

The Union of Balls and Its Dual Shape*

H. Edelsbrunner

Department of Computer Science, University of Illinois at Urbana-Champaign,
Urbana, IL 61801, USA
edels@cs.uiuc.edu

Abstract. Efficient algorithms are described for computing topological, combinatorial, and metric properties of the union of finitely many spherical balls in \mathbb{R}^d . These algorithms are based on a simplicial complex dual to a decomposition of the union of balls using Voronoi cells, and on short inclusion–exclusion formulas derived from this complex. The algorithms are most relevant in \mathbb{R}^3 where unions of finitely many balls are commonly used as models of molecules.

1. Introduction

The primary object studied in this paper is the union of finitely many spherical d -balls in \mathbb{R}^d . One of the motivations for our considerations is their widespread use in computational biology, where a molecule is frequently modeled as the union of 3-balls in \mathbb{R}^3 [3], [22]. Each atom is represented by a ball whose size is determined by its van der Waals radius. This model is referred to as the *space-filling diagram* of the molecule. As is seen later, this diagram is related to a certain polytope, called the *dual shape* of the diagram. This paper is part of a project that studies such shapes and their applications to problems in science. A declared goal of the project is the implementation of shapes and some of their useful functions. It is therefore essential to find simple algorithms so that the implementation produces a compact system of programs. At the same time, efficiency is essential because typical applications involve thousands of balls.

* This work is supported by the National Science Foundation, under Grant ASC-9200301, and the Alan T. Waterman award, Grant CCR-9118874. Any opinions, findings, conclusions, or recommendations expressed in this publication are those of the author and do not necessarily reflect the view of the National Science Foundation.

Our study of d -balls requires a variety of concepts whose origins lie in the areas of convex geometry, geometric algorithms, and algebraic topology. We make essential use of Voronoi diagrams and Delaunay simplicial complexes [5], [7], [24] and of polytopes arising as underlying spaces of certain subcomplexes of Delaunay simplicial complexes [9], [11], [13]. Topological concepts such as homotopy equivalence and homology groups [18], [23] are instrumental in uncovering the close relationship between these geometric diagrams.

The outline of this paper follows. Section 2 introduces the basic geometric diagrams used in our study. Section 3 establishes the homotopy equivalence of the union of balls and its dual shape; it implies effective algorithms for computing the homology groups of the union. Section 4 shows how the topological insights lead to an efficient algorithm for counting the faces of the union of a set of balls. Section 5 studies the Euler relation for convex polyhedra and viewpoints. Based on these relations, Section 6 gives short inclusion–exclusion formulas for measuring the union of balls. Section 7 considers simplices defined by independent sets of d -balls. Section 8 derives another set of inclusion–exclusion formulas which are decomposable and, among other things, can measure voids formed by the union. Section 9 concludes the paper.

2. The Union of Balls and Related Diagrams

This section introduces various geometric concepts defined for a finite collection of balls, with the aim to develop tools that can enhance our understanding of the union of these balls.

Basic Definitions

Let $|xz|$ denote the Euclidean distance between two points $x, z \in \mathbb{R}^d$. A subset $b \subseteq \mathbb{R}^d$ is a d -ball if there is a point $z \in \mathbb{R}^d$ and a real $\rho > 0$ so that $b = \{x \in \mathbb{R}^d \mid |xz| \leq \rho\}$; z is the *center* and ρ is the *radius* of b . For $0 \leq k \leq d - 1$, a k -ball is the intersection of a $(k + 1)$ -ball b with a hyperplane that contains its center but not b itself. A k -sphere is the (relative) boundary of a $(k + 1)$ -ball b . The center and radius are inherited from b . For example, a 0-ball is a point, a 1-ball is a line segment, and a 2-ball is a disk. A 0-sphere is a pair of points, a 1-sphere is a circle, and a 2-sphere is what in \mathbb{R}^3 is commonly called a sphere.

Besides balls and spheres we consider simplices in \mathbb{R}^d . For $0 \leq k + 1 \leq d + 1$, a k -simplex, σ , in \mathbb{R}^d is the convex hull of $k + 1$ affinely independent points. The *dimension* of σ is $\dim \sigma = k$. The convex hull of any $0 \leq l + 1 \leq k + 1$ of these points is an l -simplex and a *face* of σ . For example, the only (-1) -simplex is \emptyset , a 0-simplex is a point, a 1-simplex is an edge, a 2-simplex is a triangle, and a 3-simplex is a tetrahedron. A tetrahedron has $2^4 = 16$ faces, namely \emptyset , four vertices, six edges, four triangles, and itself.

Abstract and geometric simplicial complexes play an important role in this paper. An *abstract simplicial complex* is a finite collection A of sets so that $X \in A$ and

$Y \subseteq X$ implies $Y \in A$. The *vertex set* of A is $\text{vert } A = \{x \in X \mid X \in A\}$. A (*geometric simplicial complex*) is a finite collection G of simplices that satisfy the following two conditions. First, if $\sigma \in G$ and σ' is a face of σ , then $\sigma' \in G$. Second, if $\sigma_1, \sigma_2 \in G$, then $\sigma_1 \cap \sigma_2$ is a face of both. As a general policy, \emptyset is considered a face of every simplex and is included in all simplicial complexes. The *underlying space* or *polytope* of G is $|G| = \bigcup_{\sigma \in G} \sigma$. A subset $H \subseteq G$ is a *subcomplex* of G if it is a simplicial complex itself, that is, it satisfies the first condition. A special subcomplex is the *k-skeleton* $G^{(k)} = \{\sigma \in G \mid \dim \sigma \leq k\}$. G is a *geometric realization* of A if there is a bijection $\phi: \text{vert } A \rightarrow G^{(0)}$ so that $X \in A$ iff the convex hull of $\phi(X)$ is a simplex in G .

Primal Diagrams

Let B be a set of n d -balls in \mathbb{R}^d , see Fig. 2.1. To simplify the forthcoming discussion we assume the d -balls are in general position. An algorithmic justification of this assumption can be found in [12]. For a subset $T \subseteq B$ with $k + 1 = \text{card } T \leq d + 1$, the centers of the d -balls in T are affinely independent, by assumption, and therefore define a k -simplex, denoted σ_T . Consider a d -ball b , with center z and radius ρ , and a point x . The *power distance* of x from b is $\pi_b(x) = |xz|^2 - \rho^2$. The (*weighted*) *Voronoi cell* of $b \in B$ is $V_b = \{x \in \mathbb{R}^d \mid \pi_b(x) \leq \pi_{b'}(x), b' \in B\}$. It is not difficult to see that within its own Voronoi cell a ball b contains all other balls of B . We state this observation explicitly for later reference.

Lemma 2.1. $V_b \cap b' \subseteq V_b \cap b$ for all $b, b' \in B$.

The collection of Voronoi cells, $\mathcal{V} = \mathcal{V}(B) = \{V_b \mid b \in B\}$, defines a decomposition of \mathbb{R}^d known under a variety of different names, including (weighted) Voronoi diagram, power diagram, and Dirichlet tessellation, see Fig. 2.2(a). The collection of cells $\mathcal{Q} = \mathcal{Q}(B) = \{V_b \cap b \mid b \in B\}$ defines a decomposition of the ball union, $\bigcup B = \bigcup_{b \in B} b$, see Fig. 2.2(b) and (c).




Fig. 2.1. This set of 12 disks (2-balls) is used as a running example to illustrate forthcoming definitions.




Fig. 2.2. (a) The Voronoi cells of the 12 disks in Fig. 2.1 cover the entire plane. (b) Each Voronoi cell is restricted to within the disk that defines it. The collection of such cells covers the union of the disks shown in (c).

The boundary of $\cup B$ consists of pieces of spheres of various dimensions. Consider the sphere $\text{bd } b$ and define $f_b = V_b \cap \text{bd } b$ for each $b \in B$. Intuitively, f_b is the spherical part of the boundary of $V_b \cap b$. The components of the f_b are the $(d - 1)$ -faces of $\cup B$. The l -dimensional faces of $\cup B$ can be defined by considering subsets $T \subseteq B$, $\text{card } T = d - l$, and intersections of the form $\bigcap_{b \in T} f_b$. By the general position assumption, these intersections are l -dimensional, and the components are the l -faces of $\cup B$.

Dual Diagrams

The *nerve* of a collection A of sets is $N(A) = \{X \subseteq A \mid \bigcap_{x \in X} x \neq \emptyset\}$. We always have $\emptyset \in N(A)$. The nerve is an abstract simplicial complex because $X \in N(A)$ and $Y \subseteq X$ implies $Y \in N(A)$. For example, the nerve of B , $N(B)$, is the collection of subsets of d -balls with nonempty common intersection. We define two of the three dual diagrams as geometric realizations of nerves. Let $T \subseteq B$ be a set of balls with affinely independent centers. As before, we denote by σ_T the convex hull of the centers. Then

$$\mathcal{D} = \mathcal{D}(B) = \{\sigma_T \mid \{V_b \mid b \in T\} \in N(\mathcal{V})\}$$

is the (weighted) *Delaunay simplicial complex* of B ,

$$\mathcal{X} = \mathcal{X}(B) = \{\sigma_T \mid \{V_b \cap b \mid b \in T\} \in N(\mathcal{Q})\} \quad \text{is the dual complex of } \mathcal{Q}, \text{ and}$$

$$\mathcal{S} = \mathcal{S}(B) = |\mathcal{X}| \quad \text{is the dual shape of } \cup B.$$

Examples of the three diagrams are shown in Fig. 2.3.

