

Inclusion–Exclusion Formulas from Independent Complexes*

Dominique Attali¹ and Herbert Edelsbrunner²

¹LIS laboratory, Domaine Universitaire,
BP 46, 38402 Saint Martin d’Hères, France
Dominique.Attali@lis.inpg.fr

²Department of Computer Science, Duke University,
Durham, NC 27708, USA
edels@cs.duke.edu
and
Geomagic,
Research Triangle Park, NC 27709, USA

Abstract. Using inclusion–exclusion, we can write the indicator function of a union of finitely many balls as an alternating sum of indicator functions of common intersections of balls. We exhibit abstract simplicial complexes that correspond to minimal inclusion–exclusion formulas. They include the dual complex, as defined in [3], and are characterized by the independence of their simplices and by geometric realizations with the same underlying space as the dual complex.

1. Introduction

In this paper we study inclusion–exclusion formulas for unions of finitely many balls in \mathbb{R}^d , generalizing previous results that derive such formulas from Delaunay triangulations and dual complexes.

Motivation. It is common in biochemistry to identify a molecule with the portion of space it occupies. This portion is sometimes referred to as the *space-filling diagram*, and its simplest and most common form is a union of balls in \mathbb{R}^3 , each ball representing an atom of the molecule. The volume and surface area of this union are fundamental concepts that relate to physical forces acting on the molecules. We refer to [4] for a

* The authors were partially supported by the EU under Contract IST-2002-506766 (Aim@Shape), by the NSF under Grant CCR-00-86013 (Bio-geometry), and by the CNRS under Grant PICS-3416.

recent survey that describes this connection and also discusses derivatives of the volume and surface area, which are needed in simulating the motion of molecules.

We consider a finite set of balls in \mathbb{R}^3 and focus on the volume of the union. Generally, there are many inclusion–exclusion formulas that give the correct volume, even if we limit our attention to minimal formulas. The starting point of the work reported in this paper is the idea that this ambiguity could be useful in maintaining a formula for a moving set of balls. If we understand how long a formula remains valid, we can save time by delaying any changes until they become necessary. As a first step towards such an understanding, we study the family of minimal inclusion–exclusion formulas for a given set of balls.

Prior Work. The principle of inclusion–exclusion is perhaps the most natural approach to measuring a union of measurable sets. Letting B be a finite set of balls, the volume of the union is the alternating sum of volumes of intersections:

$$\text{vol} \bigcup B = \sum_{\emptyset \neq X \subseteq B} (-1)^{\text{card } X - 1} \text{vol} \bigcap X, \quad (1)$$

where $\text{card } X$ is the number of balls in X . Writing n for the number of balls in B , we have $2^n - 1$ terms, each plus or minus the volume of the intersection of at most n balls. It seems the formula is only useful if all sets X with non-empty common intersection are small. More than a quarter of a century ago, Kratky [7] pointed out that even if this is not the case, one can substitute lower-order for higher-order terms and thus reduce the complexity of the formula. The software of Scheraga and collaborators [9] is based on this observation, but it is sometimes difficult to do the reduction correctly. In 1992 Naiman and Wynn [8] proved that (1) is correct even if we limit the sum to sets X that correspond to simplices in the weighted Delaunay triangulation of B . By definition, this is the dual of the weighted Voronoi diagram of the balls, sometimes called the power diagram or the Dirichlet tessellation [1]. In the geometry literature, this dual is also known as the regular triangulation and the coherent triangulation of B [5]. In agreement with Kratky, this result implies that in \mathbb{R}^3 we only need sets X of cardinality at most four. Taking all such sets would lead to an incorrect formula, and Naiman and Wynn’s result is a recipe for selecting sets that give a correct formula. In 1995 Edelsbrunner [3] further reduced the formula by proving that (1) remains correct if we limit the sum to simplices in the dual complex, which is a subcomplex of the weighted Delaunay triangulation of B . Besides giving a shorter formula, the terms obtained from the dual complex consist of balls that intersect in a unique pattern, which allows a simple implementation without case analysis [6].

Results. We refer to the specific intersection pattern exhibited by the balls in the dual complex formula as independent, a term whose technical definition will be given shortly. Our first result generalizes this formula to a family of formulas whose terms exhibit the same pattern. Specifically, *if K is an abstract simplicial complex whose simplices are independent sets of balls and whose canonical geometric realization has the same boundary complex and underlying space as the dual complex then the corresponding inclusion–exclusion formula is correct.* To prove that this is a proper generalization, we show in Fig. 1 that even already for four disks in the plane we can have more than




Fig. 1. Four disks that permit two correct, minimal inclusion–exclusion formulas. The upper complex on the right is the dual complex of the disks and corresponds to the formula $a + b + c + d - ab - ad - bc - bd - cd + abd + bcd$, in which we write a for the area of disk a , ab for the area of the intersection of a and b , etc. To get the formula of the lower complex, we substitute $-ac + abc + acd$ for $-bd + abd + bcd$.

one such formula. Our second result states that *the inclusion–exclusion formulas in the family specified in our first result are minimal and exhaust all minimal formulas that correspond to simplicial complexes.*

Outline. Section 2 presents definitions and the formal statements of our two results. Section 3 proves the first result and Section 4 proves the second. Section 5 concludes this paper.

2. Statement of Results

In this paper a simplex may be abstract (a collection of balls) or geometric (the convex hull of affinely independent points). We use both interpretations interchangeably and introduce notation that does not distinguish between them.

Independent Simplices. Let B be a finite set of closed balls in \mathbb{R}^d . Throughout this paper we assume that the balls are in general position, which includes that every ball has positive radius and the common intersection of any $k + 1$ bounding spheres is either empty or a (non-degenerate) sphere of dimension $d - k - 1$. In particular, this prevents the common intersection degenerating to a single point. An *abstract simplex* is a non-empty subset $\alpha \subseteq B$ and its *dimension* is one less than its cardinality, $\dim \alpha = \text{card } \alpha - 1$. A k -*simplex* is an abstract simplex of dimension k . If $\beta \subseteq \alpha$ is a non-empty subset, we call β a *face* of α and α a *coface* of β . A simplex α is *independent* if for every subset $\gamma \subseteq \alpha$, including $\gamma = \emptyset$, there is a point that belongs to all balls in γ but not to any ball not in γ :

$$\bigcap \gamma - \bigcup (\alpha - \gamma) \neq \emptyset.$$

By assumption of general position, each set of the form $\bigcap \gamma - \bigcup (\alpha - \gamma)$ contains an open subset of \mathbb{R}^d . Clearly, every face of an independent simplex is independent. In \mathbb{R}^2 there are only three types of independent simplices, one each for one, two, and three disks, as shown in Fig. 2. Four disks cannot be independent because the four bounding circles decompose the plane into at most 14 regions but we need 16, one each for the 2^4 subsets. Similarly, in \mathbb{R}^d we have $d + 1$ types of independent simplices. To construct an




Fig. 2. From left to right: an independent simplex of dimension $k = 0, 1, 2$.

independent k -simplex, for $k \leq d$, we consider the standard k -simplex whose vertices are the endpoints of the Cartesian coordinate frame of \mathbb{R}^{k+1} . This is a regular simplex with edges of length $\sqrt{2}$ and vertices at distance 1 from the origin. Placing the centers of $k + 1$ unit balls at these vertices, it is not difficult to see that the centroid of any subset of $j + 1$ centers lies inside the corresponding $j + 1$ balls but outside the other $k - j$ balls. There are no independent simplices of dimension $k > d$ because in this case $k + 1$ $(d - 1)$ -spheres decompose \mathbb{R}^d into fewer than 2^{k+1} open regions (see the Chamber Lemma below).

