Topics in Combinatorial Differential Topology and Geometry

Robin Forman

IAS/Park City Mathematics Series
Volume 14, 2004

Topics in Combinatorial Differential Topology and Geometry

Robin Forman

Many questions from a variety of areas of mathematics lead one to the problem of analyzing the topology or the combinatorial geometry of a simplicial complex. We will see a number of examples in these notes. Some very general theories have been developed for the investigation of similar questions for smooth manifolds. Our goal in these lectures is to show that there is much to be gained in the world of combinatorics from borrowing questions, tools, motivation, and even inspiration from the theory of smooth manifolds.

These lectures center on two main topics which illustrate the dramatic impact that ideas from the study of smooth manifolds have had on the study of combinatorial spaces. The first topic has its origins in differential topology, and the second in differential geometry.

One of the most powerful and useful tools in the study of the topology of smooth manifolds is Morse theory. In the first three lectures we present a combinatorial Morse theory that posesses many of the desirable properties of the smooth theory, and which can be usefully applied to the study of very general combinatorial spaces. In the first two lectures we present the basic theory along with numerous examples. In the third lecture, we show that discrete Morse theory is a very natural tool for the study of some questions in complexity theory.

Much of the study of global differential geometry is concerned with the relationship between the geometry of a Riemannian manifold and its topology. One long conjectured, still unproved, relationship is the Hopf conjecture, which states that if a manifold has nonpositive sectional curvature, then the sign of its Euler characteristic depends only on its dimension. (See Lecture 4 for a more precise statement.) In [**15**] Charney and Davis formulated a combinatorial analogue of

[1]The Department of Mathematics, Rice University, Houston, TX, USA 77251.
E-mail address: `forman@rice.edu`.
The author was supported in part by the National Science Foundation. The author would also like to thank Carsten Lange, who served as the TA for these lectures, created most of the figures in these notes, and whose enthusiasm and attention to detail dramatically increased the comprehensibility of the text. The author expresses his gratitude to the organizers of the IAS/PC Summer Institute for their tireless dedication and enthusiasm for all things organizational and mathematical. Their support greatly improved the lectures and these notes.

this conjecture, and then observed that their conjecture is related to some of the central questions in geometric combinatorics. There has been some fascinating recent work on this subject, which has resulted in some very tantalizing, more general conjectures. In Lectures 4 and 5 we present an introduction to the conjectures of Charney and Davis, discuss some of the known partial results, and survey the most recent developments.

LECTURE 1
Discrete Morse Theory

1. Introduction

There is a very close relationship between the topology of a smooth manifold M and the critical points of a smooth function f on M. For example, if M is compact, then f must achieve a maximum and a minimum. Morse theory is a far-reaching extension of this fact. Milnor's beautiful book [**71**] is the standard reference on this subject. In these notes we present an adaptation of Morse theory that may be applied to any simplicial complex (or more general cell complex). There have been other adaptations of Morse Theory that can be applied to combinatorial spaces. For example, a Morse Theory of piecewise linear functions appears in [**59**] and the very powerful "Stratified Morse Theory" was developed by Goresky and MacPherson [**46**],[**47**]. These theories, especially the latter, have each been successfully applied to prove some very striking results.

We take a slightly different approach than that taken in these references. Rather than choosing a suitable class of continuous functions on our spaces to play the role of Morse functions, we will be working entirely with discrete structures. Hence, we have chosen the name *discrete Morse theory* for the ideas we will describe. Moreover, in these notes, we will describe the theory entirely in terms of the (discrete) gradient vector field, rather than an underlying function. We show that even without introducing any continuity, one can recreate, in the category of combinatorial spaces, a complete theory that captures many of the intricacies of the smooth theory, and can be used as an effective tool for a wide variety of combinatorial and topological problems.

The goal of these lectures is to present an overview of the subject of discrete Morse theory that is sufficient both to understand the major applications of the theory to combinatorics, and to apply the the theory to new problems. We will not be presenting theorems in their most recent or most general form, and simple examples will often take the place of proofs. Those interested in a more complete presentation of the theory can consult the reference [**32**]. Earlier surveys of this work have appeared in [**31**] and [**35**], and earlier, and similar, versions of some of the sections in these notes appeared in [**39**] and [**40**].

2. Cell Complexes and CW Complexes

The main theorems of discrete (and smooth) Morse theory are best stated in the language of CW complexes, so we begin with an overview of the basics of such complexes. J. H. C. Whitehead introduced CW complexes in his foundational work on homotopy theory, and all of the results in this section are due to him. The reader should consult [**68**] for a very complete introduction to this topic. In these notes we will consider only finite CW complexes, so many of the subtleties of the subject will not appear.

The building blocks of cell complexes are cells. Let B^d denote the closed unit ball in d-dimensional Euclidean space. That is, $B^d = \{x \in \mathbb{E}^d \text{ s.t. } |x| \leq 1\}$. The boundary of B^d is the unit $(d-1)$-sphere $S^{(d-1)} = \{x \in \mathbb{E}^d \text{ s.t. } |x| = 1\}$. A *$d$-cell* is a topological space which is homeomorphic to B^d. If σ is d-cell, then we denote by $\dot{\sigma}$ the subset of σ corresponding to $S^{(d-1)} \subset B^d$ under any homeomorphism between B^d and σ. A *cell* is a topological space which is a d-cell for some d.

The basic operation of cell complexes is the notion of *attaching a cell.* Let X be a topological space, σ a d-cell and $f : \dot{\sigma} \to X$ a continuous map. We let $X \cup_f \sigma$ denote the disjoint union of X and σ quotiented out by the equivalence relation that each point $s \in \dot{\sigma}$ is identified with $f(s) \in X$. We refer to this operation by saying that $X \cup_f \sigma$ is the result of attaching the cell σ to X. The map f is called *the attaching map.*

We emphasize that the attaching map must be defined on all of $\dot{\sigma}$. That is, the entire boundary of σ must be "glued" to X. For example, if X is a circle, then Figure 1(i) shows one possible result of attaching a 1-cell to X. Attaching a 1-cell to X cannot lead to the space illustrated in Figure 1(ii) since the entire boundary of the 1-cell has not been "glued" to X.

We are now ready for our main definition. A *finite cell complex* is any topological space X such that there exists a finite nested sequence

$$\emptyset \subset X_0 \subset X_1 \subset \cdots \subset X_n = X \tag{1}$$

such that for each $i = 0, 1, 2, \ldots, n$, X_i is the result of attaching a cell to $X_{(i-1)}$.

Note that this definition requires that X_0 be a 0-cell. If X is a cell complex, we refer to any sequence of spaces as in (1) as a *cell decomposition of X*. Suppose that

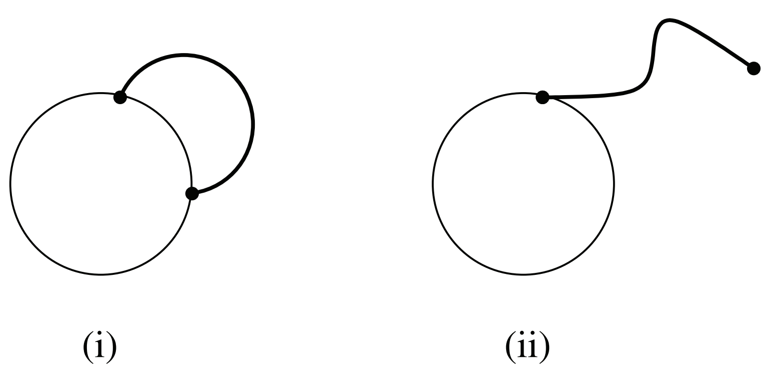

Figure 1. On the left a 1-cell is attached to a circle. This is not true for the picture on the right.

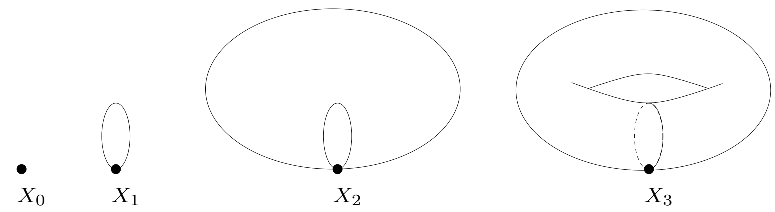

Figure 2. A cell decomposition of the torus.

in the cell decomposition (1), of the $n+1$ cells that are attached, exactly c_d are d-cells. Then we say that the cell complex X has a cell decomposition consisting of c_d d-cells for every d.

We note that a (closed) d-simplex is a d-cell. Thus a finite simplicial complex is a cell complex, and has a cell decomposition in which the cells are precisely the closed simplices.

In Figure 2 we demonstrate a cell decomposition of a 2-dimensional torus which, beginning with the 0-cell, requires attaching two 1-cells and then one 2-cell. Here we can see one of the most compelling reasons for expanding our view from simplicial complexes to more general cell complexes. Every simplicial decomposition of the 2-torus has at least 7 vertices, 21 edges and 14 triangles.

It may seem that quite a bit has been lost in the transition from simplicial complexes to general cell complexes. After all, a simplicial complex is completely described by a finite amount of combinatorial data. On the other hand, the construction of a cell decomposition requires the choice of a number of continuous maps. However, if one is only concerned with the homotopy type of the resulting cell complex, then things begin to look a bit more manageable. Namely, the homotopy type of $X \cup_f \sigma$ depends only on the homotopy type of X and the homotopy class of f.

Theorem 1. *Let $h : X \to X'$ denote a homotopy equivalence, σ a cell, and $f_1 : \dot{\sigma} \to X$, $f_2 : \dot{\sigma} \to X'$ two continuous maps. If $h \circ f_1$ is homotopic to f_2, then $X \cup_{f_1} \sigma$ and $X' \cup_{f_2} \sigma$ are homotopy equivalent.*

(See Theorem 2.3 on page 120 of [**68**].) An important special case is when h is the identity map. We state this case separately for future reference.

Corollary 2. *Let X be a topological space, σ a cell, and $f_1, f_2 : \dot{\sigma} \to X$ two continuous maps. If f_1 and f_2 are homotopic, then $X \cup_{f_1} \sigma$ and $X \cup_{f_2} \sigma$ are homotopy equivalent.*

Therefore, the homotopy type of a cell complex is determined by the homotopy classes of the attaching maps. Since homotopy clases are discrete objects, we have now recaptured a bit of the combinatorial atmosphere that we seemingly lost when generalizing from simplicial complexes to cell complexes.

Let us now present some examples.

1) Suppose X is a topological space which has a cell decomposition consisting of exactly one 0-cell and one d-cell. Then X has a cell decomposition $\emptyset \subset X_0 \subset X_1 = X$. The space X_0 must be the 0-cell, and $X = X_1$ is the result of attaching the d-cell to X_0. Since X_0 consists of a single point, the only possible attaching map is the constant map. Thus X is constructed from taking a closed d-ball and

identifying all of the points on its boundary. One can easily see that this implies that the resulting space is a d-sphere.

2) Suppose X is a topological space which has a cell decomposition consisting of exactly one 0-cell and n d-cells. Then X has a cell decomposition as in (1) such that X_0 is the 0-cell, and for each $i = 1, 2, \ldots, n$ the space X_i is the result of attaching a d-cell to $X_{(i-1)}$. From the previous example, we know that X_1 is a d-sphere. The space X_2 is constructed by attaching a d-cell to X_1. The attaching map is a continuous map from a $(d-1)$-sphere to X_1. Every map of the $(d-1)$-sphere into X_1 is homotopic to a constant map (since $\pi_{(d-1)}(X_1) \cong \pi_{(d-1)}(S^d) \cong 0$). If the attaching map is actually a constant map, then it is easy to see that the space X_2 is the wedge of two d-spheres, denoted by $S^d \vee S^d$. (The wedge of a collection of topological spaces is the space resulting from choosing a point in each space, taking the disjoint union of the spaces, and identifying all of the chosen points.) Since the attaching map must be homotopic to a constant map, Corollary 2 implies that X_2 is homotopy equivalent to a wedge of two d-spheres.

When constructing X_3 by attaching a d-cell to X_2, the relevant information is a map from S^{d-1} to X_2, and the homotopy type of the resulting space is determined by the homotopy class of this map. All such maps are homotopic to a constant map (since $\pi_{d-1}(X_2) \cong \pi_{d-1}(S^d \vee S^d) \cong 0$). Since X_2 is homotopy equivalent to a wedge of two d-spheres, and the attaching map is homotopic to a constant map, it follows from Theorem 1 that X_3 is homotopy equivalent to the space that would result from attaching a d-cell to $S^d \vee S^d$ via a constant map, i.e. X_3 is homotopy equivalent to a wedge of three d-spheres.

Continuing in this fashion, we can see that X must be homotopy equivalent to a wedge of n d-spheres.

The reader should not get the impression that the homotopy type of a cell complex is determined by the number of cells of each dimension. This is true only for very few spaces (and the reader might enjoy coming up with some other examples). The fact that wedges of spheres can, in fact, be identified by this numerical data partly explains why the main theorem of many papers in combinatorial topology is that a certain simplicial complex is homotopy equivalent to a wedge of spheres. Namely such complexes are the easiest to recognize. However, that does not explain why so many simplicial complexes that arise in combinatorics are homotopy equivalent to a wedge of spheres. I have often wondered if perhaps there is some deeper explanation for this.

3) Suppose that X is a cell complex which has a cell decomposition consisting of exactly one 0-cell, one 1-cell and one 2-cell. Let us consider a cell decomposition for X with these cells: $\emptyset \subset X_0 \subset X_1 \subset X_2 = X$. We know that X_0 is the 0-cell. Suppose that X_1 is the result of attaching the 1-cell to X_0. Then X_1 must be a circle, and X_2 arises from attaching a 2-cell to X_1. The attaching map is a map from the boundary of the 2-cell, i.e. a circle, to X_1 which is also a circle. Up to homotopy, such a map is determined by its winding number, which can be taken to be a nonnegative integer. If the winding number is 0, then without altering the homotopy type of X we may assume that the attaching map is a constant map, which yields that $X \sim S^1 \vee S^2$ (where $\sim$ denotes homotopy equivalence). If the winding number is 1 then without altering the homotopy type of X we may assume that the attaching map is a homeomorphism, in which case X is a 2-dimensional disc. If the winding number is 2, then without altering the homotopy type of X

we may assume that the attaching map is a standard degree 2 mapping (i.e. that wraps one circle around the other twice, with no backtracking). The reader should convince him/herself that the result in this case is that X is the 2-dimensional projective space $\mathbb{P}^2$. In fact, each winding number results in a homotopically distinct space. These spaces can be distinguished by their homology, since $H_1(X, \mathbb{Z})$ for the space X resulting from an attaching map with winding number n is isomorphic to $\mathbb{Z}/n\mathbb{Z}$.

It seems that we are not quite done with this example, because we assumed that the 1-cell was attached before the 2-cell, and we must consider the alternative order, in which X_1 is the result of attaching a 2-cell to X_0. In this case, X_1 is a 2-sphere, and $X = X_2$ is the result of attaching a 1-cell to X_1. The attaching map is a map of S^0 into S^2. Since S^2 is connected (i.e. $\pi_0(S^2) = 0$) all such maps are homotopic to a constant map. Taking the attaching map to be a constant map yields that $X = S^1 \vee S^2$. Thus adding the cells in this order merely resulted in fewer possibilities for the homotopy type of X. This is a general phenomenon. Generalizing the argument we just presented, using the fact that $\pi_i(S^d) = 0$ for $i < d$, yields the following statement.

Proposition 3. *Let*

$$\emptyset \subset X_0 \subset X_1 \subset \cdots \subset X_n = X \tag{2}$$

be a cell decomposition of a finite cell complex X. Then X is homotopy equivalent to a finite cell decomposition with precisely the same number of cells of each dimension as in (2), and with the cells attached so that their dimensions form a nondecreasing sequence.

A CW complex is one that can be constructed in this fashion. In fact, even more is required.

Definition 4. A *CW complex* is a cell complex with the property that the boundary of each cell is mapped into the union of the cells of lower dimension.

In some sense, this is a merely technical requirement, as every cell complex is homotopy equivalent to a CW complex. However, there are certain advantages to working with CW complexes, and all of the cell complexes which arise in these notes will be CW complexes.

I first learned of simplicial complexes in a course on algebraic topology. They were introduced as a category of topological spaces for which it was rather easy to define homology and cohomology, i.e. in terms of the simplical chain- and cochain-complexes. One might be concerned that in the transition from simplicial complexes to cell complexes we have lost this ability to easily compute these topological invariants. In fact, much of this computability remains. Let X be a cell complex with a fixed cell decomposition. Suppose that in this decomposition X is constructed from exactly c_d cells of dimension d for each $d = 0, 1, 2, \ldots, n = \dim(K)$, and let $C_d(X, \mathbb{Z})$ denote the space $\mathbb{Z}^{c_d}$ (more precisely, $C_d(X, \mathbb{Z})$ denotes the free abelian group generated by the d-cells of X, each endowed with an orientation). The following is one of the fundamental results in the theory of cell complexes.

Theorem 5. *There are boundary maps $\partial_d : C_d(X, \mathbb{Z}) \to C_{d-1}(X, \mathbb{Z})$, for each d, so that*

$$\partial_{d-1} \circ \partial_d = 0$$

and such that the resulting differential complex

$$0 \longrightarrow C_n(X,\mathbb{Z}) \xrightarrow{\partial_n} \dots \xrightarrow{\partial_1} C_0(X,\mathbb{Z}) \longrightarrow 0$$

calculates the homology of X. That is, if we define

$$H_d(C,\partial) = \frac{\mathrm{Ker}(\partial_d)}{\mathrm{Im}(\partial_{d+1})},$$

then for each d

$$H_d(C,\partial) \cong H_d(X,\mathbb{Z}),$$

where $H_d(X,\mathbb{Z})$ denotes the singular homology of X.

The actual definition of the boundary map ∂ is slightly nontrivial and we will not go into it here (see [**68**, Ch. V, Sec. 2] for the details). In fact, it is here that we see the main distinction between general cell complexes and CW complexes. There may exist multiple choices for the boundary map for a general cell complex, but the boundary map is canonical for a CW complex. At first it may seem that without knowing this boundary map, there is little to be gained from Theorem 5. In fact, much can be learned from just knowing of the existence of such a boundary map. For example, let us choose a coefficient field $\mathbb{F}$, and tensor everything with $\mathbb{F}$ to get a differential complex

$$0 \longrightarrow C_n(X,\mathbb{F}) \xrightarrow{\partial_n} \dots \xrightarrow{\partial_1} C_0(X,\mathbb{F}) \longrightarrow 0$$

which calculates $H_*(X,\mathbb{F})$, where now $C_d(X,\mathbb{F}) \cong \mathbb{F}^{c_d}$.

From basic linear algebra we can deduce the following inequalities.

