^{n} \frac{\tau_1^{(i-1)}(j)}{\binom{j}{d-1}}\right) \tag{3}$$

Consider the case $i = 3$ first. From Theorem 2.1, $\tau_1^{(2)}(n) = O(n^{d-1} \log n)$. We thus obtain

$$\begin{aligned}
\varphi_2^{(3)}(n) &= O\left(n + \sum_{j=d-1}^{n} \frac{\tau_1^{(2)}(j)}{(j-d+2)^{d-1}}\right) \\
&= O\left(n + \sum_{j=1}^{n} \log j\right) \\
&= O(n \log n).
\end{aligned}$$

For $i > 3$, using the definition of φ and Proposition 3.1, we rewrite (3) as

$$\begin{aligned}
\varphi_2^{(i)}(n) &= O\left(n + \sum_{j=d-1}^{n} \frac{\tau_2^{(i-1)}(j) + j^{d-1}}{\binom{j}{d-1}}\right) \\
&= O\left(n + \sum_{j=1}^{n} \frac{\varphi_2^{(i-1)}(j)}{j-d+2}\right).
\end{aligned} \tag{4}$$

By inductive hypothesis, $\varphi_2^{(i-1)}(j) = O(j \log j)$. Substituting into (4) yields $\varphi_2^{(i)}(n) = O(n \log n)$, or $\tau_2^{(i)}(n) = O(n^{d-1} \log n)$, as claimed. This completes the proof of our main theorem. □

Before continuing, let us reflect for a moment on the differences between the analysis techniques in the proofs of this section and the preceding one. Roughly speaking, the analysis in Section 2 is a static one—we look at a fixed arrangement, remove a simplex, add it back, estimate the increase in complexity caused by the re-insertion of the simplex, and average these increments over all simplices. The analysis given here is, in contrast, an incremental one—we start with a small set of simplices, fix the order in which the remaining simplices are to be inserted, one by one, and estimate the increase in the initial complexity produced during the entire insertion process. Such a process fares well with a charging scheme where each unit of the increase is charged to a unique feature of the arrangement being constructed, so that, after having incrementally constructed the entire arrangement, it can be concluded that the increase in the complexity under analysis does not exceed the overall number of charged features in the complete arrangement. This approach is known as a "combination-lemma" approach, and has been successfully applied to many combinatorial problems concerning arrangements; see [AS], [EGS1], [EGS2], and [GSS].

The reflex-corner count r is a measure of this type—insertion of a new simplex increases the number of reflex corner charges made by a 2-face f by the number of reflex corners that were on the boundary of f and incident to popular i-faces before the insertion, but are no longer incident to such faces after the insertion. Hence, if we insert simplices one by one until the entire arrangement is constructed, each reflex corner in the arrangement is charged only a constant number of times, so the total contribution of reflex corners to the increase in $\tau_2^{(i)}$ over all faces f and sides R is proportional to the number of such corners in the entire arrangement, which is $O(n^{d-1})$. Notice that this analysis is independent of the order of insertion of simplices.

In contrast, the q-count (refer to the definition in the proof) is difficult to bound in this manner, because, in a single insertion step of some simplex σ, the charge covers only a portion of $\tau_1^{(i-1)}$, namely, only those popular $(1, i-1)$-borders that lie in σ get charged. When we sum these charges over the entire incremental construction process, for a fixed order of inserting simplices, those portions do not add up to anything meaningful. Intuitively, the difficulty is that the quantity $\tau_1^{(i-1)}$ is evanescent—it strongly depends on the current subarrangement and does not seem to relate well to its value at the end of the process.

The idea that has enabled us to estimate these quantities in a meaningful fashion is to average the increases over all possible insertion orders, and to rearrange the averaging in such a way that it will consist of subaverages, in each of which the entire quantity $\tau_1^{(i-1)}$ of some subarrangement will be charged. This averaging process does not affect the r-count, because it is a worst-case count and applies to any insertion order. We thus obtain a mixed charging scheme that can handle both static and incremental charges. We expect that this technique will find additional applications in other contexts.

Corollary 3.3. *The number of faces of all dimensions bounding the* zone *of a $(d-1)$-simplex σ (or a hyperplane) in an arrangement of n $(d-1)$-simplices in $\mathbb{R}^d$ (that is, the collection of all cells crossed by σ) is $O(n^{d-1}\log n)$.*

Proof. All the cells in such a zone can be made into a single cell by cutting each of the given simplices into two subsets by the hyperplane containing σ, leaving a tiny gap between the two pieces, and decomposing each piece, if necessary, into a constant number of simplices. The claim is now immediate from Theorem 1.1. This argument is similar to that made in [EGP$^+$] for two-dimensional arrangements of curves. □

Remarks. (1) The constants of proportionality in the analysis given above are probably much too large, because each step of induction on i multiplies the previous bound by a fairly large constant. It would be interesting to refine the analysis so as to reduce the constants.

(2) The case $i = 3$ was treated separately, because Proposition 3.1 does not apply in this case. Thus our bound follows from the bound $\tau_1^{(2)}(n) = O(n^{d-1} \log n)$ of Theorem 2.1; this was the only use of it in our proof. In particular, to improve Theorem 1.1, it suffices to improve the bound on $\tau_1^{(2)}(n)$. For example, showing that $\tau_1^{(2)}(n) = O(n^{d-1}\alpha(n))$ would close the gap between the upper and lower bounds in Theorem 1.1. The simplest instance of this challenging problem is for $d = 3$; namely, the problem is to bound the number of edges bounding popular faces in a single cell in an arrangement of n triangles in 3-space.