The definition of \mathcal{X} as a geometric realization of the nerve of \mathcal{Q} is different although equivalent to the definition of weighted alpha shapes in [9]. It should be clear that \mathcal{X} is a subcomplex of \mathcal{D} . Indeed, $\sigma_T \in \mathcal{X}$ only if $\sigma_T \in \mathcal{D}$ and $T \in N(B)$, but no necessarily vice versa. Another interesting simplicial complex is the *boundary complex* of \mathcal{S} . It consists of all simplices $\sigma_T \in \mathcal{X}$ contained in $\text{bd } \mathcal{S}$. Call such a




Fig. 2.3. The diagrams \mathcal{D} , \mathcal{A} , and \mathcal{S} of the 12 disks in Fig. 2.1 are shown from left to right. By the definition of nerve, \mathcal{D} is dual to the Voronoi diagram in Figure 2.2(a) and \mathcal{A} is dual to the decomposition of the union defined by \mathcal{C} .

simplex a *face* of \mathcal{S} . The faces of \mathcal{S} correspond to the faces of $\cup B$ in the following manner, see also [9].

Lemma 2.2. For each $T \subseteq B$ with $1 \leq \text{card } T \leq d$, σ_T is a face of \mathcal{S} iff

$$\bigcap_{b \in T} f_b \neq \emptyset.$$

3. An Explicit Deformation Retraction

The nerve theorem of algebraic topology [23] implies that $\cup B = \cup_{b \in B} (V_b \cap b)$ and $\mathcal{S} = |\mathcal{N}|$ are homotopy equivalent. We prefer to give a direct proof of this result, which has also been observed by Naiman and Wynn [19]. It reveals some detailed relations between the diagrams used in Sections 4 and 8. We begin with some definitions and then prove homotopy equivalence results between $\cup B$ and \mathcal{S} .

Homotopy Equivalence and Deformation Retractions

It is not necessary to define homotopy equivalence in its full generality. A more restrictive notion is the following. Let $X \subseteq Y$ be two topological spaces. A *retraction* of Y onto X is a continuous map $\phi: Y \rightarrow X$ so that $\phi(x) = x$ for all $x \in X$. A *deformation retraction* of Y onto X is a continuous map $\Phi: Y \times [0, 1] \rightarrow Y$ so that $\Phi(x, t) = x$ for all $x \in X$ and $t \in [0, 1]$, Φ is the identity on Y for $t = 0$, and Φ is a retraction of Y onto X for $t = 1$. If such a Φ exists, then X is a *deformation retract* of Y . If X is a deformation retract of Y , then X and Y are homotopy equivalent. The reverse is not true, although to show that X and Y are homotopy equivalent it suffices to find a topological space Z and embeddings, $\epsilon: X \rightarrow Z$ and $\varepsilon: Y \rightarrow Z$ so that both $\epsilon(X)$ and $\varepsilon(Y)$ are deformation retracts of Z . As proved in [15] the existence of Z , ϵ , and ε is also a necessary condition for the homotopy equivalence of X and Y .

A basic property necessary for our construction is $\mathcal{S} \subseteq \cup B$. Indeed, assuming general position we get $\mathcal{S} \subseteq \text{int } \cup B$. It suffices to show the following result.

Lemma 3.1. *If $\sigma_T \in \mathcal{X}$, then $\sigma_T \subseteq \text{int } \cup T$.*

Proof. The assertion is obviously true for vertices. So let $\text{card } T = k + 1 \geq 2$ and assume inductively that the assertion holds for simplices of dimension less than k . In particular, $\sigma_U \subseteq \text{int } \cup U \subseteq \text{int } \cup T$ for each proper face σ_U of σ_T . The only possibility for $\sigma_T \not\subseteq \text{int } \cup T$ is therefore that the complement of $\text{int } \cup T \cap \text{aff } \sigma_T$ be disconnected. Consider $\cap T = \cup_{b \in T} b$ and note that $\sigma_T \in \mathcal{X}$ implies that $\text{int } \cap T = \cap_{b \in T} \text{int } b \neq \emptyset$. Because $\cap T$ lies symmetric with respect to $\text{aff } \sigma_T$, a point $x \in \text{int } \cap T \cap \text{aff } \sigma_T$ exists. Now, $\text{int } \cup T \cap \text{aff } \sigma_T$ is star-convex with respect to x , which implies that the complement within $\text{aff } \sigma_T$ is connected. Therefore $\sigma_T \subseteq \text{int } \cup T$. \square

Covering with Joins

We construct a deformation retraction of $\text{int } \cup B$ onto \mathcal{S} based on a natural covering of $\cup B$. Because of general position, the omission of $\text{bd } \cup B$ does not affect the final result. The sets of this covering are joins of simplices of \mathcal{X} and faces of $\cup B$. In general, the join of two sets $U, V \subseteq \mathbb{R}^d$ exists provided any two edges $u_1v_1 \neq u_2v_2$, with $u_1, u_2 \in U$ and $v_1, v_2 \in V$, are either disjoint or meet at a common endpoint. Then the *join* of U and V is $U * V = \cup_{u \in U, v \in V} uv$. For convenience, $U * \emptyset = \emptyset * U = U$.

Consider a subset $T \subseteq B$, with $k + 1 = \text{card } T \leq d$, with $\sigma_T \in \mathcal{X}$. Note that $s_T = \cap_{b \in T} \text{bd } b$ is a $(d - k - 1)$ -sphere. By Lemma 2.2, σ_T is a face of \mathcal{S} iff $s_T \cap \text{bd } \cup B$ is nonempty. In this case the affine hull of σ_T is a k -flat, and that of s_T is a $(d - k)$ -flat. These two flats are orthogonal and meet in the center of s_T . This implies that the join of σ_T and s_T exists, and therefore also the join of σ_T and any component f of $s_T \cap \text{bd } \cup B$. Now define $\mathcal{F} = \mathcal{A}(B) = \{\sigma_T * f\}$, where $\emptyset \neq \sigma_T \in \mathcal{X}$ and $f = \emptyset$ if σ_T is not a face of \mathcal{S} and f is a component of $s_T \cap \text{bd } \cup B$ if σ_T is a face of \mathcal{S} . A two-dimensional example is shown in Fig. 3.1. It is not difficult although tedious to prove that \mathcal{F} is indeed a covering of $\cup B$ and that the interiors




Fig. 3.1. The covering of $\cup B$ using joins between simplices of \mathcal{X} and faces of $\cup B$.




Fig. 3.2. This is $\Phi(\cup B, \frac{1}{2})$: at time $t = \frac{1}{2}$ the fringe is narrowed to half the original width.

of the sets in \mathcal{J} are pairwise disjoint. If we ignore the d -simplices of \mathcal{X} in \mathcal{J} we get a covering of the *fringe*, $\cup B - \mathcal{S}$. Alternatively, we can take the set of edges that make up the joins and get a covering of the fringe by an infinite set of edges without common interior points. Indeed, if two joins in $\mathcal{J} - \mathcal{X}$ share a common face, then the face is again a join and the union of edges common to both joins.

The Deformation Retraction

We construct a map $\Phi: \cup B \times [0, 1] \rightarrow \cup B$ whose restriction to $\text{int } \cup B$ satisfies the requirements of a deformation retraction onto \mathcal{S} . We specify Φ for each individual join in \mathcal{J} . Let $\sigma \in \mathcal{X}$ be a face of \mathcal{S} , and let f be a face of $\cup B$ with $\sigma * f \in \mathcal{J}$. For every point $y \in \sigma * f$ there are unique $u \in \sigma$, $v \in f$, and $\lambda \in [0, 1]$ so that $y = \lambda u + (1 - \lambda)v$. For each $t \in [0, 1]$ we define $\Phi(y, t) = (\lambda - \lambda t + t)u + (1 - \lambda)(1 - t)v$. Intuitively, this means that uv is continuously shortened at v so that at time t its length is $(1 - t)|uv|$, see Fig. 3.2. If $f = \emptyset$ we set $\Phi(y, t) = y$ for all $t \in [0, 1]$. The map Φ restricted to $\text{int } \cup B$ is continuous because it is continuous within each join, except possibly at points of $\text{bd } \cup B$. Clearly, Φ is the identity of $\text{int } \cup B$ for $t = 0$, its restriction to \mathcal{S} is the identity for all $t \in [0, 1]$, and $\Phi(\text{Int } \cup B, 1): \text{int } \cup B \rightarrow \mathcal{S}$ is a retraction.

Remark. The construction of Φ can be modified to get a deformation retraction Φ' of $\mathbb{R}^d - \mathcal{S}$ onto $\mathbb{R}^d - \text{int } \cup B$. If $y \in \cup B - \mathcal{S}$, then $y = \lambda u + (1 - \lambda)v$, and we define $\Phi'(y, t) = \lambda(1 - t)u + (1 - \lambda + \lambda t)v$. For $y \in \mathbb{R}^d - \text{int } \cup B$ we set $\Phi'(y, t) = y$ for all t .

Links and Unions of Caps

A relationship like the one between $\cup B$ and \mathcal{S} can be shown between some of their substructures. Consider a subset $T \subseteq B$ so that $\sigma_T \in \mathcal{X}$. The *link* of σ_T in \mathcal{X} is $\text{lk}_{\mathcal{X}}(\sigma_T) = \{\sigma \in \mathcal{X} | \sigma_T * \sigma \in \mathcal{X}\}$. For example, the link of vertex x in Fig. 2.3 is $\{\emptyset, a, b, c, d, bc, cd\}$, and the link of edge xc is $\{\emptyset, b, d\}$. Let $1 \leq k + 1 = \text{card } T$.

