COMPUTATIONAL TOPOLOGY

Tamal K. Dey *, Herbert Edelsbrunner † and Sumanta Guha ‡

Abstract

The authors of this article believe there is or should be a research area appropriately referred to as computational topology. Its agenda includes the identification and formalization of topological questions in computer applications and the study of algorithms for topological problems. It is hoped this article can contribute to the creation of a computational branch of topology with a unifying influence on computing and computer applications.

Keywords. Survey; topology, geometry, algorithms, computer applications.

INTRODUCTION

The title of this article combines *computation* with *topology*, suggesting a general research activity that studies the computational aspects of problems with topological flavor. What we have in mind is distinctly different from studying the *topology of computing* or the *computer animation of topology*. Computational studies of topological questions can be found in the mathematics and the computer science literature, but no concerted effort is apparent. The authors hope that together with the like-minded survey paper by Vegter [116] this article can contribute to the general awareness of the pervasive presence of topological notions in computer applications.

What is it about a problem that makes it topological? The standard answer mentions a type of geometry devoid of concrete spatial notions such as straightness, convexity, distance, and the like. The emphasis lies on connectivity, continuity, on space, and on maps. This does not mean that the more geometric notions have to be absent but rather that their role is deemphasized benefiting the focused attention on topological properties. The authors believe in a synergy between geometry and topology, and in the value of intuition and visualization as a general path to understanding. They also believe in the general scientific method of studying aspects in deliberate isolation. In other words, it is worthwhile to occasionally focus on topological problems. The connection to geometry can be forgotten temporarily but needs to be remembered eventually.

Many computer application areas deal with geometric data and problems, and there we find a rich collection of questions with topological flavor. Typically, these questions are not well defined and it is part of the task to find the most meaningful formalization in the given context. Once a concrete problem has been formulated we can study its computational complexity. It is important to notice this is an over-simplification of the actual situation: the problem formalization cannot be independent of studying the computational complexity, or else we are likely to generate many computationally infeasible problems. First we need to be convinced there really are essentially topological questions in computer application areas. Part I of this article provides sufficient evidence for their existence by discussing problems in image processing, cartography, computer graphics, solid modeling, mesh generation, and molecular modeling. Second, we need the necessary background in topology. Part II surveys topological methods categorized under decompositions, fixed-points, surfaces, embedding, three-manifolds, and homology computation.

It has been observed that a good fraction of the literature in topology is not or only barely intelligible to readers with interest but without proper topological education. The first barrier is the occasionally complex notation and the usual assumption of accumulated concepts and definitions. Indeed, theory building is the main purpose of topology, and as a consequence def-

*Department of Computer Science and Engineering, Indian Institute of Technology, Kharagpur 721302, India.

†Department of Computer Science, University of Illinois at Urbana-Champaign, Urbana, Illinois 61801, USA.

‡Department of Electrical Engineering and Computer Science, University of Wisconsin at Milwaukee, Milwaukee, Wisconsin 53201, USA.

initions are more important than in some other areas of mathematics. The authors attempt to produce a document accessible to non-specialists by collecting formal and standard definitions in an appendix. The main parts of the paper replace technical detail by intuitive examples and arguments. However, it seems impossible to avoid technical definitions altogether and we hope the appendix will be helpful in recalling their meaning.

I. APPLICATION PROBLEMS

We consider problems in six applications areas of computer science: image processing, cartography, computer graphics, solid modeling, mesh generation, and molecular modeling.

1 Image Processing

A *pixel* is a unit square and an *image* is a rectangle decomposed into pixels. Images are the most common means for representing pictorial information in a computer, which explains the enormous literature in the area of image processing. The papers collected in [88] cover a fairly wide range of the topics. The related area of mathematical morphology is treated in Serra [101].

We restrict our attention to *binary images* where each pixel is either black or white. The *foreground* consists of all black pixels, and the *background* consists of all white pixels.

Connectivity. Each pixel shares an edge each with 4 other pixels, and it shares a vertex each with 4 additional pixels. The numbers are smaller for pixels at the image boundary. Questions of proximity and connectivity can be based on the two corresponding notions of adjacency: pixels $p \neq q$ are 4-*adjacent* if they share an edge and they are 8-*adjacent* if they share an edge or a vertex. 4-adjacency implies 8-adjacency. A *path* is a sequence of pixels so that any two contiguous ones are adjacent. It is *simple* if two pixels are adjacent only if they are contiguous, and it is a *closed curve* if the first is adjacent to the last pixel. Depending on the notion of adjacency we get 4-*connected* and 8-*connected* paths and curves. For $i = 4, 8$ two pixels are i-*connected* if they belong to a common i-connected path.

A basic result in plane topology is the Jordan curve theorem which states that every simple closed curve in $\mathbb{R}^2$ decomposes the plane into two disjoint open connected components. Neither notion of connectivity among pixels supports an analoguous result for images. The following version of the Jordan curve theorem uses different notions of connectivity for the curve and the components [68]:

> a simple closed 8-connected curve partitions an image into the path and two 4-connected components,

see figure 1. Symmetrically, if the closed curve is 4-connected then the two components are 8-connected. A necessary assumption for both results is that the curve touches the image boundary at most once.

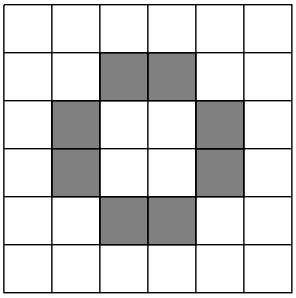

Figure 1: A simple closed 8-connected curve with two 4-connected components in the complement.

The image version of the Jordan curve theorem can be generalized to 3- and higher-dimensional arrays of unit cubes and hypercubes, see Herman [60]. This should be compared to the situation in Euclidean space. Already $\mathbb{R}^3$ permits a counterexample to the straightforward extension of the Jordan curve theorem: there is a *wild* 2-sphere in $\mathbb{R}^3$ so that the bounded component of $\mathbb{R}^3$ minus the 2-sphere is not homeomorphic to the open 3-ball, see e.g. [10, chapter 2.9]. Such counterexamples do not exist for *tame* embeddings and the Jordan curve theorem generalizes as expected.

Skeletonizing. A common operation in image processing replaces the foreground by its skeleton. There is no commonly accepted definition of what exactly this is. One expects a thin subset that reflects the connectivity of the foreground. Often the skeleton is stored together with additional information that permits the exact or approximate reconstruction of the foreground.

Algorithms for constructing skeletons are abundant in the image processing literature. One class of algorithms is based on the medial axis transform. For each foreground pixel p consider a shortest path that connects p to a background pixel q. The length of this path is $\delta(p)$. Boundary pixel are characterized by $\delta = 1$. The medial axis consists of all pixels p with at least two shortest paths to the background that intersect in no pixel other than p.

Another class of algorithms constructs a skeleton by repeatedly removing pixels of the boundary [86], compare this with the discussion on collapsing in section 7.

The foreground is eroded layer by layer, and pixels are removed as long as the connectivity of the foreground does not change. The skeleton is the collection of foreground pixels that remain. Although this collection retains the connectivity of the foreground it often does not preserve its overall shape.

Inner and Outer Boundaries. Consider a collection of disjoint 8-connected simple closed curves that bound the foreground of a binary image. Each curve C in this collection defines two 4-connected components: an outer and an inner region. C is an *outer boundary* if the foreground component bounded by C is part of its inner region. Otherwise, C is an *inner boundary*. How can we distinguish outer from inner boundaries?

Lee, Poston and Rosenfeld [72] propose an algorithm that can be explained using the topological concept of winding number, which counts for an oriented closed curve how often the normal vector winds around the circle of directions. The number is signed counting counterclockwise order positive and clockwise order negative. Orient each curve C so that the foreground lies to its left. Then the winding number of the outer boundary is $+1$ and that of any inner boundary is -1.

The algorithm traverses C and adds the supplementary angles at the pixel centers. This angle is positive if C makes a left turn and negative if C makes a right turn. Because C is simple the sum of angles can only be 360 or -360 degrees, which is 360 times the winding number.

2 Cartography

The purposeful deformation of a geographic map is a common operation in cartography. We consider two types of deformations, one bringing two maps into correspondence and the other reflecting quantities other than geographic distance and area.

Rubber Sheets. As a common practice in cartography, different information about a geographic region is represented in different maps. We can imagine a map M showing mineral distribution as obtained from a geological survey, and another map N displaying agricultural land distribution as obtained from satellite images. The two geographic maps need to be brought into correspondence so that mineral and agricultural land distributions can be related.

To model this problem let $P \subseteq M$ and $Q \subseteq N$ be two sets of n points each together with a bijection $b : P \to Q$. The construction of a homeomorphism $h : M \to N$ that agrees with b at all points of P is popularly known as *rubber sheeting*, see for example [48, 121].

Suppose $\mathcal{K}$ and $\mathcal{L}$ are simplicial complexes whose simplices cover M and N: $M = |\mathcal{K}|$ and $N = |\mathcal{L}|$. Suppose also that the points in P and Q are vertices of $\mathcal{K}$ and $\mathcal{L}$: $P \subseteq \mathrm{Vert}\,\mathcal{K}$ and $Q \subseteq \mathrm{Vert}\,\mathcal{L}$, and that there is a vertex map

$$v : \mathrm{Vert}\,\mathcal{K} \to \mathrm{Vert}\,\mathcal{L}$$

that agrees with b at all points of P. The extension of v to a simplicial map $f : M \to N$ is a simplicial homeomorphism effectively solving the rubber sheet problem.

Several variations of the construction of such complexes $\mathcal{K}$ and $\mathcal{L}$ have been tackled in the recent past. Aronov, Seidel and Souvaine [1] consider simply-connected polygons M and N with n vertices each. They show there are always isomorphic complexes $|\mathcal{K}| = M$ and $|\mathcal{L}| = N$ with at most $O(n^2)$ vertices each. They also prove that sometimes $\Omega(n^2)$ vertices are necessary and they show how to construct the complexes in $O(n^2)$ time. Gupta and Wenger [54] solve the same problem with at most $O(n + m \log n)$ vertices, where m is the minimum number of extra points required in any particular problem instance.

Cartograms. A *cartogram* is a geographic map that is deliberately deformed so that area in the map reflects a quantity other than geographic area, such as for example population density [30]. Figure 2 shows a deformation of the United States so that the area of each state is proportional to the number of electoral votes in the 1992 presidential election. The cartogram

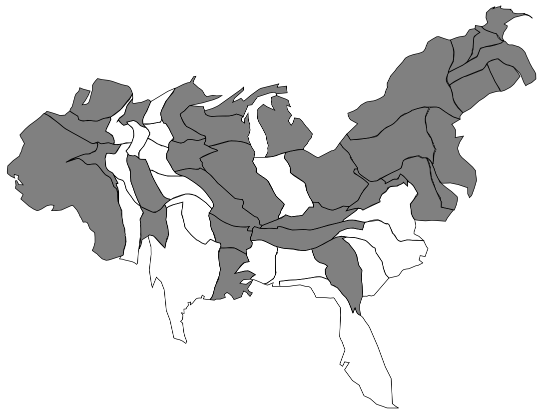

Figure 2: The shaded states have a majority of votes for the elected president.

is related to the usual geographic representation of the United States through a homeomorphism displayed in figure 3. This particular example has been generated by an iterative algorithm described in [40]. It deforms the

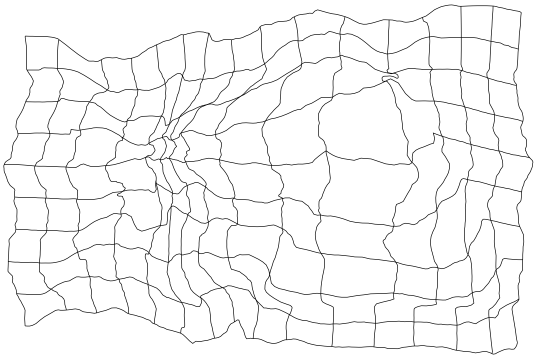

Figure 3: The image of a regular square grid under the homeomorphism that generates the cartogram of the United States shown in figure 2.

Euclidean plane through a series of elementary piecewise linear homeomorphisms. Each elementary homeomorphism leaves the area measure invariant everywhere except in a local region. The approach is based on a regular triangular grid tiling and can be extended to three dimensions using the Delaunay tiling of the body-centered cube lattice.

3 Computer Graphics

In computer graphics we find an unexpected connection between computer hardware and topology. A *texture map* is the hardware implementation of a continuous function, which is a basic concept in topology. It requires a triangulated surface and a range or texture space, which is a k-dimensional unit cube of colored pixels that can be accessed through k texture coordinates. Typical values of k are 1, 2 and 3. The map is explicitly stored at vertices and computed through barycentric combination at all other points of the surface.

Warping. *Warping* refers to the distortion or deformation of a two-dimensional image. A popular approach is based on Coons patches which are bivariate parametric functions [58]. They warp a source image bounded by four arbitrary curves to a rectangular target image. Given the left, right, bottom, top boundary curves $L(v) = C(0, v)$, $R(v) = C(1, v)$, $B(u) = C(u, 0)$, $T(u) = C(u, 1)$, the Coons patch adds an interpolation between L and R to an interpolation between B and T and subtracts an interpolation between the four corner points:

$$\begin{aligned} C(u,v) &= (1-u)\cdot L(v) + u \cdot R(v) \\ &+ (1-v)\cdot B(u) + v \cdot T(u) \\ &- (1-u)(1-v)\cdot C(0,0) - uv \cdot C(1,1) \\ &- u(1-v)\cdot C(1,0) - (1-u)v \cdot C(0,1). \end{aligned}$$

The inverse, C^{-1}, maps the source image bounded by L, R, B, T to the unit square.

A more general but little explored approach to warping can be based on simplicial maps carrying the contents of one triangulated image to another. This is similar to the approach to cartograms mentioned in section 2. Similarly we can use a simplicial map to warp a rectangular image onto a torus or any other surface. An important question in this context is the formulation of algorithms that automatically generate interesting simplicial maps between two given 2-dimensional domains.

Morphing. The transformation of one geometrical or graphical object into another is referred to as *metamorphosis* or *morphing*. The general objective is to compute a transformation that is smooth, in a vague and intuitive sense, and visually pleasing.