General Position. We think of \mathbb{R}^d as the subspace of \mathbb{R}^{d+1} spanned by the first d coordinate axes. For each d -dimensional ball b_i with center $z_i \in \mathbb{R}^d$ and radius $r_i > 0$, we construct the corresponding *ghost sphere*,

$$s_i = \{y \in \mathbb{R}^{d+1} \mid \|y - z_i\|^2 - r_i^2 = 0\},$$

which is a d -sphere in \mathbb{R}^{d+1} . Using this concept, we can now define what exactly we mean by a set of balls to be in *general position*, namely that the common intersection of any $k + 1$ ghost spheres is either empty or a sphere of dimension $d - k$. We note that for $0 \leq k < d$, this is equivalent to the condition we mentioned earlier: all radii are positive and the common intersection of any $k + 1$ bounding spheres is either empty or a sphere of dimension $d - k - 1$. For $k \geq d$ we get new conditions. We need some definitions before we can explain them in terms of concepts intrinsic to \mathbb{R}^d .

Call $\pi_i(x) = \|x - z_i\|^2 - r_i^2$ the *power distance* of the point $x \in \mathbb{R}^d$ from b_i and note that $b_i = \pi_i^{-1}(-\infty, 0]$ and the zero-set of π_i is the $(d - 1)$ -sphere bounding b_i . Using the power distance, we decompose $\bigcup B$ into convex cells, one for each ball. Specifically, the cell of b_i consists of all points $x \in b_i$ with $\pi_i(x) \leq \pi_j(x)$ for all $b_j \in B$. It is not difficult to see that the cell of b_i is the intersection of b_i with a convex polyhedron, keeping in mind that this polyhedron or its intersection with b_i may be empty. To describe the relation between this cell and the ghost sphere of b_i , we extend π_i to \mathbb{R}^{d+1} by defining $\sigma_i(y) = \|y - x\|^2 + \pi_i(x)$, where x is the orthogonal projection of $y \in \mathbb{R}^{d+1}$ to \mathbb{R}^d . The ghost sphere itself is the zero-set, $s_i = \sigma_i^{-1}(0)$. For a given point $x \in b_i$, define $y_i(x) \in \mathbb{R}^{d+1}$ above \mathbb{R}^d such that $\|y_i(x) - x\|^2 + \pi_i(x) = 0$; it is the point on the upper hemi-sphere of s_i whose orthogonal projection to \mathbb{R}^d is x . The condition for x to belong to the cell of b_i now translates to $\|y_i(x) - x\|^2 \geq \|y_j(x) - x\|^2$ whenever $y_j(x)$ is defined. In words, the cell of b_i is the orthogonal projection of s_i 's contribution to the upper envelope of the ghost spheres, as illustrated in Fig. 3.

Let us return to the case $k \geq d$ of our general position assumption. It says that the common intersection of any $k + 1 \geq d + 1$ ghost spheres is either empty or a 0-sphere, and the latter case can only happen if $k = d$. Equivalently, the common intersection of




Fig. 3. Upper envelope of the ghost spheres of three independent disks and the corresponding decomposition of the union of disks into convex cells.

the cells decomposing the union of any $k + 1 \geq d + 1$ balls is either empty or a point in the interior of the union, and the latter case can happen only if $k = d$.

Characterizing Independence. Besides for expressing our general position assumption, ghost spheres can be used for characterizing independent simplices. This characterization will be important in establishing the Non-Nesting Lemma in Section 3, a crucial step in the proof of our first result. We make use of stereographic projections to flatten d -spheres into d -planes. Assuming a d -sphere s passing through the origin of some Euclidean space, the stereographic projection is the restriction of the inversion map to the sphere minus the origin, N , sending every point $x \in s - \{N\}$ to $x/\|x\|^2$. The image of $s - \{N\}$ is a d -dimensional plane, and the image of a k -sphere in $s - \{N\}$ is a k -sphere in that image plane.

To prepare the characterization of independent simplices, we prove a technical result on $(d - 1)$ -spheres in \mathbb{R}^d . Any finite collection of such spheres decomposes this space into open cells. We refer to the (connected) cells of dimension d as *chambers*, but for technical reasons we modify the notion of connectivity by first compactify \mathbb{R}^d with a single point at infinity. The only case in which the compactification makes a difference is in dimension 1 where the two unbounded intervals in the decomposition defined by 0-spheres form a single chamber.

Chamber Lemma. *The number of chambers in an arrangement of $k + 1$ $(d - 1)$ -spheres in \mathbb{R}^d is at most 2^{k+1} and it is strictly less than 2^{k+1} unless $k \leq d$.*

Proof. We use induction over the number of spheres as well as the dimension. The induction basis consists of $k + 1 = 1$ $(d - 1)$ -sphere in \mathbb{R}^d forming two chambers and of $k + 1$ 0-spheres in \mathbb{R}^1 forming at most $2k + 2$ chambers. To prove the upper bound in the general case, let s be one of the $k + 1$ $(d - 1)$ -spheres. By induction hypothesis, the remaining k $(d - 1)$ -spheres define an arrangement with at most 2^k chambers. These spheres also decompose s into $(d - 1)$ -dimensional patches, and we complete the arrangement by adding one patch at a time. To get an upper bound on the number of patches, we use stereographic projection from a point N on s , mapping $s - \{N\}$ to a $(d - 1)$ -dimensional plane and its patches to the chambers in the arrangement of k $(d - 2)$ -spheres, the image of the intersection between s and the other $(d - 1)$ -spheres.

By induction hypothesis, there are at most $2^k (d - 1)$ -dimensional chambers and thus at most 2^k patches on s . Each such patch cuts a single d -dimensional chamber and either leaves it connected or decomposes it into two. This implies the claimed upper bound of 2^{k+1} for the number of chambers.