Because $\sigma_T \in \mathcal{X}$, s_T is nonempty and thus a $(d - k - 1)$ -sphere in \mathbb{R}^d . Define $\mathcal{X}_T = \text{lk}_{\mathcal{X}}(\sigma_T)$, $\mathcal{S}_T = |\mathcal{X}_T|$, and $B_T = \{s_T \cap b \mid b \in B - T\}$. By the definition of link and by Lemma 2.2, the spheres $s_{T \cup U}$ in s_T that contain faces of $\text{bd } \cup B_T$ correspond to simplices $\sigma_{T \cup U}$ that are faces of \mathcal{S} . Of course, such simplices exist only if σ_T is a face of \mathcal{S} .

Unlike \mathcal{S} , which is a subset of $\text{int } \cup B$, \mathcal{S}_T is usually not contained in $\text{int } \cup B_T$. However, it is possible to embed \mathcal{S}_T in $\text{int } \cup B_T$ using a projection map ψ_T . Let x be a point not contained in the k -flat $\text{aff } \sigma_T$. Hence, $\text{aff } \sigma_{T \cup \{x\}}$ is a $(k + 1)$ -flat and $\text{aff } \sigma_T$ decomposes it into halves. The half that contains x intersects s_T in a point $\psi_T(x)$. Intuitively, ψ_T projects x into s_T ; the center of the projection is $\text{aff } \sigma_T$. The restriction of ψ_T to $|\text{lk}_{\mathcal{S}}(\sigma_T)|$ is continuous and one-to-one. Hence, ψ_T embeds \mathcal{S}_T in s_T , and using an argument as in Lemma 3.1 we see that $\psi_T(\mathcal{S}_T) \subseteq \text{int } \cup B_T$. Similarly, ψ_T embeds the joins $\sigma' * f'$, where $\sigma' = \sigma_U \in \mathcal{X}_T$ and f' is a face of a component of $s_{T \cup U} \cap \text{bd } \cup B$. The embedded joins define a covering \mathcal{F}_T of $\cup B_T$, analogous to the covering \mathcal{F} of $\cup B$. The composition $\psi_T \circ \Phi$ restricted to the joins mentioned above describes a deformation retraction of $\text{int } \cup B_T$ onto $\psi_T(\mathcal{S}_T)$. It follows that $\cup B_T$ and \mathcal{S}_T are homotopy equivalent. We summarize the above results.

Theorem 3.2.

- (i) \mathcal{S} is homotopy equivalent to $\cup B$.
- (ii) For each $T \subseteq B$ with $\sigma_T \in \mathcal{X}$, \mathcal{S}_T is homotopy equivalent to $\cup B_T$.

Remark. Recall that Φ' is a deformation retraction of $\mathbb{R}^d - \mathcal{S}$ onto $\mathbb{R}^d - \text{int } \cup B$. The composition $\psi_T \circ \Phi'$ thus defines a deformation retraction of $s_T - \psi_T(\mathcal{S}_T)$ onto $s_T - \text{int } \cup B_T$. Intuitively, this means that also the complements of $\cup B_T$ and \mathcal{S}_T are homotopy equivalent. This is used in the next section.

Algorithmic Implications

Theorem 3.2(i) has algorithmic consequences concerning the homology groups of $\cup B$. We refer to [18] and [23] for an introduction to homology groups of a topological space Y . For each integer k , the *kth homology group*, $H_k = H_k(Y)$, is an abelian group expressing the k -dimensional connectivity of Y . If the dimension of Y is d , then the possibly nontrivial homology groups are H_0 through H_d . An important related numerical value is the *kth betti number* of Y , which is the rank of H_k . There is a general algorithm for computing H_k , provided Y is given as a finite simplicial complex. Since $\mathcal{S} = |\mathcal{X}|$, this algorithm computes the homology groups of \mathcal{S} . Two homotopy equivalent topological spaces have isomorphic homology groups, and thus the algorithm just mentioned also computes the homology groups of $\cup B$.

Before we say more about this algorithm, let us briefly discuss the complement spaces, $\mathbb{R}^d - \cup B$ and $\mathbb{R}^d - \mathcal{S}$. We have seen that both spaces are homotopy equivalent and thus have isomorphic homology groups. However, since the underlying space of \mathcal{S} is only a bounded subset of \mathbb{R}^d , we do not have a simplicial

representation of $\mathbb{R}^d - \mathcal{S}$. This deficiency can be remedied as follows. Call a simplex $\sigma_T \in \mathcal{D}$ a *hull simplex* if $\sigma_T \subseteq \text{bd} |\mathcal{D}|$. Add a point ω as a new 0-simplex “at infinity” to \mathcal{D} , and for each hull simplex σ_T add $\sigma_{T \cup \{\omega\}}$ to \mathcal{D} . Now, \mathcal{D} is a triangulation of \mathbb{S}^d and no further distinction between hull and other simplices is necessary.

The general algorithm for computing homology groups of simplicial complexes is based on computing Smith normal forms of integer matrices, see, e.g., [18]. Improvements of the original Smith normal form algorithm with polynomial behavior can be found in [8] and [17]. A fast combinatorial algorithm that works for simplicial complexes embedded in \mathbb{S}^3 is described in [6]. For a large problem size, which could mean thousands of balls defining $\cup B$ or similar numbers of simplices constituting \mathcal{X} , only the algorithm in [6] performs satisfactorily. This leaves us with the open problem of finding faster algorithms for computing homology groups of simplicial complexes embedded in dimensions higher than three.

4. Counting Faces

In this section we consider the algorithmic problem of counting the faces of $\cup B$. The assumption is that \mathcal{X} is given as a subcomplex of \mathcal{D} , and we seek an algorithm that computes the number of l -faces of $\cup B$, for each $0 \leq l \leq d - 1$. This problem is related to determining the betti numbers of links in \mathcal{X} because the faces of $\cup B$ are typically not simply connected. The basic strategy is to consider all l -spheres of the form $s_T = \bigcap_{b \in T} \text{bd } b$, $\text{card } T = d - l$, with $\sigma_T \in \mathcal{X}$. For each such l -sphere we compute the number of l -faces of $\cup B$ it contains, and we take the sum of these numbers. The result is n_l , the number of l -faces of $\cup B$.

Components of Link Complements

Recall the definition of $B_T = \{s_T \cap b | b \in B - T\}$. The complement of $\cup B_T$, $s_T - \cup B_T$, is the interior of the union of l -faces contained in the l -sphere s_T . Since we assume general position of the d -balls, the connectivity of the interior is the same as that of its closure. Hence, each component of the complement is the interior of an l -face and is to be counted. For each $\sigma_T \in \mathcal{X}$, define $\tilde{\mathcal{X}}_T = \text{lk}_{\mathcal{D}}(\sigma_T) - \text{lk}_{\mathcal{X}}(\sigma_T)$, where we assume that \mathcal{D} is extended to a triangulation of \mathbb{S}^d as described at the end of Section 3. By the remark after Theorem 3.2, the number of components of the complement is the same as that of $\tilde{\mathcal{X}}_T$. For each $\sigma_T \in \mathcal{X}$ let n_T be the number of components of $\tilde{\mathcal{X}}_T$. Then we have the following result.

Lemma 4.1. *For each $0 \leq l \leq d - 1$,*

$$n_l = \sum_{\sigma_T \in \mathcal{X}, \text{card } T = d - l} n_T.$$

Remark. If $\sigma_T \in \mathcal{X}$ is not a face of \mathcal{S} , then $\text{lk}_{\mathcal{X}}(\sigma_T) = \text{lk}_{\mathcal{D}}(\sigma_T)$ is a complete

triangulation of S^l . Hence, $n_T = 0$, which implies that the equation in Lemma 4.1 remains valid if the sum extends only over the simplices of \mathcal{K} that are faces of \mathcal{S} .

The Algorithm

We assume the following graph representation of \mathcal{D} and \mathcal{K} . Algorithms for constructing \mathcal{D} and \mathcal{K} can be found in [9] and [14]. The nodes of the graph \mathcal{D}^* are the d -simplices of \mathcal{D} , and the arcs of \mathcal{D}^* are the $(d - 1)$ -simplices of \mathcal{D} ; this includes the d - and $(d - 1)$ -simplices incident to ω . Each node and arc is labeled whether or not it belongs to \mathcal{K} . The subgraph that consists of the nodes and arcs in $\mathcal{D} - \mathcal{K}$ is denoted $\bar{\mathcal{K}}^*$. Since \mathcal{K} is a proper complex, $\bar{\mathcal{K}}^*$ is a proper graph. Given an arc of \mathcal{D}^* , we have access to the two incident nodes in constant time. Similarly, given a node we have access to the incident arcs in constant time. This is a reasonable assumption if d , which is the number of dimensions as well as one less than the node degree, is considered a constant. Furthermore, we assume that given a simplex $\sigma \in \mathcal{D}$ of dimension less than d , we can find an incident node in \mathcal{D}^* in constant time. Starting at this node, all other nodes incident to σ can be enumerated in constant time per node.