The morphing of a digital image $\mathbb{X}$ into another such image $\mathbb{Y}$ has found broad commercial applications. A common step in the generation of the morphing sequence is the construction of isomorphic simplicial complexes $|\mathcal{K}| = \mathbb{X}$ and $|\mathcal{L}| = \mathbb{Y}$. The transformation can be performed by slowly moving each vertex in $\mathcal{K}$ to the corresponding vertex in $\mathcal{L}$ and letting this process carry the entire image along. The construction of $\mathcal{K}$ and $\mathcal{L}$ is made difficult by external requirements that have to do with features that need to be matched across the two images.

A more challenging version of the morphing problem concerns geometric shapes and in particular shapes in $\mathbb{R}^3$. An early paper on this subject is Kent, Carlson and Parent [66]. It presents a method that takes a topological 2-sphere $\mathbb{X}$ to another such sphere $\mathbb{Y}$ by first constructing isomorphic triangulations and then moving parts of $\mathbb{X}$ towards the corresponding parts of $\mathbb{Y}$. This approach is limited to surfaces that share the same topology. The morphing of a sphere to a torus is difficult because it is not clear when and how the sphere surface should be pinched and punctured to change its topology.

An entirely different approach based on the representation of shapes as envelopes of spheres is described in [17]. Both $\mathbb{X}$ and $\mathbb{Y}$ are given in terms of finitely many spheres, and the deformation is created by moving the spheres of $\mathbb{X}$ to the spheres of $\mathbb{Y}$. Duplication and merging of spheres is permitted. At any moment in time we have a finite set of spheres, and the shape at that time is the envelope of an infinite family of shrunken spheres [35]. The infinite family makes the envelope smooth and guarantees that any motion translates into a smooth deformation. This method changes topology automatically and extends to morphs between three or more given shapes.

4 Solid Modeling

This section discusses three topological questions which arise in the area of solid modeling: how to represent a model in the computer, how to construct or reconstruct a surface from discrete point data, and how to extract features of a model.

Model Representation. Models are continuous domains and they need finite representations in the computer. Most commonly we find complexes of one kind or another used for that purpose. The model or its boundary is decomposed into simple pieces or cells, and each cell is stored along with adjacencies to neighboring cells.

The *boundary representation*, for short *b-rep*, describes a 3-dimensional model through its boundary complex, which consists of vertices, edges, and 2-dimensional faces. The winged-edge structure of Baumgart [5] and the quad-edge structure of Guibas and Stolfi [52] are examples of b-rep data structures that assume the boundary is a manifold. This means every point of the boundary has a local neighborhood homeomorphic to $\mathbb{R}^2$. If the boundary is triangulated then every edge belongs to two triangles and every vertex belongs to a simple cycle of triangles and edges. Non-manifold b-reps have been proposed by Wesley [120] and a survey of b-rep schemes can be found in Weiler [119].

A data structure for 3-dimensional complexes is described in Dobkin and Laszlo [29]. The basic element is the oriented edge-face pair. All edge-faces sharing the same edge belong to a cycle or ring around the edge. Symmetrically, all edge-faces sharing the same face are arranged in a cycle describing the face boundary. The two types of rings connect all edge-faces into a single topological data structure. More general complexes of arbitrary fixed dimension are discussed in Brisson [12].

A common concern is the topological validity of a given complex data structure. Topologically invalid structures may be caused by numerical errors inevitable in algorithms that perform geometric calculations in floating-point arithmetic. To check a 2-manifold we would include tests that verify it is orientable and that it satisfies the Euler relation, see section 7. A popular approach to building b-reps manipulates the data structure exclusively through Euler operators, see Mäntylä [77]. The purpose of the restriction is to avoid any side-effects that may compromise the topological validity of the data structure.

Surface Reconstruction. Given a finite set of points in $\mathbb{R}^3$, the *surface reconstruction* problem aims at producing a surface that passes through the points. Versions of the problem differ in their assumptions on the point distribution and their requirements on the surface. In all cases the surface is supposed to adapt to the general shape of the point set and surfaces with fewer wrinkles and less area are generally preferred.

The most common version of the problem assumes the points represent slices through the object and thus lie on a series of parallel planes in space. Catscans in medical imaging are a prime example of data that falls into this category. Within each plane the points are to be connected by edges forming a closed polygon, and polygons in contiguous planes are to be connected by triangles forming a cylindrical surface, see figure 4. The collection of cylinders makes up the reconstructed surface.

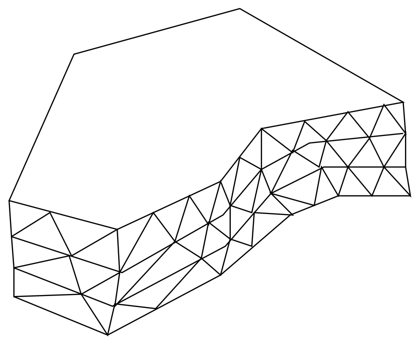

Figure 4: A triangulated cylinder connecting polygons in four parallel planes.

The operation that connects polygons in two parallel planes is often referred to as *lofting*. Let P and Q be two polygons and let n be the total number of vertices. In one of the first papers on the topic, Fuchs, Kedem and Uselton [46] describe a dynamic programming algorithm that finds the minimum area cylinder connecting P and Q in time $O(n^2 \log n)$. The cylinder is not guaranteed to be free of self-intersections, and it is generally not clear how to modify the construction so that a guarantee can be given. In fact, there are cases where all lofting cylinders that avoid self-intersections have at least one vertex in addition to the n vertices of P and Q [49].

The surface reconstruction problem seems to get considerably more difficult if no assumptions on the point distribution are made. An approach to this general problem based on 3-dimensional Delaunay complexes is described in Boissonnat [9]. The initial surface is the convex hull, which is the boundary of the Delaunay complex. The surface is moved inwards by successively removing tetrahedra from the complex. The removal heuristic of Boissonnat is replaced in [34] by a rational and unambiguous collapsing rule sensitive to local and global distribution patters. Figure 5 shows a sample

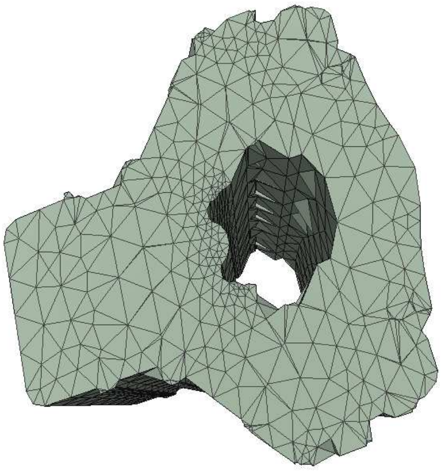

Figure 5: An engine block bounded by a 2-manifold of genus 1.

surface reconstructed with this method.

Feature Extraction. There is no common agreement on what exactly a feature of a geometric model is or should be. If the model is a solid shape in $\mathbb{R}^3$ then the three types of holes characterizing its connectivity belong to almost everybody's list: the *gaps* that separate components, the *tunnels* that pass through the shape, and the *voids* that are components of the complement space inaccessible from the outside. These features are counted by the Betti numbers $\beta_0, \beta_1, \beta_2$.

It is easier to define components than gaps between them and they are counted by β_0. The voids are similarly unambiguous and are counted by β_2. The concept of a tunnel and how to count tunnels is somewhat more delicate. For example consider the 1-skeleton of a tetrahedron, which consists of 4 vertices and 6 edges. There are 4 triangle entries connecting the outside with the inside, but β_1 is only 3. The reason is that only 3 of the tunnels are independent, and all other paths around and through the 1-skeleton are combinations of the independent tunnels: the 3 independent tunnels form a basis of the group of all tunnels. The basis is not unique and in general there is no canonical or best choice. See section 12 for more information about Betti numbers and how to compute them.

Another type of commonly sought features are pronounced protrusions and intrusions. Intrusions are similar to the three types of holes except that they are not holes in any topological sense and require a geometric characterization. A protrusion of the shape is an intrusion of the complement space, and an intrusion of the shape is a protrusion of the complement space. A more detailed discussion of these features in the context of macromolecules can be found in section 6.

5 Mesh Generation

Finite element and other numerical methods for scientific analysis rely on the availability of complexes, called meshes or grids, that decompose shapes and space of interest [65]. The two dominant types of complexes in $\mathbb{R}^3$ are the structured mesh [21, 113] and the tetrahedral mesh [6].

Structured meshes. The prototypical structured mesh in $\mathbb{R}^2$ is the decomposition of a recangular region into smaller rectangles by drawing vertical and horizontal lines. Each rectangle is referred to as an *element*. Each element has 4 neighbor elements and each vertex is surrounded by 4 elements. Exceptions to this rule are found only on the boundary of the region where the numbers are less than 4. There is flexibility in the shape of elements but not in the structure of adjacencies, which is rigid and motivates the name of this kind of mesh. For example the drawing in figure 3 qualifies as a structured mesh. One of the reasons for insisting on the regular adjacency structure is the convenient representation with two-dimensional matrices. Each matrix entry stores an element and adjacencies are represented by index increments and decrements.

A structured mesh in $\mathbb{R}^3$ consists of hexahedral elements adjacent to 6 other elements. Each edge is surrounded by 4 elements and each vertex is surrounded by 8 elements. Exceptions are found only at the boundary where the numbers are smaller. Such a mesh can again be stored in a matrix with adjacencies represented through index increments and decrements. Even for only moderately interesting 3-dimensional shapes the problem of constructing a structured mesh can be challenging. In many cases it amounts to constructing a homeomorphism between the shape and the unit cube. Meshes with hexahedral elements that relax the degree requirements for edges and vertices are considered in [41, 99].

Tetrahedral Meshes. The most common unstructured mesh is the triangular mesh in $\mathbb{R}^2$ and the tetrahedral mesh in $\mathbb{R}^3$, see figure 5. By far the most popular such mesh is the Delaunay complex named after Boris Delaunay [24]. For a finite point set $S \subseteq \mathbb{R}^3$

the Delaunay complex is unique and decomposes the convex hull of the set. In the non-degenerate case all elements are tetrahedra, and a tetrahedron spanned by 4 points in S belongs to the complex iff its circumsphere encloses no point of S.

An important but seemingly difficult problem is the generation of quality meshes. The precise meaning of this term depends on the application and sometimes is not clear altogether. In many cases it means the angles of the triangles and tetrahedra are not too small and not too large, the edges are not too long and not too short, etc. Ideally the shape and size of the elements should adapt to the solution of whatever numerical problem is considered. The resulting chicken-and-egg paradox suggests that both the mesh and the solution are to be refined simultaneously or at least in alternating steps.

It has been observed that Delaunay complexes in $\mathbb{R}^3$ tend to contain many "slivers", which are fairly flat tetrahedra of quadrilateral shape [16]. In contrast to a triangle shaped one, a quadrilateral shaped flat tetrahedron can have a fairly small circumsphere and can therefore exist in the middle of the Delaunay complex. It would be interesting to find a point insertion strategy that can eliminate slivers:

> is there always a point p so that the Delaunay complex of $S \cup \{p\}$ contains fewer slivers than the Delaunay complex of S?

Slivers close to the boundary can just be taken off from the outside, but for slivers in the middle it would be important to have other means of elimination, which an answer to the above question would provide.

Mesh Partitioning. The primary motivation for partitioning a mesh are parallel analysis algorithms that assign processors to work concurrently in different parts of the mesh. The partition defines *blocks* which should be roughly the same size and there should be few adjacencies across blocks.

For a planar graph such a partition can be constructed by removing a small number of vertices. Let the *size* of a graph be the number of vertices. Lipton and Tarjan [74] prove that

> every planar graph of size n contains a subset of at most $\sqrt{8n}$ vertices whose removal decomposes the graph into components of size at most $2n/3$ each.

Call the removed vertex set a *separator*. Lipton and Tarjan also give an algorithm that constructs a separator in time $O(n)$. For any given size bound the graph can be decomposed into blocks of at most that size by repeated removal of separators.

A geometric proof of the existence of separators for planar graphs is described in [81]. It is based on Koebe's theorem, see section 10, and the existence of centerpoints in $\mathbb{R}^3$, see section 8. The geometric proof extends to tetrahedral meshes in $\mathbb{R}^3$ provided the aspect ratios are bounded from above by some constant. In this case the mesh contains $O(n^{2/3})$ vertices whose removal leaves components of size at most $c \cdot n$, for any positive constant $c < 1$.

6 Molecular Modeling

We are interested in the study of molecules as 3-dimensional shapes and conformations. A big motivation for research in this area is the protein folding problem that aims at a computer simulation of the natural folding process. That process takes an amorphous chain of amino acid components into a folded state where shape and geometry are determining factors of the protein's function. We quote from a recent article by Rose [94]:

> "What role a protein takes in the grand biological opera depends on exactly one thing: its shape. For a protein molecule, function follows form."

At the current state of the research no successful computer simulation of the folding process is in sight.

Molecular Surfaces. A basic question concerns the geometric definition of a folded protein. Even if we make the unrealistic assumption that the protein is a static assembly of its atoms, it is not clear which part of space exactly the protein occupies. What is its interior and where is its boundary?

A rigorous but computationally demanding answer to these questions is described in a recent book by Bader [3]. Based on Morse theoretic ideas of flow and critical points, the 3-dimensional space is partitioned into catchment regions or unstable manifolds of various dimensions. The protein itself is a union of such regions and its shape changes sometimes catastrophically as small changes in the flow field imply big changes in the catchment regions.

A more pragmatic answer to the shape definition question starts from the geometric idea of modeling an atom as a spherical ball. The atom kernel is at the center and the electrons populate the region between the center and the sphere boundary, see [71, 92]. Chemically bound atoms correspond to overlapping balls. The protein itself is the union of all the balls.

This geometric idea can be extended to modeling the shape of a protein as it appears to a solvent such as water. If we model the solvent as an omnipresent sphere of radius ϱ, we can just increase the atom balls by that radius and shrink the solvent to a point. The resulting model is sometimes referred to as the *solvent accessible* model of the protein.

Dual Complexes. The union of balls model has a simplicial complex as its dual. The complex is defined in geometric terms and topological language is useful in understanding its relation to the model.

Let B be the finite set of balls defining the model $\bigcup B$ of some protein. We decompose $\bigcup B$ into convex cells using weighted Voronoi cells. The *weighted distance* of a point $x \in \mathbb{R}^3$ from a ball $b \in B$ is the square distance to the center minus the square radius. The *weighted Voronoi cell* V_b of b is the set of points x for which b minimizes the weighted distance, see figure 6.