To see that 2^{k+1} chambers cannot be attained for $k > d$ we note that the above argument reduces the problem for $k + 1 (d - 1)$ -spheres in \mathbb{R}^d to $k (d - 1)$ -spheres in \mathbb{R}^d and to $k (d - 2)$ -spheres in \mathbb{R}^{d-1} . Furthermore, a maximum number of chambers for both subproblems is necessary to get 2^{k+1} chambers for the original problem. If $k > d$ then the second subproblem consists of $k > d (d - 2)$ -spheres in \mathbb{R}^{d-1} which, by induction hypothesis, form strictly fewer than 2^k chambers. \square

We are now ready to characterize the independent simplices.

Ghost Sphere Lemma. *A k -simplex of $k + 1$ balls in general position in \mathbb{R}^d is independent iff the common intersection of its $k + 1$ ghost spheres is a (non-degenerate) sphere of dimension $d - k \geq 0$.*

Proof. We prove the lemma by establishing five equivalences, $\mathbf{A} \Leftrightarrow \mathbf{B} \Leftrightarrow \mathbf{C} \Leftrightarrow \mathbf{D} \Leftrightarrow \mathbf{E} \Leftrightarrow \mathbf{F}$, connecting statement \mathbf{A} : “a k -simplex of $k + 1$ balls in general position in \mathbb{R}^d is independent” with statement \mathbf{F} : “the common intersection of $k + 1$ ghost spheres is a (non-degenerate) sphere of dimension $d - k \geq 0$ ”. For the fourth equivalence, we use induction over the dimension. By definition, \mathbf{A} is equivalent to \mathbf{B} : “the $(d - 1)$ -spheres bounding the $k + 1$ balls form an arrangement with 2^{k+1} chambers and any two of the chambers are separated by at least one of the $(d - 1)$ -spheres”. Since we will encounter the latter property again, we call such a set of chambers *pairwise separated*.

From \mathbf{B} we prove statement \mathbf{C} : “the $k + 1$ ghost spheres form an arrangement of 2^{k+1} pairwise separated chambers in \mathbb{R}^{d+1} ”. By the Chamber Lemma there are at most 2^{k+1} chambers in \mathbb{R}^{d+1} . They intersect \mathbb{R}^d in 2^{k+1} (d -dimensional) chambers, and because the latter are pairwise separated, each chamber in \mathbb{R}^{d+1} meets \mathbb{R}^d in a unique chamber and the chambers in \mathbb{R}^{d+1} are also pairwise separated. To prove $\mathbf{C} \Rightarrow \mathbf{B}$, we use the fact that the arrangement of ghost spheres is symmetric with respect to \mathbb{R}^d . The fact that its chambers are pairwise separated then implies that each intersects \mathbb{R}^d in a single (d -dimensional) chamber and that the latter are also pairwise separated.

For the next step, let s be one of the $k + 1$ ghost spheres, let N be a point of s , and map $s - \{N\}$ by stereographic projection to a d -dimensional plane, which we refer to as \mathbb{R}^d . The other k ghost spheres decompose s into patches, which map to (d -dimensional) chambers in an arrangement of $k (d - 1)$ -spheres. Since the chambers in the ghost sphere arrangement are pairwise separated, this implies \mathbf{D} : “there are 2^k pairwise separated chambers in \mathbb{R}^d ”. The reverse implication, $\mathbf{D} \Rightarrow \mathbf{C}$, can be seen by remembering that each patch of s separates two chambers in \mathbb{R}^{d+1} , one inside and the other outside s . The inside chambers are pairwise separated by the k ghost spheres because the patches are pairwise separated. Similarly, the outside chambers are pairwise separated. Finally, s separates all inside from all outside chambers implying that the chambers in \mathbb{R}^{d+1} are pairwise separated.

By induction hypothesis, \mathbf{D} is equivalent to \mathbf{E} : “the $k (d - 1)$ -spheres meet in a common sphere of dimension $d - k \geq 0$ ”. Mapping the $(d - k)$ -sphere by the inverse of




Fig. 4. A finite set of disks and the canonical realization of an abstract simplicial complex over that set. The vertices, edges, and triangles correspond to abstract simplices of dimension 0, 1, and 2. Take a moment to verify that all simplices are independent.

the stereographic projection gives another $(d - k)$ -sphere that belongs to all $k + 1$ ghost spheres and thus implies **F**. It is clear that **F** implies **E**, which completes the proof. \square

Simplicial Complexes. An *abstract simplicial complex* is a collection of non-empty abstract simplices, K , that contains, with every simplex, the faces of that simplex. If B is the set of vertices then K is a subset of the power set, $K \subseteq 2^B$. Figure 4 illustrates the definitions. A *geometric realization* maps every abstract simplex to a geometric simplex of the same dimension such that the intersection of the images of two abstract simplices α and β is the image of $\alpha \cap \beta$, which is either empty or a face of both. In this paper the vertices are closed balls and we map every abstract simplex to its *canonical image*, defined as the convex hull of the centers of its balls. We call K *canonically realizable* if this map is a geometric realization. We use the letters α and β to denote the abstract simplices as well as their images, which are geometric simplices. Similarly, we use the letter K to denote the abstract simplicial complex as well as its geometric realization, which is a geometric simplicial complex. Its *underlying space* is the set of points covered by the geometric simplices, which we denote as $|K|$. The *star* of an abstract simplex β is the set of cofaces $\alpha \in K$, and the *link* of β is the set of simplices $\alpha - \beta$ with $\alpha \supset \beta$. Assuming K is geometrically realized in \mathbb{R}^d , the link of every k -simplex is a triangulation of the sphere of dimension $d - k - 1$ or a proper subcomplex of such a triangulation. We define the *boundary complex* of K as the subset of simplices in the latter category. This is also the subcomplex of simplices contained in the boundary of $|K|$.

Dual and Other Independent Complexes. Let B be a set of closed balls and recall the decomposition of $\bigcup B$ into convex cells described above. The nerve of this collection of cells is particularly important for the results described in this paper. The *dual complex* of B is the canonical realization of this nerve, obtained by mapping every $k + 1$ cells with non-empty intersection to the k -simplex spanned by the centers of the corresponding balls. This construction is illustrated in Fig. 5, where we see the dual complex superimposed on the decomposition of the union into convex cells. It is perhaps not obvious but true




Fig. 5. The dual complex of the disks in Figure 4. Its simplices record the overlap pattern of the cells in the decomposition of the union. In this example, the dual complex has the same boundary complex and underlying space as the independent complex in Figure 4 but differs from it in six edges and twelve triangles.

that the canonical mapping of abstract simplices defines a geometric realization of the nerve, provided the balls in B are in general position [3].

Given a finite set of balls in general position, B , we are primarily interested in abstract simplicial complexes K of B that satisfy the following three conditions:

Independence: all simplices in K are independent.

Realizability: K is canonically realizable in \mathbb{R}^d .

Boundary: the boundary complex and underlying space of K are the same as those of the dual complex.