The algorithm relies on the fact that the number of components of $\bar{\mathcal{K}}_T$ is also the number of components of the subgraph of $\bar{\mathcal{K}}^*$ induced by the d -simplices $\sigma_U \in \mathcal{D} - \mathcal{K}$ with $T \subseteq U$. Denote this induced subgraph by $\bar{\mathcal{K}}_T^*$. This is because $\sigma_V \in \bar{\mathcal{K}}_T^*$ iff $T \cap V = \emptyset$ and $\sigma_{T \cup V} \in \mathcal{D} - \mathcal{K}$. In particular, σ_V is an l - or $(l - 1)$ -simplex of $\bar{\mathcal{K}}_T^*$ iff $\sigma_{T \cup V}$ is a d - or $(d - 1)$ -simplex of $\mathcal{D} - \mathcal{K}$. The faces of $\cup B$ can thus be counted by finding components of various induced subgraphs of $\bar{\mathcal{K}}^*$. A more detailed formulation of the algorithm that computes n_l follows. Initially, all nodes of \mathcal{D}^* are unmarked.

```

 $n_l := 0;$ 
for each  $(d - l - 1)$ -simplex  $\sigma_T \in \mathcal{K}$  do
  for each node  $\sigma_U \in \mathcal{D}^*$  incident to  $\sigma_T$  do
    if  $\sigma_U$  is not marked and  $\sigma_U \notin \mathcal{K}$  then
      mark  $\sigma_U$ ;  $n_l := n_l + 1$ ;
      start a graph search to mark all nodes  $\sigma_U$  that
        belong to the same component of  $\bar{\mathcal{K}}_T^*$  as  $\sigma_U$ 
      endif
    endfor;
  unmark all marked nodes
endfor.
```

As remarked earlier, it is actually sufficient to run the outer for-loop only over all faces of \mathcal{S} . In any case, each simplex of \mathcal{D} is touched only a constant number of times, so the entire algorithm runs in time at most proportional to the number of simplices in \mathcal{D} . Indeed, the step that employs graph searching also takes only constant time per node it marks, see, e.g., Chapter VI,23 [4]. We summarize the results of this section.

Theorem 4.2. *Given a suitable representation of \mathcal{X} as a subcomplex of \mathcal{D} , for $0 \leq l \leq d - 1$, the number of l -faces of $\cup B$ can be computed in time proportional to the number of simplices in \mathcal{D} .*

5. Euler Relation from a Viewpoint

This section derives Euler relations for convex polyhedra and viewpoints in \mathbb{R}^d . For each point $x \in \mathbb{R}^d$ we specify an alternating sum for the faces of a polyhedron visible from x . This sum will be 1 inside the polyhedron and 0 outside. These sums are used in Section 6 to derive short formulas for measuring a polyhedron or its intersection with another body.

Inclusion–Exclusion for Convex Polyhedra

Let H be a finite set of closed half-spaces in \mathbb{R}^d that defines a nonempty convex polyhedron $\cap H = \cap_{h \in H} h$. For simplicity we assume general position of the half-spaces. For every $x \in \mathbb{R}^d$ and every $I \in 2^H$ define the characteristic function

$$\gamma_I(x) = \begin{cases} 1 & \text{if } x \notin h \text{ for all } h \in I, \\ 0 & \text{otherwise.} \end{cases}$$

For $L \subseteq 2^H$ define $\Gamma_L(x) = \sum_{I \in L} (-1)^{\text{card } I} \gamma_I(x)$. The general inclusion–exclusion principle implies that

$$\Gamma_{2^H}(x) = \begin{cases} 1 & \text{if } x \in \cap H, \\ 0 & \text{if } x \notin \cap H. \end{cases}$$

A direct proof can easily be given. Define $G = \{h \in H \mid x \notin h\}$. Clearly, $\Gamma_{2^H}(x) = \Gamma_{2^G}(x)$. If $x \in \cap H$, then $G = \emptyset$ and $\Gamma_{2^G}(x) = \gamma_{\emptyset}(x) = 1$. If $x \notin \cap H$, then $G \neq \emptyset$ and $\Gamma_{2^G}(x) = \sum_{I \in 2^G} (-1)^{\text{card } I} = (1 - 1)^{\text{card } G} = 0$.

It should be clear that redundant half-spaces can be eliminated, that is, if $\cap H = \cap G$, for some $G \subseteq H$, then $\Gamma_{2^H}(x) = \Gamma_{2^G}(x)$. We claim that a more dramatic reduction of the set 2^H is possible. For $I \in 2^H$ define $f_I = \cap H \cap \cap_{h \in I} \text{bd } h$. If $f_I \neq \emptyset$, then it is a unique face of $\cap H$. This includes the case $I = \emptyset$ where $f_I = f_{\emptyset} = \cap H$. Define $D = D(H) = \{I \in 2^H \mid f_I \neq \emptyset\}$. Observe that D is an abstract simplicial complex, and because of general position any $I \in D$ has cardinality at most d . Figure 5.1 illustrates the following result. It can be proved using the Euler relation for convex polyhedra [16, Chapter 8]. We prefer to give an explicit proof using induction over the number of half-spaces. We also indicate how this proof extends to an elementary inductive proof of the classic Euler relation for convex polyhedra and polytopes.




Fig. 5.1. The faces f_I of $\cap H$ for which $\gamma_I(x) = 1$ are the ones visible from x . The point x sees the cube itself, three facets, three edges, and one vertex. Hence $\Gamma_D(x) = 1 - 3 + 3 - 1 = 0$ as claimed.

Lemma 5.1.

$$\Gamma_D(x) = \begin{cases} 1 & \text{if } x \in \cap H, \\ 0 & \text{if } x \notin \cap H. \end{cases}$$

Proof. We use induction over the size of $G = \{h \in H \mid x \notin h\}$. G is empty iff $x \in \cap H$, and indeed we have $\Gamma_D(x) = \gamma_{\emptyset}(x) = 1$ in this case. So assume $x \notin \cap H$. There is at least one half-space $g \in G$: let \bar{g} be the other half-space bounded by the same hyperplane. Define $H' = H - \{g\}$ and $H'' = H' \cup \{\bar{g}\}$. By the induction hypothesis the assertion applies to $\cap H'$ and to $\cap H''$, see Fig. 5.2. Define $D' = D(H')$ and $D'' = D(H'')$. We express D , D' , and D'' as disjoint unions of smaller sets. By definition, this translates to addition for Γ . For a fixed point $x \notin \cap H$ we have

$$\Gamma_{D'} = \Gamma_{L'} + \Gamma_{X'} + \Gamma_{U'},$$

where $L' = \{I \in D' \mid f_I \subseteq g\}$, $X' = \{I \in D' \mid f_I \cap g \neq \emptyset \text{ and } f_I \cap \bar{g} \neq \emptyset\}$, and $U' = \{I \in D' \mid f_I \subseteq \bar{g}\}$. Similarly,

$$\Gamma_D = \Gamma_{L'} + \Gamma_{X'} + \Gamma_X,$$

where $X = \{I \cup \{g\} \mid I \in X'\}$. Indeed X' represents all faces of $\cap H'$ that intersect the hyperplane bounding g , and at the same time it represents all faces of $\cap H$ and




Fig. 5.2. $\cap H'$ is the union of $\cap H$ and $\cap H''$. Since $x \notin \cap H$, by assumption, we have $x \in \cap H'$ iff $x \in \cap H''$.

of $\cap H''$ that touch this hyperplane. Finally,

$$\Gamma_{D''} = \Gamma_{X'} + \Gamma_{X''} + \Gamma_{U'},$$

where $X'' = \{I \cup \{\bar{g}\} | I \in X'\}$. Now we express Γ_D in terms of the other sets:

$$\begin{aligned} \Gamma_D &= \Gamma_{D'} - \Gamma_{U'} + \Gamma_X \\ &= \Gamma_{D'} - \Gamma_{D''} + \Gamma_{X''} + \Gamma_{X'} + \Gamma_X. \end{aligned}$$

We have $x \in \cap H'$ iff $x \in \cap H''$ and therefore $\Gamma_{D'} - \Gamma_{D''} = 0$ by the induction hypothesis. Furthermore, $\Gamma_{X''} = 0$ because $x \in \bar{g}$ and each $I \in X''$ contains \bar{g} . finally, $\Gamma_{X'} + \Gamma_X = 0$ because $\gamma_I = \gamma_{I \cup \{g\}}$ and therefore $(-1)^{\text{card } I} \gamma_I + (-1)^{\text{card } I+1} \gamma_{I \cup \{g\}} = 0$ for each $I \in X'$. Therefore $\Gamma_D = 0$ as required. \square

Remarks. (1) The following modifications generalize Lemma 5.1 and its proof to cover degenerate positions of the half-spaces. First, D is defined so it contains only maximal sets defining faces of $\cap H$: $D = \{I \in 2^H | f_I \neq \emptyset \text{ and } f_I \neq f_J \text{ if } I \subset J\}$. Second, the signs in Γ_D alternate with the codimensions of the faces: $\Gamma_D(x) = \sum_{I \in D} (-1)^{\text{codim } f_I} \gamma_I(x)$, where $\text{codim } f_I = d - \dim f_I$. Third, in the proof we have three additional sets,

$$\begin{aligned} C' &= \{I \in D' | f_I \subseteq \text{bd } g\}, \\ C &= \{I \cup \{g\} | I \in C'\}, \\ C'' &= \{I \cup \{\bar{g}\} | I \in C'\}, \end{aligned}$$

which are subsets of D' , D , and D'' disjoint from L', U', X', X, X'' . In the final expression for Γ_D we get $\Gamma_{C''} - \Gamma_{C'} + \Gamma_C$ as an additional term. It vanishes just as $\Gamma_{X''} + \Gamma_{X'} + \Gamma_X$ does.

(2) Lemma 5.1 implies the Euler relation for unbounded convex polyhedra. To see this take x outside all half-spaces. The following standard decomposition of the boundary of a bounded convex polyhedron can be used to extend this result to a proof of the Euler relation for convex polytopes. Choose a generic direction classifying each facet either as a *front* or *back* facet. The collection of front facets forms a $(d - 1)$ -ball, and so does the collection of back facets. The intersection of the two balls projected along the chosen direction forms the boundary of a $(d - 1)$ -dimensional convex polytope, which is decomposed the same way.

Intersection with a Convex Body

We generalize Lemma 5.1 so it makes a statement about points x of a compact convex set A . Define $K = K(A, H) \subseteq D(H)$ so that $I \in K$ iff $f_I \cap \text{int } A \neq \emptyset$.