The weighted Voronoi cells decompose $\bigcup B$ into convex cells of the form $b \cap V_b$. Any two cells have disjoint interiors but they may overlap along their boundary. The boundary overlap figures in the definition of the dual complex. To describe this let T' be the set of centers of the balls in T. We are interested in the subsets $T \subseteq B$ whose cells have non-empty common intersection. The *dual complex* of B consists of all simplices spanned by the centers of balls in such subsets:

$$\mathcal{K} \;=\; \{\sigma = \operatorname{conv} T' \mid \bigcap_{b \in T} (b \cap V_b) \neq \emptyset\}.$$

Assuming general position, at most 4 Voronoi cells can meet in a point. In this case all elements of $\mathcal{K}$ are simplices in $\mathbb{R}^3$ and altogether they form a simplicial complex, see figure 6.

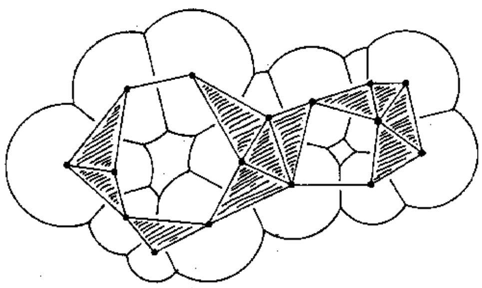

Figure 6: The union of disks is decomposed by the weighted Voronoi cells. The dual complex represents the overlap structure of the convex cells in the decomposition.

The original definition of the dual complex dates back to [37] where unit-disks in the plane are considered. Algorithms for the 3-dimensional case are described in [38]. The general definition for mixed radii and arbitrary fixed dimension can be found in [32]. The latter reference also contains a proof of the homotopy equivalence between the ball union and the dual complex and short inclusion-exclusion formulas for measuring the volume and surfaces area of a ball union model.

Pockets and Protrusions. The dual complex of a ball union model is a combinatorial structure consisting of vertices and of pairs, triplets, and quadruplets of vertices. It allows fast combinatorial algorithms for computing the connectivity in terms of Betti numbers, see section 12 of this paper. Because the ball union and the dual complex are homotopy equivalent, we know that the two have the same Betti numbers. The second Betti number counts the *voids*, which are the bounded components of the complement of $\bigcup B$. These are interesting to biologists because they may contain water molecules caught at some moment during the folding process and influencing the further evolution of that process.

Of even greater significance than voids are cavities or depressions that extend to the surface of a protein. These are not holes in a topological sense although common language refers to them often as such [15]. To avoid any confusion we refer to such cavities as *pockets*, see figure 7. The significance of pockets lies in their prominent

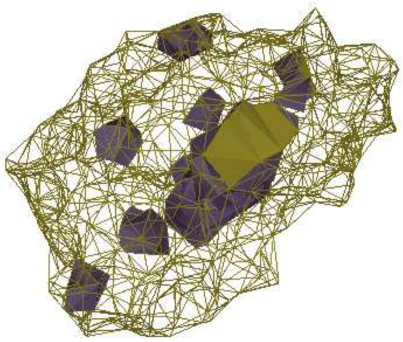

Figure 7: Visualization of voids and pockets of the dual complex of HIV-1 proteise.

role in molecular interactions. Proteins communicate with other molecules via surface contact along partial albeit imperfect shape complementarity. We wish to give a geometrically rigorous definition of a pocket as a part of space that can be computed and analyzed. Such a definition using concepts in Morse theory is given in [36]. We present a slightly simplified version of that definition.

For each point $x \in \mathbb{R}^3$ let $f(x)$ be the minimum weighted distance to any ball $b \in B$. We can make f smooth using convenient cut-off functions close to the boundary of Voronoi cells. Assuming smoothness we take the gradient of f,

$$\nabla f = \left(\frac{\partial f}{\partial x_1}, \frac{\partial f}{\partial x_2}, \frac{\partial f}{\partial x_3} \right),$$

which is a smooth vector field $\nabla f : \mathbb{R}^3 \to \mathbb{R}^3$. From a point x we follow the flow of ∇f, which either takes us to infinity or we get caught at a local maximum. A *pocket* is defined as a maximal connected piece of space outside $\bigcup B$ whose points do not flow to infinity.

It appears that a pocket is an inherently continuous concept whose computation requires following flow lines for as many points as the desired accuracy of the construction requires. Fortunately, this is not the case and a close relationship between pockets and Delaunay simplices can be used to obtain a fast combinatorial algorithm, see [36].

II. TOPOLOGICAL METHODS

We consider six topics in topology: decompositions, fixed points, surfaces, embedding, three-manifolds, and homology computation. This leaves large domains untouched, such as the entire fields of Morse theory [83] and differential topology [53]. The deliberate omission of differentiable topics is motivated by respect for the size and depth of that area and the attempt to preserve the elementary character of this paper.

7 Decompositions

Questions about decomposing a shape into simple pieces can have geometric as well as topological flavor. Restrictions on the type of pieces that are allowed and how they can overlap give rise to different types of complexes. This section considers a few decomposition questions for the simplest kind of complex where each cell is convex and any two cells intersect in another cell.

Euler Characteristic. The Euler relation for the boundary complex of a polyhedron is possibly *the* oldest topological result in all of mathematics. If the polyhedron is convex and bounded in $\mathbb{R}^3$ the relation states that the number of vertices and 2-dimensional faces exceeds the number of edges by 2:

$$v - e + f = 2.$$

This result was know already to René Descartes. Initial attempts to extend this result to polyhedra with holes and to polyhedra bounded by faces with holes were only partially successful, see Lakatos [70]. Eventually, Poincaré settled the issue [89, 90]. He observes that the result of the alternating face sum is a property of the space the faces cover, and that it equals the alternating sum of Betti numbers counting the holes in that space, see section 12. For example, the boundary of a 3-dimensional convex bounded polyhedron is homeomorphic to the 2-sphere, $\mathbb{S}^2$, and 2 is the characteristic number for $\mathbb{S}^2$.

The *Euler characteristic* of a space $\mathbb{X}$ is the alternating simplex count of any triangulation $\mathcal{K}$ of $\mathbb{X}$:

$$\chi(\mathbb{X}) = \sum_{\emptyset \neq \sigma \in \mathcal{K}} (-1)^{\dim \sigma}.$$

For the torus we have $\chi = 0$, and more generally for orientable closed surfaces with genus g we have $\chi = 2 - 2g$. For non-orientable surfaces with genus g we have $\chi = 2 - g$, see section 9. It is thus easy to determine the type of a closed surface given as a triangulation. First check whether or not the triangles can be consistently oriented and then computing the Euler characteristic by counting vertices, edges, and triangles. The same approach fails for compact 3-manifolds because they all have vanishing Euler characteristic.

The Euler characteristic has found a multitude of applications in geometry and topology, including volume formulas for unions of spherical balls [32, 87], see also section 6. A survey of extensions and generalizations can be found in the booklet by Shashkin [103].

Covers and Nerves. A *cover* of a topological space $\mathbb{X}$ is a collection $\mathcal{C}$ of subsets whose union is $\mathbb{X}$: $\mathbb{X} = \bigcup \mathcal{C}$. For example the set of d-faces of a convex polyhedron of dimension $d+1$ is a cover of the polyhedron boundary. The *nerve* of a cover $\mathcal{C}$ is the system of subcollections with non-empty common intersection:

$$\mathrm{Nrv}\, \mathcal{C} = \{ X \subseteq \mathcal{C} \mid \bigcap X \neq \emptyset \}.$$

The nerve is an abstract simplicial complex because it is closed under taking subsets: $X \in \mathrm{Nrv}\, \mathcal{C}$ and $Y \subseteq X$ implies $Y \in \mathrm{Nrv}\, \mathcal{C}$. We get a geometric realization of the abstract complex by mapping each $C \in \mathcal{C}$ to a point $p(C)$ in $\mathbb{R}^k$. Every collection $X \in \mathrm{Nrv}\, \mathcal{C}$ is represented by the convex hull of the points $p(C)$, $C \in X$. If $\mathcal{C}$ is finite then $\mathrm{Nrv}\, \mathcal{C}$ is finite and if k is sufficiently large then the points can be chosen so that the resulting simplices form a genuine simplicial complex. Call this complex $\mathcal{K}$ and consider its underlying space, $|\mathcal{K}| = \bigcup_{\sigma \in \mathcal{K}} \sigma$, which is again a topological space. The nerve theorem originally due to Leray [73] states that under reasonable assumptions about the sets in the cover and about their common intersections the underlying space of $\mathcal{K}$ has the same homotopy type as $\mathbb{X}$:

> if $\mathcal{C}$ is finite and $\bigcap X$ is either empty or contractible for every $X \subseteq \mathcal{C}$ then $|\mathcal{K}|$ and $\mathbb{X} = \bigcup \mathcal{C}$ are homotopy equivalent.

The literature contains versions of this theorem that differ in the requirements on the sets in the cover, the allowed size of the cover, and the resulting topological relation between $|\mathcal{K}|$ and $\mathbb{X}$, see e.g. [8, 122].

The homotopy equivalence implies that the Euler characteristics of $\mathcal{K}$ and $\mathbb{X}$ are the same. This fact has been used in [33] to derive boolean formulas for non-convex polyhedra. The nerve theorem is also useful in computational settings because covers are relatively easy to get and through the nerve operation large and meaningful complexes can be automatically generated. An example of this idea is the dual complex of a union of balls discussed in section 6. Another example is the automatic triangulation of a topological space developed in [18, 39, 79].

Collapsing. If a simplicial complex can be collapsed then its underlying space is contractible. Unfortunately, the reverse is not correct: there are contractible spaces that have non-collapsible triangulations. An example of such a space is the dunce cap [85]. Although the cap is contractible every edge in its triangulation belongs to at least two triangles.

What exactly is a collapse? Given a simplicial complex $\mathcal{K}$, an *elementary collapse* removes a pair of simplices $\tau \subseteq \sigma$ provided they are the only simplices in the star of τ: $\mathrm{St}\,\tau = \{\tau, \sigma\}$. For example the right complex in figure 8 is obtained from the left complex by removing a vertex τ and an edge σ. More generally, a *collapse*

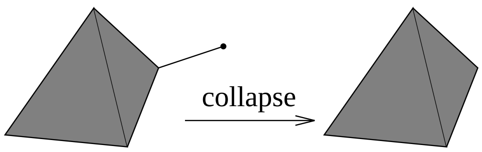

Figure 8: Collapsing a vertex together with the only edge that contains it.

removes the entire star of a simplex τ provided there is a unique highest-dimensional simplex in this star. For example the rightmost vertex of the right complex in figure 8 can be removed together with the two edges and one triangle that contains it. The topmost vertex cannot be removed through a single collapse. A collapse removing a star of 2^k simplices can be expressed as the composition of 2^{k-1} elementary collapses.

It is fairly easy to see that for each collapse there is a corresponding deformation retraction that modifies the underlying space of the complex. This implies that a collapse retains the homotopy type. If $\mathcal{K}$ can be reduced by a sequence of collapses all the way down to a single vertex then $|\mathcal{K}|$ is contractible.

An interesting result due to Wegner [118] relates the collapsibility of a simplicial complex with its representability by convex sets. Wegner calls an abstract simplicial complex $\mathcal{A}$ *d-representable* if it is the nerve of a collection of convex sets in $\mathbb{R}^d$, and he shows that $\mathcal{A}$ is d-representable iff it is d-collapsible. The notion of d-collapsibility used here is slightly different from collapsibility as it allows the removal of principal simplices of dimension less than d.

Shelling. A concept similar to collapsing that also applies to non-simplicial complexes is shelling. A superficial difference is that shelling is applied in forward direction constructing a complex by adding one cell at a time. Suppose $\mathcal{K}$ is a complex of convex cells so the union of cells is a topological ball: $|\mathcal{K}| \approx \mathbb{B}^d$. A *shelling* is an ordering of the d-dimensional cells, $\sigma_1, \sigma_2, \ldots, \sigma_n$, so that every prefix defines a d-ball:

$$\bigcup_{i=1}^{j} \sigma_i \quad \approx \quad \mathbb{B}^d$$

for every $1 \leq j \leq n$. If $|\mathcal{K}| \approx \mathbb{S}^d$, as is the case for the boundary complex of every $(d+1)$-dimensional convex polytope, then the requirement is that all *proper* prefixes define a d-ball. $\mathcal{K}$ is *shellable* if it has a shelling.

Shelling is important in the theory of convex polytopes. Schläfli [98] proves the Euler relation for convex polytopes of any dimension assuming their boundary complexes are shellable. Many years later, Bruggesser and Mani [14] establish that these complexes are indeed shellable. Seidel [100] uses the same shelling order to compute the convex hull of a set of n points in time $\mathrm{O}(\log n)$ per face of the constructed polytope. The algorithm requires an initial preprocessing phase of time $\mathrm{O}(n^2)$ before the faces can be enumerated.

Given a complex $\mathcal{K}$, Danaraj and Klee [22] study algorithms that compute a shelling, if it exists. If $|\mathcal{K}| \approx \mathbb{B}^2$ there always is a shelling and it can be constructed in time proportional to the number of cells. Already for $|\mathcal{K}| \approx \mathbb{B}^3$ it is possible that $\mathcal{K}$ is not shellable [96] and no polynomial time algorithm deciding the case is known. The problem is made difficult by the existence of partial shellings that cannot be extended to a complete shelling, see Ziegler [124].

8 Fixed Points

A point is *fixed* by a function if the point is its own image. In many circumstances, the continuity of the function implies the existence of a fixed point. A nice and elementary treatment of this subject can be found in the booklet by Shashkin [104].

A Metric Fixed Point Theorem. A map $f : \mathbb{R}^d \to \mathbb{R}^d$ *contracts* if there exists a constant γ, $0 < \gamma < 1$, such that the Euclidean distance between image points is less than γ times the Euclidean distance between arguments:

$$\|f(x) - f(y)\| \quad \leq \quad \gamma \cdot \|x - y\|.$$

For example if you consult a geographic map, you can think of each point on the map as the image of a point in the real world. The geographic map is the image of a map or continuous function, and usually this map is contracting. A classical result on contracting maps is that they have a unique fixed point. In the case of a posted geographic map this fixed point is typically indicated by the words "you are here" or a similar label.