An *independent complex* is an abstract simplicial complex that satisfies the independence condition. We note that there is an alternative way to express the boundary condition, without references to the dual complex, by comparing the boundaries of K and $\bigcup B$. In particular, a simplex α belongs to the boundary complex of K iff there is a point on the boundary of $\bigcup B$ that belongs to all balls in α and to no others.

First Result: Indication. The *indicator function* of a subset $A \subseteq \mathbb{R}^d$ is the map $\mathbf{1}_A: \mathbb{R}^d \rightarrow \{0, 1\}$ defined by

$$\mathbf{1}_A(x) = \begin{cases} 1 & \text{if } x \in A, \\ 0 & \text{if } x \notin A. \end{cases}$$

Given a finite set of balls in \mathbb{R}^d , our first result states that the inclusion–exclusion formula defined by a simplicial complex that satisfies the above three conditions gives the correct indicator function of the union.

Theorem A. *Let B be a finite set of closed balls in general position in \mathbb{R}^d and let K be an independent complex that is canonically realizable in \mathbb{R}^d and satisfies the boundary condition. Then $\mathbf{1}_{\bigcup B} = \sum_{\alpha \in K} (-1)^{\dim \alpha} \mathbf{1}_{\cap \alpha}$.*

Using Theorem A, we obtain formulas for the volume or other measures of the union by integrating the density function, $\mu: \mathbb{R}^d \rightarrow \mathbb{R}$:

$$\begin{aligned} \text{meas} \bigcup B &= \int_{x \in \bigcup B} \mu(x) \, dx \\ &= \int_{x \in \mathbb{R}^d} \mu(x) \mathbf{1}_{\bigcup B}(x) \, dx \\ &= \sum_{\alpha \in K} (-1)^{\dim \alpha} \int_{x \in \mathbb{R}^d} \mu(x) \mathbf{1}_{\cap \alpha}(x) \, dx \\ &= \sum_{\alpha \in K} (-1)^{\dim \alpha} \text{meas} \cap \alpha. \end{aligned}$$

For $d = 2$, the edge skeleton of K is a planar graph implying that the number of terms in the above formula is less than six times $n = \text{card } B$. More generally, the number of terms is bounded from above by some constant times $n^{\lceil d/2 \rceil}$ [2].

Second Result: Minimality. The inclusion–exclusion formula corresponding to an abstract simplicial complex $K \subseteq 2^B$ gives a map $\text{IEF}_K: \mathbb{R}^d \rightarrow \mathbb{Z}$ defined by

$$\text{IEF}_K(x) = \sum_{\alpha \in K} (-1)^{\dim \alpha} \mathbf{1}_{\cap \alpha}(x).$$

The formula is *minimal* if $\text{IEF}_L \neq \text{IEF}_K$ for all proper subsets $L \subset K$. By Theorem A we have $\text{IEF}_K = \mathbf{1}_{\bigcup B}$ if K is an independent complex that is canonically realizable in \mathbb{R}^d and satisfies the boundary condition. Our second result states that such complexes have minimal formulas and that they exhaust the family of complexes with minimal formulas.

Theorem B. *Let B be a finite set of closed balls in general position in \mathbb{R}^d and let $K \subseteq 2^B$ be an abstract simplicial complex with $\text{IEF}_K = \mathbf{1}_{\bigcup B}$. This formula is minimal iff K is independent, canonically realizable in \mathbb{R}^d , and satisfies the boundary condition.*

3. Proof of Theorem A

In this section we present our proof of Theorem A. Starting with a finite set of balls, first we add small balls covering the rest of \mathbb{R}^d to get an infinite but discrete set, and second we use this discrete set as the basis for a continuous set. Both steps are instrumental in obtaining the technical results that imply Theorem A.

Induced Subcomplexes. Given an abstract simplicial complex $K \subseteq 2^B$, a subset $B_0 \subseteq B$ induces the subcomplex $K_0 = K \cap 2^{B_0}$. To establish our first result, we associate to each point $x \in \mathbb{R}^d$ the subset $B_x \subseteq B$ of balls that contain x and the subcomplex $K_x \subseteq K$ induced by B_x . We have

$$\begin{aligned} \text{IEF}_K(x) &= \sum_{\alpha \in K} (-1)^{\dim \alpha} \mathbf{1}_{\cap \alpha}(x) \\ &= \sum_{\alpha \in K_x} (-1)^{\dim \alpha}. \end{aligned}$$

The latter sum is the Euler characteristic of K_x , which we denote as $\chi(K_x)$. For all points $x \notin \bigcup B$, $K_x = \emptyset$ and $\text{IEF}_K(x) = \chi(K_x) = 0$. To tackle the points inside the union, we recall that the Euler characteristic of every contractible set is 1. As explained later, such a set has the homotopy type of a point, which in the plane includes trees and closed disks. We will show that $|K_x|$ is contractible, for every point $x \in \bigcup B$, which will then imply Theorem A.

The union of the balls in B_x is star-convex, which implies that $\bigcup B_x$ is contractible. In spite of the fact that the underlying space of the subcomplex K_x induced by B_x is not necessarily star-convex, we will be able to prove that $|K_x|$ is also contractible. Before embarking on this proof, we introduce the anticipated discrete and continuous sets of balls. Using the continuous set, we will find a set between $|K_x|$ and $\bigcup B_x$, which we will show is star-convex and of the same homotopy type as $|K_x|$.

From Finite to Discrete Sets of Balls. A simplicial complex is *locally finite* if the star of every vertex is finite. We extend the finite set of balls B to a discrete set $\bar{B} = B \cup B_\varepsilon$. Simultaneously, we construct a locally finite independent complex $\bar{K} \supseteq K$ whose vertices are the balls in \bar{B} and whose underlying space is \mathbb{R}^d . The construction depends on a positive number ε , the radius of the balls in B_ε added to B . We require that \bar{B} covers \mathbb{R}^d while the center of every ball in B_ε lies outside all other balls in \bar{B} , as illustrated in Fig. 6. Choosing $\varepsilon > 0$ sufficiently small, we construct B_ε one ball at a time, picking the center outside all previous balls, until \bar{B} covers \mathbb{R}^d . Assuming the balls in B are in general position, it is clear that we can construct \bar{B} such that its balls are also in general position. To see which ε is sufficiently small, we consider the cells in the decomposition of $\bigcup B$. As we add balls of radius ε , these cells give up territory to the new balls, but not more than what is covered by the new balls. By shrinking ε , we can make the loss of territory as narrow as we like. By assumption of general position, we can therefore guarantee that any non-empty common intersection of cells in the decomposition of $\bigcup B$ remains non-empty in the decomposition of $\bigcup \bar{B}$. It follows that the dual complex of B is a subcomplex of the dual complex of \bar{B} . The boundary complex of the dual complex




Fig. 6. Extension of the independent complex in Fig. 4 by adding disks of radius ε . The rectangular frame delimits the portion of the configuration reused in Fig. 7.

of B is the same as the boundary complex of K . We can therefore construct \bar{K} equal to K inside and equal to the dual complex of \bar{B} outside that boundary, as illustrated in Fig. 6. We finally note the choice of balls implies that \bar{K} is locally finite.