Lemma 5.2.

$$\Gamma_K(x) = \begin{cases} 1 & \text{if } x \in A \cap \bigcap H, \\ 0 & \text{if } x \in A - \bigcap H. \end{cases}$$

Proof. We choose a suitable polyhedral approximation of A . Let H_A be a finite collection of closed half-spaces so that $A \subseteq \bigcap H_A$ and $K(A, H) = K(\bigcap H_A, H)$. Such a finite set H_A can be constructed during the following iteration. Find a face f_j of $\bigcap H$ disjoint from $\text{int } A$ with $f_j \cap \text{int } \bigcap H_A \neq \emptyset$. By convexity of f_j and A there is a separating hyperplane. Add the half-space bounded by this hyperplane that contains A to H_A . The process is finite because $\bigcap H$ has only finitely many faces.

Define $D_{\bigcap} = D(H \cup H_A)$ and use Lemma 5.1 to get

$$\Gamma_{D_{\bigcap}}(x) = \begin{cases} 1 & \text{if } x \in \bigcap H_A \cap \bigcap H, \\ 0 & \text{if } x \notin \bigcap H_A \cap \bigcap H. \end{cases}$$

Note that $D_{\bigcap} = K \dot{\cup} L$, where $K = K(\bigcap H_A, H)$ and $L = \{I \in D_{\bigcap} \mid I \cap H_A \neq \emptyset\}$. For all $x \in \bigcap H_A$ and each $I \in L$ we have $\gamma_I(x) = 0$. Hence,

$$\Gamma_K(x) = \Gamma_{D_{\bigcap}}(x) - \Gamma_L(x) = \Gamma_{D_{\bigcap}}(x)$$

for all $x \in \bigcap H_A$. The assertion follows because $A \cap \bigcap H \subseteq \bigcap H_A \cap \bigcap H$ and $A - \bigcap H \subseteq \bigcap H_A - \bigcap H$. \square

Remark. The relation in Lemma 5.2 also applies to compact nonconvex sets A if the collection of faces considered is determined by the convex hull of A , that is, $K = K(\text{conv } A, H)$. For example, it applies to the boundary of a compact convex set, and the intersection of this boundary with the boundary of $\bigcap H$, etc.

6. Measuring the Union of Balls

This section and the one after the next simplify, improve, and generalize earlier work on algorithms for measuring the union of balls [1], [2]. Based on the correspondences between the various diagrams introduced in Section 2, this section derives short inclusion-exclusion formulas for the d -dimensional volume or Lebesgue measure of $\bigcup B$ and the total l -dimensional Lebesgue measure of its l -dimensional faces. We begin by studying inclusion-exclusion formulas for convex polyhedra.

Measuring by Integration

We measure $A \cap \bigcap H$ using Lemma 5.2. Consider a compact convex set A with the nonempty interior in \mathbb{R}^d . The d -dimensional measure of $A \cap \bigcap H$,

$\mu_d(A \cap \cap H)$, is the integral of $\Gamma_K(x)$ over all points $x \in A$, where $K = K(A, H)$. We get

$$\begin{aligned} \mu_d\left(A \cap \cap H\right) &= \int_{x \in A} \Gamma_K(x) \, dx \\ &= \int_{x \in A} \sum_{I \in K} (-1)^{\text{card } I} \gamma_I(x) \, dx \\ &= \sum_{I \in K} (-1)^{\text{card } I} \int_{x \in A} \gamma_I(x) \, dx. \end{aligned}$$

The integral of $\gamma_I(x)$ over all $x \in A$ is the d -dimensional measure of $A \cap Q_I$, where $Q_I = \cap_{h \in I} \bar{h}$.

The same calculation can be done for lower-dimensional sets. We are interested in the sets $\text{bd } A \cap \cap H^{(l+1)}$, where $\cap H^{(l+1)}$ is the union of all $(l + 1)$ -dimensional faces of $\cap H$. Assuming general position, these sets are l -dimensional. We state the results.

Lemma 6.1.

- (i) $\mu_d(A \cap \cap H) = \sum_{I \in K} (-1)^{\text{card } I} \mu_d(A \cap Q_I)$.
- (ii) For $0 \leq l \leq d - 1$,

$$\mu_l\left(\text{bd } A \cap \cap H^{(l+1)}\right) = \sum_{I \in K} (-1)^{\text{card } I} \mu_l(\text{bd } A \cap Q_I^{(l+1)}).$$

Remarks. (1) The lowest dimension of any face of Q_I is $d - \text{card } I$. This implies that in Lemma 6.1(ii) all terms for sets I with $\text{card } I \leq d - l - 2$ vanish and can therefore be omitted.

(2) The relations in Lemma 6.1 are based on the assumption of a uniform density distribution. All results hold without change for any other reasonable density function. To see this redefine $\gamma_I(x)$ equal to the density at x , provided $x \notin h$ for all $h \in I$. Otherwise, $\gamma_I(x) = 0$ as before.

(3) Consider a simple special case of Lemma 6.1(ii): $\cap H$ is a triangular cone with apex y in \mathbb{R}^3 , and A is a 3-ball with the unit surface area centered at y . By Lemma 6.1(ii) the size, μ_2 , of the spherical triangle $\text{bd } A \cap \cap H$ is $1 - \frac{3}{2} + (\alpha + \beta + \gamma) - \mu_2$, where α, β, γ are the three dihedral angles of $\cap H$ normalized between 0 and 1. This implies the famous formula

$$\mu_2 = \frac{\alpha + \beta + \gamma}{2} - \frac{1}{4}$$

for spherical triangles. In a similar vein it is possible to derive Gram's formulas for convex polyhedra, see, e.g., Chapter 14 of [16].

Inclusion–Exclusion with Intersections of Balls

We write $\mu_d = \mu_d(\cup B)$ for the d -dimensional Lebesgue measure of $\cup B = \cup_{b \in B} b$, and $\mu_l = \mu_l(\cup B)$ for the total l -dimensional Lebesgue measure of all l -faces of $\cup B$, for $0 \leq l \leq d - 1$. In particular, μ_{d-1} is the size of $\text{bd } \cup B$, and μ_0 is the number of vertices or corners of $\cup B$. We derive formulas that express μ_l in terms of l -dimensional measures of intersections of at most $d + 1$ balls from B . These formulas are shorter than similar formulas in [20] because they take K as the index set rather than D . This difference turns out to be essential for the derivation of the decomposable formulas in Section 8. We note also that the weighted Voronoi cells used to handle varying radii, see Section 2, are different from the ones suggested in [20].

Call $T \subseteq B$ independent if for each subset $U \subseteq T$ we have $\text{int}(\cap U - \cup T - U) \neq \emptyset$. For example, if σ_T is a simplex in \mathcal{X} , then T is independent. If T is independent, then $\cap T \neq \emptyset$ and its face structure is dual to that of σ_T , see Figure 6.1. We write $\mu_l(\cap T)$ for the total l -dimensional measure of all l -faces of $\cap T$. Clearly, the lowest dimension of any face of $\cap T$ is $d - \text{card } T$, so $\mu_l(\cap T) = 0$ if $l < d - \text{card } T - 1$. We are ready to state the first set of inclusion–exclusion formulas for $\cup B$.

Theorem 6.2.

- (i) $\mu_d(\cup B) = \sum_{\emptyset \neq \sigma_T \in \mathcal{X}} (-1)^{\text{card } T - 1} \mu_d(\cap T)$.
- (ii) For $0 \leq l \leq d - 1$,

$$\mu_l(\cup B) = \sum_{\sigma_T \in \mathcal{X}} (-1)^{\text{card } T - d + l} \mu_l(\cap T).$$




Fig. 6.1. The intersection of one, two, three, and four 3-balls. The face structure is dual to that of a vertex, an edge, a triangle, and a tetrahedron.




Fig. 6.2. Under inversion in $d + 1$ dimensions, the image of a hyperplane is a d -sphere, and every $(d - 1)$ -sphere in the hyperplane maps to a $(d - 1)$ -sphere on the d -sphere.

Remark. The sets T of size $\text{card } T \leq d - l - 1$ can be omitted from the sums in Theorem 6.2(ii) because their contribution to μ_l is zero anyway. Note that the thus simplified formula for $l = 0$ counts 2 for each $(d - 1)$ -simplex $\sigma_T \in \mathcal{X}$ and then subtracts the number of incident d -simplices that are in \mathcal{X} . The result is the same as derived in Lemma 4.1.

Proof. We present the detailed argument for (i) using an embedding of \mathbb{R}^d as a hyperplane Π in \mathbb{R}^{d+1} . Let z be a point in $\mathbb{R}^{d+1} - \Pi$ and consider the inversion transform with center z . It maps every point $x \neq z$ to a point x^0 so that x and x^0 lie on the same half-line with endpoint z and $|zx^0| = 1/|zx|$. The image of Π under inversion is a d -sphere Π^0 that contains z . Furthermore, each d -ball b in Π maps to a spherical cap b^0 on Π^0 , see Fig. 6.2. Let h_b be the half-space in \mathbb{R}^{d+1} so that $b^0 = \Pi^0 \cap \bar{h}_b$.

With an eye on Lemma 6.1 we define $A = \text{conv } \Pi^0$, $\text{bd } A = \Pi^0$, and $H = \{h_b | b \in B\}$. Inversion maps $\text{bd } A - \cap H$ to $\text{int } \cup B$. Note that for points on Π^0 , inversion is the same as stereographic projection into Π centrally from z . The same projection maps the facets of $\cap H$ to the Voronoi cells of $\mathcal{V} = \mathcal{V}(B)$. Moreover, a face f_I of $\cap H$ maps to the intersection of Voronoi cells $\cap_{h_b \in I} V_b$. Furthermore $f_I \cap \text{int } A \neq \emptyset$ iff this intersection of Voronoi cells has a common point with $\text{int } \cup B$. By assumption of general position the same is true if we replace $\text{int } A$ by A and $\text{int } \cup B$ by $\cup B$. Hence, $I \in K(A, H)$ iff the simplex spanned by the centers of the corresponding balls belongs to $\mathcal{X} = \mathcal{X}(B)$.