Brouwer's Theorem. Brouwer's fixed-point theorem is one of the most basic facts about topological spaces and generalizes the metric fixed point theorem. Originally formulated for the d-dimensional ball, the theorem is stated here for a d-simplex σ^d homeomorphic to $\mathbb{B}^d$:

> every map $f : \sigma^d \to \sigma^d$ has a fixed point.

Observe that the only assumption on f is that it be continuous. Brouwer [13] provides an existential proof for this theorem. Cohen [19] gives a constructive proof based on a combinatorial result of Sperner [108]. We describe the combinatorial result and comment on its connection to Brouwer's theorem.

Let σ^d be a d-dimensional simplex with vertex set $T = \{v_0, v_1, \ldots, v_d\}$ and triangulation $\mathcal{K}$. A *labeling* is a function $\ell : \text{Vert}\,\mathcal{K} \to T$. Each vertex of $\mathcal{K}$ either belongs to the interior of σ or to the interior of a proper face of σ. The labeling is *proper* if $\ell(u)$ is one of the vertices of the face containing u in its interior. A d-simplex in $\mathcal{K}$ is *complete* if its vertices have $d+1$ different labels. Sperner's lemma guarantees the existence of a complete d-simplex in every properly labeled triangulation. More precisely:

> every properly labeled triangulation of σ has an odd number of complete d-simplices.

How does this combinatorial result relate to the continuous Brouwer theorem? Given a map $f : \sigma^d \to \sigma^d$, we can define a labeling of a triangulation based on the vectors $f(x) - x$ so that a complete d-simplex indicates a possible fixed point in its interior. If the triangulation is chosen sufficiently fine then the possibility becomes certainty and the existence of complete d-simplices implies the existence of fixed points.

Borsuk-Ulam Theorem. Intuitively it seems clear that the d-dimensional sphere, $\mathbb{S}^d$, cannot be embedded in $\mathbb{R}^d$ without overlap. A theorem conjectured by Ulam and proved by Borsuk [11] is a stronger version of this statement.

Two points $x, y \in \mathbb{S}^d$ with coordinates x_i and y_i in $\mathbb{R}^{d+1}$ are *antipodal* if $x_i = -y_i$ for $1 \leq i \leq d+1$. We write $y = -x$ if x and y are antipodal. The Borsuk-Ulam theorem states that if $\mathbb{S}^d$ is mapped continuously to $\mathbb{R}^d$ then there are two antipodal points with identical image:

> every map $f : \mathbb{S}^d \to \mathbb{R}^d$ has a point $x \in \mathbb{S}^d$ with $f(x) = f(-x)$.

Ramos [91] extends the result to maps from a Cartesian products of spheres to the real space of same dimension. Not all such maps have antipodal points with identical images and Ramos describes a combinatorial criterion that determines when the map has such a point.

Geometric Applications. The Bosuk-Ulam theorem has many interesting geometric consequences, see for example the survey by Bárány [4]. Consider d finite sets, $S_1, S_2, \ldots, S_d$, in $\mathbb{R}^d$. A hyperplane h *bisects* S_i if at most half of the points in S_i lie in each open half-space defined by h. A *ham sandwich cut* is a hyperplane that bisects all d sets. Using the Borsuk-Ulam theorem it is fairly straightforward to prove that a ham sandwich cut always exists. A linear time algorithm for finding ham sandwich cut in two dimensions can be found in [75].

Consider a single set $S \subseteq \mathbb{R}^d$ and several hyperplanes cutting S. $k \leq d$ hyperplanes $h_1, h_2, \ldots, h_k$ form an *equipartition* if each of the 2^k open d-dimensional cells contains at most $n/2^k$ points, $n = \text{card}\, S$. Using the ham sandwich cut theorem it is easy to prove that every $S \subseteq \mathbb{R}^2$ admits an equipartion formed by 2 lines. It is also true that every $S \subseteq \mathbb{R}^3$ admits an equipartition formed by 3 planes, but the argument requires more than just the 3-dimensional ham sandwich cut theorem, see Hadwiger [55] for a proof and [123] for applications to range searching. Five hyperplanes cut the moment curve in $\mathbb{R}^5$ into at most $5 \cdot 5 + 1 < 2^5$ pieces, which implies that for $d \geq 5$ not every $S \subseteq \mathbb{R}^d$

admits a d-partition [2]. The question for $d = 4$ is currently undecided. Ramos [91] studies extensions of these equipartition results to values $k \leq d$.

A related geometric concept is a point so all hyperplanes passing through this point cut a given set in reasonably balanced fractions. A point $x \in \mathbb{R}^d$ is a *centerpoint* of $S \subseteq \mathbb{R}^d$ if every closed halfspace that contains x contains at least $\frac{n}{d+1}$ points of S. The centerpoint theorem states that every finite set has a centerpoint. The classic proof of this theorem is based on Helly's result on overlapping convex sets, see e.g. [31]. Using a new approach using Brouwer's fixed point theorem, Živaljević and Vrećica [125] prove a more general result that contains the ham sandwich and the centerpoint theorem as special cases. A linear time algorithm for finding a centerpoint of a set of n points in $\mathbb{R}^2$ is given in [62].

9 Surfaces

The notion of homotopy formalizes the idea of a continuous deformation of a topological space. Since the fundamental group of a space is invariant under such deformations, this group can sometimes be used to prove the non-existence of a continuous deformation of one space to another.

Classification. A classic result due to Brahana and to Dehn and Heegard says that two surfaces or 2-manifolds are topologically different iff their fundamental groups are non-isomorphic. It follows that surfaces can be classified just by considering their fundamental groups. Descriptions of this result can be found in many textbooks in topology [80, 111].

The fundamental group $\pi(\mathbb{M})$ of a surface $\mathbb{M}$ can be finitely presented as $\pi(\mathbb{M}) = (G : R)$, where G is a set of generators and R is a set of relations. Specifically, for an orientable surface of genus g, there are $2g$ generators in G, which are typically represented as g pairs where one member is called 'latitudinal' and the other is 'longitudinal'. For a non-orientable surface of genus g there are g generators. Each relation identifies a product consisting of generators and inverses to the trivial path. For example, $(a, b : aba^{-1}b^{-1})$ is a finite representation of the fundamental group of the torus, see figure 9.

Alternatively, each surface can be combinatorially represented as a polygon P with an even number of edges. Each edge is oriented and labeled. Such a representation is called a *polygonal schema.* The surface can be obtained from P by gluing edges with equal labels consistent with their orientations. Among all polygonal schema of a given surface it is possible to identify a canonical one. We use a superscribed '$-$' to indicate orientation reversal. The canonical polygonal schema of an orientable surface of genus $g \geq 1$ has the form

$$x_1 y_1 x_1^- y_1^- \ldots x_g y_g x_g^- y_g^- .$$

Similarly, the canonical polygonal schema of a non-orientable surface of genus $g \geq 1$ has the form

$$x_1 x_1 x_2 x_2 \ldots x_g x_g .$$

The only remaining case is the 2-sphere, which has genus 0 and canonical schema xx^-.

Algorithms. The algorithmic problem of computing polygonal schemas has been considered in the computational geometry literature. Let $\mathcal{K}$ be a simplicial complex triangulating a surface $\mathbb{M}$, and let $n = \operatorname{card} \mathcal{K}$ be the number of simplices. Vegter and Yap [117] describe an algorithm that takes time $\mathrm{O}(n \log n)$ to construct the canonical schema of $\mathbb{M}$. A not necessarily canonical polygonal schema that has the same number of edges can be constructed in optimal $\mathrm{O}(n)$ time [27].

At the beginning of this century, Dehn posed two problems related to homotopy on surfaces. Recall that $\mathbb{M}$ is a surface and $\mathcal{K}$ is a triangulation of $\mathbb{M}$.

(i) Given a closed curve in terms of a cycle of edges in $\mathcal{K}$, determine whether or not the curve is contractible.

(ii) Given two closed curves, determine whether or not they are homotopic.

Optimal algorithms for both problems are given in [25]. The algorithms are based on observing that deciding the homotopy equivalence of two closed curves is equivalent to deciding whether their representatives in the fundamental group of the surface are conjugate. The conjugacy problem for this fundamental group is then solved by methods from combinatorial group theory [50, 76]. The group theoretic results allow a reduction to Dehn-like algorithms in all cases except for orientable surfaces of genus 2 and for non-orientable surfaces of genus 3 and 4. The problem of finding an optimal algorithm that decides the homotopy between two curves on such surfaces is still open.

Algorithms for deciding the contractibility of a closed curve can also be based on universal covering spaces, see figure 9. Let $\mathbb{Y}$ be a universal covering space of $\mathbb{X}$ with projection $p : \mathbb{Y} \to \mathbb{X}$, and let $y_0 \in \mathbb{Y}$ and $x_0 = p(y_0)$. Every path on $\mathbb{X}$ with initial point x_0 has a unique lifting on $\mathbb{Y}$ with initial point y_0. Furthermore, a closed curve is contractible in $\mathbb{X}$ if and only if its lift on $\mathbb{Y}$ is closed. This was first discovered by Dehn [23].

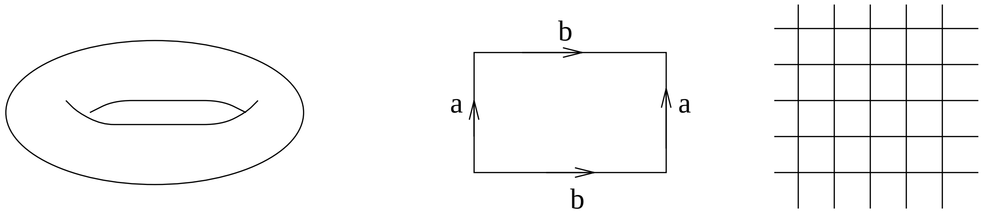

Figure 9: The polygonal schema of a torus is a rectangle with opposite edges identified. The torus can be constructed from a rectangular piece of paper by gluing edges as shown. The universal covering space of the torus is the plane decomposed into rectangles each a copy of the polygonal schema.

Let n be the number of simplices in the given triangulation of $\mathbb{M}$ and m the number of edges forming the closed curve. The algorithm of Schipper [97] takes time $O(gn + g^2 m)$ to construct a finite portion of the covering space to check whether the lifted curve is indeed closed. The algorithm has been improved by Dey and Schipper [27] to run in time $O(n + m \log g)$ and memory $O(n + m)$.

Surface Homology. It may be remarked that by the Hurewicz isomorphism theorem [107] the generators for the homotopy group of a surface are also generators for its homology group. In fact, for an orientable surface of genus g the 1-st homology group is a free abelian goup on the set of g latitudinal and g longitudinal generators. In the case of a non-orientable surface of genus g the 1-st homology group is a direct sum of a $(g-1)$-dimensional free abelian group and the torsion group $\mathbb{Z}_2$.

10 Embedding

An embedding of one topological space in another is an injection whose restriction to the image is a homeomorphism. Using a geometric realization as intermediate step we can stretch the concept and talk about embeddings of an abstract simplicial complex $\mathcal{A}$. To do this we first get a simplicial complex $\mathcal{K}$ that geometrically realizes $\mathcal{A}$ and then embed $|\mathcal{K}|$.

Geometric Realization. Every abstract simplicial complex $\mathcal{A}$ has a geometric realization in $\mathbb{R}^d$, provided d is large enough. We can for example use one dimension per vertex of $\mathcal{A}$ and place that vertex at unit distance from the origin on the corresponding coordinate axis. This is clearly a waste of real space and we should expect that the minimum d depends on $k = \dim \mathcal{A}$ but not on the cardinality of $\mathrm{Vert}\,\mathcal{A}$.

How large does d have to be? As it turns out $d = 2k + 1$ always suffices. To see this map the vertices $u \in \mathrm{Vert}\,\mathcal{A}$ to points $p(u)$ in general position in $\mathbb{R}^d$. Every subset of $d + 1 = 2k + 2$ or fewer points is affinely independent. It follows that the intersection of two geometric simplices is the realization of the intersection of the two corresponding abstract simplices:

$$\mathrm{conv}\, p(\alpha) \cap \mathrm{conv}\, p(\beta) = \mathrm{conv}\, p(\alpha \cap \beta),$$

which is exactly what is required from a geometric realization. To prove that the bound is tight, Flores [43] and van Kampen [63] independently construct k-dimensional abstract simplicial complexes without geometric realization in $\mathbb{R}^{2k}$.

One of their examples generalizes K_5, the complete graph of 5 vertices. Think of this graph as the 1-skeleton of a 4-simplex; it is the case $k = 1$ of the following construction. Take a $(2k+2)$-simplex σ^{2k+2}, which has $2k+3$ vertices and a face for every subset of the vertices. The k-skeleton of σ^{2k+2} is obtained after removing all faces of dimension larger than k. As proved in [43, 63] the k-skeleton of σ^{2k+2} cannot be embedded in $\mathbb{R}^{2k}$.

Planar Graphs. The question of embedding a graph in the Euclidean plane is one of the very few embedding problems that has been studied from a mathematical as well as computational point of view. A graph is a 1-dimensional abstract simplicial complex and can always be embedded in $\mathbb{R}^3$. A graph is *planar* if it can be embedded in $\mathbb{R}^2$. Planar graphs are completely characterized by a classic result due to Kuratowski [69]:

> a graph is planar iff it does not contain a subgraph homeomorphic to K_5 or to $K_{3,3}$.

K_5 is the complete graph with 5 vertices and $K_{3,3}$ is the complete bipartite graph with $3 + 3$ vertices. It is not entirely clear how to turn this result into an efficient algorithm for testing whether or not a given graph is planar. Indeed, the fastest algorithms for planarity testing are based on rather different principles and construct a concrete embedding, if it exists. By Euler's relation, a graph with $n \geq 3$ vertices can only be planar if it has $3n - 6$ or fewer edges. It thus seems possible to test the planarity of such a graph in time $O(n)$. The first such algorithm due to Hopcroft and Tarjan [61] is based on fast graph traversal methods.

Fáry [42] proves that every planar graph has a straight line embedding in $\mathbb{R}^2$. In other words, every planar graph has a geometric realization in $\mathbb{R}^2$, see also Tutte [115]. This result also follows from an older result by Steinitz which states that every 3-connected planar graph is the 1-skeleton of the boundary complex of a 3-dimensional convex polytope [109]. The existence of a straight line embedding also follows from a considerably stronger result by Koebe [67]:

> for every planar graph with n vertices there are n non-overlapping disk in the plane and a bijection between the vertices and the disks so that two vertices are connected by an edge iff the two corresponding disks touch.