Theorem 6. *Let X be a cell complex with a fixed cell decomposition with c_d cells of dimension d for each d. Fix a coefficient field $\mathbb{F}$ and let b_* denote the Betti numbers of X with respect to $\mathbb{F}$, i.e. $b_d = \dim(H_d(X,\mathbb{F}))$.*

(i) (The Weak Morse Inequalities) For each d

$$c_d \geq b_d.$$

(ii) Let $\chi(X)$ denote the Euler characteristic of X, i.e.

$$\chi(X) = b_0 - b_1 + b_2 - \dots.$$

Then we also have

$$\chi(X) = c_0 - c_1 + c_2 - \dots.$$

As the name "Weak Morse Inequalities" implies, this theorem can be strengthened. The following inequalities, known as the "Strong Morse Inequalities", also follow from standard linear algebra.

Theorem 7 (The Strong Morse Inequalities)**.** *With all notation as in Theorem 6, for each $d = 0, 1, 2, \dots$*

$$c_d - c_{d-1} + c_{d-2} - \dots + (-1)^d c_0 \geq b_d - b_{d-1} + b_{d-2} - \dots + (-1)^d b_0.$$

As the names imply, Theorem 7 does directly imply Theorem 6, as one can see by comparing Strong Morse Inequalities for consecutive values of d, and using the fact that $b_i = 0$ for i larger than the dimension of K.

We mentioned earlier that a great benefit of passing from simplicial complexes to the more general cell complexes is that one often can use many fewer cells. Let us take another look at this phenomenon in light of the Morse inequalities. Consider the case where X is a two-dimensional torus, so that with respect to any coefficient field $b_0 = 1, b_1 = 2, b_2 = 1$. From the weak Morse inequalities, we have that for any cell decomposition,

$$c_0 \geq b_0 = 1$$

$$c_1 \geq b_1 = 2$$

$$c_2 \geq b_2 = 1.$$

A simplicial decomposition is a special case of a CW decomposition, so these inequalities are satisfied when c_d denotes the number of d-simplices in a fixed simplicial decomposition. However, every simplicial decomposition has at least seven 0-simplices, twenty-one 1-simplices and fourteen 2-simplices, so these inequalities are far from equality. It is generally the case that for a simplicial decomposition these inequalities are very far from optimal, and hence are generally of little interest. On the other hand, earlier we demonstrated a CW decomposition of the two-torus with exactly one 0-cell, two 1-cells and one 2-cell. The inequalities tell us, in particular, that one cannot build the torus using fewer cells.

3. The Morse Theory

In this section we introduce the main topic of the first three lectures, namely discrete Morse theory. Morse theory, in the standard setting of smooth manifolds, is usually described in the language of smooth functions on smooth manifolds (e.g. [**71**]). In practice, though, it is often useful to work with gradient vector fields rather than functions (e.g. [**72**], [**82**]). In the discrete setting, too, one can follow either path. In these notes, we will focus on the notion of a (discrete) gradient vector field. To see how discrete Morse theory can be presented from the function point of view, see [**31**] or [**32**],

Let K be a CW complex. (Most of our examples will be simplicial complexes, but in a few places, even when our object of study is a simplicial complex, it will be convenient to allow more general cell complexes.)

Definition 8. Let β be a $(p+1)$-cell of K, with attaching map $h : S^p \to K_p$, where K_p denotes the union of the cells of dimension $\leq p$.

(i) A cell α is a *face* of β, denoted by $\alpha < \beta$ (or $\beta > \alpha$) if $\beta \neq \alpha \subset \beta$ (where here we are identifying a cell with its image in K).

(ii) A face α of β is said to be *regular* if
   (a) $h^{-1}(\alpha)$ is homeomorphic to a ball, and
   (b) h restricted to $h^{-1}(\alpha)$ is a homeomorphism onto α.

(iii) A *regular CW complex* is a CW complex in which every face is regular. We note that every simplicial complex or polyhedron is a regular CW complex.

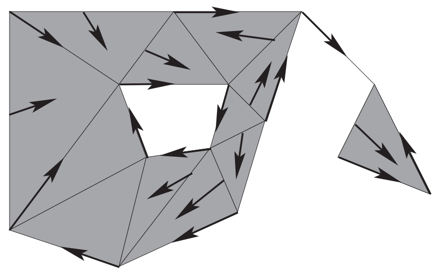

Figure 3. A discrete vector field.

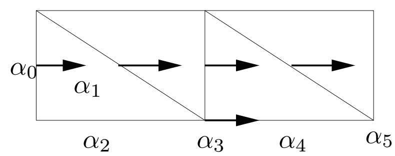

Figure 4. A V-path.

Definition 9. A *discrete vector field* V on K is a collection of pairs $\{\alpha^{(p)} < \beta^{(p+1)}\}$ of cells of K such that each cell is in at most one pair of V, and such that if $\{\alpha^{(p)} < \beta^{(p+1)}\}$ is in V then α is a regular face of β.

We picture such vector fields by drawing, for each pair $\{\alpha^{(p)} < \beta^{(p+1)}\} \in V$, an arrow whose tail lies in α and whose head lies in β (Figure 3). Such pairings were studied in the case of a simplicial complex in [**85**] and [**27**] as a tool for investigating the possible f-vectors for a such complexes. Here we take a different point of view. Our first step is to introduce a special class of vector fields which will play the role of gradient vector fields.

Definition 10.

(1) Given a discrete vector field V on a cell complex K, a *V-path* is a sequence of cells $\alpha_0, \alpha_1, \alpha_2, \ldots, \alpha_r$ such that for each $i = 1, 2, \ldots, r$, either $\{\alpha_{i-1} < \alpha_i\} \in V$ or α_i is a codimesion-one face of α_{i-1} and $\{\alpha_i < \alpha_{i-1}\} \notin V$ (Figure 4). We say such a path is a *non-trivial closed path* if $r > 0$ and $\alpha_0 = \alpha_r$.

(2) A discrete vector field V is a *gradient vector field* if there are no non-trivial closed V-paths.

(3) If V is a gradient vector field on a cell complex K and α is a cell of K which is not contained in any pair in V, then we say that α is a *critical cell of V*.

The main theorem of discrete Morse theory is the following.

Theorem 11. *Let K be a CW complex with a discrete gradient field V. Then K is (simple-)homotopy equivalent to a CW complex with precisely one cell of dimension p for each critical cell of V of dimension p.*

Before presenting the very simple proof, we will recall the notion of simple-homotopy. This idea was introduced by J.H.C. Whitehead in an effort to establish a combinatorial basis for homotopy theory. Let K be a CW complex.

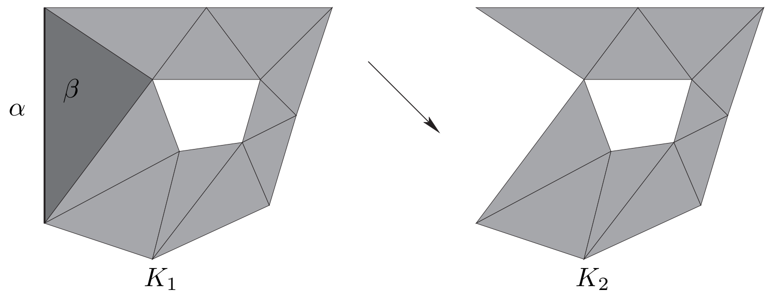

Figure 5. An elementary collapse.

Definition 12. Let β be a $(p+1)$-cell of K, and α a regular face of β. We say that α is a *free face* of β if α is not the face of any other cell of K. (This implies that β is maximal, i.e. is not the face of any cell in K, and that $\dim(\alpha) = p$.)

If α is a free face of β then $K - (int(\alpha) \cup int(\beta))$ is a deformation retract of K. Such a deformation retract is called an *elementary collapse* (and in the category of simplicial complexes, an *elementary simplicial collapse*). See Figure 5. Simple-homotopy is the equivalence relation generated by elementary collapse.

We are now ready to present the proof of Theorem 11. (Many essentially equivalent proofs have appeared since the original proof in [**32**]. Here we present the very short proof that can be found in [**59**].)

Proof. Since V has no closed paths, we can find a cell α of K which has no predecessors, i.e. such that there is no cell β such that β, α is a V-path. There are two possibilities, either (i) α is a maximal face, and is critical for V, or (ii) α is a free face of a cell β, and $\{\alpha < \beta\} \in V$, see Figure 6. In case (i), K is the result of attaching the cell α to $K' = K - int(\alpha)$. In case (ii), K collapses onto $K' = K - (int(\alpha) \cup int(\beta))$. The proof now follows by induction. □

Combining Theorems 11, 6, and 7, and the fact that homotopy equivalent spaces have isomorphic homology, yields the following theorem.

Theorem 13. *Let K be a simplicial complex with a discrete gradient vector field. Let m_p denote the number of critical simplices of dimension p. Let $\mathbb{F}$ be any field, and $b_p = \dim H_p(K, \mathbb{F})$ the p^{th} Betti number with respect to $\mathbb{F}$. Then we have the following relationships.*

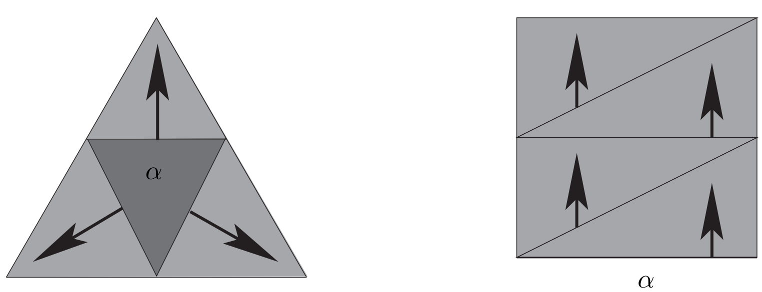

Figure 6. Two possibilities for a cell α with no predecessors.

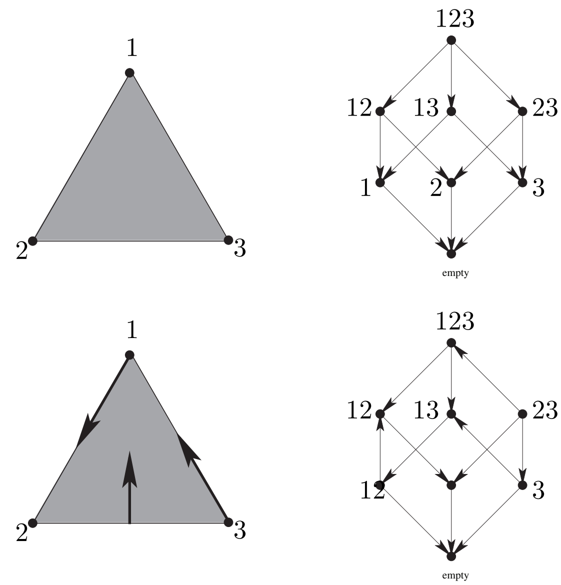

Figure 7. A 2-simplex and its Hasse diagram as a directed graph. A discrete vector field defines a modified Hasse diagram.

I. The Weak Morse Inequalities.

(i) For each $p = 0, 1, 2, \ldots, n$ *(where* n *is the dimension of* K*)*

$$m_p \geq b_p.$$

(ii) $m_0 - m_1 + m_2 - \cdots + (-1)^n m_n = b_0 - b_1 + b_2 - \cdots + (-1)^n b_n \ \ [= \chi(K)].$

II. The Strong Morse Inequalities.

For each $p = 0, 1, 2, \ldots, n, n+1,$

$$m_p - m_{p-1} + \cdots \pm m_0 \geq b_p - b_{p-1} + \cdots \pm b_0.$$

4. A More Combinatorial Language

The notion of a gradient vector field has a very nice purely combinatorial description due to Chari [**14**], with which we can recast the Morse theory in an appealing form. Let K be a regular CW complex. The *Hasse diagram* of K is defined to be the partially ordered set of cells of K ordered by the face relation. Consider the Hasse diagram as a directed graph, directed downward. That is, the vertices of the graph are in 1-1 correspondence with the cells of K, and there is a directed edge from β to α if and only if α is a codimension-one face of β. Now let V be a combinatorial vector field. We modify the directed graph as follows. If $\{\alpha < \beta\} \in V$ then reverse the orientation of the edge between α and β, so that it now goes from α to β. A V-path is precisely a directed path in this modified graph.

Thus, in this combinatorial language, a discrete vector field is a partial matching of the Hasse diagram, and a discrete vector field is a gradient vector field if the partial matching is acyclic, in the sense that the resulting directed graph has no directed loops.

When using this language, there is one possible minor source of confusion. When working with a simplicial complex, one usually includes the empty set as an

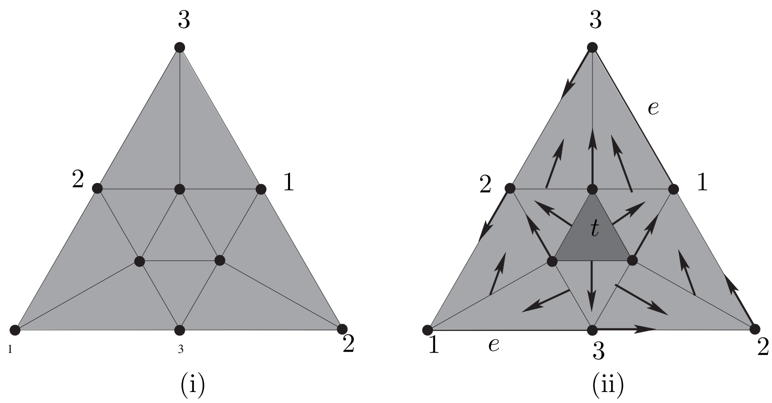

Figure 8. (i) A triangulation of the projective plane. (ii) A discrete vector field on the projective plane.

element of the Hasse diagram (considered as a simplex of dimension -1), while we have not considered the empty set previously. This issue will appear repeatedly in these lectures.

5. Our First Example: The Real Projective Plane

Figure 8(i) shows a triangulation of the real projective plane $\mathbb{P}^2$. The vertices along the boundary with the same labels are to be identified, as are the edges whose endpoints have the same labels. In Figure 8(ii) we illustrate a discrete vector field V on this simplicial complex. One can easily see that there are no closed V-paths (since all V-paths go to the boundary of the figure and there are no closed V-paths on the boundary), and hence is a gradient vector field. The only cells which are neither the head nor the tail of an arrow are the vertex label 1, the edge e, and the triangle t. Thus, by Theorem 11, the projective plane is homotopy equivalent to a CW complex with exactly one 0-cell, one 1-cell and one 2-cell. (Of course, we already knew this from our discussion of Example 3 in Section 2.)

This example gives rise to two potential concerns. The first is that from the main theorem we learn only a statement about "homotopy equivalence". This is sufficient if one is only interested in calculating homology or homotopy groups. However, one might be interested in determining the (PL-)homeomorphism type of the complex. This is possible, in some cases, using deep results of J. H. C. Whitehead. We revisit this topic briefly in the next section.

The second potential point of concern is that as we saw in Section 2 there are an infinite number of different homotopy types of CW complexes which can be built from exactly one 0-cell, one 1-cell and one 2-cell. One might wonder if Morse theory can give us any additional information as to how the cells are attached. In fact, one can deduce much of this information if one has enough information about the gradient paths of the gradient vector field. This point is discussed further in Section 3 of Lecture 2, where we will return to this example of the triangulated projective plane.

6. Sphere Theorems

As mentioned in our discussion at the end of Section 5, one can sometimes use discrete Morse theory to make statements about more than just the homotopy type of the simplicial complex. One can sometimes classify the complex up to homeomorphism or combinatorial equivalence. In this section we give some examples of such arguments. An interesting application of these ideas is presented in the next section. So far, we have not placed any restrictions on the simplicial complexes under consideration. The main idea of this section is that if our simplicial complex has some additional structure, then one may be able to strengthen the conclusion. This idea rests on some very deep work of J. H. C. Whitehead [**95**].

A simplicial complex K is a *combinatorial d-ball* if K and the standard d-simplex σ_d have isomorphic subdivisions. A simplicial complex K is a *combinatorial $(d-1)$-sphere* if K and $\dot{\sigma}_d$ have isomorphic subdivisions (where $\dot{\sigma}_d$ denotes the boundary of σ_d with its induced simplicial structure). A simplicial complex K is a *combinatorial d-manifold with boundary* if the link of every vertex is either a combinatorial $(d-1)$-sphere or a combinatorial $(d-1)$-ball. The following is a special case of the powerful main theorem of [**95**].

Theorem 14. *Let K be a combinatorial d-manifold with boundary which simplicially collapses to a vertex. (That is, K can be a reduced to a vertex by a sequence of elementary simplicial collapses.) Then K is a combinatorial d-ball.*

With this theorem, and its generalizations, one can sometimes strengthen the conclusion of Theorem 11 beyond homotopy equivalence. We present just one example.

Theorem 15. *Let X be a combinatorial d-manifold with a discrete gradient vector field with exactly two critical simplices. Then X is a combinatorial d-sphere.*

The proof is quite simple (given Theorem 14). The statement is trivial for $d = 0$, so we assume that $d \geq 1$. Suppose that X is a combinatorial d-manifold with a discrete gradient vector field V with exactly two critical simplices. Let x_0 be a vertex of X. If x_0 is not critical, then $\{x_0 < e\}$ is an element of V, for some edge e. Let x_1 be the other endpoint of e. Then x_0, e, x_1 is a V-path. If x_1 is not critical, we can follow the V-path to the next vertex x_2, etc. Since there are only a finite number of vertices, and there are no loops, we must eventually reach a critical vertex. We can run this argument in reverse for d-simplices. That is, if α_0 is a d-simplex, and α_0 is not critical, then $\{\beta < \alpha_0\}$ is an element of V for some $(d-1)$-simplex β. Let α_1 denote the other d-simplex incident to β. Then $\alpha_1, \beta, \alpha_0$ is a V-path, and we can follow this path backwards until reaching a critical d-simplex. Thus, there must be precisely one critical vertex x, and one critical d-simplex α. Then $X - \alpha$ is a combinatorial d-manifold with boundary with a discrete gradient vector field with only a single critical simplex, namely the vertex x. It follows that $X - \alpha$ collapses to x. Whitehead's theorem now implies that $X - \alpha$ is a combinatorial d-ball, which implies that X is a combinatorial d-sphere.

7. Our Second Example

In this section we demonstrate some of the ideas of the previous sections with a simple example from algebra. Fix a positive integer n, and consider the following

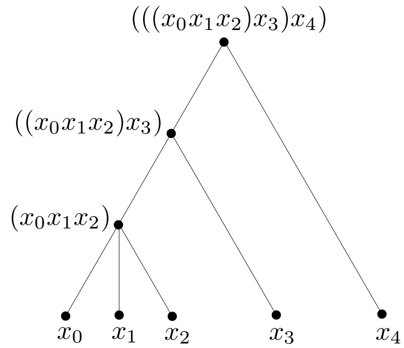

Figure 9. The planar rooted tree corresponding to $(((x_0x_1x_2)x_3)x_4)$.

$(n-2)$-dimensional simplicial complex, which we denote M_n. Starting with the following expression

$$(x_0x_1x_2\dots x_n)$$

consider all ways of adding legal pairs of parentheses. An expression resulting from adding $p+1$ pairs of parentheses will be a p-simplex in our complex. The faces of this p-simplex are all expressions that result from removing corresponding pairs of parentheses.