(3) In our argument we used a single point, p, to mark the popular cell. Suppose that, instead, we fix an algebraic surface σ of small constant degree, or a convex surface σ (the boundary of an arbitrary convex set), and consider the collection of all cells constituting the zone of σ in $\mathscr{A}(T)$. Extend the definition of popularity, so that all these cells are now popular. A face would be popular if all of its sides border popular cells, in other words, if all the incident cells are zone cells. The definition is naturally extended to (k, i)-borders, etc. As noted in Section 2, Corollary 2.3 and Lemma 2.4, which provide the "bootstrapping" of our argument, apply to the zone of an algebraic or convex surface as well. An examination of Sections 2 and 3 shows that the entire argument goes through unchanged. A crucial property that makes the analysis work is that popularity can never be regained, in the sense that inroducing new simplices can never change an unpopular feature into a popular one. Put in a different way, if a face g is popular and we remove a simplex not containing g, g becomes (a portion of) a popular face in the reduced arrangement, and this holds for both the case of a single cell and the case of a zone, as is easily checked. This facilitates the derivation of the same recurrence relationship as above, and thus yields the following theorem (see also [APS] for a more detailed analysis that uses a popularity-based argument for estimating the zone complexity of an algebraic or convex surface in the simpler case of an arrangement of hyperplanes):

Theorem 3.4. *The number of faces of all dimensions bounding the zone of an algebraic surface in an arrangement of n $(d-1)$-simplices in $\mathbb{R}^d$ is $O(n^{d-1} \log n)$, where the constant of proportionality depends on d and the degree of the surface. The same statement holds when the algebraic surface is replaced by any convex surface, i.e., the boundary of an arbitrary convex set; with the constant of proportionality depending only on d.*

4. Popularity Extended...

We next extend the definition of popularity and obtain an interesting variant of Theorem 1.1. Instead of concentrating on a single cell (or the zone of some surface), we consider the entire arrangement. For $k < d$, we define an inner k-face to be *popular* if every cell of $\mathscr{A}(T)$ contains an *even* number of the 2^{d-k} sides of the face on its boundary. As an illustration, an edge e in an arrangement of triangles in 3-space is popular if two (not necessarily distinct) cells C_1, C_2 of $\mathscr{A}(T)$ exist so that, as we turn around e, the four sides of e lie in C_1, C_1, C_2, and C_2, respectively, or in C_1, C_2, C_1, and C_2, respectively. This extended definition of popularity is quite natural for facets (a facet is popular if both of its sides lie in the same cell of $\mathscr{A}(T)$), but becomes somewhat less natural as the dimension of the face decreases. Note that this is indeed a generalization of the "old" notion of popularity, as long as $k < d$. Also observe that if a k-face f is popular, for $k < d$, then *all* cells incident to f are nonconvex.

Define *popular borders* in terms of popular faces, as before. Let $\hat{\tau}_k^{(i)}(T)$ be the number of popular (k, i)-borders in $\mathscr{A}(T)$ and $\hat{\tau}_k^{(i)}(n) = \max_{|T|=n} \hat{\tau}_k^{(i)}(T)$. Interestingly, the analysis of Sections 2 and 3 can be adapted, with relatively few modifications, to yield a similar bound on $\hat{\tau}_k^{(i)}$, as long as $i < d$. We briefly sketch the new analysis, focusing on the modifications that have to be made.

Consider first the "bootstrapping" Lemma 2.4. Its proof is easily modified to handle the new kind of popularity. All we need to do is to observe, as in Corollary 2.3, that the union of all nonconvex cells can be decomposed into $O(n^{d-1})$ convex polyhedra with pairwise disjoint interiors, and that the charge each popular face makes in the proof is to a vertex which is locally lowest in some nonconvex cell, and is thus the lowest vertex of one of these convex polyhedra.

Consider next the recurrence derived in Section 2, and recall that we only need it to establish a bound on $\hat{\tau}_1^{(2)}(n)$. Thus let $((f, R), (g, Q))$ be a popular $(1, 2)$-border in $\mathscr{A}(T)$, in the extended sense, and let σ be a simplex of T not containing g. It is easily checked that, when σ is removed, f and g may grow larger, say to edge f' and 2-face g', respectively, but f' and g' still form a popular $(1, 2)$-border in $\mathscr{A}(T \setminus \{\sigma\})$. Indeed, removal of σ may cause some of the cells on which g borders to merge together, but clearly each of these new cells still contains an even number of sides of g'. We can therefore apply the analysis of Section 2, observing that, if σ splits f into two subedges so that, together with an appropriate side and an appropriate piece of g, each of them forms a popular $(1, 2)$-border in the full arrangement, then the point $f \cap \sigma$ and an appropriate edge of $g \cap \sigma$ adjacent to it form a popular $(0, 1)$-border, as follows from the new definition of popularity. Thus the analysis of Section 2 goes through, and we can conclude that $\hat{\tau}_1^{(2)}(n) = O(n^{d-1} \log n)$.

The analysis of Section 3 also applies to the new kind of popularity. Proposition 3.1 applies essentially to any collection of faces of dimension 3 and their borders. The main analysis accounts for the increase in the number of popular $(2, i)$-borders that occurs when a simplex σ is inserted into the arrangement. As argued in Section 3, such an increase can occur only when σ splits a currently popular $(2, i)$-border into several pieces, and then we can charge this increase either to edges of $\sigma \cap f$ that bound popular features on both sides (with respect to some side R of f), or to reflex vertices in subfaces of f which "disappear," in the sense that they no longer bound popular i-faces (again, with respect to some side R). In the extended context, both types of charges are still appropriate. The first type of charge will be to popular $(1, i-1)$-borders, as follows from the new definition of popularity. The second type of charge will be to reflex vertices in subfaces of f that no longer bound popular i-faces (on certain sides of f), and, clearly, no such vertex will be charged more than once for each side of f during the incremental process.