We derive a formula for the d -dimensional measure of $\text{bd } A - \cap H$ using Lemma 6.1(ii). For the intersection, $\text{bd } A \cap \cap H$, we get

$$\mu_d(\text{bd } A \cap \cap H) = \sum_{I \in K(A, H)} (-1)^{\text{card } I} \mu_d(\text{bd } A \cap Q_I),$$

where $Q_I = \cap_{h \in I} \bar{h}$, as usual. For $\text{bd } A - \cap H$ we therefore get

$$\begin{aligned} \mu_d(\text{bd } A - \cap H) &= \mu_d(\text{bd } A) - \mu_d(\text{bd } A \cap \cap H) \\ &= \sum_{\emptyset \neq I \in K(A, H)} (-1)^{\text{card } I - 1} \mu_d(\text{bd } A \cap Q_I). \end{aligned}$$

To get (i) note that for each $I \in K(A, H)$ the set $\text{bd } A \cap Q_I$ is the image under inversion of the corresponding intersection of balls, $\cap T$. Finally, to get the d -dimensional measure of $\cup B$ we put a density function on $\text{bd } A = \Pi^0$ whose image under inversion is the uniform density in Π .

The relations in (ii) are obtained by similar arguments using Lemma 6.1(ii) for values of l less than d . The main difference to the above proof for $l = d$ is that for $l < d$ we compute $\text{bd } A \cap \cap H^{(l+1)}$ directly, without considering any complement. This explains the inconsistency in sign between the formulas in (i) and (ii). \square

7. Results on Independent Simplices

This section proves several results on simplices, which are used in Section 8 where another set of formulas for the measure of $\cup B$ is derived. The main result is Theorem 7.3. It expresses the common intersection of $d + 1$ d -balls in terms of the simplex spanned by their centers and common intersections of d or fewer of the d -balls. The theorem applies only if the $d + 1$ balls are independent, which is assumed throughout this section.

Inclusion–Exclusion for Simplices

Let H be a set of $d + 1$ closed half-spaces in \mathbb{R}^d defining a d -simplex $\cap H$. Each proper subset $I \subseteq H$ defines a proper face, f_I , of $\cap H$. Except for the vanishing term $\gamma_H(x)$, Lemma 5.1 for $\cap H$ coincides with the trivial inclusion–exclusion formula,

$$\Gamma_{2^H}(x) = \begin{cases} 1 & \text{if } x \in \cap H, \\ 0 & \text{if } x \notin \cap H. \end{cases}$$

Let D_I contain all sets in 2^H that contain I , so $D_I = \{J \cup I \mid J \in 2^{H-I}\}$. Recall that $Q_I = \cap_{h \in I} \bar{h}$ and define $P_I = Q_I \cap \cap (H - I)$, see Fig. 7.1. We are interested in $\Gamma_{D_I}(x)$ for points $x \in Q_I$. Because Γ_{D_I} coincides with $\Gamma_{2^{H-I}}$ for such points we get the following result.

Lemma 7.1.

$$\Gamma_{D_I}(x) = \begin{cases} 1 & \text{if } x \in P_I, \\ 0 & \text{if } x \in Q_I - P_I. \end{cases}$$

Intuitively, this means that with respect to inclusion–exclusion P_I behaves in Q_I the same way as $\cap H$ behaves in \mathbb{R}^d .




Fig. 7.1. Regions P_I of a 3-simplex, with card $I = 1$ to the left, card $I = 2$ in the middle, and card $I = 3$ to the right.

Independent Sets of Balls

Recall the definition of an independent set T of $d + 1$ d -balls in \mathbb{R}^d . T has 2^{d+1} subsets each defining a nonempty set in \mathbb{R}^d consisting of all points inside balls of the subset and outside balls not in the subset. The $(d - 1)$ -spheres bounding $d + 1$ d -balls decompose \mathbb{R}^d into at most 2^{d+1} cells. In the case of $d + 1$ independent d -balls the numbers are the same, so each cell must belong to a unique subset of T . This implies each set defined by a subset of T is connected. Notice that T is independent iff $\mathcal{R}(T)$ consists of σ_T and all its faces.

Let H be a set of $d + 1$ half-spaces so that $\cap H = \sigma_T$, as before. Each hyperplane $\text{bd } h, h \in H$, contains the centers of d d -balls in T . For each $I \subseteq H$ let $X = X_I \subseteq T$ contain the d -balls whose centers lie in all hyperplanes bounding half-spaces in I . Define $Y = Y_I = T - X$ and note that $\text{card } I = \text{card } Y$. For a choice of I we are interested in P_I , see Fig. 7.1. In particular, we claim that within P_I the intersection of the d -balls in Y is contained in the union of the d -balls in X . To help the discussion we call P_I the *focus* of Y in T . For example, $\cap H$ is the focus of \emptyset in T , and \emptyset is the focus of T in T . See Fig. 7.2 for an illustration. The disks around a and b at the right intersect their focus outside the disk around c , but the intersection of the two disks does not. The claim is now formally stated and proved.

Lemma 7.2. *For each $I \subseteq H$, we have*

$$\cap Y_I \cap P_I \subseteq \cup X_I.$$

Proof. Note that the assertion holds in \mathbb{R}^1 , where we have two intersecting 1-balls (intervals) that do not nest. They define a 1-simplex connecting the midpoints of the 1-balls. Assume the assertion inductively for dimensions less than d . Take a subset $I \subseteq H$ and consider the focus, P_I , of Y in T . If $I \neq H, \emptyset$, then P_I is a proper convex polyhedron which shares the face σ_X with $\cap H$. By Lemma 3.1, this face is contained in $\cup X$. All other proper faces of P_I are lower-dimensional foci, namely, foci of Y' in T' , where $Y' \subseteq Y, X' \subseteq X$, and $Y' \cup X' = T' \subset T$. For each choice of Y' and X' the assertion holds by induction hypothesis. It follows that $\cap Y \cap \text{bd } P_I \subseteq \cup X$. We just need to extend this result from $\text{bd } P_I$ to P_I itself.

To get a contradiction, assume $\cap Y$ is not contained in $\cup X$. Choose a point $x \in \cap Y \cap P_I$ not contained in $\cup X$. Note that $\cap Y$ is symmetric with respect to

aff σ_Y . Similarly, $\cup X$ is symmetric with respect to aff σ_X . Let h_Y be the hyperplane that contains aff σ_Y and is parallel to aff σ_X . h_Y is unique because the dimensions of the two affine hulls add to $d - 1$. Let y be the reflection of x with respect to h_Y , and observe that $y \in \cap Y$. By construction, $h_Y \cap P_I = \emptyset$, and thus $y \notin P_I$. Since y is further away from all $b \in X$ than x , we also have $y \notin \cup X$. This implies that x and y belong to the region of the same subset Y of T , that is, $x \in b$ iff $y \in b$ for all $b \in T$. Furthermore, $x \in P_I$, $y \notin P_I$, and this region does not intersect $\text{bd } P_I$. This implies that the region is disconnected, which contradicts the independence of T . \square

Measuring Simplices and Balls

Using Lemmas 7.1 and 7.2 we derive a relation for the measure of $\cap T$. We still suppose that T is independent. Hence, $\sigma_T \in \mathcal{K}(T)$ and, by Lemma 3.1, $\sigma_T \subseteq \cup T$. For each face σ_U , $U \subseteq T$, let $\varphi_{U,T}$ be the angle at σ_U inside σ_T . We normalize angles between 0 and 1, so all angles can be interpreted as follows. Take a point $x \in \text{int } \sigma_U$ and a sufficiently small $(d - 1)$ -sphere s with center x . Then $\varphi_{U,T}$ is the fraction of s inside σ_T , that is,

$$\varphi_{U,T} = \frac{\mu_{d-1}(s \cap \sigma_T)}{\mu_{d-1}(s)}.$$

For example, $\varphi_{T,T} = 1$, and $\varphi_{U,T} = \frac{1}{2}$ if $\text{card } U = d$. It is convenient to set $\varphi_{\emptyset,T} = 0$. In \mathbb{R}^3 an angle at a vertex is usually referred to as a solid angle, and an angle at an edge as a dihedral angle.

Theorem 7.3.

- (i) $\sum_{U \subseteq T} (-1)^{\text{card } U - 1} \varphi_{U,T} \cdot \mu_d(\cap U) = \mu_d(\sigma_T)$.
- (ii) For $0 \leq l \leq d - 1$,

$$\sum_{U \subseteq T} (-1)^{\text{card } U - 1} \varphi_{U,T} \cdot \mu_l(\cap U) = 0.$$

Before proving these relations, let us consider a two-dimensional example, see, e.g., Fig. 7.2. There are three disks satisfying the assumptions of independence.




Fig. 7.2. The three disks to the left define a favorable case, whereas the disks to the right do not. Indeed, the intersection of the disk around b with the half-plane opposite ac is not covered by the disks around a and c . The same is true for the disk around a and the half-plane opposite bc .

Relation (i) expresses the area of the triangle spanned by the centers of the three disks in terms of angular pieces of the disks and their intersections. The fraction of each disk is defined by the angle at its vertex inside the triangle. Half the area of each pairwise intersection is subtracted, and the area of the common intersection of all three disks is finally added. The validity of relation (i) is obvious for the three disks shown on the left in Fig. 7.2, and it is less obvious for the disks on the right in Fig. 7.2.