For example, consider the case $n=3$. The vertices of M_3 are the expressions

$$v_1 = ((x_0x_1)x_2x_3), \qquad v_2 = ((x_0x_1x_2)x_3), \qquad v_3 = (x_0(x_1x_2)x_3),$$
$$v_4 = (x_0(x_1x_2x_3)), \qquad v_5 = (x_0x_1(x_2x_3)),$$

and the edges are the expressions

$$e_1 = (((x_0x_1)x_2)x_3), \qquad e_2 = ((x_0(x_1x_2))x_3), \qquad e_3 = (x_0((x_1x_2)x_3)),$$
$$e_4 = (x_0(x_1(x_2x_3))), \qquad e_5 = ((x_0x_1)(x_2x_3)).$$

One can easily check the relations

$$e_1 = \{v_1, v_2\}, \qquad e_2 = \{v_2, v_3\}, \qquad e_3 = \{v_3, v_4\},$$
$$e_4 = \{v_4, v_5\}, \qquad e_5 = \{v_5, v_1\},$$

so that M_3 is a circle triangulated with 5 edges and 5 vertices.

These complexes arise in a number of different settings. For example, they arise in the study of planar rooted trees. To illustrate by an example, the edge $(((x_0x_1x_2)x_3)x_4)$ of M_4 can naturally be associated with the planar rooted tree shown in Figure 9. From this point of view, the top dimensional simplices correspond to binary trees. (See [**10**] and the references therein for an extensive discussion of such issues.) Moreover, the complexes M_n arise in geometry, as they are closely related to the simplicial complex of subdivisions of an $(n+1)$-gon into subpolygons (see, e.g. [**60**]). In the study of homotopy associative algebras ([**86**], [**87**]) one studies an algebra which is associative only up to homotopy. In that case, M_2, for example, arises from studying all ways of multiplying 3 elements, with $(x_0x_1x_2)$ representing a homotopy between $((x_0x_1)x_2)$ and $(x_0(x_1x_2))$. Note that here we see a slight difference. From this point of view, one would like to think of $((x_0x_1)x_2)$ and $(x_0(x_1x_2))$ as vertices, and $(x_0x_1x_2)$ as an edge between them. Thus, in this context, one is essentially working with the dual of the complex we have defined. We will say more about this a bit later (see the remarks following Theorem 17).

The main goal of this section is to use discrete Morse theory to give a simple proof of the following result.

Theorem 16. *The complex M_n is homotopy equivalent to an $(n-2)$-sphere.*

This result is well known, and it is only our proof that is new. We will prove this theorem by showing that one can easily construct a discrete gradient vector field on M_n which has precisely two critical simplices, namely one critical vertex and one critical $(n-2)$-simplex. The theorem then follows from Theorem 11. In fact, one can deduce more. We saw above that M_3 is not just a homotopy circle, but rather it is an actual combinatorial circle. One can easily see that the link of every vertex of M_n is isomorphic to a complex of the form $M_p * M_{n-p}$ (where $*$ denotes join). By induction, M_p and M_{n-p} are combinatorial spheres of dimension $p-2$ and $n-p-2$, respectively, so the link is a combinatorial sphere of dimension $n-3$ (see Proposition II.1 of [**45**]). Since the link of every vertex of M_n is a combinatorial $(n-3)$-sphere, it follows that M_n is a combinatorial $(n-2)$-manifold (see page 19 of [**45**]). Therefore we can apply Theorem 6 to learn the following stronger result.

Theorem 17. *The complex M_n is a combinatorial $(n-2)$-sphere.*

Before beginning our proof, we return to our earlier comments about the complex arising in the study of homotopy associative algebras. As remarked above, in that case one considers what is essentially the dual of the complex M_n. However, there is a slight modification. Let M_n^* denote a combinatorial $(n-2)$-sphere endowed with the cell decomposition which is dual to that of M_n. In M_n the trivial expression $(x_0x_1\dots x_n)$ corresponds to a simplex of dimension -1, i.e. the empty set. In the dual setting, $(x_0x_1,\dots x_n)$ corresponds to a cell of dimension $n-1$, whose boundary sphere is identified with all of M_n^*. Adding in this cell to form the cone on M_n^* results in a complex, introduced in [**86**] (see also [**87**]) called the associahedron (or Stasheff polytope), and which is often denoted A_{n+1}. Thus we learn

Corollary 18. *The associahedron A_{n+1} is a combinatorial* $(n-1)$*-ball.*

A proof of this appears in [**86**], by very different methods, and numerous alternative proofs have also been presented. In fact, A_{n+1} is a polytope ([**60**]). For more about the associahedron, from many points of view, one should certainly consult Fomin and Reading's wonderful lecture notes in this volume [**30**].

Let us now describe the construction of the desired gradient vector field V on M_n. Let s be a simplex of M_n. Suppose that there is not a pair of parenthesis around x_0 and x_1. If it is possible to legally add a pair of parentheses around x_0 and x_1 do so and call the resulting simplex t. We then add the pair $\{s \prec t\}$ to V. For example, in M_4 the expression $((x_0x_1x_2)(x_3x_4))$ is paired with $(((x_0x_1)x_2)(x_3x_4))$. After this step, the expressions which have not been paired with any other expression are those that have at least one parenthesis between x_0 and x_1, and it is simple to see that any such parenthesis must be a left parenthesis. There is one additional unpaired expression, namely the expression $s^* = ((x_0x_1)x_2x_3\dots x_n)$. According to our rule, this should be paired with the original expression $(x_0,x_1\dots x_n)$ with no added parentheses, but this is not permitted.

If s is any expression other than s^* that is currently unpaired, and a pair of parentheses can legally be added around the elements x_1 and x_2, do so and call the resulting simplex t. We then add the pair $\{s \prec t\}$ to V. After this step, the expressions which have not been paired with any other expression are s^* and those that have at least one left parenthesis between x_0 and x_1, and at least one left parenthesis between x_1 and x_2. Pair such an expression with the one resulting from adding a pair of parentheses around x_2 and x_3 if possible. Continue this process as long as possible. When it has terminated, the only expressions that have not been paired up with any other expression are s^* and the one that has a left parenthesis between every consecutive pair x_1 and x_{i+1} for $i = 0, 1, \dots, n-1$, i.e. the expression $t^* = (x_0(x_1(x_2(\dots(x_{n-2}(x_{n-1}x_n)))) \dots)$. Note that t^* is an $(n-2)$-simplex of the complex M_n.

This completes our construction of the vector field V. All that needs to be checked is there are no closed V-paths. Denote by V_k the discrete vector field that has been constructed after the k^{th} step in the construction, i.e. after consideration of the pair x_{k-1}, x_k. It is simple to check that V_1 has no closed orbits. Let $s_0^{(p)}, t_0^{(p+1)}, s_1^{(p)}$ denote a V-path. This requires that s_0 and t_0 be paired in V. Suppose that s_0 and t_0 are paired in V_k The reader can check that this implies that either s_1 is the head of an arrow in V_k (and hence the V-path cannot be continued) or s_1 is paired in V_{k-1}. Thus, by induction, there can be no closed V-paths.

8. Exercises for Lecture 1

(1) (a) Prove the strong Morse inequalities. That is, suppose that

$$V:\ 0 \to V_n \overset{\partial_n}{\to} V_{n-1} \overset{\partial_{n-1}}{\to} V_{n-2} \overset{\partial_{n-2}}{\to} \cdots \overset{\partial_1}{\to} V_0 \to 0$$

is a differential complex (i.e. $\partial_{i+1} \circ \partial_i = 0$ for all i). Let m_i denote the dimension of V_i, and b_i the dimension of the i^{th} homology ($=\mathrm{Ker}(\partial_i)/\mathrm{Im}(\partial_{i+1})$). Prove that for each i

$$m_i - m_{i-1} + m_{i-2} - \cdots \pm m_0 \geq b_i - b_{i-1} + b_{i-2} - \cdots \pm b_0.$$

Make sure you see how these inequalities imply the Weak Morse Inequalities.

(b) Now prove the converse of the Morse inequalities. That is, suppose that we are given finite lists of nonnegative integers $m_0, \dots, m_n$, and $b_0, \dots, b_n$ which satisfy the above inequality for each i. Prove that there is a complex V as above with $m_i = \dim(V_i)$ for each i, and such that b_i is the dimension of the i^{th} homology. [This shows that one cannot deduce anything stronger than the strong Morse inequalities using only the abstract existence of a complex which calculates the desired homology.]

(2) Prove that every triangulated disc is collapsible (i.e. collapses to a vertex).

(3) Triangulate a torus (more precisely, construct a simplicial complex which is homeomorphic to the torus) and find a discrete gradient field on the resulting simplicial complex with as few critical simplices as possible.

(4) Prove that every triangulated surface has a perfect gradient vector field. That is, let M be a connected simplicial complex which is homeomorphic to a compact surface. Prove that there is a gradient vector field on M with precisely 1 critical vertex, 1 critical 2-simplex, and g critical edges, where g denotes the genus of M. (Hint: Use the Morse inequalities to see that it is sufficient to find a discrete gradient vector field with exactly one critical vertex, and exactly one critical triangle.)

(5) One can also present discrete Morse theory using the language of functions, rather than gradient vector fields. Let K be a finite simplicial complex.

A function $f : K \to \mathbb{R}$ (i.e. f assigns a single real number to each simplex) is called a *discrete Morse function* if for each p-simplex α

$$\#\{\beta^{(p+1)} > \alpha \ s.t. \ f(\beta) \leq f(\alpha)\} \leq 1$$

and

$$\#\{\gamma^{(p-1)} < \alpha \ s.t. \ f(\gamma) \geq f(\alpha)\} \leq 1.$$

Given such a function f, define a set of pairs V_f by declaring that $\{\alpha < \beta\} \in V_f$ if α is a codimension-one face of β and $f(\beta) \leq f(\alpha)$.

(a) Show that V_f is actually a discrete vector field (i.e. that each simplex is contained in at most one pair in V).
(b) Show that V_f is a gradient vector field.
(c) Show that every gradient vector field arises in this way. That is, if V is a gradient vector field, then there is a discrete Morse function f such that $V = V_f$.

LECTURE 2
Discrete Morse Theory, continued

1. Suspensions and Discrete Morse Theory

Let K be a simplicial complex, and let x and y be two points not in K. Then the *suspension* of K is defined to be the join of K and the set $\{x, y\}$. More geometrically, embed K in some $\mathbb{R}^d$, and embed $\mathbb{R}^d$ in $\mathbb{R}^{d+1}$ by adding a final coordinate. Let x be the point $(0, \ldots, 0, 1)$ and y the point $(0, \ldots, 0, -1)$. Then the suspension of K is the union of all of the closed line segments connecting x to a point in K and all of the closed line segments connecting the point y to a point in K. This space comes with a natural simplicial decomposition induced from that of K.

Let S be a simplex, and M a nonempty proper subcomplex of S. There are two interesting topological spaces to consider in this setting. One is M itself, and the other is S/M, the result of identifying all of the points in M to a single point. While S/M is not a simplicial complex, it does have a canonical cell decomposition giving S/M the structure of a CW complex. Moreover, if $\alpha < \beta$ are two faces of S which are not in M, and α^* and β^* are their images in S/M, then $\alpha^* < \beta^*$, and moreover, α^* is a regular face of β^*.

In fact, the two spaces M and S/M are closely related, and one can deduce essentially the entire topological structure of either one from a knowledge of the other. More precisely, we have the following statement.

Theorem 19. *S/M is homotopy equivalent to the suspension of M.*

Of particular interest to us is the following result.

Corollary 20. *For any p, $\widetilde{H}_{p+1}(S/M, \mathbb{Z}) \cong \widetilde{H}_p(M, \mathbb{Z})$.*

These results are not hard to prove using standard methods, but we present a discrete Morse theory proof of Corollary 20, as the technique (more than the result) will prove useful later (see the next section). In fact, a more careful analysis of this proof allows one to deduce Theorem 19, but we will leave that to the reader. Our approach is to simultaneously construct gradient vector fields U and V on M and S/M, respectively. Let v be any vertex of M. If α is a nonempty simplex of M which does not contain v and which has the property that $v * \alpha$ is also in M, then pair α with $v * \alpha$. Let U_1 denote this collection of pairs. That is

$$U_1 = \{\{\alpha < v * \alpha\} \text{ s.t. } \emptyset \neq \alpha \text{ and } v * \alpha \subset M\}.$$

It is a simple observation that U_1 is a gradient vector field. Similarly, define a gradient vector field V_1 on S/M by setting

$$V_1 = \{\{\alpha < v * \alpha\} \text{ s.t. } \emptyset \neq \alpha \text{ and } \alpha \not\subseteq M\}.$$

(We are now identifying a simplex α in S, $\alpha \not\subseteq M$, with its image in S/M.) The simplices of M which are critical for U_1 are the vertex v and any nonempty simplex α of M with the property that $v * \alpha \not\subseteq M$. Let C_U denote this collection of critical simplices of U_1. The cells of S/M which are critical for V_1 are the special 0-cell m in S/M resulting from identifying all of the points in M, along with any nonempty simplex $\beta \not\subseteq M$ which has the property that $v \in \beta$, and $\beta - v \subset M$. Let C_V denote this collection of critical simplices of V_1. We observe that there is a canonical identification of the elements of C_U with those of C_V. Namely, identify $v \in M$ with $m \in S/M$, and identify α with $v * \alpha$ whenever $\alpha \subset M$ and $v * \alpha \not\subseteq M$. Let U denote any vector field on M which is an extension of U_1, and let $U_2 = U - U_1$ (so that U_2 consists of pairs of elements in C_U). Define $V_2 = \{\{v * \alpha < v * \beta\}$ *s.t.* $\{\alpha < \beta\} \in U_2\}$, and let $V = V_1 \cup V_2$.

Lemma 21.

(i) Let $A = \alpha_1, \alpha_2, \alpha_3, \ldots, \alpha_k$ be a sequence of elements in C_U, and let $B = v * \alpha_1, v * \alpha_2, v * \alpha_3, \ldots, v * \alpha_k$ be the corresponding sequence of elements in C_V. Then A is a U-path if and only if B is a V-path

(ii) V is a gradient vector field if and only if U is.

Proof. Part (i) follows immediately from the construction of V. To prove part (ii) let $M' = M - C_U$, and $S' = S/M - C_V$. (These are the cells that are paired in U_1 and V_1, respectively.) It is easy to see that any U-path that begins in M' stays in M', and hence is a U_1-path. Since U_1 is a gradient vector field, none of these U-paths are closed. Similarly, any V-path that begins in S' stays in S', and none of these are closed. Hence any closed U-path must lie entirely in C_U. Now the result follows from part (i). □

If U is a vector field on M which contains U_1, we say that U *collapses towards* v. If V is a vector field on S/M which contains V_1, then we say that V *collapses towards* v^*. Then Lemma 21 leads to the following result.

Theorem 22. *For any vertex v of M, there is a canonical identification of gradient vector fields of M which collapse towards v, and those of S/M which collapse towards v^*. If U is a gradient vector field on M which collapses towards v, and V is the corresponding gradient vector field on S/M, then v is critical for U, and m is critical for V. For every additional critical simplex α of U, $v * \alpha$ is critical for V, and every critical cell of V arises in this manner.*

In Section 3 we will introduce the Morse complex, a method of calculating the homology of a cell complex exactly using a knowledge of the critical cells and the gradient paths. The preceeding discussion is sufficient, modulo some minor details which can be supplied by the reader, to deduce that the Morse complex for the relative pair (M, v), which computes the reduced homology of M, is isomorphic to the Morse complex of the relative pair $(S/M, m)$, which computes the reduced homology of S/M, with the isomporphism shifting all degrees up by 1. This suffices to prove Theorem 20. A more careful consideration of the implications of Theorem 22 yields Theorem 19.

2. Monotone Graph Properties

A number of fascinating simplicial complexes arise from the study of monotone graph properties. Let K_n denote the complete graph on n vertices, and suppose we have label the vertices 1,2,...,n. Let $\mathcal{G}_n$ denote the set of spanning subgraphs of K_n, that is, the subgraphs of K_n that contain all n vertices. (Elements of $\mathcal{G}_n$ are permitted to be disconnected and to have isolated vertices.) A subset $\mathcal{P} \subset \mathcal{G}_n$ is called a *graph property* of graphs with n vertices if inclusion in $\mathcal{P}$ only depends on the isomorphism type of the graph. That is, $\mathcal{P}$ is a graph property if for all pairs of graphs $G_1, G_2 \in \mathcal{G}_n$, if G_1 and G_2 are isomorphic (ignoring the labelings on the vertices) then $G_1 \in \mathcal{P}$ if and only if $G_2 \in \mathcal{P}$. A graph property $\mathcal{P}$ of graphs with n vertices is said to be *monotone decreasing* if for any graphs $G_1 \subset G_2 \in \mathcal{G}_n$, if $G_2 \in \mathcal{P}$ then $G_1 \in \mathcal{P}$.

Monotone decreasing properties abound in the study of graph theory. Here are some typical examples: graphs having no more than k edges (for any fixed k), graphs such that the degree of every vertex is less that δ (for any fixed δ), graphs which are not connected, graphs which are not i-connected (for any fixed i), graphs which do not have a Hamiltonian cycle, graphs which do not contain a minor isomorphic to H (for any fixed graph H), graphs which are r-colorable (for any fixed r), and bipartite graphs.

Any monotone decreasing graph property $\mathcal{P}$ gives rise to a simplicial complex $\mathcal{K}$ where the d-simplices of $\mathcal{K}$ are the graphs $G \in \mathcal{P}$ which have $d+1$ edges. In particular, if G is a d-simplex in $\mathcal{K}$, then the faces of G are all of the nontrivial spanning subgraphs of G (the monotonicity of $\mathcal{P}$ implies that each of these graphs is in $\mathcal{K}$). Said in another way, if $\mathcal{P}$ is nonempty, then the vertices of $\mathcal{K}$ are the edges of K_n (more precisely, the spanning subgraphs of K_n which include all n vertices and precisely one edge), and a collection of vertices in $\mathcal{K}$ span a simplex if the spanning subgraph of K_n consisting of all edges which correspond to these vertices lies in $\mathcal{P}$.

The simplicial complexes induced by many of the above-mentioned monotone decreasing graph properties have been studied using the techniques of these notes. See for example [**14**], [**25**], [**52**], [**53**], [**65**], and [**79**]. These papers contain some beautiful mathematics in which the authors construct "by hand" explicit discrete gradient vector fields, along the way illuminating some of the intricate finer structures of the graph properties.