From Discrete to Continuous Sets of Balls. An abstract simplex is a finite set of balls, $\alpha = \{b_0, b_1, \dots, b_k\}$. We extend α to an infinite set by considering convex combinations of balls in α . Recall that $\pi_i: \mathbb{R}^d \rightarrow \mathbb{R}$ maps each point $x \in \mathbb{R}^d$ to its power distance from b_i and that $b_i = \pi_i^{-1}(-\infty, 0]$. An *affine combination* of α is a ball $b = \pi^{-1}(-\infty, 0]$ for which there are real numbers λ_i , summing to 1, such that $\pi = \sum_{i=0}^k \lambda_i \pi_i$. To determine the center z of b observe that the gradient of π at a point $x \in \mathbb{R}^d$ is

$$\begin{aligned} \nabla \pi(x) &= \sum_{i=0}^k \lambda_i \nabla \pi_i(x) \\ &= \sum_{i=0}^k 2\lambda_i (x - z_i) \\ &= 2 \left(x - \sum_{i=0}^k \lambda_i z_i \right), \end{aligned}$$

where z_i is the center of b_i . This gradient vanishes at the center of b , which implies $z = \sum_{i=0}^k \lambda_i z_i$. When α is independent, we can use the ghost spheres of the b_i to complete the characterization of b . Since the σ_i are the same as the π_i , only extended from \mathbb{R}^d to \mathbb{R}^{d+1} , it is clear that the ghost sphere s of b is the zero-set of $\sigma = \sum_{i=0}^k \lambda_i \sigma_i$. A point $y \in \mathbb{R}^{d+1}$ belongs to the common intersection of the ghost spheres of the b_i , which is a $(d - k)$ -sphere, iff $\sigma_i(y) = 0$ for $0 \leq i \leq k$. It follows that $\sigma(y) = 0$, implying that y also belongs to the ghost sphere of b . We summarize these findings for later reference.

Affine Combination Lemma. *If b is the affine combination of the balls b_i with centers z_i determined by coefficients λ_i , for $0 \leq i \leq k$, then the center of b is $z = \sum_{i=0}^k \lambda_i z_i$. If furthermore the b_i form an independent k -simplex then the ghost sphere of b passes through the common intersection of the ghost spheres of the b_i .*

A *convex combination* b of α is an affine combination for which all λ_i are non-negative. If a point x belongs to all balls in α then $\pi_i(x) \leq 0$, for all i , which implies $\pi(x) \leq 0$. Furthermore, if $\pi(x) \leq 0$ then $\pi_i(x) \leq 0$ for at least one index i , which implies that x belongs to at least one ball in α . We rewrite this finding in set notation for later reference.

Convex Combination Lemma. *If b is a convex combination of the balls in α then $\bigcap \alpha \subseteq b \subseteq \bigcup \alpha$.*

Letting $\alpha \in K$ be the simplex whose interior contains the point $z \in \mathbb{R}^d$, we write b_z for the (unique) convex combination of balls in α whose center is z . The Affine Combination Lemma is useful when we consider a line and the balls b_z whose centers z lie on the line. These balls intersect the line in intervals. It turns out that as we move

the center monotonically along the line, the left endpoint also moves monotonically and so does the right endpoint. It is convenient to prove this for the extension \bar{K} of K for which there are balls for all points along the line. As usual, we assume the balls in \bar{B} are in general position.

Non-Nesting Lemma. *For any two points $x \neq y$ in \mathbb{R}^d , the two balls b_x and b_y are either disjoint or independent.*

Proof. Consider first the case in which x and y are points of a common d -simplex α in \bar{K} . Since α is independent, the ghost spheres of α intersect in a common 0-sphere. By the second part of the Affine Combination Lemma, the ghost spheres of b_x and b_y pass through this 0-sphere and thus meet in a $(d - 1)$ -sphere. It follows that b_x and b_y are independent.

If x and y do not belong to a common d -simplex, there is a point z on the line segment connecting x and y that lies on a $(d - 1)$ -simplex. The number of $(d - 1)$ -simplices separating x from z is strictly smaller than the number separating x from y , and similar for z and y . We can therefore use induction to show that b_x and b_z as well as b_z and b_y are either disjoint or independent. Consider the directed line that passes through x, z, y , in this sequence. It intersects b_x, b_z, b_y in three intervals. Because b_x and b_z are non-nested so are the intervals in which b_x and b_z meet the line. Similarly, the intervals in which b_z and b_y meet the line are non-nested. However, because x, z, y are the centers of these intervals, the left endpoints are ordered the same way as x, z, y and so are the right endpoints. It follows that the intervals and therefore the balls of x and y are non-nested. \square

Intermediate Set of Centers. For a point $x \in \mathbb{R}^d$, we write $\bar{B}_x \subseteq \bar{B}$ for the set of balls that contain x and $\bar{K}_x \subseteq \bar{K}$ for the subcomplex induced by \bar{B}_x . We prove that $|\bar{K}_x|$ is contractible by showing it has the same homotopy type as

$$\bar{Z}_x = \{z \in \mathbb{R}^d \mid x \in b_z\},$$

which we later prove is star-convex. Let α be the simplex whose interior contains the point $z \in \mathbb{R}^d$. By the Convex Combination Lemma, $z \in \bar{Z}_x$ if all balls in α contain x . Similarly, $z \notin \bar{Z}_x$ if none of the balls in α contains x . As illustrated in Fig. 7, the first property implies $|\bar{K}_x| \subseteq \bar{Z}_x$. Let \bar{L}_x be the subcomplex of \bar{K} induced by $\bar{B} - \bar{B}_x$. Each vertex of \bar{K} is either in \bar{K}_x or in \bar{L}_x . It follows that each abstract simplex α that is in neither induced subcomplex is the union of its largest faces $\alpha_K \in \bar{K}_x$ and $\alpha_L \in \bar{L}_x$. The corresponding geometric construction writes α as the union of line segments pq connecting points $p \in \alpha_K$ with points $q \in \alpha_L$. Because of the Convex Combination Lemma, b_p contains x and b_q does not contain x . Since p and q belong to a common independent simplex, the ghost spheres of b_p and b_q intersect in a $(d - 1)$ -sphere. It follows that the bounding spheres of b_p and b_q intersect in a $(d - 2)$ -sphere, namely the portion of the $(d - 1)$ -sphere in \mathbb{R}^d . This $(d - 2)$ -sphere is also contained in the bounding sphere of every ball along the line segment between p and q . As we move from p to q , the ball shrinks on the side of the $(d - 1)$ -plane passing through this $(d - 2)$ -sphere that contains x . This implies that there is a unique point y on pq such that $x \in b_y$ for