In other words, the graph is the nerve of a set of n possibly touching but not otherwise overlapping disks. Numerical algorithms for constructing such sets of disks are studied by Stephenson [110].

Beyond graphs. Embedding problems for abstract complexes of dimension 2 or higher are considerably more difficult than for graphs. For example the problem to decide whether or not a 2-complex can be embedded in $\mathbb{R}^3$ seems to be related to the recognition problem for the 3-sphere for which algorithms have been found only recently, see section 11.

A few results concerning the extension of Kuratowski's theorem beyond graphs are described in [102]. For dimension $k = 3$ or higher there is an obstruction so that a k-dimensional complex $\mathcal{A}$ can be embedded in $\mathbb{R}^{2k}$ iff the obstruction vanishes. The obstruction itself belongs to the cohomology groups of a complex related to $\mathcal{A}$. For $k \geq 5$ similar obstructions are known for embeddings in $\mathbb{R}^{2k-1}$.

For a graph the questions whether it has an embedding or a geometric realization in the plane are the same: every planar graph has a geometric realization in $\mathbb{R}^2$. This is no longer true for complexes of dimension higher than 1: there are 2-dimensional complexes with embeddings in $\mathbb{R}^3$ that have no geometric realization in $\mathbb{R}^3$. Examples can be found in Bing [7] and Grünbaum [51, chapter 11].

Instead of general 2-dimensional complexes consider triangulations of 2-manifolds. Such triangulations are characterized by having every edge shared by exactly two triangles and the link of every vertex being a simple cycle of edges. A 2-manifold can be embedded in $\mathbb{R}^3$ iff it is orientable. Examples of 2-manifolds that are non-orientable and thus have no embedding in $\mathbb{R}^3$ are the projective plane and the Klein bottle.

11 Three-manifolds

A complete classification of all 2-manifolds is mentioned at the beginning of section 9. Such a thing is not possible for d-manifolds if $d \geq 4$. The case of 3-manifolds is still open and Thurston [114] sketches a program that is hoped to eventually settle the problem. A central problem in this context is the Poincaré conjecture which has not been settled in spite of a considerable amount of effort in the mathematics community.

Poincaré Conjecture. Poincaré originally conjectured that every 3-manifold with betti numbers $\beta_0 = 1$, $\beta_1 = \beta_2 = 0$, and $\beta_3 = 1$ is a 3-sphere, or more accurately is homeomorphic to $\mathbb{S}^3$. He found a counterexample himself constructed from a dodecahedron whose pentagonal faces are pairwise identified with a twist. Pointcaré modified his original conjecture to

> every simply-connected 3-manifold is a 3-sphere.

It is still open whether or not this is the case. The specialization of the conjecture to 2-manifolds is correct. The generalization to d-manifolds has been settled by Smale [105] for $d \geq 5$ and by Freedman and Quinn [44] for $d = 4$. See Bing [7] for a very accessible survey of topological problems and issues related to the Poincaré conjecture.

An algorithm that decides whether or not a given complex is a triangulation of the 3-sphere has recently been discovered by Rubinstein [95], see also [112]. The existence of such an algorithm is possibly surprising considering the unsettling results of past attempts to characterize the 3-sphere algebraically.

Knots. A *knot* is an embedding of a closed curve in space:

$$K : \mathbb{S}^1 \to \mathbb{R}^3.$$

It is *polygonal* if it consists of a finite number of line segments. The *standard knot* is the unit circle in the x_1x_2-plane, and a knot is *trivial* if there is an isotopy that connects it to the standard knot. The computational complexity of deciding the triviality of a knot is as yet unresolved. For a polygonal knot with n line segments an algorithm with running time at most doubly-exponential in n is given by Haken [56].

A remarkable result due to Seifert states that every knot is the boundary of an orientable 2-manifold embedded in $\mathbb{R}^3$. Such a 2-manifold with boundary is called a *Seifert surface* for the knot. The Seifert surface can

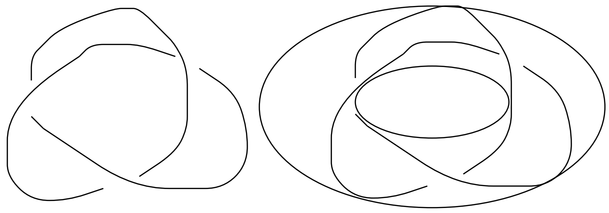

Figure 10: A trefoil knot and its embedding on a torus.

be doubled to form a 2-manifold without boundary, see figure 10. The *genus* of a knot is the least genus of any such 2-manifold. It is known that a knot is trivial iff its genus vanishes or equivalently if it has a Seifert surface that is a disk. It seems difficult to use this observation in the design of a polynomial time algorithm for testing knot triviality because there exist polygonal knots K with n edges so that every piecewise linear Seifert disk whose boundary is K has at least exponentially many vertices, see Snoeyink [106].

A useful invariant of knots is the *knot group*, that is the fundamental group $\pi(\mathbb{R}^3 - K)$ of the complement space of the knot. Although the knot group is invariant under isotopy it does not classify knots: there exist non-isotopic knots with isomorphic knot groups [93].

Knots can be classified by classifying a special class of 3-manifolds [59]. Consider a regular neighborhood R of K in $\mathbb{R}^3$. The space $\mathbb{R}^3 - R$ when compacted with a single point at infinity forms an orientable compact 3-manifold with boundary. These manifolds are often referred to as *Haken manifolds* named after Wolfgang Haken who initiated their study [56].

Computational Complexity. Markov [78] proves that the question whether two finitely represented groups are isomorphic is undecidable. Since every such group is the fundamental group of a 4-manifold, it follows that the question whether two 4-manifolds are homeomorphic is undecidable. The undecidability result extends to dimensions higher than 4. This is why we said earlier that the classification of 4- and higher-dimensional manifolds is out the question.

Manifolds can be specified in terms of triangulations. Triangulations are finite combinatorial objects and the existence of an isomorphism between two can be decided by trying all possibilities. Probably, this is the best one can do. Even for graphs or 1-complexes no polynomial time algorithm is known for deciding isomorphism, see Garey and Johnson [47].

12 Homology Computation

Homology groups offer a formal algebraic framework for studying and counting holes in a topological space. The situation is fairly intuitive for spaces or shapes $\mathbb{X} \subseteq \mathbb{R}^3$. The ranks or Betti numbers $\beta_0, \beta_1, \beta_2$ of the three possibly non-trivial homology groups count the number of gaps, tunnels, and voids. By Alexander duality these correspond to the number of components, loops, and shells in $\mathbb{X}$. The corresponding algebraic notions are non-bounding 0-, 1-, and 2-cycles. They form cosets in the homology groups counted by $\beta_0, \beta_1, \beta_2$.

This section considers different approaches to computing Betti numbers. Since Betti numbers are invariant over all triangulations of a topological space we may assume without loss of generality that $\mathbb{X}$ is given in terms of a triangulating simplicial complex $\mathcal{K}$.

Smith Normal Form. The classic algorithm for Betti numbers manipulates the incidence matrices of simplices of contiguous dimensions. For complexes in $\mathbb{R}^3$ we have vertices, edges, triangles, and tetrahedra and therefore 3 incidence matrices for dimensions 0 and 1, 1 and 2, 2 and 3. Using standard row and column operations, the matrices are brought into a normal form, from which the Betti numbers can be derived. The normal forms also encode the torsion subgroups and thus contain complete information about all homology groups of $\mathcal{K}$.

If implemented as described in textbooks like Munkres [85], the running time can be at least exponential in $n = \text{card}\,\mathcal{K}$. Part of the reason for the slow running time is that the number of row and column operations depends not only on the number of simplices but also on the size of the integer coefficients describing chains of simplices computed during the algorithm. An improvement of the worst-case running time to $\mathrm{O}(n^5)$ can be found in [64]. The complexes that arise in the study of macromolecules, see section 6, are usually too large to even keep the incidence matrices in internal memory. We need new ideas to work on these applications.

Incremental algorithm. A combinatorial approach to computing Betti numbers is described in [28]. Let $\sigma_1, \sigma_2, \ldots, \sigma_n$ be an ordering of the simplices in $\mathcal{K}$ so that every prefix $\mathcal{K}_i = \{\sigma_1, \sigma_2, \ldots, \sigma_i\}$ is a complex. For example the sequence of simplices sorted in order of non-decreasing dimension satisfies this requirement. The sequence of σ_i is a *filter* and the sequence of complexes $\mathcal{K}_i$ is a *filtration* of $\mathcal{K}$.

The Betti numbers of $\mathcal{K}_i$ can be computed from σ_i and the Betti numbers of $\mathcal{K}_{i-1}$. Let k be the dimension of σ_i. If σ_i completes a k-cycle in $\mathcal{K}_i$ then

$$\beta_k(\mathcal{K}_i) = \beta_k(\mathcal{K}_{i-1}) + 1.$$

Indeed, none of the simplices in the star of σ_i have been added yet so σ_i does not bound in $\mathcal{K}_i$. In the other case

$$\beta_{k-1}(\mathcal{K}_i) = \beta_{k-1}(\mathcal{K}_{i-1}) - 1$$

because σ_i fills a non-bounding $(k-1)$-cycle in $\mathcal{K}_{i-1}$, which now bounds in $\mathcal{K}_i$. A formal proof of this case analysis can be based on Mayer-Vietoris sequences as described in any textbook in algebraic topology.

If $\mathcal{K}$ is a complex in $\mathbb{R}^3$ then there are fast algorithms that distinguish between the two possible cases. Given the filter, all questions can be decided in time $O(n\alpha(n))$, where $\alpha(n)$ is the extremely slowly growing inverse of the Ackermann function, see [20]. In the same time the Betti numbers of all complexes in the filtration can be computed. Already in $d = 4$ dimensions it is not clear how to efficiently decide whether or not a simplex σ_i belongs to a cycle. It would be interesting to find fast data structures for the cycle test in all dimensions.

Solids. Call a connected 3-manifold with boundary embedded in $\mathbb{R}^3$ a *solid*. Suppose $\mathbb{M}$ is a solid and $\mathcal{K}$ is a triangulation of $\mathbb{M}$. By definition $\mathcal{K}$ is a pure 3-complex which means that every $\sigma \in \mathcal{K}$ is face of a tetrahedron. The boundary of $\mathbb{M}$ is a 2-manifold and the boundary complex of $\mathcal{K}$ triangulates it. Let $\mathbb{S}_1, \mathbb{S}_2, \ldots, \mathbb{S}_m$ be the components of the boundary of $\mathbb{M}$. Each $\mathbb{S}_i$ is a connected orientable closed surface. Assume $\mathbb{S}_1$ through $\mathbb{S}_{m-1}$ are 'inner' surfaces and $\mathbb{S}_m$ is the 'outer' surface. Note that the dispositions of the surfaces in $\mathbb{R}^3$ can form exotic configurations that may be knotted or linked, see figure 11.

By assumption $\mathbb{M}$ is connected and we have $\beta_0 = 1$. The number of loops and tunnels is the sum of surface tunnels:

$$\beta_1 = \sum_{i=1}^{m} g_i,$$

where g_i is the genus of $\mathbb{S}_i$. Furthermore, there is a void for each inner surface and therefore $\beta_2 = m-1$. In other words the Betti numbers of $\mathbb{M}$ can be computed directly from the genuses of the $\mathbb{S}_i$. As explained in section 7, the genus of $\mathbb{S}_i$ can be calculated from the triangulation:

$$g_i = 1 - \frac{1}{2}(v_i - e_i + f_i),$$

where v_i, e_i, f_i are the number of vertices, edges, triangles in the triangulation of $\mathbb{S}_i$.

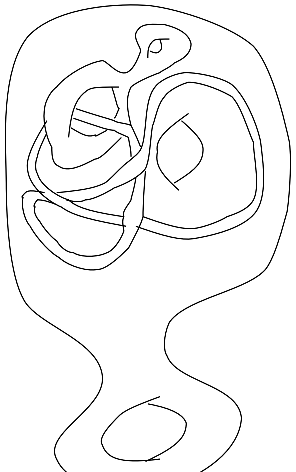

Figure 11: A solid with knotted and linked boundaries.

The above algorithm applies to all simplicial complexes in $\mathbb{R}^3$, not just to triangulations of solids. This is because every such complex $\mathcal{K}$ has a homotopy equivalent solid that can be obtained by thickening $\mathcal{K}$. Delfinado and Edelsbrunner [28] explain how the thickening process can be avoided. Dey and Guha [26] show that the latitudinal generators of the inner surfaces together with the longitudinal generators of the outer surface form a basis of the first homology group of $\mathbb{M}$. This basis can be computed in time $O(gn^2)$, where g is the maximum genus of any surface [117, 27].

Combinatorial Laplacians. Friedman [45] proposes the use of Laplacians for computing Betti numbers of simplicial complexes of arbitrary dimension. His approach appears to be more efficient than computing Smith normal forms, at least for some types of complexes.

Consider the homology groups of a simplicial complex $\mathcal{K}$ with real rather than integer coefficients; the Betti numbers remain the same. More precisely let

$$\ldots \xrightarrow{\partial_{j+2}} \mathsf{C}_{j+1} \xrightarrow{\partial_{j+1}} \mathsf{C}_j \xrightarrow{\partial_j} \mathsf{C}_{j-1} \xrightarrow{\partial_{j-1}} \ldots$$

be the chain complex where C_j is the vector space of j-chains with coefficients in $\mathbb{R}$. $\mathsf{H}_j(\mathbb{R}) = \mathrm{Ker}\,\partial_j | \mathrm{Im}\,\partial_{j+1}$ is the j-th homology group with real coefficients, and β_j is the dimension of H_j. With the usual inner product $\langle \cdot, \cdot \rangle$ on vector spaces over $\mathbb{R}$, we get maps

$$\partial_j^* : \mathsf{C}_{j-1} \to \mathsf{C}_j$$

defined by $\langle \partial_j c, c' \rangle = \langle c, \partial_j^* c' \rangle$, where $c \in \mathsf{C}_j$ and $c' \in \mathsf{C}_{j-1}$. In other words, ∂_j^* is the transpose of ∂_j. Finally for each j define the *Laplacian* $\triangle_j : \mathsf{C}_j \to \mathsf{C}_j$:

$$\triangle_j = \partial_{j+1}\partial_{j+1}^* + \partial_j^*\partial_j,$$

and the *harmonic j-forms*:

$$\mathcal{H}_j = \{c \in \mathsf{C}_j \mid \triangle_j c = 0\}.$$

By the elementary linear algebra behind Hodge theory, $\mathcal{H}_j$ is isomorphic to $\mathsf{H}_j(\mathbb{R})$ so β_j of $\mathcal{K}$ is the dimension of $\mathcal{H}_j$, see [45]. By definition, this is the dimension of the null space of $\triangle_j$, which is the multiplicity of the eigenvalue 0 of $\triangle_j$. Friedman uses matrix-theoretic methods to probabilistically estimate the eigenvalues of the Laplacian $\triangle_j$.