Proof. We present the proof of (i) in detail. Recall that $\sigma_T = \cap H$, and that every subset $I \subseteq H$ defines a face $f_I = \sigma_X$, with $X = X_I \subseteq T$. We use Lemma 6.1(i) with $\cap H = \sigma_T$ and $A = \cup T$ and get

$$\mu_d\left(\cap H\right) = \mu_d\left(A \cap \cap H\right) = \sum_{I \subseteq H} (-1)^{\text{card } I} \mu_d\left(A \cap Q_I\right),$$

where $A \cap Q_I$ is the closure of A outside all half-spaces $h \in I$, as before. We first assume the favorable case where $A \cap Q_I = \cup X \cap Q_I$ for all $I \subseteq H$, see Fig. 7.2. This case is favorable because all hyperplanes $\text{bd } h, h \in I$, contain the centers of all $b \in X$ and thus cut these d -balls into halves. We use the fact that the angle at σ_X inside $\cap H$ is that same as the opposite angle inside Q_I and get

$$\begin{aligned} \mu_d\left(A \cap Q_I\right) &= \mu_d\left(\cup X \cap Q_I\right) \\ &= \varphi_{X, T} \cdot \mu_d\left(\cup X\right) \\ &= \varphi_{X, T} \cdot \sum_{\emptyset \neq U \subseteq X} (-1)^{\text{card } U - 1} \mu_d\left(\cap U\right). \end{aligned}$$

The last line is obtained by straightforward application of the inclusion–exclusion principle. Now we plug the last relation into the earlier one for $\mu_d(\cap H)$ and get

$$\mu_d\left(\cap H\right) = \sum_{I \subseteq H} \left((-1)^{\text{card } I} \varphi_{X, T} \cdot \sum_{\emptyset \neq U \subseteq X} (-1)^{\text{card } U - 1} \mu_d\left(\cap U\right) \right).$$

Summing over all subsets I of H is the same as summing over all subsets X of T . Therefore,

$$\begin{aligned} \mu_d\left(\cap H\right) &= \sum_{X \subseteq T} \left((-1)^{d - \text{card } X + 1} \varphi_{X, T} \cdot \sum_{\emptyset \neq U \subseteq X} (-1)^{\text{card } U - 1} \mu_d\left(\cap U\right) \right) \\ &= \sum_{\emptyset \neq U \subseteq T} \left(\mu_d\left(\cap U\right) \cdot \sum_{X \supseteq U} (-1)^{d - \text{card } X + \text{card } U} \varphi_{X, T} \right) \\ &= \sum_{U \subseteq T} (-1)^{\text{card } U - 1} \varphi_{U, T} \cdot \mu_d\left(\cap U\right). \end{aligned}$$

The last line is obtained by observing that $\sum_{X \supseteq U} (-1)^{d - \text{card } X + 1} \varphi_{X,T} = \varphi_{U,T}$ for all $U \subseteq T$, see remark (3) after Lemma 6.1. This proves (i) in the favorable case.

With Lemmas 7.1 and 7.2 we can reduce the unfavorable case to the favorable one. Consider a subset $I \subseteq H$. In the unfavorable case we have $A \cap Q_I \neq \cup X \cap Q_I$. By Lemma 7.2, $\cap Y$ intersects P_I at most inside $\cup X$. By Lemma 7.1, within $Q_I - P_I$ all points of $\cap Y$ outside $\cup X$ can be ignored without penalty. After doing this for all $I \subseteq H$ we have the same derivation as in the favorable case.

The proof for (ii) is essentially the same using Lemma 6.1(ii) instead of 6.1(i). For $l \leq d - 1$ the right side vanishes because σ_T does not intersect $\text{bd } A = \text{bd } \cup T$. \square

8. Decomposable Metric Formulas

From Theorem 6.2 we derive a second set of inclusion–exclusion formulas for $\cup B$. In contrast to Theorem 6.2, the new formulas have terms that express the contribution of individual simplices of \mathcal{D} . This is useful in situations where only a part of $\cup B$ or its complement is to be measured. Another advantage of the second set of formulas is that its terms correspond to intersections of at most d d -balls, rather than $d + 1$ as in Theorem 6.2.

Inclusion–Exclusion with Angle Weights

We first make the relation in Theorem 6.2 more complicated, and then replace or eliminate large parts using Theorem 7.3. It is convenient to cover the part of $\cup B$ outside $|\mathcal{D}|$ with simplices. This can be done by adding $d + 1$ points (degenerate d -balls), whose convex hull contains $\cup B$, to B . Consider Theorem 6.2(i) and decompose $\cap T$ into the parts defined by the d -simplices incident to σ_T . That is, use

$$\mu_d(\cap T) = \sum \varphi_{T,S} \cdot \mu_d(\cap T),$$

where the sum is taken over all $S \supseteq T$, $\text{card } S = d + 1$, so that $\sigma_S \in \mathcal{D}$. We need some notation. For subsets \mathcal{L} and \mathcal{L}' of a simplicial complex in \mathbb{R}^d let $\mathcal{L}^{[d]} = \mathcal{L}^{(d)} - \mathcal{L}^{(d-1)}$ be the collection of d -simplices $\sigma_S \in \mathcal{L}$, and let $[\mathcal{L}', \mathcal{L}]$ denote the collection of pairs (σ_T, σ_S) so that $\sigma_T \in \mathcal{L}'$ is a face of $\sigma_S \in \mathcal{L}^{[d]}$. With this notation, Theorem 6.2(i) becomes

$$\mu_d(\cup B) = \sum_{(\sigma_T, \sigma_S) \in [\mathcal{X}, \mathcal{D}]} (-1)^{\text{card } T - 1} \varphi_{T,S} \cdot \mu_d(\cap T).$$

Now we make a substitution using Theorem 7.3(i) whenever $\sigma_S \in \mathcal{X}$, and get the final result stated as Theorem 8.1(i). The same derivation works also for $\mu_l(\cup B)$, $0 \leq l \leq d - 1$. In this case the substitution uses Theorem 7.3(ii) and is, in fact, an elimination. We state the resulting second set of formulas for $\cup B$ and note that the remark after Theorem 6.2 also applies to Theorem 8.1.

Theorem 8.1.

(i)

$$\mu_d\left(\bigcup B\right) = \sum_{\sigma_S \in \mathcal{X}^{[d]}} \mu_d(\sigma_S) + \sum_{(\sigma_T, \sigma_S) \in [\mathcal{X}, \mathcal{D} - \mathcal{X}]} (-1)^{\text{card } T - 1} \varphi_{T,S} \cdot \mu_d\left(\bigcap T\right).$$

(ii) For $0 \leq l \leq d - 1$,

$$\mu_l\left(\bigcup B\right) = \sum_{(\sigma_T, \sigma_S) \in [\mathcal{X}, \mathcal{D} - \mathcal{X}]} (-1)^{\text{card } T - d + l} \varphi_{T,S} \cdot \mu_l\left(\bigcap T\right).$$

How can we interpret Theorem 8.1(i) in \mathbb{R}^2 ? It says the area of $\bigcup B$ can be computed as follows. First, take the triangles in \mathcal{X} and compute their total area. Second, for each vertex σ_T of \mathcal{S} , $T = \{b\}$, compute the angle, φ_T , around σ_T outside \mathcal{S} , and add φ_T times the area of $\bigcap T = b$ to the total area. Third, for each edge σ_T of \mathcal{S} , $T = \{b, b'\}$, subtract half the area of $\bigcap T = b \cap b'$ if there is one triangle in $\mathcal{D} - \mathcal{X}$ incident to σ_T , and subtract the entire area if there are two such triangles. Similar interpretations apply to Theorem 8.1(ii) and in higher dimensions.

Measuring a Void

Note that Theorem 8.1(i) consists of two sums. The first measures the d -dimensional part of \mathcal{S} , and the second measures the *fringe*, $\bigcup B - \mathcal{S}$. This relates to the considerations in Section 3, where the fringe is deformed in a continuous manner until it disappears. We can also measure the fringe simply by dropping the first sum in Theorem 8.1(i). This suggests it should be possible to measure a *void*, that is, a bounded component of $\mathbb{R}^d - \bigcup B$. In \mathbb{R}^3 , measuring voids is of some significance in the study of proteins [3], [22].

Let V_0 be a void of $\bigcup B$. As proved in Section 3, there is a void $\bar{\mathcal{F}}_0$ of \mathcal{S} that contains V_0 . Moreover, $\bar{\mathcal{F}}_0$ contains no other void of $\bigcup B$, that is, $V_0 = \bar{\mathcal{F}}_0 - \bigcup B$. It thus seems natural to collect all simplices $\sigma \in \mathcal{D} - \mathcal{X}$ with $\text{int } \sigma \subseteq \bar{\mathcal{F}}_0$ using the ideas described in Section 4. Call this set $\bar{\mathcal{X}}_0$ and note that $\bar{\mathcal{X}}_0$ is not a simplicial complex, but $\mathcal{D} - \bar{\mathcal{X}}_0$ is one. To measure V_0 we adapt the formulas in Theorem 8.1. Recall that $\bar{\mathcal{X}}_0^{[d]}$ is the collection of d -simplices $\sigma_S \in \bar{\mathcal{X}}_0$.

Theorem 8.2.