We leave it to the reader to fill in the missing details and to verify that the analysis indeed carries over in its entirety to the new kind of popularity. In summary, we obtain:

Theorem 4.1. *The total number, $\hat{\tau}_k^{(i)}(n)$, of k-borders bounding popular i-borders in an arrangement of n $(d-1)$-simplices in $\mathbb{R}^d$, in the extended sense of popularity just defined, is $O(n^{d-1} \log n)$, for all $0 \le k \le i < d$.*

Why is this an interesting result? First, putting $i = d - 1$ and letting k vary in the above bound, we see that the total complexity (i.e., the total number of faces of all dimensions) of all popular *facets* in the entire arrangement is only $O(n^{d-1} \log n)$, while the total complexity of the arrangement is $O(n^d)$ and the total complexity of all nonconvex cells is in the worst case at least $\Omega(n^{d+1/3})$ (this is an easy generalization of a result of Aronov and Sharir [AS])—a significant gap. Thus, in a very strong sense, most of the complexity of the arrangement comes from "unpopular" facets. Another reason for investigating $\hat{\tau}$ is that it arises in the analysis of the "excess" in arrangements of simplices, as presented in the following section. To facilitate this analysis, we first derive an interesting consequence of Theorem 4.1.

Let f be an inner k-face of $\mathcal{A}(T)$, for $k < d$. We call f *semipopular* if at least two of its sides bound the same cell of $\mathcal{A}(T)$. In what follows we focus on the case $k = 0$, i.e., we consider only semipopular vertices. At the end of the analysis we comment on the easy extension of these results to arbitrary values of k.

Let v be a semipopular (inner) vertex in $\mathcal{A}(T)$, and suppose it is the intersection of d simplices $\sigma_1, \ldots, \sigma_d$. For each σ_i, the hyperplane spanning it partitions space into two half-spaces, and we regard one of them as the positive side of σ_i and the other as the negative side. Let R, R' be two distinct sides of v contained in a common cell C. We can represent R and R' by two respective *sign sequences* $(\delta_1, \delta_2, \ldots, \delta_d)$ and $(\delta'_1, \delta'_2, \ldots, \delta'_d)$, where $\delta_i, \delta'_i \in \{+1, -1\}$ for $i = 1, \ldots, d$, so that $\delta_i = +1$ (resp. -1) if R lies on the positive (resp. negative) side of σ_i, and similarly for δ'_i and R'. Let $j = j(v, R, R')$ be the number of indices i for which $\delta_i = \delta'_i$. Clearly, $0 \le j(v, R, R') \le d - 1$. Let a *semipopular triple* (v, R, R') consist of a semipopular vertex v and two of its sides, R, R', such that (v, R) and (v, R') border the same cell. Our goal is to bound the number of semipopular triples for which $j(v, R, R')$ is equal to some fixed value j, for $j = 0, \ldots, d - 1$. Let us denote the number of such triples by $\zeta_j(T)$, and its maximum value over all collections T of n simplices by $\zeta_j(n)$. Notice that the total number of semipopular triples provides an immediate upper bound on the number of semipopular vertices, since each semipopular triple (v, R, R') is a witness of v being a semipopular vertex.

First observe that $\zeta_{d-1}(n) = O(n^{d-1} \log n)$. Indeed, let (v, R, R') be a semipopular triple with $j(v, R, R') = d - 1$, and let σ be the (unique) simplex containing v so that R and R' lie on different sides of σ. It is easily verified that a facet f on σ exists which has v as a vertex and R, R' as its two sides (or, more precisely, as subsets of its two respective sides Q, Q'). Since (v, R) and (v, R') bound the same cell, it follows that f is popular in the extended sense and that, say, $((v, R), (f, Q))$ is a popular $(0, d-1)$-border. Hence $\zeta_{d-1}(n)$ is proportional to $\hat{\tau}_0^{(d-1)}(n) = O(n^{d-1} \log n)$.

Suppose next that $j < d - 1$. Let (v, R, R') be a semipopular triple with $j(v, R, R') = j$, let $\sigma_1, \ldots, \sigma_d$ be the simplices meeting at v, and let $(\delta_1, \delta_2, \ldots, \delta_d)$ and $(\delta'_1, \delta'_2, \ldots, \delta'_d)$ be the two sign sequences representing R and R', respectively, as above. Without loss of generality, suppose that $\delta_d \neq \delta'_d$. If we remove σ_d from T, the remaining simplices $\sigma_1, \ldots, \sigma_{d-1}$ intersect in a 1-flap which contains an edge g that contains v as an interior point. Since $j(v, R, R') \leq d - 2$, g has two sides, S and S', that lie in the same cell of $\mathscr{A}(T \setminus \{\sigma_d\})$; these sides can be represented by the respective truncated sign sequences $(\delta_1, \delta_2, \ldots, \delta_{d-1})$ and $(\delta'_1, \delta'_2, \ldots, \delta'_{d-1})$, in complete analogy with the representation of the original sides R, R'. Clearly, these sequences agree in j places.

Thus our goal is to bound the number of quadruples (g, S, S', σ) in $\mathscr{A}(T)$, such that:

- (T_1) g is an edge of $\mathscr{A}(T \setminus \{\sigma\})$.
- (T_2) S and S' are two sides of g that lie in the same cell of this subarrangement.
- (T_3) The two sign sequences that represent S and S' have j indices at which they agree.
- (T_4) σ intersects g at an interior point of both σ and g.

Note that such a quadruple can arise from two different semipopular triples, one in which $\delta_d = +1$ and $\delta'_d = -1$ and one in which $\delta_d = -1$ and $\delta'_d = +1$. Hence the number of semipopular triples under consideration is at most twice the number of such quadruples.

Lemma 4.2. *The number of quadruples satisfying* (T_1)–(T_4) *is* $O(n^{d-1} + \zeta_{j+1}(n/2))$.