Some monotone graph properties have recently been the focus of intense interest because of their relation to knot theory. Unfortunately this is probably not a good time for an in depth discussion of this fascinating topic. We will mention only that Vassiliev has shown how one can derive finite type knot invariants from the study of the space of "singular knots" (i.e. maps from S^1 to $\mathbb{R}^3$ which are not embeddings). The homology of the simplicial complexes of disconnected and not-2-connected graphs show up in his spectral sequence calculation of the homology of this space. This is explained in [**93**], where Vassiliev derives the homotopy type of the complex of disconnected graphs. In [**92**] and [**6**], the topology of the space of

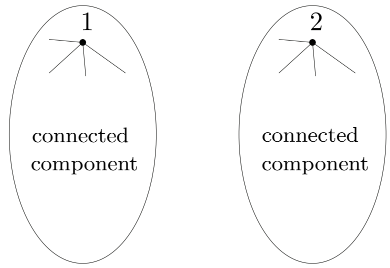

Figure 10. Graphs which are critical for V_{12} have two components.

not-2-connected graphs is determined, with discrete Morse Theory playing a minor role in the latter reference. This topic is reexamined in [**79**], in which the entire investigation is framed in the language of discrete Morse theory. We examine this topic in Section 2.2. Discrete Morse theory is used to determine the topology of not-3-connected graphs in [**52**].

2.1. The Complex of Disconnected Graphs

In this section, we will provide an introduction to this work by taking a look at the simpler case of the complex of disconnected graphs. We will show how the ideas of these lectures may be used to determine the topology of Δ_n, the simplicial compex of disconnected graphs on n vertices. Let me begin by pointing out that this complex can be well studied by more classical methods, and the answer has also been found by Vassiliev in [**93**]. The only novelty of this section is our use of discrete Morse theory.

Our goal is to construct a discrete gradient vector field V on Δ_n, the simplicial complex of all disconnected graphs with the vertex set $\{1, 2, 3, \ldots, n\}$. The construction will be in steps. Let V_{12} denote the discrete vector field consisting of all pairs $\{G, G + (1, 2)\}$, where G is any graph in Δ_n which does not contain the edge $(1, 2)$ and such that $G + (1, 2) \in \Delta_n$. Another way of describing V_{12} is that if G is any graph in Δ_n which contains the edge $(1, 2)$, then $G - (1, 2)$ and G are paired in V_{12}. Actually, there is one exception to this rule. Let g denote the graph consisting of only the single edge $(1, 2)$. Then $g - (1, 2)$ is the empty graph, which corresponds to the empty simplex in Δ_n, and may not be paired in a discrete vector field. Thus, g is unpaired in V_{12}.

The graphs in Δ_n other than g which are unpaired in V_{12} are those that do not contain the edge $(1, 2)$ and have the property that $G + (1, 2) \notin \Delta_n$. That is, those disconnected graphs G with the property that $G + (1, 2)$ is connected. Such a graph must have exactly two connected components, one of which contains the vertex labeled 1, and one which contains the vertex labeled 2. We denote these connected components by G_1 and G_2, resp. See Figure 10.

Let G be a graph other than g which is unpaired in V_{12}, and consider vertex 3. This vertex must either be in G_1 or G_2. Suppose that vertex 3 is in G_1. If G does not contain the edge $(1, 3)$ then $G + (1, 3)$ is also unpaired in V_{12}, so we can pair G with $G + (1, 3)$. If vertex 3 is in G_1, then the graph G is still unpaired if and only if G contains the edge (1,3) and $G - (1, 3)$ is the union of three connected components, one containing vertex 1, one containing vertex 2, and one containing vertex 3.

Similarly, if vertex 3 is in G_2 and G does not contain the edge $(2, 3)$, then pair G with $G + (2, 3)$. Let V_3 denote the resulting discrete vector field.

The unpaired graphs in V_3 are g and those that either contain the edge (1,3) and have the property that $G - (1, 3)$ is the union of three connected components, one containing vertex 1, one containing vertex 2, and one containing vertex 3, or contain the edge (2,3) and have the property that $G - (2, 3)$ is the union of three connected components, one containing vertex 1, one containing vertex 2, and one containing vertex 3. We illustrate these graphs in Figure 11. The circles in this figure indicate connected subgraphs.

Now consider the location of the vertex label 4, and pair any graph G which is unpaired in V_3 with $G + (1, 4)$, $G + (2, 4)$, or $G + (3, 4)$ if possible (at most one of these graphs is unpaired in V_3). Call the resulting discrete vector field V_4. We continue in this fashion, considering in turn the vertices label $5, 6, \ldots, n$. Let V_i denote the discrete vector field that has been constructed after the consideration of vertex i, and $V = V_n$ the final discrete vector field. When we are done the only unpaired graphs in V will be g and those graphs that are the union of two connected trees, one containing the vertex 1 and one containing the vertex 2. In addition, both trees have the property that the vertex labels are increasing along every ray starting from the vertex 1 or the vertex 2. There are precisely $(n-1)!$ such graphs, and they each have $n-2$ edges, and hence correspond to an $(n-3)$-simplex in Δ_n.

It remains to see that the discrete vector field V is a gradient vector field, i.e. that there are no closed V-paths. We first check that V_{12} is a gradient vector field. Let $\gamma = \alpha_0^{(p)}, \beta_0^{(p+1)}, \alpha_1^{(p)}$ denote a V_{12}-path. Then α_0 must be the "tail of an arrow", i.e. the smaller graph of some pair in V_{12}, with β_0 being the head of the arrow, i.e. $\beta_0 = \alpha_0 + (1, 2)$. The simplex α_1 is a codimension-one face of β_0 other than α_0. Thus, α_1 corresponds to a graph of the form $\alpha_0 + (1, 2) - e$, where e is an edge of α_0 other than (1,2). Since α_1 contains the edge $(1, 2)$ it is the "head of an arrow" in V_{12}, i.e. the larger graph of some pair in V_{12}, which implies that γ cannot be continued to a longer V_{12}-path. This certainly implies that there are no closed V_{12}-paths.

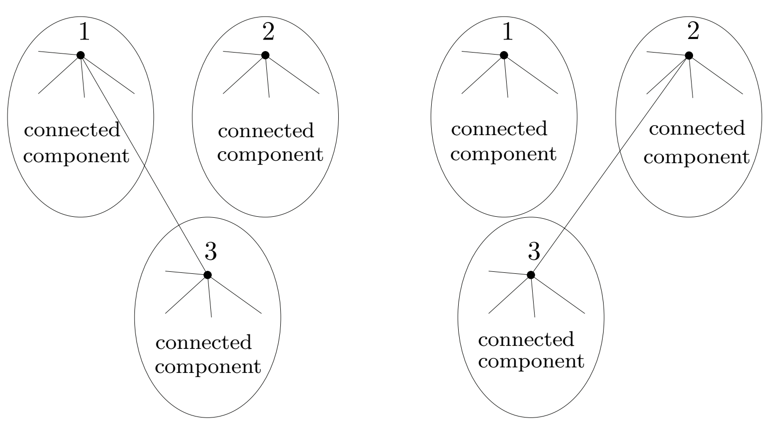

Figure 11. The two types of graphs which are critical for V_3.

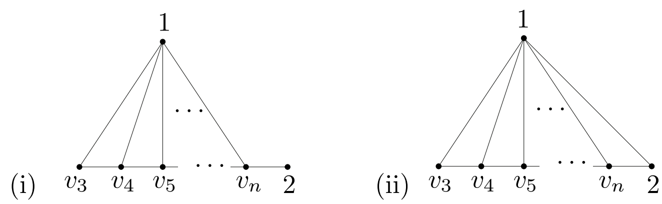

Figure 12. (i) Critical Graphs in Δ_n^2. (ii) Critical Graphs in N_n^2.

The same sort of argument will work for V. Recall that V is constructed in stages, by first considering the edge $(1,2)$ and then the vertices $3,4,5,\dots$ in order. Let $\gamma = \alpha_0, \beta_0, \alpha_1$ denote a V-path. In particular, α_0 and β_0 must be paired in V. The reader can check that if α_0 and β_0 are first paired in V_i, $i \geq 3$, then either α_1 is the head of an arrow in V_i, in which case the V-path cannot be continued, or α_1 is paired in V_{i-1}. It follows by induction that there can be no closed V-paths.

In summary, V is a discrete gradient vector field on N_n with exactly one unpaired vertex, and $(n-1)!$ unpaired $(n-3)$-simplices. We can now apply Theorem 11 to conclude

Theorem 23 ([**93**]). *The complex Δ_n of disconnected graphs on n vertices is homotopy equivalent to the wedge of $(n-1)!$ spheres of dimension $(n-3)$.*

2.2. Not-2-connected Graphs

Recall that a graph G is *2-connected* if the removal of any vertex (along with all incident edges) results in a connected graph. If G is not 2-connected, we call any vertex v a *cut vertex* if $G - v$ is not connected. Let Δ_n^2 denote the complex of not-2-connected graphs on the vertex set $\{1,2,\dots,n\}$.

In this section, we will describe a proof of the following result.

Theorem 24. *For $n \geq 3$, the space Δ_n^2 is homotopy equivalent to a wedge of $(n-2)!$ spheres of dimension $2n-5$.*

This result was first established in [**6**] and [**92**], but we will follow (with only cosmetic changes) the proof, via discrete Morse theory, presented in [**79**]. Let g denote the graph on the vertex set $\{1,2,\dots,n\}$ containing only the single edge $(1,2)$. Theorem 24 follows from the following result.

Proposition 25. *There is a discrete gradient vector field on Δ_n^2 whose critical simplices are g along with all graphs of the form shown in Figure 12(i), where $v_3, v_4, \dots, v_n$ is any permutation of $3,4,\dots,n$.*

Let $C_n^2 = \mathcal{G}_n/\Delta_n^2$. Then the cells of C_n^2, with the exception of the distinguished point, correspond to the 2-connected graphs, so we call C_n^2 the *complex of 2-connected graphs*. Our construction of the gradient vector field in Theorem 24 first begins by collapsing towards g. Hence, following the discussion in the previous section, Theorem 25 implies the following result.

Corollary 26. *There is a discrete gradient vector field on C_n^2 whose critical simplices are the special point p in N_n^2 corresponding to Δ_n^2, and all graphs of the form shown in Figure 12(ii).*

Proposition 25 and Corollary 26 will be proved simultaneously, inductively on n. For $n = 3$, the set $\mathcal{G}_3$ of graphs on the vertex set $\{1, 2, 3\}$, is a 2-dimensional simplex on the vertex set consisting of the 3 possible edges $\{(1, 2), (2, 3), (1, 3)\}$. The only graph on 3 vertices which is 2-connected is K_3, which corresponds to the maximal face of $\mathcal{G}_3$. That is, Δ_3^2 is a circle, and the gradient vector field which collapses towards $(1, 2)$, has critical vertex $\{[1, 2]\}$ and critical edge $\{(2, 3), (1, 3)\}$. The space C_n^2, resulting from collapsing Δ_3^2 to a point, is a 2-sphere consisting of the point p and the 2-cell K_3. The only possible gradient vector field in C_n^2 is empty so that both cells are critical. These gradient vector fields satisfy the conclusions of Proposition 25 and Corollary 26.

Now let us begin to construct a gradient vector field on Δ_n^2 for general n (assuming the construction of such a gradient vector field on C_{n-1}^2). First, we collapse towards g. That is, set

$$V_1 = \{\{G - (1, 2), G\}\}$$

where G ranges over all graphs which are not 2-connected and contain the edge $(1, 2)$. Let M_1 denote the graphs which remain unpaired. Then M_1 consists of all graphs G which are not 2-connected, and do not contain the edge $(1, 2)$, and have the property that $G + (1, 2)$ is 2-connected.

To describe the next step in our construction of V, we must take a closer look at such graphs. Such a graph G must be connected (as otherwise $G + (1, 2)$ cannot be 2-connected). Let us now recall the basic structure of connected, not-2-connected graphs. Let H be such a graph, and let H_1 be an induced 2-connected subgraph which is maximal among all induced 2-connected subgraphs. (A subgraph H of a graph G is said to be *induced* if H contains all edges of G which connect two vertices of H.) Let H_2 denote another maximal induced 2-connected subgraph. Then $H_1 \cap H_2$ can contain at most one vertex (as otherwise the induced graph on $V(H_1) \cup V(H_2)$ would be 2-connected, and larger than H_1 and H_2). If $H_1 \cap H_2$ contains a vertex, then that vertex must be a cut vertex of H. Conversely, any cut vertex of H is of the form $H_1 \cap H_2$ for some maximal induced 2-connected subgraphs H_1 and H_2. Now let $H(2)$ denote the graph whose vertices are the maximal induced 2-connected subgraphs of H, with the property that if H_1 and H_2 are maximal induced 2-connected subgraphs of H, then the corresponding vertices of $H(2)$ are adjacent if and only if $H_1 \cap H_2$ is not empty. Clearly every vertex of H is contained in some maximal induced 2-connected subgraph of H. Moreover, H is not 2-connected, which implies that $H(2)$ has at least 2 vertices. Lastly, we observe that every minimal loop in H is contained in some maximal induced 2-connected subgraph of H, and hence appears as a vertex in $H(2)$, from which one can deduce that $H(2)$ has no loops, i.e. $H(2)$ is a tree. Now let G be a connected, not-2-connected graph with the property that $G + (1, 2)$ is 2-connected. Note that vertices 1 and 2 cannot be contained in the same maximal induced 2-connected subgraph of G, as otherwise the blocks of $G + (1, 2)$ would be the same as those of G, and hence $G + (1, 2)$ would not be 2-connected. Let G_1 denote the maximal induced 2-connected subgraph of G that contains the vertex 1, and G_2 the maximal induced 2-connected subgraph of G that contains the vertex 2. It is easy to see that $G_1 \neq G_2$ (as otherwise $G + (1, 2)$ cannot be 2-connected). In fact, the following result is easily established.

Lemma 27. *With all notation as above, $G(2)$ is a path from G_1 to G_2.*

Now let G_3 denote the induced maximal 2-connected subgraph which is adjacent to G_1 in $G(2)$, and let $v(G)$ be the vertex $G_1 \cap G_3$. It is clear that $v(G) \neq 2$. Moreover, if $v(G) = 1$ then vertex 1 would be a cut vertex of $G+(1,2)$, contradicting the assumption that $G+(1,2)$ is 2-connected. Therefore $v(G) \notin \{1,2\}$. Suppose $G \in M_1$ and $(v(G),2) \notin G$. It is easy to see that $v(G)$ is a cut vertex of $G+(v(G),2)$, and hence $G+(v(G),2)$ is not 2-connected. Moreover, $[G+(v(G),2)]+(1,2) = [G+(1,2)]+(v(G),2)$ is 2-connected (since $[G+(1,2)]$ is), so $G+(v(G),2) \in M_1$. Now define

$$V_2 = \{\ \{G, G+(v(G),2)\}\}$$

where G ranges over all graphs in M_1 which do not contain $(v(G),2)$.

Let M_2 contain the graphs which are not paired in V_1 or V_2. Then M_2 consists of those graphs G in M_1 which contain $(v(G),2)$, and which have the property that $G-(v(G),2) \notin M_1$. First note that since $(v(G),2) \in G$, $v(G)$ and 2 are contained in an induced 2-connected subgraph, which implies that $v(G) \in G_2$, and hence G_1 and G_2 are connected in $G(2)$. From the previous lemma, we learn that $G(2)$ must consist of only the two vertices G_1 and G_2 and the edge between them. The only way $G-(v(G),2)$ could fail to be in M_1 is if $G-(v(G),2)$ failed to be connected. This can happen only if $G_2-(v(G),2)$ is not connected. However, since G_2 is 2-connected, this can happen only if G_2 consists entirely of the vertices 2 and $v(G)$ and the edge between them. Thus, the graphs G in M_2 are precisely those that can be constructed by taking a 2-connected graph G_1 on the vertex set $\{1,2,\ldots,n\}-\{2\}$, adding the vertex 2, and adding the edge $(i,2)$ for some $i \notin \{1,2\}$ (in which case $v(G) = i$).

Let $M_2(i), i = 3,4,\ldots,n$, denote those graphs G in M_2 with $v(G) = i$. Then M_2 is the disjoint union of the $M_2(i)$'s. Each $M_2(i)$ can be canonically identified with the complex Γ of 2-connected graphs on the $n-1$ vertices $\{1,3,4,\ldots,n\}$. By induction, there is a gradient vector field on Γ with precisely $(n-3)!$ critical simplices of dimension $2(n-1)-5. = 2n-7$. Using the identification, we get a gradient vector field $V_3(i)$ on $M_2(i)$ with $(n-3)!$ critical simplices of dimension $2n-6$. Let

$$V = V_1 \cup V_2 \cup (\cup_{i=3}^{n} V_3(i)).$$

Since there are $n-2$ such $M_2(i)$'s, the total number of unmatched simplices in V is $(n-2)(n-3)! = (n-2)!$, each of dimension $2n-6$. The theorem now follows once we know that V is a gradient vector field.

Lemma 28. *The vector field V constructed above is a gradient vector field.*

The proof is left as a (rather non-trivial) exercise.

It is, in fact, quite easy to identify more explicitly the critical simplices in the above gradient vector field. To find the critical graphs in $M_2(i)$, $i = 3,4,\ldots,n$, we take the critical graphs in the complex of 2-connected graphs on the vertex set $\{1,3,4,\ldots,n\}$ with respect to some optimal gradient vector field add the vertex 2 and the edge $(i,2)$ for some $i = 3,4,\ldots,n$. Fixing i, identify $\{1,3,4,\ldots,n\}$ with $\{1,2,\ldots,n-1\}$ via a correspondence that identifies 1 with 1, and identifies i with 2. By induction, there is a gradient vector field on the 2-connected graphs on the vertex set $\{1,2,\ldots,n-1\}$ whose critical simplices have the form shown in Figure 12(ii). Using the identification, we get a gradient vector field on 2-connected graphs on the

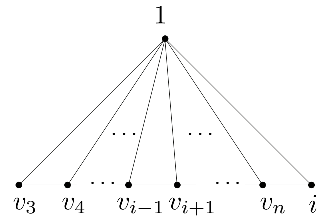

Figure 13. Critical 2-connected graphs on the vertex set $\{1, 2, 3, \ldots, n\}$.

vertex set $\{1, 3, 4, \ldots, n\}$ whose critical simplices are of the form shown in Figure 13 (where $v_3, v_4, \ldots, v_{i-1}, v_{i+1}, \ldots, v_n$ is any permutation of $3, 4, \ldots, i-1, i+1, \ldots, n$).

Adding a vertex 2 to each such graph, and adding an edge between vertex i and vertex 2 yields the desired collection of graphs shown in Figure 12(i). Corollary 26 now follows from Theorem 22.

2.3. Some further thoughts

The reader may wonder why we stopped with not-2-connected graphs. In fact, with quite a bit of hard work, it is possible to go further. In [**52**] J. Jonsson used discrete Morse theory to prove the following result.

Theorem 29. *The simplicial complex Δ_n^3 of not-3-connected graphs is homotopy equivalent to a wedge of $(n-3)\cdot(n-2)!/2$ spheres of dimension $(2n-1)$.*

Many of the gradient vector fields presented in these notes, including the two examples in this section, follow a similar pattern, in that one constructs the gradient vector field in several stages, following distinct rules for each stage. In this way, a user of discrete Morse theory generally discovers the following useful observation, which appeared implicitly earlier, but seems to have been first explicitly stated in [**52**] and [**50**].