(i)

$$\mu_d(V_0) = \sum_{\sigma_S \in \bar{\mathcal{X}}_0^{[d]}} \mu_d(\sigma_S) - \sum_{(\sigma_T, \sigma_S) \in [\mathcal{X}, \bar{\mathcal{X}}_0]} (-1)^{\text{card } T - 1} \varphi_{T,S} \cdot \mu_d\left(\bigcap T\right).$$

(ii) For $0 \leq l \leq d - 1$,

$$\mu_l(V_0) = \sum_{(\sigma_T, \sigma_S) \in [\mathcal{X}, \bar{\mathcal{X}}_0]} (-1)^{\text{card } T - d + l} \varphi_{T,S} \cdot \mu_l\left(\bigcap T\right).$$

Proof. We cover the void V_0 with finitely many d -balls and consider the difference

in measure before and after adding the d -balls. Let B' be the set of d -balls that cover V_0 , and consider $\cup(B \cup B')$ and $\mathcal{X}' = \mathcal{X}(B \cup B')$. We require that (i) B' is finite, (ii) \mathcal{X} is a subcomplex of \mathcal{X}' , and (iii) $V_0 = \cup B' - \cup B$.

We argue that such a set B' exists. Choose $\varepsilon > 0$ small enough so that $\mathcal{X} = \mathcal{X}(B) = \mathcal{X}(B_\varepsilon)$, where

$$B_\varepsilon = \left\{ b_\varepsilon = \left(z, \sqrt{\rho^2 - \varepsilon^2} \right) \mid b = (z, \rho) \in B \right\}.$$

Note that $\mathcal{X}(B) = \mathcal{X}(B_\varepsilon)$, by definition of \mathcal{X} , and therefore $\mathcal{D} = \mathcal{D}(B) = \mathcal{D}(B_\varepsilon)$. Since general position of the d -balls in B is assumed, we can find ε small enough so that also the subcomplexes $\mathcal{X}(B) \subseteq \mathcal{D}$ and $\mathcal{X}(B_\varepsilon) \subseteq \mathcal{D}$ coincide. Now let B' be a sufficiently large set of d -balls $b' = (z', \varepsilon)$, with $z' \in V_0$, so that $V_0 \subseteq \cup B'$. Since V_0 is bounded and $\varepsilon > 0$ we can certainly choose B' finite. We show that (ii) and (iii) are also satisfied. Define B'_ε the same way as B_ε before. The balls of this set are degenerate, that is, B'_ε is a finite point set. Therefore, $\mathcal{X}(B_\varepsilon \cup B'_\varepsilon)$ is just $\mathcal{X}(B_\varepsilon)$ together with finitely many isolated 0-simplices. Hence, $\mathcal{X} = \mathcal{X}(B_\varepsilon) \subseteq \mathcal{X}(B_\varepsilon \cup B'_\varepsilon)$. From this (ii) follows because $\mathcal{X}(B_\varepsilon \cup B'_\varepsilon) \subseteq \mathcal{X}(B \cup B') = \mathcal{X}'$. If σ_τ is a simplex in $\mathcal{X}' - \mathcal{X}$, then $T \cap B' \neq \emptyset$. So $\mathcal{S}' - \mathcal{S} \subseteq \bar{\mathcal{S}}_0$, where $\mathcal{S}' = |\mathcal{X}'|$. In fact, $\mathcal{S}' - \mathcal{S} = \bar{\mathcal{S}}_0$ because $V_0 \subseteq \cup B'$. Condition (iii) follows because the correspondence between \mathcal{S}' and $\cup B'$ expressed in Lemma 2.2 guarantees that $\cup B'$ and $\cup B$ coincide outside V_0 .

So we have $\mu_d(V_0) = \mu_d(\cup B') - \mu_d(\cup B)$ and $\mu_l(V_0) = \mu_l(\cup B) - \mu_l(\cup B')$ for $0 \leq l \leq d - 1$. Note that the first sum in (i) is equal to the first sum of Theorem 8.1(i) for $\cup B'$ minus the first sum of Theorem 8.1(i) for $\cup B$. Similarly, the second sum in (i) is equal to the second sum of Theorem 8.1(i) for $\cup B'$ minus the second sum of Theorem 8.1(i) for $\cup B$. The sum in (ii) is Theorem 8.1(ii) for $\cup B$ minus Theorem 8.1(ii) for $\cup B'$. □

9. Discussion

This paper studies the union of finitely many d -balls in \mathbb{R}^d . It is demonstrated that many properties can be computed without explicit construction of the union. Instead, the nerve of the balls intersected with their respective (weighted) Voronoi cells is constructed. This is a simplicial complex that can be derived directly from the (weighted) Delaunay simplicial complex of the balls. For constant d , the size of this complex is no more than some constant times $n^{\lfloor d/2 \rfloor}$, where n is the number of balls, and for most distributions it is much less than that.

Specific algorithms are discussed that compute topological, combinatorial, and metric properties of the union of balls directly from the complex. The advantage of this complex, obtained by clipping balls to within their Voronoi cells, over the nerve of the set of unclipped balls is the significantly reduced size, see also [21]. It leads to much improved running times which make computations practical also for fairly large data sets. This is relevant to computational problems in biology, where proteins are modeled as unions of hundreds or thousands of 3-balls in \mathbb{R}^3 . For further applications it would be interesting to extend the inclusion–exclusion formulas of

Theorems 6.2, 8.1, and 8.2 to physical forces associated with a molecule. The most demanding step in obtaining running implementations of the algorithms in this paper is the construction of \mathcal{N} . Software for $d = 3$ is described in [13] and for $d = 2, 3$ it is available via ftp from ftp.ncsa.uiuc.edu. Algorithms in dimensions beyond three are described [9], [14]. The time-complexity of these algorithms depends on the distribution of the balls, and is often roughly of the same order as the number of simplices in \mathcal{D} . An implementation of the formulas in Theorems 6.2, 8.1, and 8.2 for \mathbb{R}^3 is also available and described in [10].

Acknowledgments

The author thanks Daniel Naiman for insightful suggestions improving the presentation and for pointers to related literature.

References

1. F. Aurenhammer. Improved algorithms for disks and balls using power diagrams. *J. Algorithms* **9** (1988), 151–161.
2. D. Avis, B. K. Bhattacharya, and H. Imai. Computing the volume of the union of spheres. *Visual Comput.* **3** (1988), 323–328.
3. M. L. Connolly. Analytical molecular surface calculation. *J. Appl. Cryst.* **16** (1983), 548–558.
4. T. H. Cormen, Ch. E. Leiserson, and R. L. Rivest. *Introduction to Algorithms*. MIT Press, Cambridge, MA, 1990.
5. B. Delaunay. Sur la sphère vide. *Izv. Akad. Nauk SSSR Otdel. Mat. Estestv. Nauk* **7** (1934), 793–800.
6. C. J. A. Delfinado and H. Edelsbrunner. An incremental algorithm for betti numbers of simplicial complexes. In *Proc. 9th Ann. Symp. on Computational Geometry*, 1993, pp. 232–239.
7. P. G. L. Dirichlet. Über die Reduktion der positiven quadratischen Formen mit drei unbestimmten ganzen Zahlen. *J. Reine Angew. Math.* **40** (1850), 209–227.
8. P. D. Domich. Residual Hermite normal form computations. *ACM Trans. Math. Software* **15** (1989), 275–286.
9. H. Edelsbrunner. Weighted alpha shapes. Report UIUCDCS-R-92-1760, Computer Science Department, University of Illinois, Urbana, IL, 1992.
10. H. Edelsbrunner, M. Facello, P. Fu, and J. Liang. Measuring proteins and voids in proteins. In *Proc. 28th Hawaii Internat. Conf. on System Sciences*, 1995, vol. V, pp. 256–264.
11. H. Edelsbrunner, D. G. Kirkpatrick, and R. Seidel. On the shape of a set of points in the plane. *IEEE Trans. Inform. Theory* **29** (1983), 551–559.
12. H. Edelsbrunner and E. P. Mücke. Simulation of simplicity: a technique to cope with degenerate cases in geometric algorithms. *ACM Trans. Graphics* **9** (1990), 66–104.
13. H. Edelsbrunner and E. P. Mücke. Three-dimensional alpha shapes. *ACM Trans. Graphics* **13** (1994), 43–72.
14. H. Edelsbrunner and N. R. Shah. Incremental topological flipping works for regular triangulations. In *Proc. 8th Ann. Symp. on Computational Geometry*, 1992, pp. 43–52.
15. M. Fuchs. A note on mapping cylinders. *Michigan Math. J.* **18** (1971), 289–290.
16. B. Grünbaum. *Convex Polytopes*. Wiley-Interscience, London, 1967.
17. R. Kannan and A. Bachem. Polynomial algorithms for computing the Smith and Hermite normal forms of an integer matrix. *SIAM J. Comput.* **8** (1979), 499–507.
18. J. R. Munkres. *Elements of Algebraic Topology*. Addison-Wesley, Redwood City, CA, 1984.
19. D. Q. Naiman and H. P. Wynn. A topological approach to discriminant analysis. *IMS Bull.* **19** (1990), 40–41.

20. D. Q. Naiman and H. P. Wynn. Inclusion–exclusion–Bonferroni identities and inequalities for discrete tube-like problems via Euler characteristics. *Ann. Statist.* **20** (1992), 43–76.
21. D. Q. Naiman and H. P. Wynn. Abstract tubes and improved inclusion–exclusion identities and inequalities. Manuscript, 1994.
22. F. M. Richards. Areas, volumes, packing, and protein structure. *Ann. Rev. Biophys. Bioeng.* **6** (1977), 151–176.
23. J. J. Rotman. *An Introduction to Algebraic Topology*. Springer-Verlag, New York, 1988.
24. G. F. Voronoi. Nouvelles applications des paramètres continus à la théorie des formes quadratiques. *J. Reine Angew. Math.* **133** (1907), 97–178.

Received September 23, 1993, and in revised form June 22, 1994.