Fig. 7. The point x is contained in four disks, which induce the dark-gray subcomplex \bar{K}_x , consisting of two triangles and their faces. The other disks induce the light-gray subcomplex \bar{L}_x . We get \bar{Z}_x by adding initial portions of the line segments covering the white in-between simplices to the underlying space of \bar{K}_x .

all z between p and y , including y , and $x \notin b_z$ for all z between y and q , excluding y . In other words, \bar{Z}_x can be written as the union of $|\bar{K}_x|$ and all line segments py as described. These line segments can be shrunk continuously towards $|\bar{K}_x|$. Formally, we define $z(\lambda) = (1 - \lambda)z + \lambda p$ for each point z on py and each $0 \leq \lambda \leq 1$. Since p varies continuously with z , we have here a continuous map from $\bar{Z}_x \times [0, 1]$ to \bar{Z}_x that keeps the points of $|\bar{K}_x|$ fixed. We thus constructed a deformation retraction that takes \bar{Z}_x to $|\bar{K}_x|$, proving that the two have the same homotopy type.

Homotopy Type Lemma. $|\bar{K}_x| \simeq \bar{Z}_x$.

As mentioned earlier, a contractible set has the homotopy type of a point. By the above lemma, $|\bar{K}_x|$ is contractible iff \bar{Z}_x is contractible. We prove the latter by showing that \bar{Z}_x is the union of line segments emanating from a common endpoint. This implies that \bar{Z}_x is contractible because we can again exhibit a deformation retraction by shrinking the line segments, this time toward their common endpoint.

Star-Convexity Lemma. \bar{Z}_x is star-convex.

Proof. Observing that $x \in \bar{Z}_x$, we show that any line that passes through x intersects \bar{Z}_x in a single line segment. To reach a contradiction, assume there are points y and z on such a line through x such that z lies strictly between x and y and $x \in b_y$ but $x \notin b_z$. Then $b_z \subseteq b_y$, which contradicts the Non-Nesting Lemma. \square

It is not too difficult to show that the boundary of \bar{Z}_x is piecewise linear, as suggested by Fig. 7. In other words, \bar{Z}_x is a star-convex polytope.

Finale. We finally state and prove the crucial technical result that implies Theorem A.

Contractibility Lemma. Let B be a finite set of closed balls in general position in \mathbb{R}^d . For every point $x \in \bigcup B$, the underlying space of the subcomplex K_x induced by the balls that contain x is contractible.

Proof. We first establish the result for points x in the interior of $\bigcup B$. We may assume that $\varepsilon > 0$ is sufficiently small such that none of the balls in B_ε contains x . Hence, $\bar{B}_x = B_x$ and $\bar{K}_x = K_x$. By the Homotopy Type Lemma, $|K_x|$ and \bar{Z}_x have the same homotopy type, and by the Star-Convexity Lemma, \bar{Z}_x is star-convex and therefore contractible. It follows that $|K_x|$ is contractible. By assumption of general position, every point x on the boundary has a point y in the interior of $\bigcup B$ that is contained in the same balls of B as x . Therefore, $K_x = K_y$ and the claim follows by the first argument. \square

As mentioned earlier, the contractibility of $|K_x|$ implies $\text{IEF}_K(x) = \chi(K_x) = 1$ for all points $x \in \bigcup B$. Theorem A follows.

4. Proof of Theorem B

In this section we present a proof of Theorem B. We begin by establishing (2) as our main technical tool.

Witness Points. Let B be a finite set of closed balls in general position in \mathbb{R}^d , as usual. Let $\beta \subseteq B$ be an independent k -simplex, for $k < d$, and let $l(\beta)$ be the $(d - k - 1)$ -sphere common to the $(d - 1)$ -spheres bounding the balls in β . By assumption of general position, almost all points of $l(\beta)$ do not lie on any other bounding $(d - 1)$ -sphere, and we let $y \in l(\beta)$ be one such point. We consider 2^{k+1} points x_γ near y , one for each subset $\gamma \subseteq \beta$, as illustrated in Fig. 8. We require that the points witness the independence of β , that is, $x_\gamma \in \bigcap \gamma - \bigcup (\beta - \gamma)$ for all γ , and that every other ball in $B - \beta$ either contains all of the points or none of them. Supposing $\alpha \subseteq B$ is an independent simplex, we consider

$$\kappa(\alpha) = \sum_{\gamma \subseteq \beta} (-1)^{\dim \gamma} \mathbf{1}_{\bigcap \alpha}(x_\gamma),$$

where $\dim \emptyset = -1$. If $\bigcap \alpha$ does not contain y then it contains none of the points x_γ and we have $\kappa(\alpha) = 0$. Otherwise, $\bigcap \alpha$ contains some but not necessarily all points x_γ . To specify which ones it contains, we define $\delta = \beta \cap \alpha$ and note that x_γ belongs to $\bigcap \alpha$ iff $\delta \subseteq \gamma \subseteq \beta$. Writing γ as the disjoint union of δ and $\gamma' \subseteq \beta - \delta$ we have $\dim \gamma = \dim \delta + \dim \gamma' + 1$. If δ is a proper subset of β then $\beta - \delta$ is a (non-empty) simplex and

$$\kappa(\alpha) = (-1)^{\dim \delta + 1} \sum_{\gamma' \subseteq \beta - \delta} (-1)^{\dim \gamma'} = 0,$$

as before. The only remaining case is when $y \in \bigcap \alpha$ and $\delta = \beta$ or, equivalently, $y \in \bigcap \alpha$ and $\beta \subseteq \alpha$. This is also the only case in which κ does not vanish, namely $\kappa(\alpha) = (-1)^{\dim \beta}$.

We now consider a set $L \subseteq 2^B$ of independent simplices, and the inclusion–exclusion formulas this set defines for the points x_γ . The alternating sum of these formulas is

$$\begin{aligned}\chi &= \sum_{\gamma \subseteq \beta} (-1)^{\dim \gamma} \text{IEF}_L(x_\gamma) \\ &= \sum_{\gamma \subseteq \beta} (-1)^{\dim \gamma} \sum_{\alpha \in L} (-1)^{\dim \alpha} \mathbf{1}_{\cap \alpha}(x_\gamma) \\ &= \sum_{\alpha \in L} (-1)^{\dim \alpha} \kappa(\alpha).\end{aligned}$$

Recalling that $\kappa(\alpha)$ is non-zero only if $\beta \subseteq \alpha$, we let $L_\beta \subseteq L$ be the subset of simplices α that contain β . The contribution of such a simplex α is $(-1)^{\dim \beta}$ if $y \in \cap \alpha$, and zero if $y \notin \cap \alpha$. Hence,

$$\chi = (-1)^{\dim \beta} \text{IEF}_{L_\beta}(y). \quad (2)$$

We are interested in two special cases. The first case is characterized by IEF_L being constant in a neighborhood of y . By the choice of points x_γ , we have $\text{IEF}_L(x_\gamma) = \text{IEF}_L(y)$ for all γ . Plugging the common value into the definition, we get $\chi = 0$, and using (2), we get $\text{IEF}_{L_\beta}(y) = 0$. We state this result in words, letting β be an independent simplex of dimension $k < d$, and $y \in l(\beta)$ be a point not on the bounding sphere of any ball in $B - \beta$, as before.