Acknowledgements

The authors of this paper are grateful to the advise and suggestions offered by three referees and many friends and collegues in the mathematics and computer science communities.

References

[1] B. Aronov, R. Seidel and D. Souvaine. On compatible triangulations of simple polygons. *Comput. Geom. Theory Appl.* **3** (1993), 27–35.

[2] D. Avis. On the partitionability of point sets in space. *In* "Proc. 1st Ann. Sympos. Comput. Geom. 1985", 116–120.

[3] R. F. W. Bader. *Atoms in Molecules. A Quantum Theory.* Clanrendon Press, Oxford, England, 1994.

[4] I. Bárány. Geometric and combinatorial applications of Borsuk's theorem. Chapter IX in *New Trends in Discrete Geometry*, J. Pach (ed.), Springer-Verlag, Berlin, 1993, 235–249.

[5] B. Baumgart. A polyhedron representation for computer vision. *In* "Proc. AFIPS Conf. 1975", 589–596.

[6] M. Bern and D. Eppstein. Mesh generation and optimal triangulation. In *Computing in Euclidean Geometry*, F. K. Hwang and D.-Z. Du (eds.), World Scientific, 1992, 47–123.

[7] R. Bing. Some aspects of the topology of 3-manifolds related to the Poincaré conjecture. In *Lectures on Modern Mathematics, vol. II*, T. L. Saaty (ed.), Wiley, New York, 1964, 93–128.

[8] A. Björner. Topological methods. In *Handbook of Combinatorics, vol. I and II*, R. L. Graham, M. Grötschel and L. Lovász (ed.), Elsevier, Amsterdam, 1995, 1819–1872.

[9] J.-D. Boissonnat. Geometric structures for three-dimensional shape representation. *ACM Trans. Graphics* **3** (1984), 266–286.

[10] V. G. Boltjanskij and V. A. Efermovič. *Anschauliche kombinatorische Topologie.* Vieweg, Braunschweig/Wiesbaden, Germany, 1986.

[11] K. Borsuk. Drei Sätze über die n-dimensionale euklidische Sphäre. *Fund. Math.* **20** (1933), 177–190.

[12] E. Brisson. Representing geometric structures in d dimensions: topology and order. *Discr. Comput. Geom.* **9** (1993), 387–426.

[13] L. E. J. Brouwer. Über Abbildung von Mannigfaltigkeiten. *Math. Ann.* **71** (1912), 97–115.

[14] H. Bruggesser and P. Mani. Shellable decompositions of cells and spheres. *Math. Scand.* **29** (1971), 197–205.

[15] R. Casati and A. C. Varzi. *Holes and Other Superficialities.* MIT Press, Cambridge, Massachusetts, 1994.

[16] J. C. Cavendish, D. A. Field and W. H. Frey. An approach to automatic three-dimensional finite element mesh generation. *Internat. J. Numer. Methods Engrg.* **21** (1985), 329–347.

[17] H.-L. Cheng, H. Edelsbrunner and P. Fu. Geometric deformation with skin and alpha shape. Manuscript, 1996.

[18] L. P. Chew. Guaranteed-quality mesh generation for curved surfaces. *In* "Proc. 9th Ann. Sympos. Comput. Geom. 1993", 274–280.

[19] D. I. A. Cohen. On the Sperner lemma. *J. Combin. Theory* **2** (1967), 585–587.

[20] T. H. Cormen, Ch. E. Leiserson and R. L. Rivest. *Introduction to Algorithms.* MIT Press, Cambridge, Massachusetts, 1990.

[21] J. E. Costillo (ed.) *Mathematical Aspects of Numerical Grid Generation.* Frontiers in Applied Math. **8**, SIAM, Philadelphia, 1991.

[22] G. Danaraj and V. Klee. Which spheres are shellable? *Ann. Discrete Math.* **2** (1978), 33–52.

[23] M. Dehn. Über unendliche diskontinuierliche Gruppen. *Math. Ann.* **71** (1912), 116–144.

[24] B. Delaunay. Sur la sphère vide. *Izv. Akad. Nauk SSSR, Otdelenie Matematicheskii i Estestvennyka Nauk* **7** (1934), 793–800.

[25] T. K. Dey and S. Guha. Optimal algorithms for curves on surfaces. *In* "Proc. 35th IEEE Ann. Sympos. Found. Comput. Sci. 1995", 266–274.

[26] T. K. Dey and S. Guha. Algorithms for manifolds and simplicial complexes in Euclidean 3-space. *In* "Proc. 28th ACM Sympos. Theory Comput. 1996", 398–407.

[27] T. K. Dey and H. Schipper. A new technique to compute polygonal schema for 2-manifolds with application to null-homotopy detection. *Discrete Comput. Geom.* **14** (1995), 93–110.

[28] C. J. A. Delfinado and H. Edelsbrunner. An incremental algorithm for Betti numbers of simplicial complexes on the 3-sphere. *Comput. Aided Geom. Design* **12** (1995), 771–784.

[29] D. P. Dobkin and M. J. Laszlo. Primitives for the manipulation of three-dimensional subdivisions. *Algorithmica* **4** (1989), 3–32.

[30] J. A. Dougenik, N. R. Chrisman and D. R. Niemeyer. An algorithm to construct continuous area cartograms. *Professional Geographer* **37** (1985), 75–81.

[31] H. Edelsbrunner. *Algorithms in Combinatorial Geometry.* Springer-Verlag, Heidelberg, Germany, 1987.

[32] H. Edelsbrunner. The union of balls and its dual shape. *Discrete Comput. Geom.* **13** (1995), 415–440.

[33] H. Edelsbrunner. Algebraic decomposition of non-convex polyhedra. *In* "Proc. 36th Ann. IEEE Sympos. Found. Comput. Sci. 1995", 248–257.

[34] H. Edelsbrunner. Surface reconstruction by wrapping finite sets in space. Rept. rgi-tech-96-001, Raindrop Geomagic, Urbana, Illinois, 1996.

[35] H. Edelsbrunner. Deformable smooth surface design. Rept. rgi-tech-96-002, Raindrop Geomagic, Urbana, Illinois, 1996.

[36] H. Edelsbrunner, M. A. Facello and J. Liang. On the definition and the construction of pockets in macromolecules. *In* "Proc. Pacific Sympos. Biocomputing 1996", World Scientific.

[37] H. Edelsbrunner, D. G. Kirkpatrick and R. Seidel. On the shape of a set of points in the plane. *IEEE Trans. Inform. Theory* **IT-29** (1983), 551–559.

[38] H. Edelsbrunner and E. P. Mücke. Three-dimensional alpha shapes. *ACM Trans. Graphics* **13** (1994), 43–72.

[39] H. Edelsbrunner and N. R. Shah. Triangulating topological spaces. *Internat. J. Comput. Geom. Appl.* **7** (1997), 365–378.

[40] H. Edelsbrunner and R. Waupotitsch. A combinatorial approach to cartograms. *J. Comput. Geom. Theory Appl.* **7** (1997), 343–360.

[41] D. Eppstein. Linear complexity hexahedral mesh generation. *In* "Proc. 12th Ann. Sympos. Comput. Geom. 1996", 58–67.

[42] I. Fáry. On straight line representation of planar graphs. *Acta Sci. Math. (Szeged)* **11** (1948), 229–233.

[43] A. Flores. Über n-dimensionale Komplexe die in R_{2n+1} selbstverschlungen sind. *Ergeb. Math. Koll.* **6** (1933/34), 4–7.

[44] M. H. Freedman and F. Quinn. *Topology of 4-manifolds.* Princeton Univ. Press, Princeton, New Jersey, 1990.

[45] J. Friedman. Computing Betti numbers via combinatorial Laplacians. *In* "Proc. 28th ACM Sympos. Theory Comput. 1996", 386–391.

[46] H. Fuchs, Z. M. Kedem and S. P. Uselton. Optimal surface reconstruction from planar contours. *Commun. ACM* **20** (1977), 693–702.

[47] M. R. Garey and D. S. Johnson. *Computers and Intractibility: A Guide to the Theory of NP-Completeness.* Freeman, New York, 1979.

[48] D. Gillman. Triangulations for rubber-sheeting. *Auto Cartography* **7** (1985), 191–197.

[49] C. Gitlin, J. O'Rourke and V. Subramanian. On reconstructing polyhedra from parallel slices. *Intern. J. Comput. Geom. Appl.* **6** (1996), 103–122.

[50] M. D. Greendlinger. On Dehn's algorithm for the conjugacy and word problems with applications. *Comm. Pure Appl. Math.* **13** (1960), 641–677.

[51] B. Grünbaum. *Convex Polytopes.* John Wiley & Sons, London, 1967.

[52] L. Guibas and J. Stolfi. Primitives for the manipulation of general subdivisions and the computation of Voronoi diagrams. *ACM Trans. Graphics* **4** (1985), 74–123.

[53] V. Guillemin and A. Pollack. *Differential Topology.* Prentice Hall, Englewood Cliffs, New Jersey, 1974.

[54] H. Gupta and R. Wenger. Constructing piecewise linear homeomorphisms of simple polygons. *J. Algorithms*, to appear.

[55] H. Hadwiger. Simultane Vierteilung zweier Körper. *Arch. Math. (Basel)* **17** (1966), 274–278.

[56] W. Haken. Theorie der Normalflächen. *Acta Math.* **105** (1961), 245–375.

[57] W. Haken. Various aspects of the three-dimensional Poincaré problem.

[58] P. S. Heckbert. Bilinear Coons patch image warping. In *Graphics Gems IV*, P. S. Heckbert (ed.), Academic Press, Boston, 1994, 438–446.

[59] G. Hemion. *The Classification of Knots and 3-dimensional Spaces.* Oxford Univ. Press, Oxford, 1992.

[60] G. T. Herman. Oriented surfaces in digital spaces. *CVGIP: Graphical Models Image Process.* **55** (1993), 381–396.

[61] J. E. Hopcroft and R. E. Tarjan. Efficient planarity testing. *J. Assoc. Mach. Soc.* **21** (1974), 549–568.

[62] S. Jadhav and A. Mukhopadhyay. Computing a centerpoint of a finite planar set of points in linear time. *In* "Proc. 9th Ann. Sympos. Comput. Geom. 1993", 83–90.

[63] E. R. van Kampen. Komplexe in euklidischen Räumen. *Abh. Math. Sem.* **9** (1932), 72–78 and 152–153.

[64] 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.

[65] W. J. Kaufmann III and L. L. Smarr. *Supercomputing and the Transformation of Science.* Scientific American Library, New York, 1993.

[66] J. R. Kent, W. E. Carlson and R. E. Parent. Shape transformation for polyhedral objects. *Computer Graphics* **26** (1992), 47–54.

[67] P. Koebe. Kontaktprobleme der konformen Abbildung. *Ber. Verh. Sächs. Akad. Wiss., Leipzig, Math.-Phys. Kl.* **88** (1936), 141–164.

[68] T. Y. Kong and A. Rosenfeld. Digital topology: introduction and survey. *Comput. Vision Graph. Image Process.* **48** (1989), 357–393.

[69] K. Kuratowski. Sur le problème des courbes en topologie. *Fund. Math.* **15** (1930), 271–283.

[70] I. Lakatos. *Proofs and Refutations. The Logic of Mathematical Discovery.* Cambridge Univ. Press, England, 1976.

[71] B. Lee and F. M. Richards. The interpretation of protein structures: estimation of static accessibility. *J. Mol. Biol.* **55** (1971), 379–400.

[72] C. N. Lee, T. Poston and A. Rosenfeld. Winding and Euler numbers for 2D and 3D digital images. *CVGIP: Graphical Models Image Process.* **53** (1991), 522–537.

[73] J. Leray. Sur la forme des espaces topologiques et sur les point fixes des représentations. *J. Math. Pures Appl.* **24** (1945), 95–167.

[74] R. J. Lipton and R. E. Tarjan. A separator theorem for planar graphs. *SIAM J. Appl. Math.* **36** (1979), 177–189.

[75] C.-Y. Lo, J. Matoušek and W. Steiger. Algorithms for ham-sandwich cuts. *Discrete Comput. Geom.* **11** (1994), 433–452..

[76] R. C. Lyndon and P. E. Schupp. *Combinatorial Group Theory.* Springer-Verlag, New York, 1977.

[77] M. Mäntylä. *An Introduction to Solid Modeling.* Comput. Sci. Press, Rockville, Maryland, 1988.

[78] A. A. Markov. Insolvability of the problem of homeomorphy. *In* "Proc. Intern. Cong. Math. 1958", Cambridge, 300-306.

[79] T. Martinetz and K. Schulten. Topology representing networks. *Neural Networks* **7** (1994), 507–522.

[80] W. S. Massey. *Algebraic Topology: An Introduction.* Springer-Verlag, New York, 1967.

[81] G. L. Miller, S.-H. Teng, W. P. Thurston and S. A. Vavasis. Separators for sphere-packings and nearest neighbor graphs. *J. Assoc. Comput. Mach.*, to appear.

[82] J. Milnor. Two complexes which are homeomorphic but combinatorially distinct. *Ann. Math.* (1961), 575–590.

[83] J. Milnor. *Morse Theory.* Annals Math. Studies, Princeton Univ. Press, New Jersey, 1963.

[84] E. Moise. *Geometric Topology in Dimensions 2 and 3.* Springer-Verlag, New York, 1977.

[85] J. R. Munkres. *Elements of Algebraic Topology.* Addison-Wesley, Redwood City, 1984.

[86] N. J. Naccache and R. Shinghal. SPTA: A proposed algorithm for thinning binary patterns. *IEEE Trans. System, Man, Cybernet.* **14** (1984), 409–418.

[87] 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.