Proof. We apply an argument adapted from Clarkson and Shor [CS]. Note that a quadruple (g, S, S', σ) is defined in terms of at most $d + 2$ simplices, $d - 1$ of which define the 1-flap containing g, at most two define the endpoints of g (by their intersection with that 1-flap, unless the respective endpoint is an outer vertex), and the $(d + 2)$nd simplex is σ. Draw a random sample R of $r = n/2$ simplices of T, and consider the number of triples (g, S, S') such that:

- (R_1) g is an edge of $\mathscr{A}(R)$, contained in the intersection of $d - 1$ simplices of R.
- (R_2) S, S' are two sides of g, represented by a pair of sign sequences $(\delta_1, \ldots, \delta_{d-1})$, $(\delta'_1, \ldots, \delta'_{d-1})$, as above, such that both sides lie in the same cell of $\mathscr{A}(R)$.
- (R_3) The number of indices i for which $\delta_i = \delta'_i$ is j.

We can charge each such triple to a triple (v, Q, Q'), where v is one of the endpoints of g, which we assume for now to be an inner vertex formed by the intersection of the 1-flap containing g with some simplex $\sigma^* \in R$, and Q, Q' are the two sides of v obtained by intersecting the half-space containing g and bounded by the hyperplane spanning σ^*, with S and S', respectively. Clearly, Q and Q' lie in the same cell of $\mathscr{A}(R)$, and their corresponding sign sequences agree in $j + 1$ components—j of which are "inherited" from the sequences of S and S', and the $(j + 1)$st one corresponds to σ^* and signifies the fact that Q and Q' lie on the same side of σ^*. Clearly, no triple (v, Q, Q') is charged more than a constant number of times, which implies that the number of triples (g, S, S') as above is bounded by $\zeta_{j+1}(n/2)$, by definition. If v is an outer vertex, we simply charge (g, S, S') to v. Since there are only $O(r^{d-1})$ outer vertices in $\mathscr{A}(R)$, it follows that the number of triples (g, S, S') charged to outer vertices is also $O(r^{d-1}) = O(n^{d-1})$.

Next consider the probability that a quadruple (g, S, S', σ) in $\mathscr{A}(T)$, satisfying properties (T_1)–(T_4), will give rise to the corresponding triple (g, S, S') in $\mathscr{A}(R)$, so that the triple satisfies properties (R_1)–(R_3). We claim that for this to happen it is necessary and sufficient that the at most $d + 1$ simplices defining g are chosen in R and that σ is not chosen in R. This condition is clearly necessary. It is sufficient because the two sides S, S' lie in the same cell of $\mathscr{A}(T \backslash \{\sigma\})$, and so they clearly do the same in the smaller arrangement $\mathscr{A}(R)$. If we denote by t the number of simplices determining (g, S, S', σ) (where $d \leq t \leq d + 2$), then the probability for this condition to occur is easily seen to be

$$\binom{n-t}{r-t+1} \bigg/ \binom{n}{r} \geq c,$$

for some constant $c > 0$ depending on d (see also [CS]). Hence the expected number τ of triples (g, S, S') in $\mathscr{A}(R)$ satisfying (R_1)–(R_3) and induced by the quadruples (g, S, S', σ) satisfying (T_1)–(T_4) is at least $c\xi$, where ξ is the number of such quadruples. Hence, $c\xi \leq \tau \leq O(n^{d-1}) + \zeta_{j+1}(n/2)$, so $\xi = O(n^{d-1} + \zeta_{j+1}(n/2))$, as asserted. □

The preceding lemma yields the chain of inequalities

$$\zeta_0(n) \leq a\left(n^{d-1} + \zeta_1\left(\frac{n}{2}\right)\right), \qquad \zeta_1(n) \leq a\left(n^{d-1} + \zeta_2\left(\frac{n}{2}\right)\right), \qquad \ldots$$

$$\zeta_{d-2}(n) \leq a\left(n^{d-1} + \zeta_{d-1}\left(\frac{n}{2}\right)\right) = O(n^{d-1} \log n),$$

for some constant a depending on d, which implies:

Theorem 4.3. *The total number $\zeta(T)$ of semipopular triples (v, R, R') in an arrangement of n $(d-1)$-simplices in d-space is $O(n^{d-1} \log n)$. Thus the number of semipopular vertices is $O(n^{d-1} \log n)$.*

Theorem 4.3 can be extended to yield a similar bound on the number of k-faces that have at least two sides lying in a common cell of $\mathscr{A}(T)$, for any $2 \leq k < d$. Let f be such a semipopular k-face, and let v be any vertex of f. If v is an outer vertex, we charge f to v and observe that only $O(n^{d-1})$ such charges will be made. If v is an inner vertex, then it is clearly also semipopular, as witnessed by its two sides that are contained in the respective two sides of f that lie in a common cell. Hence we have:

Theorem 4.4. *The total complexity of semipopular faces of all dimensions in an arrangement of* n $(d-1)$*-simplices in* d*-space is* $O(n^{d-1}\log n)$.

5. ... Leading to Excesses

In this section we introduce the notion of the *excess* of an arrangement $\mathscr{A}$ of a set T of n $(d-1)$-simplices in $\mathbb{R}^d$, $d \geq 2$, in general position. Fix $0 \leq k \leq d-1$, and let f be a k-flap formed by the intersection of the relative interiors of some $d-k$ simplices of T. If some cell C in $\mathscr{A}(T)$ has a k-border of the form (f', R), where $f' \subseteq f$ and R is any one of the 2^{d-k} sides of f, we put $p(f, C) = 1$ and say that f is *present* on the boundary of C; otherwise put $p(f, C) = 0$. In either case we define $\kappa(f, C)$ to be the number of k-borders (f', R) of C, with $f' \subseteq f$. The *excess* $\varepsilon(f, C)$ of C relative to f is defined as $\kappa(f, c) - p(f, c)$. The k-excess of a cell C is $\varepsilon_k(C) = \sum_f \varepsilon(f, C)$, where the sum extends over all k-flaps f, and the k-excess of the entire arrangement is $\varepsilon_k(T) = \sum_C \varepsilon_k(C)$.