Even Corollary. *If IEF_L is constant in a neighborhood of y then the number of cofaces $\alpha \in L$ of β with $y \in \cap \alpha$ that have even dimension is the same as the number of such cofaces that have odd dimension.*

The name of the claim is motivated by the weaker implication that the number of cofaces α of β with $y \in \cap \alpha$ is even. The second special case is characterized by $\text{IEF}_L(x_\gamma) = \text{IEF}_L(y)$ for all $\gamma \neq \emptyset$ and $\text{IEF}_L(x_\emptyset) = \text{IEF}_L(y) - 1$. This arises, for example, when the inclusion–exclusion formula of L is the indicator function of $\bigcup B$ and y lies on the boundary of the union. Plugging the values into the definition, we get $\chi = 1$, and using (2), we get $\text{IEF}_{L_\beta}(y) = \pm 1$. We state a weaker implication in words.

Odd Corollary. *If IEF_L is constant around y , except in the orthant of x_\emptyset where it is one less, then the number of cofaces $\alpha \in L$ of β with $y \in \cap \alpha$ is odd.*

Redundant Subsets. A subset L of an abstract simplicial complex K is *redundant* if $\text{IEF}_K = \text{IEF}_{K-L}$. Equivalently,

$$\text{IEF}_L = \sum_{\alpha \in L} (-1)^{\dim \alpha} \mathbf{1}_{\cap \alpha}$$

vanishes everywhere. We use the Even Corollary to derive structural properties of redundant subsets.

Redundant Subset Lemma. *Let B be a finite set of closed balls in general position in \mathbb{R}^d , let K be an independent complex over B , and let L be a redundant subset of K .*

- (i) *If L contains a k -simplex β , with $k < d$, then L contains at least one proper coface $\alpha \supset \beta$.*
- (ii) *If L contains a $(d - 1)$ -simplex β , then L contains two d -simplices whose canonical images in \mathbb{R}^d intersect in the canonical image of β .*
- (iii) *If L contains a d -simplex α , then L contains all $d + 1$ $(d - 1)$ -faces of α .*

Proof. To get (i), let $y \in l(\beta)$. Since IEF_L vanishes everywhere, and therefore also in a neighborhood of y , the Even Corollary implies that L contains an even number of cofaces α of β with $y \in \bigcap \alpha$. One such coface is β itself, which implies the number is at least two and therefore includes at least one proper coface.

To get (ii), observe that $l(\beta)$ consists of two points, y and z . Applying the above argument to y we obtain a d -simplex $\alpha \supset \beta$ in L . Since $d - 1$ and d are the only dimensions to consider, and for trivial reasons β is the only $(d - 1)$ -simplex that contains β , the d -simplex α is unique. Since α is independent, the extra ball in α contains y and does not contain z . Symmetrically, we get a unique d -simplex whose extra ball contains z and does not contain y . The centers of the two extra balls lie on opposite sides of the $(d - 1)$ -dimensional plane spanned by β . It follows that the two d -simplices lie on opposite sides of the $(d - 1)$ -simplex, as illustrated in Fig. 8 on the right.

To get (iii), we consider a $(d - 1)$ -face β of α . Since β is independent, $l(\beta)$ consists of two points and we let y be the one contained in the single ball in $\alpha - \beta$. The Even Corollary implies that L contains at least one coface of β , besides α whose common intersection contains y . As proved above, α is the only proper coface of β with $y \in \bigcap \alpha$, leaving β itself as the only remaining possibility. \square

Sufficiency. We are ready to prove one direction of Theorem B. Specifically, we show that an abstract simplicial complex K that is independent, canonically realizable in \mathbb{R}^d , and satisfies the boundary condition has a minimal inclusion–exclusion formula. Equivalently, such a complex K contains no redundant subset.

To obtain a contradiction, we assume K has a non-empty redundant subset L . Because of (i) in the Redundant Subset Lemma, we may assume that L contains at least one d -simplex. Using (iii) of the same lemma, we see that L also contains the $(d - 1)$ -faces of that d -simplex. By iterating (ii) and (iii), we conclude that L contains all d -simplices of a component formed by connecting the d -simplices across shared $(d - 1)$ -faces. However, then L also contains the boundary $(d - 1)$ -simplices of that component, which exist because K is finite and geometrically realized in \mathbb{R}^d . Now we have arrived at a contradiction because a boundary $(d - 1)$ -simplex lacks the d -simplex on its other side which, by (ii) of the Redundant Subset Lemma, ought to be in L .

Boundary and Interior. Having established one direction of Theorem B, we now prepare the other. Let B be a finite set of balls in \mathbb{R}^d and let $K \subseteq 2^B$ be an abstract simplicial complex. The only properties we assume are that the balls are in general position and that $\text{IEF}_K = \mathbf{1}_{\bigcup B}$.

Inside–Outside Lemma. *Let $\beta \subseteq B$ be an independent simplex of dimension $k < d$ and not necessarily in K , and let $y \in l(\beta)$ be a point not on the $(d - 1)$ -sphere bounding any ball in $B - \beta$.*

- (i) *If y lies in the interior of $\bigcup B$ then $\beta \in K$ implies that K contains a proper coface $\alpha \supset \beta$.*
- (ii) *If y lies on the boundary of $\bigcup B$ then $\beta \in K$.*
- (iii) *If $\dim \beta = d - 1$, $\beta \in K$, and K is independent then the number of d -cofaces of β in K is equal to the number of points in $l(\beta)$ that lie in the interior of $\bigcup B$. If there are two then the centers of the two extra balls lie on opposite sides of the $(d - 1)$ -plane spanned by the centers of the balls in β .*

Proof. To get (i), we note that IEF_K is equal to 1 in a neighborhood of y . The Even Corollary implies that K contains an even number of cofaces of β whose common intersections contain y . If β is in K then this number is at least 2 so there is also a proper coface $\alpha \supset \beta$ in K .

To get (ii), we note that IEF_K is equal to 1 in a neighborhood of y except outside $\bigcup B$, where it is 0. The Odd Corollary implies that K contains an odd number of cofaces of β whose common intersections contain y . This odd number is at least 1, and since K is a complex, this implies that K also contains β .