Lemma 30. *Let $K = \sqcup_{i\in I} K_i$ be a partition of the faces of K, where I is some partially ordered set. Suppose that for every $i \in I$, $\cup_{j\le i} K_j$ is a subcomplex of K. Now suppose we have a discrete vector field V_i on each K_i (that is, a partial pairing of the simplices in K_i) with the property that there are no closed V_i-paths in K_i. Then $V = \cup_{i\in I} V_i$ is a gradient vector field on K.*

3. The Morse Complex

In this section we will see how more precise knowledge of the gradient vector field on a simplicial complex K allows one to strengthen the conclusions of the main theorems of discrete Morse theory. In particular, rather than just knowing the number of cells in a CW decomposition for K, one can calculate the homology exactly.

Let K be a simplicial complex with a gradient vector field V. In keeping with the standard terminology in the smooth category, we will refer to V-paths (see Section 3) as gradient paths. Let $C_p(K,\mathbb{Z})$ denote the space of simplicial p-chains, and $M_p \subseteq C_p(K,\mathbb{Z})$ the span of the critical p-simplices of V. We refer to M_* as the space of Morse chains. If we let m_p denote the number of critical p-simplices, then we obviously have

$$\mathcal{M}_p \cong \mathbb{Z}^{m_p}.$$

Since homotopy equivalent spaces have isomorphic homology, the following theorem follows from Theorems 11 and 5.

Theorem 31. *There are boundary maps* $\tilde{\partial}_d : \mathcal{M}_d \to \mathcal{M}_{d-1}$*, for each d, so that*

$$\tilde{\partial}_{d-1} \circ \tilde{\partial}_d = 0$$

and such that the resulting differential complex

$$0 \longrightarrow \mathcal{M}_n \xrightarrow{\tilde{\partial}_n} \dots \xrightarrow{\tilde{\partial}_1} \mathcal{M}_0 \longrightarrow 0 \tag{3}$$

calculates the homology of K. That is, if we define

$$H_d(\mathcal{M}, \tilde{\partial}) = \frac{\mathrm{Ker}(\tilde{\partial}_d)}{\mathrm{Im}(\tilde{\partial}_{d+1})}$$

then for each d

$$H_d(\mathcal{M}, \tilde{\partial}) \cong H_d(K, \mathbb{Z}).$$

In fact, this statement is equivalent to the Strong Morse inequalities (see Exercise 1 of Lecture 1). The main goal of this section is to present an explicit formula for the boundary operator $\tilde{\partial}$. This requires a closer look at the notion of a gradient path. Let β be a critical $(p+1)$ simplex, and and α a critical p-simplex. Then it is easy to check that any gradient-path from β to α has the form

$$\beta = \beta_0^{(p+1)}, \alpha_1^{(p)}, \beta_1^{(p+1)}, \alpha_2^{(p)}, \dots, \beta_r^{(p+1)}, \alpha_{r+1}^{(p)} = \alpha$$

such that for each $i = 0, 1, 2, \dots, r$, $\{\alpha_{i+1} < \beta_{i+1}\} \in V$, and $\alpha_{i+1} < \beta_i$, but $\{\alpha_{i+1} < \beta_i\} \notin V$. In Figure 14 we show a single gradient path from the boundary of a critical 2-simplex β to a critical edge α, where the arrows pointing from an edge to a 2-cell indicate the gradient vector field V.

Given a gradient path as shown in Figure 14, an orientation on β induces an orientation on α. We will not state the precise definition here. The idea is that one "slides" the orientation from β along the gradient path to α. For example, for the gradient path shown in Figure 14, the indicated orientation on β induces the indicated orientation on α.

We are now ready to state the desired formula.

Theorem 32. *Choose an orientation for each simplex. Then for any critical* $(p+1)$*-simplex* β *set*

$$\widetilde{\partial}\beta = \sum_{critical\ \alpha^{(p)}} c_{\alpha,\beta}\alpha \tag{4}$$

where

$$c_{\alpha,\beta} = \sum_{\gamma \in \Gamma(\beta,\alpha)} m(\gamma)$$

where $\Gamma(\beta, \alpha)$ *is the set of gradient paths which go from* β *to* α*. The multiplicity* $m(\gamma)$ *of any gradient path* γ *is equal to* ± 1*, depending on whether, given* γ*, the orientation on* β *induces the chosen orientation on* α*, or the opposite orientation. With this differential, the complex (3) computes the homology of K.*

We refer to the complex (3) with the differential (4) as the Morse complex (it goes by many different names in the literature). An extensive study of the Morse complex in the smooth category appears in [**78**]. In is section, we have focused our attention on simplicial complexes. However, it is worth noting that this entire

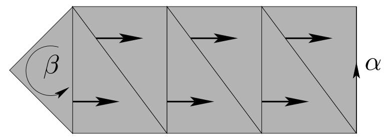

Figure 14. The flow of the edge e.

discussion applies, without any change, to any regular CW-complex, and, after some refinement of the notion of the multiplicity $m(\gamma)$, to all CW complexes. See [**32**] for details.

We only have time to present the main ideas the proof of Theorem 32. For the details, consult Sections 7 and 8 of [**32**]. The key ingredient in the proof is the notion of a (discrete time) flow associated to a discrete vector field V. In the case of smooth manifolds, the gradient vector field defines a dynamical system, namely the flow along the vector field. Viewing the Morse function from the point of view of this dynamical system leads to important new insights [**83**]. The same is true in the combinatorial category.

Up to this point in the notes, we have been thinking of V as a collection of pairs of simplices. Now it is better to think of V as a map of oriented simplices. Namely, choose an orientation for each simplex of M. If $\{\beta^{(p)} < \alpha^{(p+1)}\}$ is an element of V, then we set $V(\beta) = -i\alpha$ where $i = \pm 1$ is the incidence number of β and α (i.e $i = 1$ if the orientations agree, and -1 otherwise). Set $V(\beta^{(p)}) = 0$ if there is no such $\alpha^{(p+1)}$, i.e. if β is not the tail of any arrow in V. Now extend V linearly to a map

$$V : C_p(M, \mathbb{Z}) \to C_{p+1}(M, \mathbb{Z}),$$

and do this for each p.

The flow Φ along the gradient vector field V is a map

$$\Phi : C_p(M, \mathbb{Z}) \to C_p(M, \mathbb{Z}),$$

for each p, defined by the formula

$$\Phi \;=\; 1 + \partial V + V\partial.$$

See Figure 15 for the flow of an oriented edge e. In this figure, we indicate the orientation of e, and just enough of the vector field V in order to determine $\Phi(e)$. We observe that the map Φ commutes with the boundary operator. The other main fact is that for a finite simplicial complex, the map Φ stabilizes in finite time. That is, there is an N such that $\Phi^N = \Phi^{N+1} = \Phi^{N+2} = \ldots$ (it is only here that it is necessary that the vector field V be a gradient vector field), and we denote this map by Φ^∞.

Now let us return to the analysis of the Morse complex. Let

$$C_* : 0 \longrightarrow C_n(K, \mathbb{Z}) \xrightarrow{\partial_n} \ldots \xrightarrow{\partial_1} C_0(K, \mathbb{Z}) \longrightarrow 0$$

denote the usual simplicial chain complex of K. Let $C_p^\Phi(K, \mathbb{Z}) \subset C_p(K, \mathbb{Z})$ denote the subspace of Φ-invariant chains (i.e. the chains c such that $\Phi(c) = c$). Then, since Φ commutes with the boundary operator ∂, the boundary map takes $C_p^\Phi(K, \mathbb{Z})$ to $C_{p-1}^\Phi(K, \mathbb{Z})$. Now consider the complex of Φ-invariant chains.

$$C_*^\Phi : 0 \longrightarrow C_n^\Phi(K, \mathbb{Z}) \xrightarrow{\partial_n} \ldots \xrightarrow{\partial_1} C_0^\Phi(K, \mathbb{Z}) \longrightarrow 0.$$

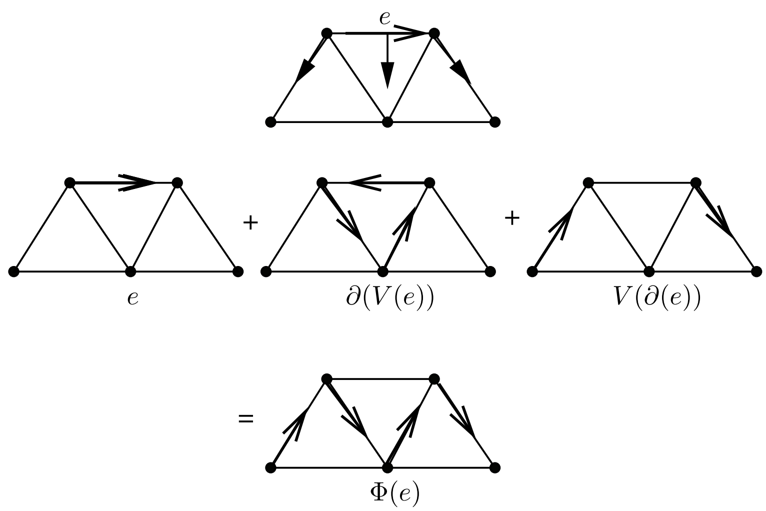

Figure 15. The flow of an oriented edge.

The first step is to see that this complex has the same homology as C_*. There are obvious maps between the two complexes, since C_*^Φ injects into C_*, and Φ^∞ maps C_* onto C_*^Φ. The composition yields the identity map on C_*^Φ. Thus, it is sufficient to show that the map $\Phi^\infty : C_* \to C_*$ induces an isomorphism on homology. For this, it is sufficient to find a homotopy operator. That is, an operator $L : C_*(K, \mathbb{F}) \to C_{*-1}(K, \mathbb{F})$ with the property that $\Phi^\infty - 1 = L\partial + \partial L$. If $\Phi^\infty = \Phi$, then one could take $L = V$. The general case of $\Phi^\infty = \Phi^N$ is similar.

To make the transition to critical simplices, one can establish that

$$\Phi^\infty : \mathcal{M}_p \to C_p(K, \mathbb{Z})$$

is an isomorphism for each p, with inverse the restriction map $r : C_p(K, \mathbb{Z}) \to \mathcal{M}_p$. Theorem 32 now follows if we take $\widetilde{\partial} = r \circ \partial \circ \Phi^\infty$. One must then calculate that this is precisely the operator defined in the statement of the theorem.

A different proof of Theorem 32 is suggested in the exercises.

Example 33. We end this section with a demonstration of how the ideas of this section may be applied to the example of the real projective plane $\mathbb{P}^2$ as illustrated in Figure 8(ii). We saw in Section 11 how discrete Morse Theory can help us see that $\mathbb{P}^2$ has a CW decomposition with exactly one 0-cell, one 1-cell and one 2-cell. Here we will see how Morse theory can distinguish between the spaces which have such a CW decomposition. Let us now calculate the boundary map in the Morse complex corresponding to the gradient vector field illustrated in Figure 8(ii). Choose an orientation for the edge e. To calculate $\widetilde{\partial}(e)$, we must count all of the gradient paths from e to v. There are precisely two such paths, since the unique gradient path beginning at each endpoint of e leads to v. (The gradient path beginning at vertex 1 is the trivial path of 0 steps.) Since the orientation of e induces a $+$ on one endpoint of e, and a $-$ orientation on the other, adding these two paths with their corresponding signs leads us to the formula that $\widetilde{\partial}(e) = 0$. Now choose an orientation for t. It can be seen from Figure 8(ii) that there are precisely two gradient paths from t to e, and both induce the same orientation on

e, so that $\widetilde{\partial}(t) = \pm 2e$. By reversing the chosen orientation on t if necessary, we may assume that $\tilde{\partial}(t) = 2e$. Therefore the homology of the real projective plane can be calculated from the following differential complex.

$$\mathbb{Z} \xrightarrow{\times 2} \mathbb{Z} \xrightarrow{0} \mathbb{Z} \longrightarrow 0.$$

Thus we see that

$$H_0(\mathbb{P}^2, \mathbb{Z}) \cong \mathbb{Z}, \qquad H_1(\mathbb{P}^2, \mathbb{Z}) \cong \mathbb{Z}/2\mathbb{Z}, \qquad H_2(\mathbb{P}^2, \mathbb{Z}) \cong 0.$$

4. Canceling Critical Points

One of the main problems in Morse theory, whether in the combinatorial or smooth setting, is to find a Morse function, or equivalently a gradient vector field, for a given space with the fewest possible critical points (much of the book [**80**] is devoted to this topic). In general this is a very difficult problem, since, in particular, it contains the Poincaré conjecture – spheres can be recognized as those spaces which have a Morse function with precisely 2 critical points. In [**72**], Milnor presents Smale's proof [**83**] of the higher dimensional Poincaré conjecture (in fact, a proof is presented of the more general h-cobordism theorem) completely in the language of Morse theory. Drastically oversimplifying matters, the proof of the higher Poincaré conjecture can be described as follows. Let M be a smooth manifold of dimension ≥ 5 which is homotopy equivalent to a sphere. Endow M with a (smooth) Morse function f. One then proceeds to show that the critical points of f can be canceled out in pairs until one is left with a Morse function with exactly two critical points, which implies that M is a (topological) sphere.

A key step in this proof is the "cancellation theorem" which provides a sufficient condition for two critical points to be canceled (see Theorem 5.4 in [**72**], which Milnor calls "The First Cancellation Theorem", or the original proof in [**74**]). In this section we will see that the analogous theorem holds for discrete Morse functions. Moreover, in the combinatorial setting the proof is much simpler. The main result is that if $\alpha^{(p)}$ and $\beta^{(p+1)}$ are 2 critical simplices, and if there is exactly 1 gradient path from β to α, then α and β can be canceled. More precisely,

Theorem 34. *Suppose V is a discrete gradient vector field on M such that $\beta^{(p+1)}$ and $\alpha^{(p)}$ are critical, and there is exactly one V-path from β to α. Then there is another gradient vector field W on M with the same critical simplices except that α and β are no longer critical. Moreover, W is equal to V except along the unique gradient path from β to α.*

In the smooth case, the proof, either as presented originally by Morse in [**74**] or as presented in [**72**], is rather technical. In our discrete case the proof is simple. If, in the top drawing in Figure 16, the indicated gradient path is the only V-path from β to α, then we can reverse the gradient vector field along this path, replacing V by the vector field W shown in the bottom drawing in Figure 14.

The uniqueness of the V-path implies that the resulting discrete vector field has no closed orbits, and hence, by Theorem 2, is a gradient vector field. Moreover, α and β are not critical for this new gradient vector field, while the criticality of all other simplices is unchanged. This completes the proof.

The proof in the smooth case proceeds along the same lines. However, in addition to turning around those vectors along the unique gradient path from β to α, one must also adjust all nearby vectors so that the resulting vector field is

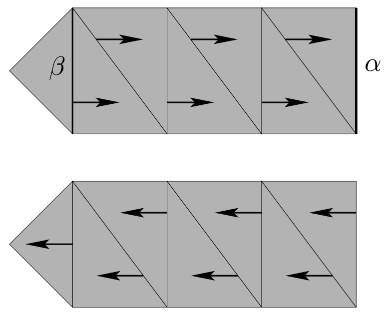

Figure 16. Canceling critical simplices.

smooth. Moreover, one must check that the new vector field is a gradient vector field, so that, in particular, modifying the vectors did not result in the creation of a closed orbit. This is an example of the sort of complications which arise in the smooth setting, but which do not make an appearance in the discrete theory.

This theorem was recently put to very good use in [**7**], in which discrete Morse theory is used to determine the homotopy type of some simplicial complexes arising in the study of partitions. It is fascinating, and quite pleasing, to see the same idea play a central role in two subjects, the Poincaré conjecture and the study of partitions, which seem to have so little to do with one another. In [**50**], Hersh generalizes this cancellation technique and investigates, among other ideas, when families of pairs of critical simplices can be canceled simultaneously. The main theorem of this section is also used extensively in [**54**] as a basic computational tool for searching for optimal gradient vector fields. To see other computational approaches to finding optimal gradient vector fields, the reader can take a look at [**62**], [**63**], [**64**].

5. Exercises for Lecture 2

(1) In the lecture we found a perfect gradient vector field on the complex Δ_n of disconnected graphs on n vertices. Since our construction began by collapsing everything towards the graph g containing only the edge $[1,2]$ we saw that this is equivalent to the construction of a perfect gradient vector field on the complex of connected graphs on n vertices. The critical connected graphs are precisely the critical disconnected graphs with the edge $[1,2]$ added. The result is the set of increasing trees on n-vertices. That is, the trees with vertex set $\{1,2,\dots,n\}$ with the property that the labels increase along every ray starting at vertex 1. Note that there are $(n-1)!$ of these, and each contains $(n-1)$ edges (and thus corresponds to a simplex of dimension $(n-2)$). Consider the collection $\mathcal{P}_n$ of graphs on n-vertices which are paths with one endpoint labeled 1. That is, graphs of the form

$$1 - v_2 - v_3 - \cdots - v_n$$

where $\{v_2, v_3, \dots, v_n\} = \{2,3,\dots,n\}$. We observe that there are precisely $(n-1)!$ of these graphs, and each has $(n-1)$ edges. Your job is to construct a gradient vector field on the simplicial complex of connected graphs on

n vertices for which the critical graphs are precisely the graphs in $\mathcal{P}_n$. (In Vassiliev's original work on this complex, this is the form in which he presented the answer.)

(2) Let G be any graph, and let P be any monotone decreasing graph property. Then we can consider the simplicial complex of all spanning subgraphs of G that satisfy P.

In the lecture we only considered the case where G is a complete graph. For other graphs G, these complexes are quite interesting and largely unexplored.

(a) Examine this in the case where P is the property of being disconnected. What is the homotopy type of the resulting complex? That is, given a graph G, what is the homotopy type of the simplicial complex of disconnected spanning subgraphs of G?

(b) Pick your favorite monotone graph property and your favorite graph and examine the resulting simplicial complex.

(3) Let M be a simplicial complex with a gradient vector field V. Prove that the homology of the Morse complex (as defined in this lecture) is isomorphic to the homology of M by following these steps:

(a) Suppose that V is the empty gradient vector field. Then the Morse complex is just the standard simplicial chain complex of M.

(b) Now prove that the homology of the Morse complex of V does not change if one pair is removed from V (i.e. if one arrow is erased). Do this by showing that if

$$M:\ 0 \to M_n \xrightarrow{d} M_{n-1} \xrightarrow{d} M_{n-2} \xrightarrow{d} \cdots \xrightarrow{d} M_0 \to 0$$

and

$$M':\ 0 \to M'_n \xrightarrow{d'} M'_{n-1} \xrightarrow{d'} M'_{n-2} \xrightarrow{d'} c \ldots \xrightarrow{d'} M'_0 \to 0$$

are the Morse complexes corresponding to gradient vector fields V and V' on M which differ by a single arrow, then there is a map $\Phi : M_i \to M'_i$ which induces an isomorphism on homology. Try to construct the map Φ as explicitly as possible.