First, notice that, by definition, $\varepsilon_k(C_p(T)) < c(p, T)$, for any point p not on a simplex of T. Therefore, $\varepsilon_k(C) = O(n^{d-1}\log n)$, for any cell C and any k. Moreover, $\varepsilon_k(C)$ is smaller than the number of k-borders of C by at most $\binom{n}{d-k}$ and, as is easily seen, families of n simplices exist in which, say, $\varepsilon_1(C) = \Omega(n^{d-1}\alpha(n))$ for some cell C in their arrangement. This more or less settles the question of the worst-case excess of a single cell in an arrangement. A more challenging problem, though, is to bound the overall excess of the entire arrangement. Let $\varepsilon_k^{(d)}(n) = \max_T \varepsilon_k(T)$, where the maximum is taken over all collections T of n $(d-1)$-simplices in general position in d-space. The case $k = 0$ is easy, because the excess of a vertex is nonzero if and only if the vertex is semipopular. Hence $\varepsilon_0^{(d)}(n) = O(n^{d-1}\log n)$. We next show:

Lemma 5.1. *For* $d \geq 3$, $\varepsilon_1^{(d)}(n) = O(n^{d-1}\log^2 n)$. *For* $d = 2$, $\varepsilon_1^{(2)}(n) = O(n \log n)$.

Proof. Fix any $\sigma \in T$, and consider the 1-flap f formed by some $d-1$ simplices of $T\setminus\{\sigma\}$. Suppose σ is inserted back into the arrangement. The increase in the contribution of f to $\varepsilon_1(T)$ can be bounded as follows: Fix some cell C of $\mathscr{A}(T\setminus\{\sigma\})$ which is bounded by at least one 1-border (g, R), with $g \subseteq f$. When σ is re-inserted, C may be split into several subcells in $\mathscr{A}(T)$. Let $g \subseteq f$ be an edge that is present on the boundary of C (in $\mathscr{A}(T\setminus\{\sigma\})$), and let $t \geq 1$ be the number of sides R such that (g, R) is a 1-border of C. It is easily verified that if $\sigma \cap g = \varnothing$, then g does not contribute anything to the increase in 1-excess. If $\sigma \cap g \neq \varnothing$, then each of the t former 1-borders (g, R) is now split into two subborders (g^+, R), (g^-, R), yielding a total of $2t$ subborders, each bounding some subcell of C. Let q be the number of subcells into which C has been split in $\mathscr{A}(T)$ and which are still bordered by pieces of g. It is easily checked that the increase in the 1-excess generated by f within C, created when σ is re-inserted, is at most $(2t - q) - (t - 1) = t - q + 1$, so we are only interested in cases where $t \geq q$, for only then does the 1-excess go up. Note that the vertex $g^* = \sigma \cap f$ has $2t$ sides that it "inherits" from the borders (g^+, R), (g^-, R) as above, and that these sides lie in at most q distinct subcells of C. It follows that we can generate at least $2t - q$ distinct semipopular triples of the form (g^*, Q, Q'), where Q and Q' are two sides of g^* that lie in the same subcell of C. Since $2t - q \geq t - q + 1$, it follows that we can charge these semipopular triples for the increase in 1-excess under consideration. Moreover, when we repeat this argument over all simplices σ (and all corresponding cells C), it is easily seen that no semipopular triple gets charged more than d times, once for every simplex containing its vertex. Hence, as in the proof of Theorem 2.1, we obtain the following recurrence for ε_1:

$$\begin{aligned}(n - d + 1)\varepsilon_1(T) &\leq \sum_{\sigma \in T} \varepsilon_1(T \backslash \{\sigma\}) + d\zeta(T) \\ &= \sum_{\sigma \in T} \varepsilon_1(T \backslash \{\sigma\}) + O(n^{d-1} \log n).\end{aligned}$$

Applying the same analysis to this recurrence, we obtain $\varepsilon_1^{(d)}(n) = O(n^{d-1} \log^2 n)$. In the plane the analysis gives the improved bound $\varepsilon_1^{(2)}(n) = O(n \log n)$, because in this case $\zeta(T) = O(\hat{\tau}_0^{(1)}(T)) = O(n)$.

Remark. In the plane we thus have $\varepsilon_1^{(2)}(n) = O(n \log n)$ and $\Omega(n\alpha(n))$. We conjecture that in fact $\varepsilon_1^{(2)}(n) = \Theta(n\alpha(n))$. More generally, we would like to know if the 1-excess of the entire arrangement can be asymptotically larger than the 1-excess of a single cell. The same question applies to k-excesses, which we consider next:

Theorem 5.2. *For $d \geq 3$ and $1 \leq k \leq d - 1$, $\varepsilon_k^{(d)}(n) = O(n^{d-1} \log^2 n)$.*