To get (iii), we first note that $l(\beta)$ is a 0-sphere consisting of only two points, y and z , as illustrated in Fig. 8. Furthermore, β is its own only coface of dimension $d - 1$ and the assumed independence of K implies all other cofaces are of dimension d . The Even Corollary therefore permits only one d -coface for each point of $l(\beta)$ in the interior of $\bigcup B$. If there are two such d -simplices then the respective extra ball contains one point of the 0-sphere but not the other, which implies that the two centers lie on opposite sides of the $(d - 1)$ -plane spanned by the centers of the balls in β . Finally, if b is the extra ball of a d -coface of β then b either contains y or z in its interior. We therefore get no d -coface if both y and z lie on the boundary of $\bigcup B$ and only one d -coface if one of the two points lies on the boundary and the other in the interior. \square

Necessity. We are finally ready to prove the second direction of Theorem B. Specifically, we show that an abstract simplicial complex K with minimal inclusion–exclusion




Fig. 8. The edge belongs to none, one, or two triangles depending on whether two, one, or none of the points y and z lie on the boundary of the union of disks. The four points near y are the points x_y used in the derivation of (2).

formula $\mathbf{IEF}_K = \mathbf{1}_{\cup B}$ is independent, canonically realizable, and has the same boundary complex and underlying space as the dual complex.

First independence. Supposing K is not independent, we let $\alpha \in K$ be a non-independent simplex. By definition, α has a face β such that $\cap \beta - \cup (\alpha - \beta) = \emptyset$ or, equivalently, $\cap \beta \subseteq \cup (\alpha - \beta)$. If $\beta = \alpha$ then $\cap \alpha = \emptyset$ so $L = \{\alpha\}$ is redundant, which contradicts the minimality of \mathbf{IEF}_K . Otherwise, $\alpha - \beta \neq \emptyset$ and we can rewrite the indicator function of the common intersection of balls in β as

$$\begin{aligned} \mathbf{1}_{\cap \beta} &= \mathbf{1}_{\cap \beta} \cdot \mathbf{1}_{\cup (\alpha - \beta)} \\ &= \mathbf{1}_{\cap \beta} \cdot \sum_{\emptyset \neq \gamma \subseteq \alpha - \beta} (-1)^{\dim \gamma} \mathbf{1}_{\cap \gamma} \\ &= \sum_{\beta \subsetneq \delta \subseteq \alpha} (-1)^{\dim \delta - \dim \beta - 1} \mathbf{1}_{\cap \delta}, \end{aligned}$$

where $\delta = \beta \cup \gamma$, and $\dim \delta = \dim \beta + \dim \gamma + 1$ because $\beta \cap \gamma = \emptyset$. After moving $\mathbf{1}_{\cap \beta}$ to the other side and multiplying with $(-1)^{\dim \beta + 1}$ we get

$$\sum_{\beta \subsetneq \delta \subseteq \alpha} (-1)^{\dim \delta} \mathbf{1}_{\cap \delta} = 0, \quad (3)$$

which implies that the set of faces of α that are cofaces of β is redundant. We again get a contradiction to the minimality of \mathbf{IEF}_K , which shows that the minimality of the inclusion–exclusion formula implies the independence of K , as claimed.

Second realizability and boundary. Recall that a simplex β belongs to the boundary complex of the dual complex of B iff there is a point $y \in l(\beta)$ on the boundary of $\cup B$. All simplices of the dual complex are independent so the Inside–Outside Lemma (ii) applies showing that β also belongs to K . By (i) of the same lemma, every simplex in K for which there is no such point y is the face of a d -simplex. By (iii) of the same lemma, every such $(d - 1)$ -simplex belongs to two d -simplices, one on each side. Intersect the (canonical images of the) simplices with an oriented line that avoids all simplices of dimension $d - 2$ or less. It meets the boundary $(d - 1)$ -simplices in some order, alternating between entering and exiting the underlying space. After entering and before exiting, the line may encounter a sequence of interior $(d - 1)$ -simplices, alternating between entering and exiting a d -simplex. Since this is true for almost all oriented lines, it is not possible that two d -simplices overlap in anything other than their shared faces. It follows the canonical mapping gives properly intersecting simplices on the boundary as well as in the interior. Equivalently, K is canonically realizable. Furthermore, the boundary complex and the underlying space of K are equal to those of the dual complex. This completes the proof of Theorem B.

5. Conclusion

The main result of this paper is a characterization of the minimal inclusion–exclusion formulas of a union of closed balls B in \mathbb{R}^d that correspond to simplicial complexes. What about inclusion–exclusion formulas that correspond to sets of simplices that do

not form complexes? The central concept is that of an independent set of balls in \mathbb{R}^d , and our results rest on the observation that the maximum size of such a set is $d + 1$. There are other classes of geometric shapes with bounds on the size of independent sets. For example, the number of independent ovals (each bounded by an ellipse in \mathbb{R}^2) is at most five. Does an upper bound of $k + 1$ on the maximum number of independent shapes imply the existence of an abstract simplicial complex of dimension at most k that gives a correct inclusion–exclusion formula? The argument leading up to (3) gives a weaker result, namely the existence of inclusion–exclusion formulas in which each term is an integer multiple of the indicator function of the intersection of at most $k + 1$ shapes.

Acknowledgments

The authors thank three anonymous referees for their constructive comments on the presentation of the results.

References

1. F. Aurenhammer. Voronoi diagrams—a survey of a fundamental geometric data structure. *ACM Comput. Surv.* **23** (1991), 345–405.
2. T. K. Dey and J. Pach. Extremal problems for geometric hypergraphs. *Discrete Comput. Geom.* **19** (1998), 473–484.
3. H. Edelsbrunner. The union of balls and its dual shape. *Discrete Comput. Geom.* **13** (1995), 415–440.
4. H. Edelsbrunner and P. Koehl. The geometry of biomolecular solvation. In *Combinatorial and Computational Geometry*, J. E. Goodman, J. Pach and E. Welzl (eds.), pp. 243–275. MSRI Publ. 52. Cambridge University Press, Cambridge, 2005.
5. I. Gelfand, M. Kapranov, and A. Zelevinsky. *Discriminants, Resultants and Multidimensional Determinants*. Birkhäuser, Boston, 1994.
6. P. Koehl. ProShape: understanding the shape of protein structures. biogeometry.duke.edu/software/proshape, 2004.
7. K. Kratky. The area of intersection of n equal circular disks. *J. Phys. A* **11** (1978), 1017–1024.
8. D. Naiman and H. Wynn. Inclusion–exclusion Bonferroni identities and inequalities for discrete tube-like problems via Euler characteristics. *Ann. Statist.* **20** (1992), 43–76.
9. G. Perrot, B. Cheng, K. D. Gibson, J. Vila, A. Palmer, A. Nayeem, B. Maigret, and H. A. Scheraga. MSEED: a program for rapid determination of accessible surface areas and their derivatives. *J. Comput. Chem.* **13** (1992), 1–11.

Received July 10, 2005, and in revised form February 17, 2006, and March 17, 2006.

Online publication December 5, 2006.