Together (a) and (b) prove the desired result.

(4) In the exercises to Lecture 1 we proved that every triangulated surface has a perfect gradient vector field. Consider the Morse complex corresponding to such a vector field. Prove that all of the differentials vanish (that is, each differential is the zero map). Can you understand this directly from the definition of the differential – that is by counting gradient paths?

LECTURE 3
Discrete Morse Theory and Evasiveness

1. The Main Results

So far, we have indicated some applications of discrete Morse theory to combinatorics and topology. We now present an application to computer science. The reader should see the reference [**36**] for a more complete treatment of the content of this section.

There is a wide variety of situations in which one has the ability to quickly ask a series of yes/no questions, with the goal of answering a more difficult question. For example, when one goes to the doctor with an illness, the doctor usually asks a series of yes/no questions, such as "Do you have a headache?", "Do you have a fever?", etc., using the information from the previous questions to decide what to ask next, with the goal of answering the more difficult question "What illness does my patient have?". When one takes a malfunctioning car to the mechanic, the mechanic often attempts to analyze the problem by testing the individual components one at a time, using the appropriate tools to ask "Is this component working properly?", with the goal of answering the more difficult question "What is wrong with this car?". The mathematical study of such questions began with the following sort of problem. Suppose that one has a network G of phones, or computers, or... which we think of as a collection of nodes, some of which are connected by arcs. We assume that the nework G is connected. That is, one can get from any node to any other node by a series of arcs. Now we suppose that there is an electrical storm, or a terrorist attack, or..., and some of the arcs are disabled. At that time, our first concern may not be "What is the precise network that remains?" but rather, we may primarily concerned with questions such as "Is the remaining network still connected?". This is the difficult question we wish to answer. We suppose further that we have the capability of testing each arc in the original network, in order to answer the question "Is this arc still working?". Of course, we can answer any question about the network if we simply test every arc in the original network, and determine precisely which of these arcs are still working, as that completely determines the remaining network. The precise question we want to analyse here is "Can one do any better?". That is, is there any strategy for testing the arcs such that we are guaranteed that we can answer the desired question before having tested each of the original arcs.

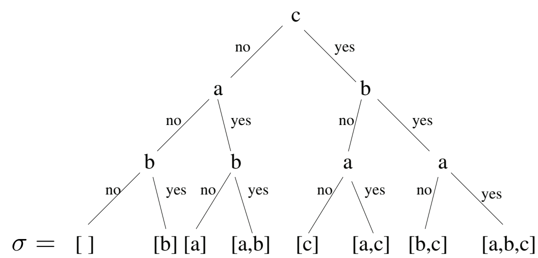

Figure 17. A search algorithm.

Let us begin with a simple example of the sort of thing we wish to study. Suppose there are three yes/no questions that we can easily ask. We label these questions $\{a, b, c\}$.

Assumption 1: We suppose that these questions have the property that their answers are independent of the order in which they are asked. (We will make this assumption for the rest of these notes.)

Then there are eight possible outcomes resulting from asking these three questions. We label these outcomes by listing the questions that yield the answer "yes" for that outcome. The possibilities are: [ ], $[a]$, $[b]$, $[c]$, $[ab]$, $[ac]$, $[bc]$, $[abc]$.

Assumption 2: We assume that every set of answers is possible.

That is, one can easily imagine a set of questions with the property that questions b and c can not both be answered "yes", but we will not consider this possibility in these lectures. We make this assumption only for reasons of simplicity. The general situation is considered in [**41**]. Suppose that the following four outcomes are good: $[a]$,$[b]$,$[c]$,$[ab]$, and the remaining outcomes are bad. By asking these three questions, our goal is to determine whether the outcome is good or bad. We can, of course, accomplish this goal by asking all three questions. We are considered to have won this game if we achieve the goal before we ask the third question. A winning strategy, then, is one which guarantees that no matter what the outcome is, we can determine whether or not it is good or bad before asking the third question.

For example, consider the search algorithm shown in Figure 17, in which case we have listed the question to be asked next, given the answers to the previous questions. For example, we ask question c first, and if we get the answer "yes" we ask question b, but if we get the answer "no", we ask the question a. We observe that, asking questions in the indicated order, if the outcome is in the set $\{[\ ], [b], [c], [ac]\}$, then we must ask the third question. Outcomes which require us to ask the third question are called *evaders* of the search algorithm, so the algorithm has 4 evaders. In fact, this is the best one can do. The following proposition is fairly easy to check by straightforward means.

Proposition 35. *Every search algorithm for the problem of determining membership in the set of good outcomes* $\{[a], [b], [c], [ab]\}$ *has at least 4 evaders. The number of evaders which are good equals the number of evaders which are bad, and hence there must be at least two of each.*

If we assume that each outcome is equally likely, then this proposition implies that no matter which search algorithm we choose, we will have to ask the third question at least half of the time. Note that this theorem does not say that every search algorithm has exactly 4 evaders, and it is rather easy to find search algorithms

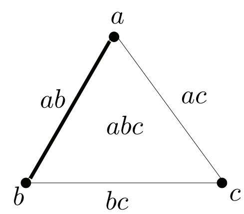

Figure 18. A topological approach to the problem.

with more than 4 evaders. If every search algorithm has some evaders, so that we have no winning strategy, then we say that the problem is *evasive*.

It is probably not at all clear to the reader what this topic is doing in a series of lectures on discrete Morse theory, but we will show that in fact these topics are intimately related. In particular, we will show that algebraic topology gives a way of understanding why some problems of this form are easy, and others are hard. First we observe that the problem can easily be stated in a more topological way. Consider the 2-dimensional simplex S with vertices labeled $\{a, b, c\}$. Then the faces of S can be indentified with the subsets of $\{a, b, c\}$, and hence with the 8 possible outcomes (see Figure 18). Then the good and bad outcomes partition the faces of S into 2 sets. In this setting we are given a partition of the set of faces, the outcome is a face σ of the simplex, and our goal is to determine which block of the partition contains σ. We are permitted to ask questions of the form "Is vertex v in σ?".

In this way, we can convert binary search problems (which satisfy Assumption 1) into the language of simplices. If we also require Assumption 2, then the sort of search problems we are considering lead to problems of the following form. Let S be an n-dimensional simplex, with vertices $v_0, v_1, \ldots, v_n$, F the set of faces of S, and

$$P: \ F = P_1 \sqcup P_2 \sqcup \ldots \sqcup P_k$$

a partition of F, which is known to you. Let σ be a face of S which is not known to you. Your goal is to determine which block of the partition P contains σ. In particular, you need not determine the face σ. You are permitted to ask questions of the form "Is v_i in σ?". You may use the answers to the questions you have already asked in determining which vertex to ask about next. Of course, you can determine which block contains σ by asking $n+1$ questions, since by asking about all $n+1$ vertices you can completely determine σ. You win this game if you answer the given question after asking fewer than $n+1$ questions.

Say that P is nonevasive if there is a winning strategy for this game, i.e there is a search algorithm that determines which block contains σ in fewer than $n+1$ questions, no matter what σ is. Say P is evasive otherwise.

One of the main issues we will have to deal with is that a block P_i of the partition need not be a subcomplex or have any other nice structure. Hence, the notion of the homology of such a set is problematic. Let P be any set of faces of a simplex S, and let $\mathbb{F}$ be a field. One of the main contributions of this and the following sections is a definition of the $\mathbb{F}$-Betti numbers of P. More precisely, for each $i = -1, 0, 1, \ldots$, we will define $B_i(P, \mathbb{F})$, the i^{th} Betti number of P with respect to the field $\mathbb{F}$. We will also define the even and odd Betti numbers, denoted $B_e(P, \mathbb{F})$ and $B_o(P, \mathbb{F})$, respectively, and the total Betti number $B(P, \mathbb{F})$. For ease of notation, we will assume that the field $\mathbb{F}$ is fixed, and refer to $B_i(P), B_e(P), B_o(P)$

and $B(P)$. We will present the precise definition of these numbers in the next section. The basic idea is that the Betti number $B_i(P)$ is defined by restricting the chain complex of S over the field $\mathbb{F}$ to the faces in P. The result need not be a complex, since the composition of its consecutive differentials might be nonzero, but the dimension of its i^{th} "homology" can still be defined as the dimension of its i^{th} kernel minus the dimension of its $(i+1)^{st}$ image (if that is nonnegative, and 0 otherwise); see Proposition 47. At this point, we will state the main properties of these Betti numbers.

Theorem 36. *For any set of faces P*

(i)
$$B_e(P) \geq \sum_{i\ even} B_i(P),$$
$$B_o(P) \geq \sum_{i\ odd} B_i(P),$$
$$B(P) = B_e(P) + B_o(P).$$

(ii) $B_i(P) = 0$ *for i larger than the dimension of P.*

(iii) $B_e(P) - B_o(P) = \chi(P) = \sum_i (-1)^i \#\{i\text{-simplices in } P\}$

Our notion of a Betti number is equal to a standard notion in a number of settings.

Theorem 37.

(1) If P is a subcomplex of S, and the empty set (considered as a face of S) is an element of P, then for each i
$$B_i(P) = \dim \widetilde{H}_i(P, \mathbb{F}).$$
where the tilde denotes reduced homology. Moreover,
$$B_e(P) = \dim \widetilde{H}_{even}(P, \mathbb{F})$$
$$B_o(P) = \dim \widetilde{H}_{odd}(P, \mathbb{F})$$
$$B(P) = \dim \widetilde{H}_*(P, \mathbb{F}).$$

(2) If P is a subcomplex of S, except that the empty set is not element of P, then for each i
$$B_i(P) = \dim H_i(P, \mathbb{F}).$$
Moreover,
$$B_e(P) = \dim H_{even}(P, \mathbb{F})$$
$$B_o(P) = \dim H_{odd}(P, \mathbb{F})$$
$$B(P) = \dim H_*(P, \mathbb{F}).$$

(3) Let $\overline{P}$ denote the closure of P (i.e. the set consisting of the faces of P along with all of their faces), and let $\dot{P} = \overline{P} - P$. If $\dot{P}$ is a subcomplex of S (which contains the empty set) then for each i
$$B_i(P) = \dim H_i(\overline{P}, P, \mathbb{F}).$$
Moreover
$$B_e(P) = \dim H_{even}(\overline{P}, P, \mathbb{F})$$
$$B_o(P) = \dim H_{odd}(\overline{P}, P, \mathbb{F})$$
$$B(P) = \dim H_*(\overline{P}, P, \mathbb{F}).$$

Assuming these results for now, as well as the still undefined notion of Betti number, we present the main theorem of this section.

Theorem 38. *With all notation as above, for any search algorithm A the number of evaders of A which lie in any block P_j of the partition P is at least $B(P_j)$. Hence the total number of evaders is at least $\sum_{j=1}^{k} B(P_j)$.*

In fact, we can make this statement much more precise. Define the dimension of an evader to be the dimension of the face of S to which it corresponds. That is, if σ is any possible outcome, $\dim(\sigma)$ is

$$\text{(the number of questions answered "yes" if the outcome is } \sigma) - 1.$$

Theorem 39. *With all notation as above, for any search algorithm A the number of evaders of A of dimension i which lie in any block P_j of the partition P is at least $B_i(P_j)$. The number of even-dimensional evaders which lie in block P_j is at least $B_e(P_j)$, and the number of odd-dimensional evaders which lie in block P_j is at least $B_o(P_j)$.*

Before discussing the proof of this result, we would like to point out that Kahn, Saks and Sturtevant [**55**] first observed the relationship between evasiveness and algebraic topology. In their setting, the partition consists of precisely two blocks, $P : S = P_1 \sqcup P_2$, in which P_1 is a subcomplex. They proved the following theorem.

Theorem 40. *If $\tilde{H}_*(P_1) \neq 0$, where $\tilde{H}_*(P_1)$ denotes the reduced homology of P_1, then P is evasive.*

In fact, they proved something stronger, and we will come back to this point later. In [**39**] we used discrete Morse theory to make some of their results more quantitative along the lines of Theorems 38 and 39. The generalization in this section to more than two blocks is relatively minor. The extension to more general sets of faces is the major value of this newer work.

We illustrate the previous theorems by returning to the example introduced at the beginning of this section. Let P_1 denote the set $\{[a], [b], [c], [ab]\}$ of good outcomes, and let P_2 denote the complement, the set of bad outcomes. We observe that P_1 is a simplicial complex which does not contain the empty face. Hence by Theorem 37, $B(P_1)$ is equal to the dimension of the (unreduced) homology of P_1, which is 2. By Theorem 38, we learn that for any search algorithm, the number of evaders which lie in P_1 is at least 2. We observe that P_2 does not satisfy any of the hypotheses presented in Theorem 37, so one can not deduce its Betti numbers from that result. However, as the reader can check (after we present the definition of Betti numbers in the next section), its total Betti number is also 2.

The link between evasiveness and algebraic topology is provided by discrete Morse theory. Morse theory comes to the fore when one observes that a search algorithm induces a discrete vector field on S. For example, the search algorithm shown in Figure 17 induces the vector field

$$V = \{\, \{[\,] < [b]\},\ \{[a] < [a,b]\},\ \{[c] < [a,c]\},\ \{[b,c] < [a,b,c]\} \,\}$$

That is, V consists of those pairs of faces of S which are not distinguished by the search algorithm until the last question. There is slight subtlety here in that a search algorithm pairs a vertex with the empty face [ ], while in our original definition, it was not permitted to pair a simplex with [ ]. Thus, to get a true discrete vector field, we must remove this pair from V. (It is precisely this subtle point that results in the reduced homology of K being the relevant measure of topological complexity

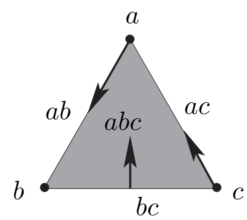

Figure 19. The induced vector field V.

in Theorem 37(1), rather than the unreduced homology.) However, for simplicity, from now on we will simply ignore this technical point.

Theorem 41. *For any search algorithm, let V denote the vector field consisting of pairs of nonempty faces of S which are not distinguished by the search algorithm until the last question. Then V is a gradient vector field.*

We will postpone the proof of this result until the end of this section.

For now, suppose that block P_j of the partition P is a subcomplex (containing the empty face). We will complete the proof in this setting before discussing the general case. Now restrict V to P_j by taking only those pairs in V such that both simplices are in P_j, and denote the resulting vector field by V_j. In our simple example, this results in the vector field

$$V_1 = \{\{[a] < [a, b]\}\}.$$

From the previous theorem, V has no closed orbits. Any discrete vector field consisting of a subset of the pairs of V has fewer paths, and hence also has no closed orbits. Therefore, V_j is a gradient vector field on P_j. Note that V pairs every face of S with another face, and hence there are no critical simplices except for the vertex which is paired with the empty set. Thus, ignoring that special vertex for the moment, the critical simplices of V_j are precisely the simplices of P_j which are paired in V with a face of S which is not in P_j. These are precisely the simplices of P_j which are the evaders of the search algorithm.

The Morse inequalities of Theorem 13 (i) immediately imply the following result.

Corollary 42. *If the block P_j of the partition P is a subcomplex (containing the empty face of S) then the number of evaders in P_j is at least* $\dim \widetilde{H}_*(P_j)$.

(We must use reduced homology here because of the minor issue surrounding the vertex paired with the empty set.) This yields Theorem 37 (in the case of a simplicial complex containing the empty face).

Suppose that P is nonevasive. Then there is some search algorithm which has no evaders. From our above discussion we have seen that this implies that P_j has a gradient vector field with no critical simplices. Actually, this is not quite true. The gradient vector field must have a critical vertex – the vertex that is paired with the empty face. These ideas lead to the following strengthening of Theorem 40.

Theorem 43. *If P is nonevasive, and if the block P_j of the partition is a subcomplex, then P_j collapses to a vertex.*

This theorem appears in [**55**], the paper that first established, and used to very good effect, a close relationship between evasiveness and topology. The interested reader can consult [**36**] for some additional refinements of this theorem. This topic has been the subject of much study, and the reader can find more information about the connection between evasiveness and topology in the references [**11**], [**56**], [**76**], [**77**], and [**94**] .

We now present a proof of Theorem 41. Let S denote an n-simplex, and fix a search algorithm. Associate to each p-simplex α of S the sequence of integers

$$n(\alpha) = n_0(\alpha) < n_1(\alpha) < \cdots < n_p(\alpha)$$

where, for each i, question number $n_i(\alpha)$ is answered "yes" if $\sigma = \alpha$, and these are the only questions answered "yes".

Let V be the vector field induced by the search algorithm and

$$\alpha_1, \alpha_2$$

be a V-path. Then either (i) α_1 is a face of α_2 and $\{\alpha_1 < \alpha_2\} \in V$, or (ii) α_2 is a face of α_1 and $\{\alpha_1 < \alpha_2\} \notin V$. Let us consider case (ii) first. In this case, α_2 has one fewer vertex than α_1, and the vertex is not the subject of the $(n+1)^{st}$ question. Suppose the the vertex is the subject of the $n_i(\alpha_1)^{st}$ question. Then this question is answered "yes" for α_1, but "no" for α_2. This implies that

$$n(\alpha_2) = n_0(\alpha_1) < n_1(\alpha_1) < \cdots < n_{i-1}(\alpha_1) < n_i(\alpha_2) < \cdots$$

for some $i < n+1$, and such that $n_i(\alpha_2) > n_i(\alpha_1)$. Thus $n(\alpha_2) > n(\alpha_1)$ in the lexicographic order.

We now consider case (i), in which $\{\alpha_1 < \alpha_2\} \in V$, and continue the V-path one more step to $\alpha_1, \alpha_2, \alpha_3$. Then α_1 and α_2 are not distinguished until the $(n+1)^{st}$ question. Thus,

$$n(\alpha_2) = n_0(\alpha_1) < n_1(\alpha_1) < \cdots < n_p(\alpha_1) < n+1.$$

We now observe that the vertices of α_3 are a subset of the vertices of α_2. Suppose the vertex of α_2 which is not in α_3 is the vertex tested in question $n_i(\alpha_2)$. Then we must have $i \neq n+1$. This demonstrates that

$$n(\alpha_3) = n_0(\alpha_1) < n_1(\alpha_1) < \cdots < n_{i-1}(\alpha_1) < n_i(\alpha_3) < \cdots$$

for some $i < n+1$, and such that $n_i(\alpha_3) > n_i(\alpha_1)$. Thus $n(\alpha_3) > n(\alpha_1)$ in the lexicographic order, which is sufficient to prove that there are no closed orbits.

In [**53**], Jonsson investigates further the question of which gradient vector fields arise from decision trees. Anyone interested in this topic should also consult [**84**].