Proof. We prove the claim by induction on k. The base case $k = 1$ has just been established. Suppose then that $k \geq 2$ and that the claim is true for $k - 1$. As above, fix any $\sigma \in T$, and consider a k-flap f formed by some $d - k$ simplices of $T \backslash \{\sigma\}$. Suppose σ is inserted back into the arrangement. The increase in the contribution of f to $\varepsilon_k(T)$ can be bounded as follows: Fix a cell C in $\mathscr{A}(T \backslash \{\sigma\})$, and suppose that there are m k-borders (g, R) on the boundary of C, where $g \subseteq f$ and R is a side of f. Then the "old" k-excess that f induces in C is $m - 1$. When σ is added back, C is split into several subcells, say $C_1, \ldots, C_q$. Similarly, any k-border (g, R) on the boundary of C, with $g \subseteq f$ and R a side of f, may be split into several subborders, which may appear on the boundaries of any new subcell C_j. Let a_j denote the number of k-borders of this kind on the boundary of C_j, for $j = 1, \ldots, q$, and put $a = \sum_{j=1}^q a_j$. Then the "new" k-excess that f induces in all subcells of C is $\sum_{j=1}^q (a_j - 1) = a - q$. (We ignore here subcells for which $a_j = 0$, since these subcells do not involve f at all; thus, in what follows, q denotes the number of subcells of C on whose boundaries f is actually present.) Let f^* be the $(k - 1)$-flap $f \cap \sigma$. Observe that, for a fixed side R, the increase in the number of k-borders (g, R) as above is controlled by the number of times the boundary ∂C of C is "cut" by f^*. More precisely, if $\partial C \cap f^*$ contains t_R $(k - 1)$-faces which split k-borders

(g, R) as above (for the fixed side R), then the number of relevant k-borders will have increased by at most t_R. Hence $a \le m + \sum_R t_R$, so the new k-excess in question is at most $m + t - q$, where $t = \sum_R t_R$.

If $t - q \le -1$, then the new excess is not larger than the old excess, so there is no increase in excess to account for. So suppose that $t - q \ge 0$. The resulting increase in k-excess is thus $t - q + 1 \ge 1$, and we need to find some way of charging for this increase. Notice that, since $q \ge 1$, we also have in this case $t \ge 1$, which implies $t - q + 1 \le 2t - q$.

With each side R of f for which $t_R > 0$, we can associate $2t_R$ $(k-1)$-borders of the form (h, R'), where h is a $(k-1)$-face on f^* which has split a k-border (g, R) as above, and R' is one of the two sides into which R is split by σ. Thus we get a total of $2t$ $(k-1)$-borders along f^*, each bounding one of the q subcells C_j of C. Hence their total contribution to the $(k-1)$-excess in the full arrangement is at least $2t - q$. In other words, we can charge the increase in k-excess within the cell C, as caused by the re-insertion of σ, to the $(k-1)$-excess within the subcells into which C is split; more precisely—to that portion of this $(k-1)$-excess generated by $(k-1)$-borders that lie on σ. Hence, summing over all cells C and simplices σ, we obtain the recurrence

$$\begin{aligned}(n - d + k)\varepsilon_k(T) &\le \sum_{\sigma \in T} \varepsilon_k(T \backslash \{\sigma\}) + (d - k + 1)\varepsilon_{k-1}(T) \\ &= \sum_{\sigma \in T} \varepsilon_k(T \backslash \{\sigma\}) + O(n^{d-1} \log^2 n),\end{aligned}$$

by induction hypothesis. Since we want to solve this recurrence only for $k \ge 2$, its solution is easily seen to be $\varepsilon_k^{(d)}(n) = O(n^{d-1} \log^2 n)$. This completes the proof of the theorem. □

The notion of excess provides a new measure for the complexity of an arrangement. The bound given above shows, informally, that the $\Omega(n^d)$ worst-case complexity of the entire arrangement cannot be accounted for by the repetition of k-flaps along the boundary of the same cell. This large complexity is thus due either to the number of cells being large, or to the number of distinct k-flaps on each cell boundary being large. Note that the extreme case, where all simplices are actually hyperplanes, gives the largest overall complexity of the arrangement and the lowest overall excess (namely 0).

We have included the above analysis of excess in this paper for three main reasons:

(a) The intrinsic interest in this new measure of complexity in arrangements, as well as its connection with the related complexity of popular and semipopular features.
(b) The fact that the analysis of excess is similar in spirit to the analysis of the complexity of a single cell.
(c) The usefulness of this notion in the analysis of other complexity problems in simplex arrangements, such as the complexity of multiple cells, which we discuss next.

6. Complexity of Multiple Cells in Arrangements of Simplices

We now apply the results concerning excesses to derive a nontrivial bound on the complexity of m distinct cells in an arrangement of n $(d-1)$-simplices in d-space. Specifically, we show:

Theorem 6.1. *The total number of $(d-1)$-borders on the boundaries of m distinct cells in an arrangement of n $(d-1)$-simplices in d-space is*

$$O(n^{d-1}\log^2 n + m^{1/2}n^{d/2}\log^{1/2} n).$$

Proof. Let T be the given collection of simplices, and let the m given cells be $C_1, \ldots, C_m$. For cell C in $\mathscr{A}(T)$, denote the number of $(d-1)$-borders on the boundary ∂C of C by $\xi(C)$, and the number of different simplices appearing on ∂C by $\mu(C)$. By definition of excess, we have

$$\sum_{i=1}^{m} \xi(C_i) \leq \varepsilon_{d-1}(T) + \sum_{i=1}^{m} \mu(C_i) = O(n^{d-1}\log^2 n) + \sum_{i=1}^{m} \mu(C_i).$$

Claim. *$\sum_C \mu(C)^2 = O(n^d \log n)$, where the sum extends over all cells C of $\mathscr{A}(T)$.*

Indeed, the sum is proportional to $\sum_C N(C)$, where we sum over all cells C, and where $N(C)$ is the number of pairs (σ, σ') of simplices both appearing on ∂C. Rearranging this sum, we obtain $\sum_C N(C) = \sum_\sigma K(\sigma)$, where now we sum over all simplices σ, and where $K(\sigma)$ is the number of pairs (C, σ') of cells C and simplices σ' such that σ appears on ∂C and σ' also appears on that boundary. Let Z_σ denote the collection of all these cells C; these are precisely the cells in the zone of σ in $\mathscr{A}(T)$. Clearly,

$$K(\sigma) \leq \sum_{C \in Z_\sigma} \mu(C) \leq \sum_{C \in Z_\sigma} \xi(C).$$

In other words, $K(\sigma)$ is bounded by the complexity of the zone of σ, measured in terms of the number of $(d-1)$-borders bounding cells in the zone. By Corollary 3.3, this complexity is $O(n^{d-1}\log n)$ and, summing over all simplices σ, we obtain:

$$\sum_C \mu(C)^2 = O(n^d \log n),$$

as claimed.