[88] Y.-L. O, A. Toet, D. Foster, H. J. A. M. Heijmans and P. Meer (eds.) *Shape in Picture. Mathematical Description of Shape in Grey-level Images.* NATO ASI Series, Springer-Verlag, Berlin, 1992.

[89] H. Poincaré. Sur la généralisation d'un théorème d'Euler relatif aux polyèdres. *C. R. Acad. Sci. Paris* **117** (1893), 144–145.

[90] H. Poincaré. Complément à l'analysis situs. *Rend. Circ. Mat. Palermo* **13** (1899), 285–343.

[91] E. A. Ramos. Equipartition of mass distributions by hyperplanes. *Discrete Comput. Geom.* **15** (1996), 147–167.

[92] F. M. Richards. Areas, volumes, packing, and protein structure. *Ann. Rev. Biophys. Bioeng.* **6** (1977), 151–176.

[93] D. Rolfsen. *Knots and Links.* Publish or Perish, Berkeley, 1976.

[94] G. D. Rose. No assembly required. *The Sciences* **36** (1996), 26–31.

[95] J. H. Rubinstein. An algorithm to recognise the 3-sphere. Manuscript, Univ. Melbourne, Australia, 1995.

[96] M. E. Rudin. An unshellable triangulation of a tetrahedron. *Bull. Amer. Math. Soc.* **64** (1958), 90–91.

[97] H. Schipper. Determining contractibility of curves. *In* "Proc. 8th Sympos. Comput. Geom. 1992", 358–367.

[98] L. Schläfli. Theorie der vielfachen Kontinuität. *Denkschr. Schweiz. naturf. Ges.* **38** (1901), 1–237.

[99] R. Schneiders and R. Bunton. Automatic generation of hexahedral finite element meshes. *Comput. Aided Geom. Design* **12** (1995).

[100] R. Seidel. Constructing higher dimensional convex hulls at logarithmic cost per face. *In* "Proc. ACM 18th Ann. Sympos. Theory Comput. 1986", 404–413.

[101] J. P. Serra. *Image Analysis and Mathematical Morphology.* Academic Press, London, 1982.

[102] A. Shapiro. Obstructions to the imbedding of a complex in a Euclidean space. I. The first obstruction. *Ann. Math.* **66** (1957), 256–269.

[103] Yu. A. Shashkin. *The Euler Characteristic.* Little Math. Lib., Mir Publ., Moscow, 1989.

[104] Yu. A. Shashkin. *Fixed Points.* Amer. Math. Soc. and Math. Assoc. Amer., 1991.

[105] S. Smale. Generalized Poincaré's conjecture in dimensions greater than 4. *Ann. Math.* **74** (1961), 391–406.

[106] J. Snoeyink. A trivial knot whose spanning disks have exponential size. *In* "Proc. 6th Sympos. Comput. Geom. 1990", 139–147.

[107] E. H. Spanier. *Algebraic Topology.* Springer-Verlag, New York, 1966.

[108] E. Sperner. Neuer Beweis für die Invarianz der Dimensionszahl und des Gebietes. *Hamburger Abhand.* **6** (1928), 265–272.

[109] E. Steinitz and H. Rademacher. *Vorlesungen über die Theorie der Polyeder.* Springer-Verlag, Berlin, 1934.

[110] K. Stephenson. Circle packings in the approximation of conformal mappings. *Bull. Amer. Math. Soc.* **23** (1990), 407–415.

[111] J. Stillwell. *Classical Topology and Combinatorial Group Theory.* Springer-Verlag, New York, 1980.

[112] A. Thompson. This position and the recognition problem for S^3. Manuscript, Dept. Math., Univ. California at Davis, 1996.

[113] J. F. Thompson, Z. U. A. Warsi and C. W. Mastin. *Numerical Grid Generation. Foundations and Applications.* North-Holland, New York, 1985.

[114] W. P. Thurston. *Three-dimensional Geometry and Topology, vol. 1.* S. Levy (ed.), Princeton Univ. Press, New Jersey, 1985.

[115] W. Tutte. Convex representation of graphs. *Proc. London Math. Soc.* **10** (1960), 304–320.

[116] G. Vegter. Computational topology. In *The Handbook of Discrete and Computational Geometry,* J. E. Goodman and J. O'Rourke (eds.), CRC Press LLC, Boca Raton, Florida, 1997, 517–536.

[117] G. Vegter and C. K. Yap. Computational complexity of combinatorial surfaces. *In* "Proc. 6th Sympos. Comput. Geom. 1990", 102–111.

[118] G. Wegner. d-collapsing and nerves of families of convex sets. *Arch. Math.* **26** (1975), 317–321.

[119] K. J. Weiler. Topological structures for geometric modeling. Ph.D. Thesis, Comput. Syst. Engin., Rensselaer Polytechnic Inst., 1986.

[120] M. A. Wesley. Construction and use of geometric modeling systems. In *Computer Aided Design - Modeling, Systems Engineering CAD Systems,* J. Encarnacao (ed.), Springer-Verlag, 1980.

[121] M. S. White Jr. and P. Griffin. Piecewise linear rubber-sheet map transformation. *Amer. Cartographer* **12** (1985), 123–131.

[122] W.-T. Wu. On a theorem of Leray. *Chinese Math.* **2** (1962), 398–410.

[123] F. F. Yao, D. P. Dobkin, H. Edelsbrunner and M. S. Paterson. Partitioning space for range queries. *SIAM J. Comput.* **18** (1989), 371–384.

[124] G. M. Ziegler. Shelling polyhedral 3-balls and 4-polytopes. Manuscript, 1996.

[125] R. T. Živaljević and S. Vrećica. An extension of the ham sandwich theorem. *Bull. London Math. Soc.* **22** (1990), 183–186.

A Definitions

The appendix presents definitions from topology grouped under topology and space, affine and convex hull, simplicial complexes, maps and equivalences, fundamental groups and covering spaces, manifolds, and homology groups and Betti numbers.

Topology and Space. The mathematical definition of a topology is exceedingly general. We present the definition and immediately specialize it assuming an embedding in Euclidean space.

A **topology** on a set $\mathbb{X}$ is a system $\mathcal{X}$ of subsets with the following properties:

(i) $\emptyset, \mathbb{X} \in \mathcal{X}$,

(ii) if $\mathcal{Z} \subseteq \mathcal{X}$ then $\bigcup \mathcal{Z} \in \mathcal{X}$, and

(iii) if $\mathcal{Z} \subseteq \mathcal{X}$ is finite then $\bigcap \mathcal{Z} \in \mathcal{X}$.

A set $\mathbb{X}$ with a topology is called a **topological space**. The sets in $\mathcal{X}$ are the **open** sets and the complements of the open sets are the **closed** sets of $\mathbb{X}$. A **neighborhood** of a point $x \in \mathbb{X}$ is an open set that contains x. A **cover** is a collection of sets whose union is $\mathbb{X}$. $\mathbb{X}$ is **compact** if every cover of $\mathbb{X}$ with open sets has a finite subcover. $\mathbb{X}$ is **connected** if the only subsets of $\mathbb{X}$ that are both open and closed are $\emptyset$ and $\mathbb{X}$. The **subspace topology** of $\mathbb{Y} \subseteq \mathbb{X}$ is the system $\mathcal{Y} = \{\mathbb{Y} \cap X \mid X \in \mathcal{X}\}$. With this topology $\mathbb{Y}$ is called a **subspace** of $\mathbb{X}$.

The d**-dimensional Euclidean space** is the set of real d-tuples, $\mathbb{R}^d = \{x = (x_1, x_2, \ldots, x_d) \mid x_i \in \mathbb{R}\}$. The **norm** of $x \in \mathbb{R}^d$ is $\|x\| = (\sum_{i=1}^{d} x_i^2)^{\frac{1}{2}}$, and the **distance** between points $x, y \in \mathbb{R}^d$ is $|xy| = \|x - y\|$. The **open ball** with center $z \in \mathbb{R}^d$ and radius $\varepsilon > 0$ is $b_{z,\varepsilon} = \{x \mid |xz| < \varepsilon\}$. The set of open balls generates what we call the **Euclidean topology** of $\mathbb{R}^d$, namely the system whose sets are unions of open balls. Every subset $\mathbb{Y} \subseteq \mathbb{R}^d$ defines a subspace topology and thus a subspace of $\mathbb{R}^d$. Particularly useful subspaces of $\mathbb{R}^d$ are the d**-halfspace**, the **(closed)** d**-ball**, and the $(d-1)$**-sphere**:

$$
\begin{aligned}
\mathbb{H}^d &= \{x \in \mathbb{R}^d \mid x_1 \geq 0\}, \\
\mathbb{B}^d &= \{x \in \mathbb{R}^d \mid \|x\| \leq 1\}, \quad \text{and} \\
\mathbb{S}^{d-1} &= \{x \in \mathbb{R}^d \mid \|x\| = 1\}.
\end{aligned}
$$

Affine and Convex Hull. We assume the usual vector space with componentwise addition of two vectors and multiplication of a vector with a scalar. The terms vector and point are used synonymously.

Given a set of points $T = \{p_0, p_1, \ldots, p_n\}$ in $\mathbb{R}^d$, the **affine hull** is

$$
\operatorname{aff} T = \{\sum_{i=0}^{n} \phi_i p_i \mid \sum_{i=0}^{n} \phi_i = 1\}.
$$

T is **affinely independent** if $\operatorname{aff} T$ is different from the affine hull of every proper subset of T. A k**-flat** is the affine hull of $k+1$ affinely independent points. The **convex hull** of T is

$$
\operatorname{conv} T = \{x \in \operatorname{aff} T \mid \phi_i \geq 0 \text{ for all } i\}.
$$

A k**-simplex** is the convex hull of $k+1$ affinely independent points. We call a 0-simplex a **vertex**, a 1-simplex an **edge**, a 2-simplex a **triangle**, and a 3-simplex a **tetrahedron**. Let $T = \{p_0, p_1, \ldots, p_k\}$ be affinely independent. Then $\sigma = \operatorname{conv} T$ is a k-simplex and its **dimension** is $\dim \sigma = k = \operatorname{card} T - 1$. A **face** of σ is a simplex $\tau = \operatorname{conv} U$ with $U \subseteq T$; it is **proper** if U is a proper subset of T. The **barycentric coordinates** of a point $x \in \sigma$ are the real numbers ϕ_i with

$$
\sum_{i=0}^{k} \phi_i p_i = x \quad \text{and} \quad \sum_{i=0}^{k} \phi_i = 1.
$$

The **barycenter** of σ is the point $b(\sigma)$ with barycentric coordinates $\phi_i = \frac{1}{k+1}$ for all i.

Simplicial Complexes. A **simplicial complex** is a finite collection $\mathcal{K}$ of simplices with the following two properties:

(i) if $\sigma \in \mathcal{K}$ and τ is a face of σ then $\tau \in \mathcal{K}$, and

(ii) if $\sigma, \sigma' \in \mathcal{K}$ then $\sigma \cap \sigma'$ is empty or a face of both.

The **vertex set** is $\operatorname{Vert} \mathcal{K} = \{\sigma \in \mathcal{K} \mid \dim \sigma = 0\}$. The **underlying space** is the part of space covered by simplices in $\mathcal{K}$: $|\mathcal{K}| = \bigcup \mathcal{K}$. The **dimension** is $\dim \mathcal{K} = \max\{\dim \sigma \mid \sigma \in \mathcal{K}\}$. A **subcomplex** is a simplicial complex $\mathcal{L} \subseteq \mathcal{K}$. The **closure** of a subset $L \subseteq \mathcal{K}$ is the smallest subcomplex that contains L: $\operatorname{Cl} L = \{\sigma \in \mathcal{K} \mid \sigma \subseteq \tau \text{ for some } \tau \in L\}$. The **star** and **link** of $\tau \in \mathcal{K}$ are

$$
\begin{aligned}
\operatorname{St} \tau &= \{\sigma \in \mathcal{K} \mid \tau \text{ is a face of } \sigma\} \quad \text{and} \\
\operatorname{Lk} \tau &= \{\sigma \in \operatorname{Cl} \operatorname{St} \tau \mid \sigma \cap \tau = \emptyset\}.
\end{aligned}
$$

A simplicial complex $\mathcal{L}$ with $|\mathcal{L}| = |\mathcal{K}|$ is a **subdivision** of $\mathcal{K}$ if every simplex in $\mathcal{L}$ is contained in a simplex in $\mathcal{K}$.

An **abstract simplicial complex** is a finite system $\mathcal{A}$ of sets so $X \in \mathcal{A}$ and $Y \subseteq X$ implies $Y \in \mathcal{A}$. The sets in $\mathcal{A}$ are called **abstract simplices**. The **dimensions** of X and $\mathcal{A}$ are $\dim X = \operatorname{card} X - 1$ and $\dim \mathcal{A} = \max\{\dim X \mid X \in \mathcal{A}\}$. The **vertex set** of $\mathcal{A}$ is $\operatorname{Vert} \mathcal{A} = \bigcup \mathcal{A}$. A **graph** is a 1-dimensional abstract simplicial complex. The **nerve** of a finite system A is

$$\operatorname{Nrv} A = \{X \subseteq A \mid \bigcap X \neq \emptyset\}.$$

$\operatorname{Nrv} A$ is an abstract simplicial complex with vertex set A. A **geometric realization** is a map $f : \operatorname{Vert} \mathcal{A} \to \mathbb{R}^d$ for which $\mathcal{K} = \{\operatorname{conv} f(X) \mid X \in \mathcal{A}\}$ is a simplicial complex. A useful result is that every abstract simplicial complex of dimension k has a geometric realization in $\mathbb{R}^d$ for some large enough d.

A **chain** is a collection of abstract simplices that can be ordered so that $X_0 \subset X_1 \subset \ldots \subset X_j$. The **order complex** is the system $\mathcal{C}$ of chains in $\mathcal{A}$; it is an abstract simplicial complex because every subset of a chain is again a chain. Let $\mathcal{K}$ be the simplicial complex defined by a geometric realization $f : \operatorname{Vert} \mathcal{A} \to \mathbb{R}^d$. Extend f to $g : \operatorname{Vert} \mathcal{C} \to \mathbb{R}^d$ by sending every abstract simplex $X \in \mathcal{A} = \operatorname{Vert} \mathcal{C}$ to its barycenter: $g(X) = b(\operatorname{conv} f(X))$. The **barycentric subdivision** of $\mathcal{K}$ is

$$\operatorname{Sd} \mathcal{K} = \{\operatorname{conv} g(C) \mid C \in \mathcal{C}\}.$$

$\operatorname{Sd} \mathcal{K}$ is indeed a simplicial complex and g is a geometric realization of the order complex.