2. Betti Numbers for General Sets of Faces

In this section we examine how to extend the results of the previous section to sets of faces of a simplex that do not form a simplicial complex or have any other special structure. A more extensive treatment of these ideas can be found in [**41**]. Let $\mathbb{F}$ be any field. Let

$$\mathcal{V}: 0 \longrightarrow V_n \xrightarrow{\partial_n} V_{n-1} \xrightarrow{\partial_{n-1}} \cdots \xrightarrow{\partial_1} V_0 \xrightarrow{\partial_0} V_{-1} \longrightarrow 0$$

be a complex of finite dimensional vector spaces over the field $\mathbb{F}$. The ∂_i's are assumed to be linear maps, but we are not assuming that $\partial_d \circ \partial_{d+1} = 0$. Our goal is to define the "homology" of this complex.

If $\partial_p \circ \partial_{p+1} = 0$ for each p, then we say that $\mathcal{V}$ is a *differential complex* and that the ∂_i's form a *differential*. We recall that one defines the homology of a differential complex by the formula

$$H_p(\mathcal{V}) := \frac{\operatorname{Ker} \partial_p}{\operatorname{Im} \partial_{p+1}}. \tag{5}$$

For each p, choose a subspace X_p which is mapped isomorphically onto $\operatorname{Im} \partial_p$. Then we have that

$$V_p = X_p \oplus \operatorname{Ker} \partial_p.$$

In the case of a differential complex, $\operatorname{Im} \partial_{p+1} \subset \operatorname{Ker} \partial_p$ we can find a $Z_p \subset V_p$ so that

$$\operatorname{Ker} \partial_p = \operatorname{Im} \partial_{p+1} \oplus Z_p,$$

which implies that

$$V_p = X_p \oplus \operatorname{Im} \partial_{p+1} \oplus Z_p,$$

and the reader can easily check that $Z_p \cong H_p(\mathcal{V})$.

We now return to the general case of a nondifferential complex. That is, we no longer assume that $\partial_p \circ \partial_{p+1} = 0$. We will use the construction of the previous paragraph to define the homology of such a complex.

Definition 44. A homological decomposition D of the complex S is a decomposition

$$V_i = X_i \oplus Y_i \oplus Z_i,$$

for each i, with the property that for each i, ∂_i maps X_i isomorphically onto Y_{i-1}.

By the notation $V_i = X_i \oplus Y_i \oplus Z_i$ we mean that X_i, Y_i and Z_i are linear subspaces of V_i, such that their pairwise intersections are $\{0\}$, and they sum to give all of V_i. Homological decompositions always exist, since one can take $X_i = 0, Y_i = 0$, and $Z_i = V_i$, for each i.

For any homological decomposition D of $\mathcal{V}$, and any i, let $B_i(\mathcal{V}, D)$ denote the dimension of Z_i. We also define the even Betti number of D

$$B_e(\mathcal{V}, \mathcal{D}) := \sum_{i \text{ even}} B_i(\mathcal{V}, \mathcal{D}),$$

the odd Betti number of D

$$B_o(\mathcal{V}, \mathcal{D}) = \sum_{i \text{ odd}} B_i(\mathcal{V}, \mathcal{D}),$$

and the total Betti number of D

$$B(\mathcal{V}, \mathcal{D}) = B_e(\mathcal{V}, \mathcal{D}) + B_o(\mathcal{V}, \mathcal{D}) = \sum_i B_i(\mathcal{V}, \mathcal{D}).$$

We now define the Betti numbers of S by

$$B_i(\mathcal{V}) := \min_{\mathcal{D}} B_i(\mathcal{V}, \mathcal{D}),$$

$$B_e(\mathcal{V}) := \min_{\mathcal{D}} B_e(\mathcal{V}, \mathcal{D}),$$

$$B_o(\mathcal{V}) := \min_{\mathcal{D}} B_o(\mathcal{V}, \mathcal{D}),$$

and

$$B(\mathcal{V}) := \min_{\mathcal{D}} B(\mathcal{V}, \mathcal{D}).$$

We observe the following facts.

Proposition 45. *Let $\mathcal{V}$ be any finite complex of finite dimensional vector spaces.*
(i) $B_e(S) \geq \sum_{i\ even} B_i(S)$ *and* $B_o(S) \geq \sum_{i\ odd} B_i(S)$.
(ii) $B(S) = B_e(S) + B_o(S)$.

Example 46. A simple example will serve to show that the inequalities in part (i) of the proposition can be strict when $\mathcal{V}$ is not a differential complex. Consider the complex $\mathcal{V}$ with $V_0 = V_1 = V_2 = \mathbb{F}$, and $V_i = 0$ for $i = -1$ and $i > 2$. Suppose that ∂_1 and ∂_0 are both the identity map.

$$\mathcal{V} : 0 \longrightarrow \mathbb{F} \xrightarrow{\partial_1} \mathbb{F} \xrightarrow{\partial_0} \mathbb{F} \longrightarrow 0$$

Let D_1 denote the homological decomposition

$$0 \longrightarrow \mathbb{F} \oplus 0 \oplus 0 \xrightarrow{\partial_1} 0 \oplus \mathbb{F} \oplus 0 \xrightarrow{\partial_0} 0 \oplus 0 \oplus \mathbb{F} \longrightarrow 0.$$

We have that $B_1(\mathcal{V}, D_1) = B_2(\mathcal{V}, D_1) = 0$, while $B_0(\mathcal{V}, D_1) = 1$, which implies that $B_1(\mathcal{V}) = B_2(\mathcal{V}) = 0$, and $B_0(\mathcal{V}) \leq 1$.

Let D_2 denote the homological decomposition

$$0 \longrightarrow 0 \oplus 0 \oplus \mathbb{F} \xrightarrow{\partial_1} \mathbb{F} \oplus 0 \oplus 0 \xrightarrow{\partial_0} 0 \oplus \mathbb{F} \oplus 0 \longrightarrow 0.$$

In this case we see that $B_0(\mathcal{V}, D_2) = B_1(\mathcal{V}, D_2) = 0$, while $B_2(\mathcal{V}, D_2) = 1$, which implies that $B_0(\mathcal{V}) = B_1(\mathcal{V}) = 0$, and $B_2(\mathcal{V}) \leq 1$.

Thus we learn that $B_i(\mathcal{V}) = 0$ for every i. On the other hand $B_o(\mathcal{V}, D_1) = B_o(\mathcal{V}, D_2) = 0$, which implies that $B_o(\mathcal{V}) = 0$. We note that $B_e(\mathcal{V}, D_1) = B_e(\mathcal{V}, D_2) = 1$, and, in fact, once can easily see that $B_e(\mathcal{V}) = 1$.

In the case that $\mathcal{V}$ is a differential complex, we have

$$\dim H_i(\mathcal{V}) = \dim(\operatorname{Ker} \partial_i) - \dim(\operatorname{Im} \partial_{i+1}).$$

A remnant of this equation holds for general complexes.

Proposition 47. *For any complex S, whether a differential complex or not,*

$$B_i(\mathcal{V}) = \max\{\dim(\operatorname{Ker} \partial_i) - \dim(\operatorname{Im} \partial_{i+1}), 0\}$$

Since the right hand side is easily algorithmically computable by standard methods, this theorem implies that the generalized Betti numbers are also readily computable. Moreover, we note that these definitions do, in fact, generalize the standard definition those for a differential complex.

Theorem 48. *Suppose that $\mathcal{V}$ is a differential complex, and let $H_i(\mathcal{V})$ denote the homology of S as defined by the standard formula (5). Then for each i,*

$$B_i(\mathcal{V}) = \dim H_i(\mathcal{V}).$$

Moreover,

$$B_e(\mathcal{V}) = \sum_{i\ even} B_i(\mathcal{V})$$

$$B_o(\mathcal{V}) = \sum_{i\ odd} B_i(\mathcal{V}),$$

and

$$B(\mathcal{V}) = \sum_i B_i(\mathcal{V}).$$

Now let S be a simplex. A *simplex space* is defined to be a set of faces of S (we consider the empty set to be a face of S). Let K be a simplex space. Fix a field $\mathbb{F}$ and let $C_d(K,\mathbb{F})$ denote the oriented d-chains in K. This is defined in the usual way, even if K is not a subcomplex, that is, for all $d \geq 2$, let $S_d(K)$ denote the set of pairs (σ, ϵ), where σ is a d-simplex in K, and ϵ is an orientation on σ (i.e. an ordering of the vertices in σ modulo even permutations – so that each σ has two distinct orientations). We let $\widetilde{C}_d(K,\mathbb{F})$ denote the vector space of all formal linear combinations of elements in $S_d(K)$, and $C_d(K,\mathbb{F})$ the result of quotienting out $\tilde{C}_d(K,\mathbb{F})$ by the relation $-1(\sigma,\epsilon) = (\sigma,-\epsilon)$, where if ϵ denotes an orientation on a d-simplex σ, then $-\epsilon$ denotes the alternate orientation. If $d = -1, 0$, then each d-simplex in K has a unique orientation, and we let $C_d(K,\mathbb{F})$ denote the vector space of all formal linear combinations of d-simplices. Note that the only possible -1-dimensional simplex is the empty simplex. Thus, if K contains the empty simplex $C_{-1}(K,\mathbb{F}) \cong \mathbb{F}$ and otherwise $C_{-1}(K,\mathbb{F}) \cong 0$.

We now define a boundary map ∂_d from $C_d(K,\mathbb{F})$ to $C_{d-1}(K,\mathbb{F})$, by setting, for any oriented d-simplex $[x_0, x_1, \ldots, x_d]$,

$$\partial_d[x_0, x_1, \ldots, x_d] = \sum_i{}' (-1)^i [x_0, x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_d].$$

where the notation $\sum'$ means that the sum is over all $i \subset \{0, 1, 2, \ldots, d\}$ such that the simplex $\{x_0, x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_d\}$ is in K. As usual, the notation $[x_0, x_1, \ldots, x_d]$ denotes the simplex $\{x_0, x_1, \ldots, x_d\}$ along with the orientation induced by ordering the vertices in the manner in which they are listed.

We let $C(K,\mathbb{F})$ denote the complex

$$0 \longrightarrow C_n(K,\mathbb{F}) \xrightarrow{\partial_n} \ldots \xrightarrow{\partial_0} C_{-1}(K,\mathbb{F}) \longrightarrow 0.$$

Using the notation of the previous section, we denote by

$$B_i(K,\mathbb{F}), \quad B_e(K,\mathbb{F}), \quad B_o(K,\mathbb{F}), \quad B(K,\mathbb{F}),$$

the numbers

$$B_i(\mathcal{C}(K,\mathbb{F})), \quad B_e(\mathcal{C}(K,\mathbb{F})), \quad B_o(\mathcal{C}(K,\mathbb{F})), \quad B(\mathcal{C}(K,\mathbb{F})).$$

We observe that if K is a simplicial complex, and contains the empty simplex, then $C(K,\mathbb{F})$ is the usual reduced chain complex. If K is a simplicial complex except that K does not contain the empty simplex, then $C(K,\mathbb{F})$ is the usual (nonreduced) chain complex. If $\dot{K} = \bar{K} - K$ is a nonempty subcomplex (where $\bar{K}$ denotes the closure of K, i.e. K along with all of the faces of the elements in K, and $\dot{K}$ denotes the elements in $\bar{K}$ which are not in K), then $C(K,\mathbb{F})$ is isomorphic to the relative chain complex $C(\bar{K},\dot{K},\mathbb{F})$. In particular, in these cases, $C(K,\mathbb{F})$ is a differential complex, so, by Theorem 48, the Betti numbers we have defined are equal to the standard Betti numbers. This implies Theorem 37.

Example 49. Let S denote the two-dimensional simplex with vertices labeled a, b, c. Let K denote the simplex space consisting of the faces $[c]$, $[b,c]$ and $[a,b,c]$ (Figure 20). The chain complex for K can easily seen to be isomorphic to the complex examined in Example 46. From the calculations done there, we see that $B_i(K) = 0$ for each $i = -1, 0, 2, \ldots$, and $B_o(K) = 0$, while $B(K) = B_e(K) = 1$.

The next goal is to indicate that these Betti numbers allow us to apply the basic notions of discrete Morse theory to general simplex spaces. Let K be a simplex

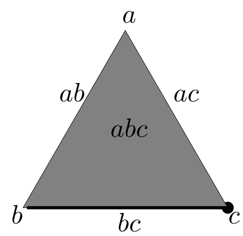

Figure 20. The complex K of example 49.

space. We define the basic combinatorial notions just as for a simplicial complex. A face a of S is said to be a *maximal element* of K if a is in K, and a is not a proper subset of any element in K. If a is a maximal element of K, we say that b is a *free face* of a in K if: b is in K, b is a codimension-one face of a, and a is the only element of K which properly contains b. Let K be a simplex space, a a maximal face of K, and b is a free face of a in K. The act of replacing K by $K - \{a, b\}$ is called a *simplicial collapse*. Say that K is *collapsible* if one can transform K into the empty simplex space by a sequence of simplical collapses.

Let K be a simplex space, and a a maximal element of K. We will call the act of replacing K by $K - a$ a *simplical removal*.

We will use the term *elementary simplicial reduction* to refer to either a simplical collapse or a simplicial removal. A *complete reduction* of K is any sequence of elementary reductions that transforms K into the empty simplex space. In particular, K is collapsible if and only if there is a complete reduction consisting solely of simplicial collapses.

Lemma 50. *Let K be a simplex space.*

(i) If $K' = K - \alpha$ for some maximal d-simplex α, then

$$B(K') \geq B(K) - 1.$$

(ii) If $K' = K - (Int(\alpha) \cup Int(\beta))$ is the result of a simplicial collapse, where α is a maximal d-simplex and β is a free face of α, then

$$B(K') \geq B(K).$$

Together, parts (i) and (ii) imply the following theorem.

Theorem 51. *Let K be a simplex space. In any complete reduction of K, the number of simplices which are taken out by a simplicial removal is at least $B(K)$.*

Corollary 52. *Let K be any simplex space, and V any gradient vector field on K. Then the number of critical cells of V is at least $B(K)$.*

Theorem 50 can be made more precise to include an understanding of how the individual $B_d(K)$ can change under simplicial collapse and simplicial removal that is sufficient to imply Theorem 39.

Example 53. We end this section with an example to illustrate that, unlike in the case of a simplicial complex, a simplicial collapse can increase the Betti numbers of a simplex space. Let S denote the two-dimensional simplex with vertices label a, b, c. Let K denote the simplex space consisting of the four faces $[a], [a, b], [b, c], [a, b, c]$. Then K is collapsible, since one can remove $[b, c]$ and $[a, b, c]$ by one simplicial collapse, and the remaining two faces with a second simplicial collapse. Thus, all Betti numbers of K are zero. One the other hand, beginning with K, one can also remove the faces $[a, b]$ and $[a, b, c]$ by a simplicial collapse, resulting in the simplex space K' consisting of the faces $[a]$ and $[b, c]$. One can easily check that $B_0(K') = B_1(K') = B_e(K') = B_o(K') = 1$.

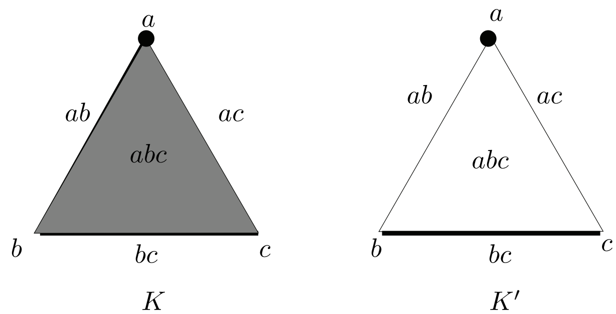

Figure 21. The simplex space K with vanishing Betti numbers collapses to K', which has nonzero Betti numbers.

Different algebraic extensions of discrete Morse theory appear in [**9**],[**52**] and [**81**]. These appoaches are quite similar in spirit to each other, and share some ideas with the work in this section.

3. Exercises for Lecture 3

(1) In Lecture 2 we constructed a perfect gradient field (i.e. one for which the Morse inequalities are equalities) on Δ_n^2, the complex of disconnected graphs on n vertices. Show that there is a decision tree which induces such a gradient vector field. (This observation is due to Jonsson, who found large classes of complexes which have perfect gradient vector fields which are induced by decision trees.)

(2) Consider the 2-dimensional simplicial complex on 5 vertices whose maximal simplices are [012], [023], [034], [045], [051], [123], [234], [345], [451], and [512]. Show that (i) this complex is collapsible and (ii) this complex is not nonevasive. (This example is due to Björner.)

(3) Returning to the first example in this lecture, prove directly from the definitions that the total Betti number of the set of bad outcomes is 2.

LECTURE 4
The Charney-Davis Conjectures

1. Introduction

These notes are intended to be an introduction to the Charney-Davis conjectures and some of their combinatorial implications. My aim is to provide a stimulating advertisement for a circle of ideas that is the subject of some fascinating recent work, most of which creates more questions than answers, and which has shed new light on some of the central questions in geometric combinatorics. The subject is a beautiful one, borrowing techniques and ideas from geometry, topology, analysis, algebra, algebraic geometry and combinatorics. My goal in these lectures is to present the topological and geometric context of these conjectures (as presented e.g. in [**15**],[**22**]), along with the most recent combinatorial understanding of them (as in Gal [**43**] and Brändén[**12**].) These notes will have been successful if some readers are inspired to consult these original sources, and to begin thinking about these conjectures.

The Charney-Davis conjectures, concerned with the relationship between geometry and topology, find their origins, as do most such questions, in the Gauss-Bonnet Theorem. Recall that the Gauss-Bonnet theorem states that if M is a compact surface with a Riemannian metric, then

$$\chi(M) = \frac{1}{2\pi}\int_M K \; d\text{area}$$

where K denotes the Gauss curvature of M. It follows that if $K \leq 0$ everywhere, then $\chi(M) \leq 0$.

Hopf conjectured the following generalization.

Conjecture 54. *If M is a compact Riemannian manifold of dimension $2n$ and the sectional curvature of M is ≤ 0 then $(-1)^n\chi(M) \geq 0$.*

[Recall that if M is an odd-dimensional manifold, then $\chi(M) = 0$.] This is not a suitable place for a primer in differential geometry, so we hope it will suffice to say that the condition that the sectional curvature is nonpositive means that every two-dimensional "orthogonal slice" of M is a surface of nonpositive Gauss curvature. This conjecture may seem a bit surprising, and perhaps unintuitive, at first glance. However, some general considerations point in this direction. Most notably, one has the following observations.

Proposition 55.

(1) Let M_1 and M_2 be compact manifolds, then
$$\chi(M_1 \times M_2) = \chi(M_1)\chi(M_2).$$
(2) If M_1 and M_2 are Riemannian manifolds with nonpositive sectional curvature, and $M_1 \times M_2$ is endowed with the product Riemannian metric, then $M_1 \times M_2$ has nonpositive sectional curvature.

Thus, if M_1 and M_2 are nonpositively curved Riemannian manifolds for which the conclusion of Hopf's conjecture holds, then the same is true for $M_1 \times M_2$. In particular, Hopf's conjecture holds for any product of arbitrarily many nonpositively curved surfaces.