Now we apply the Cauchy–Schwarz inequality to obtain

$$\sum_{i=1}^{m} \mu(C_i) \leq m^{1/2}\left[\sum_{i=1}^{m} \mu(C_i)^2\right]^{1/2} \leq m^{1/2}\left[\sum_C \mu(C)^2\right]^{1/2} = O(m^{1/2}n^{d/2}\log^{1/2} n),$$

and this clearly completes the proof of the theorem. □

Remarks. (1) Aronov *et al.* [AMS] show that, in an arrangement of n hyperplanes in $\mathbb{R}^d$, the number of faces of all dimensions bounding m distinct cells is $O(m^{1/2}n^{d/2}\log^{(\lfloor d/2 \rfloor - 1)/2} n)$. They also provide a lower bound on the maximum possible face count in m distinct cells, which is close to the upper bound, and for many values of m and n is $\Omega(m^{1/2}n^{d/2})$. Theorem 6.1 compares favorably with their upper bound (except, of course, that we are only counting facets). Their lower bound counts lower-dimensional faces, and so does not seem to say anything about the tightness of our result.

(2) An obvious open problem is thus to extend the preceding analysis to obtain similar bounds on the total number of faces of all dimensions bounding m distinct cells in an arrangement of n simplices in $\mathbb{R}^d$. What is missing is an extension of the "sum-of-squares" claim (as given in [AMS] for the case of hyperplanes), which will include faces of lower dimension as well. Unfortunately, a straightforward extension of $\mu(C)$ to include the number of different k-flaps, for all $0 \le k < d$, occurring on ∂C, does not yield a satisfactory bound for simplex arrangements—it is easy to construct an arrangement in which there is a cell C with $\mu(C)$ (in the extended sense) equal to $\Omega(n^{d-1})$, so $\sum_C \mu(C)^2 = \Omega(n^{2d-2})$ in the worst case, which is much too large for the purposes of our argument.

7. Other Applications and Discussion

Our estimate of the complexity of a single cell in an arrangement of n $(d-1)$-simplices in d-space has several applications. One such application, already mentioned in the Introduction, is to the problem of translational motion planning in three dimensions. Specifically, let B be a rigid polyhedral body free to translate in three-dimensional space bounded by a collection of polyhedral obstacles. Given an initial free placement O of B, we wish to compute the space of all free placements of B reachable from O by a collision-free translational motion. Suppose the obstacles are bounded by a total of n faces and B is bounded by k faces. For simplicity, assume that the surfaces of B and of the obstacles are triangulated. We compute the Minkowski difference of the obstacles and B. The boundary of this difference is formed by $O(kn)$ possibly intersecting triangles and parallelograms, each being the Minkowski difference of an i-face of some obstacle and of a $(2-i)$-face of B, for $i = 0, 1, 2$. The space of free placements of B reachable from O corresponds to the cell containing the point representing O in the arrangement of these $O(kn)$ triangles and parallelograms. We thus conclude:

Corollary 7.1. *The combinatorial complexity of the space of free placements of a polyhedron B with k faces, translating amidst a collection of polyhedral obstacles bounded by a total of n faces, which can be reached from a fixed placement of B by a collision-free motion, is* $O((kn)^2 \log(kn))$.

Clearly, we can extend this result to more complex instances of motion planning, involving more than three degrees of freedom, as long as all the collision-constraint surfaces (each representing placements of the system in which contact is being made between some feature of the moving system and some obstacle feature, or between two system features) are formed by appropriate simplices. For example, this happens in the case of a nonrigid polyhedral body B having any number of prismatic joints and free to translate amidst polyhedral obstacles in 3-space as above, or to any collection of such independently moving bodies. In any of these cases the above results imply that the combinatorial complexity of the portion of free configuration space, reachable from a given initial system configuration, is no more than $O(n^{d-1} \log n)$, where d is the number of degrees of freedom of the system and n is the number of simplices in a triangulation of the constraint surfaces.

Next consider the problem of efficiently computing one cell in an arrangement of triangles in 3-space:

Corollary 7.2. *Given $\delta > 0$, a set of n triangles in 3-space, and a point not on any of them, the cell containing this point in the arrangement formed by the triangles can be computed in randomized expected time $O(n^{2+\delta})$, where the constant of proportionality depends on δ.*

Proof. Aronov and Sharir [AS] present such a randomized algorithm with expected running time $O(c_3(n)n^\delta)$. The claim now follows from Theorem 1.1. □

Corollary 7.3. *Motion planning for a polyhedron translating amidst polyhedral obstacles in 3-space, as formulated above, can be performed in randomized expected time $O((kn)^{2+\delta})$, for any $\delta > 0$.*

Remarks. (1) A simplar algorithm would follow for the case of n simplices in d-dimensional space, with an expected running time close to $O(n^{d-1})$, provided a generalization of the Slicing Theorem of [AS] to higher dimension can be obtained; see the discussion following the proof of the Chopping Theorem, and the technical details of the algorithm of [AS]. There is no such generalization known at present.

(2) The algorithm in [AS] is based on random sampling. It would be interesting to obtain an efficient randomized algorithm which would add the triangles (or simplices) one by one in random order and maintain and update the cell containing the marking point p as it goes. Even more challenging is the problem of designing an efficient and simple deterministic algorithm for this problem.