Maps and Equivalences. A function $f : \mathbb{X} \to \mathbb{Y}$ is **continuous** if the preimage of every open set in $\mathbb{Y}$ is open in $\mathbb{X}$. A **map** is a continuous function. A **homeomorphism** is a bijective map whose inverse is also continuous. $\mathbb{X}$ and $\mathbb{Y}$ are **homeomorphic** or **topologically equivalent**, denoted $\mathbb{X} \approx \mathbb{Y}$, if there is a homeomorphism between them. A **triangulation** of $\mathbb{X}$ is a simplicial complex $\mathcal{K}$ together with a homeomorphism $h : |\mathcal{K}| \to \mathbb{X}$. $\mathbb{X}$ is **triangulable** if it has a triangulation. An **embedding** is an injection $e : \mathbb{X} \to \mathbb{Y}$ whose restriction to the image, $e(\mathbb{X}) \subseteq \mathbb{Y}$, is a homeomorphism. A **retraction** of a topological space $\mathbb{X}$ to a subspace $\mathbb{Y}$ is a map $r : \mathbb{X} \to \mathbb{Y}$ with $r(y) = y$ for all $y \in \mathbb{Y}$. An **immersion** is a map $f : \mathbb{X} \to \mathbb{Y}$ such that every $x \in \mathbb{X}$ has a neighborhood on which f is bijective.

A **homotopy** between two maps $f, g : \mathbb{X} \to \mathbb{Y}$ is a map

$$F : \mathbb{X} \times [0,1] \to \mathbb{Y}$$

for which $F(x,0) = f(x)$ and $F(x,1) = g(x)$ for all $x \in \mathbb{X}$. f and g are **homotopic** if there exists a homotopy F between f and g. F **fixes** x if $F(x,t)$ is the same for all $t \in [0,1]$. An **isotopy** between f and g is a homotopy F so that every $f_t : \mathbb{X} \to \mathbb{Y}$ defined by $f_t(x) = F(x,t)$ is a homeomorphism between $\mathbb{X}$ and its image. A **deformation retraction** of $\mathbb{X}$ to $\mathbb{Y} \subseteq \mathbb{X}$ is a homotopy between the identity $1_{\mathbb{X}} : \mathbb{X} \to \mathbb{X}$ and a retraction $r : \mathbb{X} \to \mathbb{Y}$ that fixes all $y \in \mathbb{Y}$.

Two topological spaces, $\mathbb{X}$ and $\mathbb{Y}$, are **homotopy equivalent**, denoted $\mathbb{X} \simeq \mathbb{Y}$, if there are maps $a : \mathbb{X} \to \mathbb{Y}$ and $b : \mathbb{Y} \to \mathbb{X}$ with $b \circ a$ homotopic to $1_{\mathbb{X}}$ and $a \circ b$ homotopic to $1_{\mathbb{Y}}$. For example, $\mathbb{X} \simeq \mathbb{Y}$ if there is a deformation retraction of $\mathbb{X}$ to $\mathbb{Y}$. Indeed, define $a = r$ and $b = 1_{\mathbb{Y}}$. Then $a \circ b = 1_{\mathbb{Y}}$ and the deformation retraction itself is a homotopy between $b \circ a$ and $1_{\mathbb{X}}$. $\mathbb{X}$ is **contractible** if it is homotopy equivalent to a point.

Topological spaces given in terms of simplicial complexes can be related via piecewise linear maps. Let $\mathcal{K}, \mathcal{L}$ be simplicial complexes. A **vertex map** is a function $v : \operatorname{Vert} \mathcal{K} \to \operatorname{Vert} \mathcal{L}$ with the property that if $\sigma = \operatorname{conv} T$ is a simplex in $\mathcal{K}$ then $\operatorname{conv} v(T)$ is a simplex in $\mathcal{L}$. A **simplicial map** $f : |\mathcal{K}| \to |\mathcal{L}|$ is the linear extension of a vertex map v:

$$f(x) = \sum_{p_i \in T} \phi_i \cdot v(p_i),$$

where $\sigma = \operatorname{conv} T$ is a simplex in $\mathcal{K}$ that contains x and ϕ_i is the barycentric coordinate of x that corresponds to $p_i \in T$. $\mathcal{K}$ and $\mathcal{L}$ are **isomorphic** or **simplicially equivalent** if they permit a bijective vertex map v. In this case the linear extension is a simplicial homeomorphism between $|\mathcal{K}|$ and $|\mathcal{L}|$. $\mathcal{K}$ and $\mathcal{L}$ are **PL-equivalent** if there are simplicially equivalent subdivisions $\mathcal{K}'$ of $\mathcal{K}$ and $\mathcal{L}'$ of $\mathcal{L}$. It has been conjectured that two simplicial complexes are PL-equivalent if they triangulate the same topological space. This is known as the **Hauptvermutung** in combinatorial topology and has since been disproved for complexes of dimension 3 and higher [82]. The conjecture is correct for 3-dimensional manifolds and 2-dimensional triangulable spaces [84].

Fundamental Group and Covering Spaces. Let $\mathbb{X}$ be a topological space. A **path** in $\mathbb{X}$ is a continuous map $p : [0,1] \to \mathbb{X}$. The path **joins** the **initial** point, $p(0)$, to the **terminal** point, $p(1)$. $\mathbb{X}$ is **path-connected** if every pair of points in $\mathbb{X}$ can be joined by a path. Two paths are **equivalent** if they are connected by a homotopy that leaves the common initial and terminal points fixed. The **inverse** of p is $p^{-1}(t) = p(1-t)$. The **product** of two paths p and q is defined if $p(1) = q(0)$:

$$(p \cdot q)(t) = \begin{cases} p(2t) & \text{if } 0 \leq t \leq \frac{1}{2}, \text{ and} \\ q(2t-1) & \text{if } \frac{1}{2} \leq t \leq 1 . \end{cases}$$

The equivalence relation and the product are compatible, that is, if p_1, q_1 and p_2, q_2 are two pairs of equivalent paths then $p_1 \cdot q_1$ and $p_2 \cdot q_2$ are equivalent. A path p is a **loop** if $p(0) = p(1) = x_0$, and the loop is **based** at x_0. For example, $p \cdot p^{-1}$ is a loop equivalent to the constant map $[0,1] \to x_0$. The equivalence classes of loops based at x_0 together with the product operation form a group called the **fundamental group** of $\mathbb{X}$ at x_0 and denoted $\pi(\mathbb{X}, x_0)$. For a path-connected space $\mathbb{X}$ any two groups $\pi(\mathbb{X}, x_0)$ and $\pi(\mathbb{X}, y_0)$ are isomorphic. Therefore, for a path-connected space $\mathbb{X}$ we have a unique fundamental group for the entire space, denoted $\pi(\mathbb{X})$. We note that for finite simplicial complexes and for manifolds the notions of path-connectedness and connectedness are the same and do not have to be distinguished. The fundamental group is invariant over homotopy equivalent spaces. $\mathbb{X}$ is **simply connected** if $\pi(\mathbb{X})$ is trivial. If $\mathbb{X}$ is contractible then it is simply connected. The reverse is not correct: for $d \geq 2$ the d-sphere is simply connected but not contractible.

Let $\mathbb{X}$ be a topological space. A **covering space** of $\mathbb{X}$ is a topological space $\mathbb{Y}$ together with a map $p : \mathbb{Y} \to \mathbb{X}$. p is called a **projection** and satisfies the following property: for each point $x \in \mathbb{X}$ there is a path-connected neighborhood U so that for each path-connected component V of $p^{-1}(U)$ the restriction of p to V is a homeomorphism. A covering space $\mathbb{Y}$ of $\mathbb{X}$ is **universal** if $\mathbb{Y}$ is connected. Any two universal covering spaces of $\mathbb{X}$ are topologically equivalent. For example, $\mathbb{R}^1$ is a universal covering space of $\mathbb{S}^1$ and $\mathbb{R}^2$ is a universal covering space of the torus. An important property of covering spaces is that they lift maps. For example, if α is a path in a space $\mathbb{X}$ with covering space $\mathbb{Y}$ and projection p, and if $y_0 \in \mathbb{Y}$ is such that $p(y_0) = \alpha(0)$, then there exists a unique path α' in $\mathbb{Y}$ with initial point $\alpha'(0) = y_0$ and such that $p \circ \alpha' = \alpha$. We say the path α' is a **lifting** of the path α.

Manifolds. A topological space $\mathbb{M}$ is a d**-manifold** if every point $x \in \mathbb{M}$ has a neighborhood homeomorphic to $\mathbb{R}^d$. A **closed surface** is a compact 2-manifold. The long-standing **Poincaré conjecture** states that every simply connected 3-manifold is homeomorphic to the 3-sphere. In spite of a tremendous amount of effort there is still no proof and no counterexample [7, 57]. $\mathbb{M}$ is a d-**manifold with boundary** if every $x \in \mathbb{M}$ has a neighborhood homeomorphic to $\mathbb{R}^d$ or to $\mathbb{H}^d$. The **boundary** of $\mathbb{M}$ is the set of points with neighborhoods homeomorphic to the d-dimensional half-space, $\mathbb{H}^d$. The boundary of $\mathbb{M}$ is either empty or a $(d-1)$-manifold. All manifolds of dimension $d \leq 3$ are triangulable.

An **ordered** k**-simplex** is a k-simplex together with a permutation of its vertices, and we write

$$\sigma \;=\; [p_0, p_1, \ldots, p_k].$$

Two orderings have the same **orientation** if they differ by an even permutation. All simplices of dimension 1 or higher have two orientations. The orientation of a $(k-1)$-face **induced** by σ is

$$\tau \;=\; (-1)^i [p_0, \ldots, p_{i-1}, p_{i+1}, \ldots, p_k],$$

where a leading minus reverses orientation. Two k-simplices sharing a $(k-1)$-face τ are **consistently oriented** if they induce *different* orientations of τ. A triangulable d-manifold is **orientable** if all d-simplices in a triangulation can be ordered so all adjacent pairs are consistently oriented. Otherwise, the d-manifold is **non-orientable**. For example, the projective plane and the Klein bottle are two non-orientable closed surfaces, and the Möbius strip is a non-orientable 2-manifold with boundary. Non-orientable closed surfaces cannot be embedded in $\mathbb{R}^3$.

Homology Groups and Betti Numbers. We consider only simplicial homology with integer coefficients. Let $\mathcal{K}$ be a simplicial complex and assign to each simplex an arbitrary but fixed ordering of its vertices. $\mathcal{K}^j \subseteq \mathcal{K}$ is the set of ordered j-simplices. A j**-chain** is a function $c : \mathcal{K}^j \to \mathbb{Z}$. We write the chain as a formal sum:

$$c \;=\; \sum_{\sigma \in \mathcal{K}^j} c(\sigma) \cdot \sigma.$$

Two j-chains are added componentwise. The j-chains together with addition form the **group of** j**-chains**, C_j. The **boundary operator** maps an ordered j-simplex $\sigma = [p_0, p_1, \ldots, p_j]$ to a $(j-1)$-chain

$$\partial_j \sigma \;=\; \sum_{i=0}^{j} (-1)^i [p_0, \ldots, p_{i-1}, p_{i+1}, \ldots, p_j].$$

The **boundary homomorphism**, $\partial_j : \mathsf{C}_j \to \mathsf{C}_{j-1}$, is defined by $\partial_j c = \sum c(\sigma) \cdot \partial_j \sigma$. The **chain complex** is the sequence of chain groups connected by boundary homomorphisms:

$$\ldots \xrightarrow{\partial_{j+2}} \mathsf{C}_{j+1} \xrightarrow{\partial_{j+1}} \mathsf{C}_j \xrightarrow{\partial_j} \mathsf{C}_{j-1} \xrightarrow{\partial_{j-1}} \ldots$$

The **image** and **kernel** of the boundary homomorphism are

$$\begin{aligned} \operatorname{Im} \partial_j &= \{\partial_j c \mid c \in \mathsf{C}_j\}, \\ \operatorname{Ker} \partial_j &= \{c \in \mathsf{C}_j \mid \partial_j c = 0\}. \end{aligned}$$

$c \in \mathsf{C}_j$ is a j**-cycle** if $c \in \operatorname{Ker} \partial_j$ and it is a j**-boundary** if $c \in \operatorname{Im} \partial_{j+1}$. A j-boundary is sometimes called a **null-homologous** j-cycle. The j-boundaries form a subgroup of the group of j-cycles which form a subgroup of the group of j-chains: $\mathsf{B}_j \subseteq \mathsf{Z}_j \subseteq \mathsf{C}_j$. The fact that every j-boundary is also a j-cycle is important and follows from $\partial_j \partial_{j+1} c = 0$ for every $c \in \mathsf{C}_{j+1}$. The **j-th homology group** is the quotient of the cycle group over the boundary group:

$$\mathsf{H}_j = \mathsf{Z}_j | \mathsf{B}_j .$$

Homology groups are finitely generated and abelian. The fundamental theorem on such groups implies H_j can be written as the direct sum of two subgroups, $\mathbb{Z}^{\beta_j}$ and T. β_j is the **rank** of H_j. T is the direct sum of finitely many cyclic groups, $\mathbb{Z}_k$, and is referred to as the **torsion subgroup** of H_j.

The above seemingly roundabout construction is justified by the fact that H_j is invariant over all simplicial complexes triangulating the same topological space $\mathbb{X} \approx |\mathcal{K}|$. It is therefore meaningful to call H_j the j-th homology group of $\mathbb{X}$. Similarly, the **j-th Betti number**, β_j, and the **Euler characteristic**,

$$\chi = \sum (-1)^j \beta_j ,$$

are invariant over all triangulations of $\mathbb{X}$. For example the Euler characteristics of the d-sphere and the d-ball are

$$\begin{aligned} \chi &= 1 + (-1)^d , \\ \chi &= 1. \end{aligned}$$

Indeed, the only non-zero Betti numbers of $\mathbb{S}^d$ are $\beta_0 = \beta_d = 1$ and the only non-zero Betti number of $\mathbb{B}^d$ is $\beta_0 = 1$. The Euler characteristic of every compact d-manifold with odd d is 0.