We next consider possible extensions of our result and related open problems. The first open problem is to tighten the remaining small gap between our upper bound of $O(n^{d-1} \log n)$ for $c_d(n)$, and the lower bound of $\Omega(n^{d-1}\alpha(n))$ established in [PS]. We conjecture, as in previous papers, that $c_d(n) = \Theta(n^{d-1}\alpha(n))$. A greater challenge is to settle the analogous conjecture for the total excess $\varepsilon_k^{(d)}(n)$ in arrangements of simplices, which is open even in the plane. Another open problem is to obtain sharp upper bounds for the total complexity of m distinct cells in such an arrangement, extending the notion of complexity to include faces of all dimensions bounding those cells.

Throughout this paper we have assumed that the input simplices are in general position. As we have already remarked, this assumption does not affect the asymptotic upper bounds on the number of faces bounding a single cell, or multiple cells, in an arrangement of simplices, as this complexity is maximized when the simplices are in general position. Other combinatorial quantities that were considered in this paper are much more sensitive to degeneracies. For example, the definition of a popular vertex has to be changed to account for the possibility of having more than 2^d "sides." Should all of these sides be required to lie in a given cell if a vertex is to be considered popular? Besides notational problems that would arise if degenerate configurations are allowed, there will be significant difficulties in carrying out some of the inductive proofs which crucially depend on the fact that a k-flap is contained in exactly $d - k$ simplices. In fact, some of the bounds derived in this paper are blatantly false in degenerate configurations, notably Theorem 4.3, which counts semipopular triples: Consider the arrangement in which all n simplices pass through a single vertex v and are otherwise in general position. The number of "sides" of v will be $\Theta(n^{d-1})$ and the number of semipopular triples $\Theta(n^{2d-2})$, which is much larger than the upper bound claimed by the theorem. In fact, we pose the following challenging open question: What are the appropriate generalizations of the combinatorial complexity measures introduced in this paper so that the bounds proven for arrangements in general position continue to hold when this assumption is removed?

Finally, a natural generalization of the problem studied above is to bound the complexity of a single cell in an arrangement of n surfaces or surface patches, assuming the surfaces are all algebraic of low degree and that their boundaries are algebraic curves or surfaces of low degree as well. It is conjectured that the complexity of such a cell is close to $O(n^{d-1})$, perhaps off by a factor that depends on the inverse Ackermann function $\alpha(n)$. This problem, even in the special case $d = 3$, appears to be substantially more difficult than the case of triangles. Even the problem of bounding the complexity of the lower envelope of such surface patches in three dimensions appears to be very difficult, and near-quadratic bounds are known only in a few special cases (including that of triangles); see [PS] and [SS2]. Solving this problem would yield improved bounds on the general motion-planning problem with nonlinear constraints, similar to what has been noted above. We mention a recent result of Halperin [H1], which obtains a near-quadratic bound for single cell complexity in three-dimensional arrangements of surfaces that arise in certain motion-planning problems with three degrees of freedom. These results are further extended in [H2], and some of these extensions apply the techniques developed in this paper.

References

[AA] P. K. Agarwal and B. Aronov, Counting facets and incidences, *Discrete Comput. Geom.* 7 (1992), 359–369.

[AMS] B. Aronov, J. Matoušek, and M. Sharir, On the sum of squares of cell complexities in hyperplane arrangements, *Proc. 7th Symp. on Computational Geometry*, 1991, pp. 307–313. To appear in *J. Combin. Theory Ser. A*.

[APS] B. Aronov, M. Pellegrini, and M. Sharir, On the zone of an algebraic surface in a hyperplane arrangement, *Discrete Comput. Geom.* **9** (1993), 177–186.

[AS] B. Aronov and M. Sharir, Triangles in space or building (and analyzing) castles in the air, *Combinatorica* **10**(2) (1990), 137–173.

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

[EGP+] H. Edelsbrunner, L. Guibas, J. Pach, R. Pollack, R. Seidel, and M. Sharir, Arrangements of curves in the plane: topology, combinatorics, and algorithms, *Theoret. Comput. Sci.* **92** (1992), 319–336.

[EGS1] H. Edelsbrunner, L. Guibas, and M. Sharir, The complexity and construction of many faces in arrangements of lines and of segments, *Discrete Comput. Geom.* **5** (1990), 161–196.

[EGS2] H. Edelsbrunner, L. Guibas, and M. Sharir, The complexity of many cells in arrangements of planes and related problems, *Discrete Comput. Geom.* **5** (1990), 197–216.

[ESS] H. Edelsbrunner, R. Seidel, and M. Sharir, On the zone theorem for hyperplane arrangements, *SIAM J. Comput.* **22** (1993), 418–429.

[GSS] L. Guibas, M. Sharir, and S. Sifrony, On the general motion planning problem with two degrees of freedom, *Discrete Comput. Geom.* **4** (1989), 491–521.

[H1] D. Halperin, On the complexity of a single cell in certain arrangements of surfaces in 3-space, *Proc. 7th ACM Symp. on Computational Geometry*, 1991, pp. 314–323.

[H2] D. Halperin, Algorithmic Motion Planning, via Arrangements of Curves and of Surfaces, Ph.D. Dissertation, Tel Aviv University, August 1992.

[HKS] D. Huttenlocher, K. Kedem, and M. Sharir, The upper envelope of Voronoi surfaces and its applications, *Discrete Comput. Geom.* **9** (1993), 267–291.

[PS] 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.

[SS1] J. T. Schwartz and M. Sharir, On the piano movers' problem: II. General techniques for computing topological properties of real algebraic manifodls, *Adv. in Appl. Math.* **4** (1983), 298–351.

[SS2] J. T. Schwartz and M. Sharir, On the 2-dimensional Davenport–Schinzel problem, *J. Symbolic Comput.* **10** (1990), 371–393.

Received November 20, 1992, and in revised form January 3, 1994.