Allendoerfer, Fenchel and Weil ([**1**],[**29**], [**2**]), and later Chern ([**17**]), proved a higher dimensional version of the Gauss-Bonnet theorem, which, for a compact Riemannian manifold, has the general form
$$\chi(M) = \int_M R\, dvol,$$
where R is a function of the curvature of M and is usually called the Chern-Gauss-Bonnet integrand. Chern [**18**] gives a proof (attributed to Milnor) that in dimension 4, if the sectional curvature is ≤ 0 everywhere, then $R \geq 0$. In particular:

Corollary 56. *If M is a compact Riemannian 4-manifold with sectional curvature ≤ 0, then $\chi(M) \geq 0$.*

However, Geroch [**44**] proved that this approach is insufficient to settle Hopf's conjecture in higher dimensions.

Theorem 57. *In even dimensions ≥ 6, there exist Riemannian metrics with sectional curvature ≤ 0 such that the Chern-Gauss-Bonnet integrand achieves both signs.*

So, in higher dimensions another approach is necessary. Before discussing alternate approaches, and partial results, we will take a detour to discuss some generalizations and extensions of Hopf's conjecture. From now on, when we say that a Riemannian manifold M has *nonpositive curvature,* we mean that all sectional curvatures are ≤ 0.

It is a theorem of Cartan and Hadamard that if M^n has nonpositive curvature, then $\widetilde{M}$, the universal cover of M, is diffeomorphic to $\mathbb{R}^n$. A manifold is said to be *aspherical* if its universal cover is contractible. Thurston generalized Hopf's conjecture to the following

Conjecture 58. *Let M^{2n} be a smooth, compact, aspherical manifold. Then*
$$(-1)^n \chi(M) \geq 0.$$

This is quite interesting, as the hypothesis has changed from a geometric condition to one that is purely topological. Our interests, however, lie in a different direction. Riemannian curvature is expressed in terms of 2^{nd} derivatives of the metric. Thus, Hopf's conjecture, as it is usually understood, is a statement about manifolds which are at least twice differentiable. However, Alexandrov showed how one could speak of nonpositive curvature for continuous, but nonsmooth, manifolds. Let M be a complete Riemannian manifold. The Hopf-Rinow theorem states that for any two points p and q in M, there exists a minimal geodesic from p to q (i.e. a curve γ

from p to q satisfying length(γ) = distance(p, q)). Let x, y and z be three points in M, and $\widetilde{x}$, $\widetilde{y}$ and $\widetilde{z}$ be three points in $\mathbb{R}^2$ such that

$$d_M(x,y) = d_{\mathbb{R}^2}(\widetilde{x},\widetilde{y}), \quad d_M(x,z) = d_{\mathbb{R}^2}(\widetilde{x},\widetilde{z}), \quad \text{and} \quad d_M(y,z) = d_{\mathbb{R}^2}(\widetilde{y},\widetilde{z}).$$

(Such $\widetilde{x}$, $\widetilde{y}$ and $\widetilde{z}$ always exist.) Let γ be a minimal geodesic from y to z, and $\widetilde{\gamma}$ the straight line from $\widetilde{y}$ to $\widetilde{z}$. Since $|\gamma| = |\widetilde{\gamma}|$ there is a natural identification between points in γ and points in $\widetilde{\gamma}$. In [**3**] Alexandrov proves the following result.

Theorem 59.

(1) *If M is simply-connected and has nonpositive curvature, then for any point p in γ, if $\tilde{p}$ is the corresponding point in $\tilde{\gamma}$, we have*

$$d_M(x,p) \le d_{\mathbb{R}^2}(\tilde{x},\tilde{p}).$$

(2) *If M is not simply-connected then the above inequality is still true if one restricts to triples x, y and z which are sufficiently close to one another.*

(3) *The converse is also true: If the above inequality holds for all sufficiently close x, y and z then M has nonpositive curvature.*

Alexandrov's theorem shows that the property of nonpositive curvature is equivalent to a property which can be expressed in terms of the distance function without reference to any derivatives. Hence we can use this theorem to make sense of the notion of nonpositive curvature for spaces without a smooth structure. By replacing $\mathbb{R}^2$ with other constant curvature surfaces, we can also make sense of the notion of having curvature bounded above by any real number. While a smooth structure is not necessary, this approach does require the existence of geodesics, so we must restrict attention to spaces for which this notion makes sense.

Definition 60. Let M be a metric space.

(1) Let x and y be two points in M and $d = dist(x,y)$. Then a *(minimal) geodesic* between x and y is an isometry $\gamma : [0,d] \to M$ with $\gamma(0) = x$ and $\gamma(d) = y$.

(2) M is said to be a *length space* (or a *geodesic space*) if every pair of points can be joined by a geodesic.

Motivated by Alexandrov's theorem, a length space M is said to be CAT(0) (C for Comparison or Cartan, A for Alexandrov, and T for Toponogov, who proved related comparison theorems) if the following condition holds: For any three points x, y and z in M, let $\widetilde{x}$, $\widetilde{y}$ and $\widetilde{z}$ be three points in $\mathbb{R}^2$ such that

$$d_M(x,y) = d_{\mathbb{R}^2}(\widetilde{x},\widetilde{y}), \quad d_M(x,z) = d_{\mathbb{R}^2}(\widetilde{x},\widetilde{z}), \quad \text{and} \quad d_M(y,z) = d_{\mathbb{R}^2}(\widetilde{y},\widetilde{z}).$$

Let γ be any minimal geodesic from y to z, and γ' the straight line from $\widetilde{y}$ to $\widetilde{z}$. Let p be any point in γ, and $\widetilde{p}$ the corresponding point in $\widetilde{\gamma}$. Then we require, for all choices of x, y, z, γ and p that

$$d_M(x,p) \le d_{\mathbb{R}^2}(\widetilde{x},\widetilde{p}).$$

Some basic facts about CAT(0) spaces are left as exercises (see the end of this section).

Parts (2) and (3) of Theorem 59 lead to the following definition (which first appeared, along with many far-reaching implications and applications of this idea, in [**48**]).

Definition 61. Say a length space M is nonpositively curved (NP) if the CAT(0) inequality is true for all sufficiently close triples x, p, q.

The basic relationship between these notions is the following.

Theorem 62. *M is CAT (0) if and only if M is simply connected and nonpositively curved.*

We now specialize to a subclass of length spaces, namely polyhedra. A *polyhedron* is defined by a set $\{p_1, p_2, p_3, \dots\}$ of convex polytopes together with a collection of isometric identifications of some faces of the polytopes. Quotienting out by these identifications yields a topological space M which has a given cell decomposition. We require that in the resulting space, if two polytopes meet, they do so along a face of each. In these notes we also require that there be a uniform upper bound on the diameter of the polytopes, and that the resulting cell complex be locally finite.

It is useful now to introduce the notion of the link of a vertex. Let M be a polyhedron, and v a vertex of M. For any ϵ, let $S(v,\epsilon)$ denote the points in M which lie in a polytope incident to v and whose distance from v is precisely ϵ. For ϵ small enough, $S(v,\epsilon)$ is a topological space with an induced cell decomposition which is, up to isomorphism, independent of ϵ. This space, along with its cell decomposition, is defined to be the *link of* v, and is denoted by $link(v)$. If M has the property that the link of every vertex is a combinatorial sphere (i.e. it has a subdivision isomorphic to a subdivision of the boundary complex of a simplex), then M is a manifold. In this case we call M a *piecewise Euclidean manifold* (or PE manifold).

It will be important later that the link of each vertex comes also with a natural geometry. If ϵ is the radius of S, we can normalize the metric on $S \cap M$ by multiplying by $1/\epsilon$, so that the cells of $S \cap M$ are now convex cells from a sphere of radius 1. When we speak of distances and lengths on $link(v)$, it is always with respect to this spherical geometry.

For any polyhedron M there is a natural notion of arc-length for curves in M, induced by the Euclidean structure in each polytope. We can then define the distance between two points in M to be the infimum of the length of the curves connecting the two points.

Theorem 63. *Any polyhedron with arclength and distance defined as above is a length space.*

Thus, using Definition 61, we can speak of a polyhedron having nonpositive curvature. We can now state the first conjecture of Charney and Davis, which is a direct analogue of Hopf's conjecture.

Conjecture 64. *If M is a nonpositively curved compact PE manifold of dimension $2n$, then $(-1)^n\chi(M) \geq 0$.*

(This conjecture, and all of the other conjectures presented in this section first appeared in [**15**], and the reader should certainly consult that paper for a more complete discussion, as well as some initial evidence for the conjectures.) Since the first positive steps towards the Hopf conjecture were proved using the generalized Gauss-Bonnet theorem, it seems reasonable to begin our examination of the Charney-Davis conjecture with a similar approach. Let us begin with the case of a PE surface. Let M be a PE surface, and v a vertex of M. Let

$$k(v) = 1 - \sum_{f>v} angle(f,v), \tag{6}$$

where the sum is over all faces f which contain v, and $angle(f,v) \in [0,1]$ denotes the normalized interior angle of f at v, i.e. the usual angle (in radians) divided by 2π. Then one can check in a straightforward manner the following very classical formula

$$\chi(M) = \sum_v k(v). \tag{7}$$

The relationship between the previous discussion and the current topic is provided by the following lemma.

Lemma 65. *A PE surface M is nonpositively curved if and only if $k(v) \leq 0$ for each vertex v.*

The Charney-Davis conjecture, in the case of PE surfaces, follows immediately.

This discussion was generalized to higher dimensions by Banchoff [**8**]. Let M be a polyhedron and v a vertex of M. Define

$$k(v) = \sum_{i=0}^{n} (-1)^i \sum_{\{\alpha^{(i)} > v\}} [v, \alpha] \tag{8}$$

where $\{\alpha^{(i)} > v\}$ denotes the set of i-dimensional cells of M which contain v, and $[v,\alpha]$ denotes the normalized exterior angle of α at v. That is, $[v,\alpha]$ is the fraction of the sphere consisting of outward pointing normals to supporting hyperplanes of α at v. Equivalently, $[v,\alpha]$ is the fraction of linear functions on α which achieve their maximum at v. Banchoff proved the following generalization of (7).

Theorem 66. *If M is a polyhedron, then*

$$\chi(M) = \sum_v k(v). \tag{9}$$

Recall that the local approach that was sufficient to prove Hopf's theorem in dimensions 2 and 4 is not sufficient in higher directions. Charney and Davis, perhaps somewhat surprisingly, conjecture that the corresponding local approach to their conjecture works in all dimensions.

Conjecture 67. *Let M be a PE manifold of dimension $2n$. If M is nonpositively curved, then for every vertex of M*

$$(-1)^n k(v) \geq 0.$$

The function $k(v)$ can, in a straightforward way, be written in terms of the link of v with its natural geometry as a complex of spherical cells. Let $\widetilde{k}$ denote this function, so that

$$k(v) = \widetilde{k}(\text{link}(v)).$$

The next step is to determine which simplicial complexes can arise as links of vertices in nonpositively curved PE manifolds. Roughly speaking, a Riemannian manifold has nonpositive curvature if and only if the boundary of each small metric ball is larger, in some sense, than the corresponding Euclidean sphere. Something similar is true for PE manifolds. That is, a PE manifold is nonpositively curved if the link is larger, in a precise sense, than a standard sphere of radius 1. More precisely, say that a complex L of spherical cells is *large* if for every pair of points x and y in L, with $dist(x,y) < \pi$, there is a unique geodesic connecting them. The following is a theorem of Gromov [**48**].

Theorem 68. *A PE manifold M is nonpositively curved if and only if the link of every vertex is large.*

The strongest version (from this point of view) of the Conjecture 67 can now be stated.

Conjecture 69. *Let L be a spherical complex (i.e. a cell complex in which each cell has the geometry of a convex cell in a sphere of radius 1) which is homeomorphic to a sphere of dimension $2n-1$. If L is large, then*

$$(-1)^n \widetilde{k}(L) \geq 0.$$

Note that this is a more general conjecture that what is needed, as our original setting required consideration only of combinatorial spheres (a more restrictive class than all topological spheres). However, it does not seem constructive at this point to quibble over such distinctions, since it is not at all clear what the right context is for this conjecture. Moreover, in just a moment we will weaken this hypothesis even further.

For the remainder of these notes, we will restrict attention to the case in which M is a cubical complex (i.e. all of the polytopal cells in M are geometric cubes). In this case, the links have a very special structure, namely every edge of each of the spherical simplices has length $\pi/2$. One can easily see that for each $\alpha^{(i)} > v$, we have $[v, \alpha] = (1/2)^i$. Since each such α corresponds to an $(i-1)$ simplex in the link of v, we have the formula

$$\widetilde{k}(L) = 1 + \sum_{i=0}^{\dim(M)-1} \left(\frac{-1}{2}\right)^{i+1} f_i(L), \tag{10}$$

where $f_i(L)$ denotes the number of i-simplices in L. Gromov showed that there is a simple combinatorial test for whether such a link is large.

Definition 70. Say that a simplicial complex L is *flag* if every clique spans a simplex. That is, is $v_1, v_2, \ldots, v_k$ are vertices in L, and they are all pairwise adjacent, then they span a simplex.

Theorem 71. *A cubical PE manifold is nonpositively curved if and only if the link of every vertex is flag.*

Thus, in this case, Conjecture 69 implies the following statement.

Conjecture 72. *Let L be a simplicial complex which is homeomorphic to a sphere of dimension $2n-1$. If L is flag, then $(-1)^n \widetilde{k}(L) \geq 0$, where $\widetilde{k}(L)$ is given by the formula (10*

This conjecture is very combinatorial in nature, but still has one topological, noncombinatorial, ingredient, namely the hypothesis that L be homeomorphic to a sphere. There is a natural generalization of triangulated spheres which has a more combinatorial flavor. A *Gorenstein** complex (or *a generalized homology sphere*) is a simplicial complex with the property that, for every $p \geq 0$, the link of every p-simplex has the homology of an $(n-p-1)$-sphere. If L is a simplicial complex which is homeomorphic to an n-sphere, or, more generally, any homology n-sphere, then L is Gorenstein*. Thus, to place these ideas in a more combinatorial setting, it is natural to consider the following generalization of Conjecture 72.

Conjecture 73. *Let L be a $(2n-1)$-dimensional Gorenstein* complex. If, in addition, L is flag, then $(-1)^n\widetilde{k}(L) \geq 0$, where $\widetilde{k}(L)$ is given by the formula (10).*

It is not at all surprising that this conjecture, which is stated completely in combinatorial terms, has received the most attention from the combinatorics community. We will discuss the connections with other combinatorial notions in Section 3. For now we note that this conjecture, or Conjecture 72, implies Conjecture 64 for cubical complexes. It is a simple, but still rather surprising, fact that the converse is true. The following result is due to Babson-Billera-Chan [**5**] and Bridson-Haefliger [**13**].

Theorem 74. *Let L be any simplicial complex. Then there is a finite cubical polyhedron M with the property that the link of every vertex of M is isomorphic to L.*

Proof. Let V denote the vertex set of L. Consider the cube $C = [0,1]^V \subset \mathbb{R}^V$ endowed with its standard cubical cell decomposition. We will find M as a subcomplex of C. For every simplex σ of L, let $\mathbb{R}^\sigma$ denote the linear subspace of $\mathbb{R}^V$ traced out by varying the coordinates corresponding to vertices in σ. Let M be the union of all faces of C which are parallel to some $\mathbb{R}^\sigma$, for some simplex σ of L. Then the vertices of M are the vertices of C, and the link of every vertex is isomorphic to L. □

Applying this result to the case when L is a combinatorial sphere yields a cubical PE manifold M. If, in addition, L is flag, it follows from Theorem 71 that M is nonpositively curved.

Corollary 75. *The Charney-Davis Conjecture 64 is true for cubical nonpositively curved PE manifolds if and only if Conjecture 72 is true for combinatorial spheres.*

2. Exercises for Lecture 4

(1) Prove that if M is a CAT(0) space then
  (a) For any points p and $q \in M$, there is a unique minimal geodesic from p to q.
  (b) M is contractible.

(2) Prove the formula (7) for any triangulated surface M.

(3) Prove the formula (9) for any finite polyhedron M.

(4) Show that the formula (8) specializes to (6) in the case of a triangulated surface.

(5) (a) Show that the barycentric subdivision of any polyhedron (in fact any CW complex) is flag.
  (b) Show that the join of any two flag simplicial complexes is flag.
  (c) Show that if L is a flag simplicial complex and v is any vertex of L, then link(v) and star(v) are both flag.

(6) Let L_1 and L_2 be simplicial spheres of dimension $2n_1 - 1$ and $2n_2 - 1$, respectively, which satisfy the conjectured inequality $(-1)^{n_i}\widetilde{k}(L_i) \geq 0$ and are flag. Show that $L_1 * L_2$, the join of L_1 and L_2, also satisfies the conjectured inequality.

LECTURE 5
From Analysis to Combinatorics

1. Hodge Theory and the Hopf-Charney-Davis Conjectures

In this section we present an overview of some of the beautiful analytical ideas that have been used to study the Hopf-Charney-Davis conjectures. Hodge theory is one of the standard ways of investigating the homological implications of geometric information, so it should not be too surprising that it has played a central role in this subject. To date, the most substantial partial results towards the Hopf and Charney-Davis conjectures have been proved using some variation of the Hodge theoretic approach we present here. To avoid some technical details, we will present the ideas in the combinatorial category. However, everything in this section can, with suitable care, be applied in the Riemannian setting.

Let X be a finite CW complex, and let $C^p(X)$ denote the space of real-valued p-cochains on X. Let

$$d_p : C^p(X) \to C^{p+1}(X)$$

denote the usual coboundary operator. Then $d^2 = 0$, and the singular cohomology of X, $H^*(X, \mathbb{R})$, is isomorphic to the cohomology of the cochain complex

$$C^*(X) : 0 \longrightarrow C^0(X) \xrightarrow{d_0} C^1(X) \xrightarrow{d_1} C^2(X) \longrightarrow \cdots$$

That is

$$H^p(X, \mathbb{R}) \cong \frac{\operatorname{Ker} d_p}{\operatorname{Im} d_{p-1}}.$$

Now endow each $C^p(X)$ with a (positive definite) inner product by declaring the canonical basis to be orthonormal. More explicitly, for each p, let $S_p(X)$ denote the set of p-cells in X. Choose an orientation for each element in $S_p(X)$. Then for α and β in $C^p(X)$ set

$$\langle \alpha, \beta \rangle = \sum_{y \in S_p(x)} \alpha(y)\beta(y). \tag{11}$$

Note that this inner product is independent of the chosen orientations.

Let d_p^* denote the adjoint of the operator d_p. That is,

$$d_p^* : C^{p+1}(X) \longrightarrow C^p(X)$$

is the unique map satisfying

$$\langle d_p\alpha, \beta\rangle = \langle \alpha, d_p^*\beta\rangle$$

for every p-cochain α and $(p+1)$-cochain β.

Define, for each p, the (combinatorial) Laplace operator

$$\Box^p := d_{p-1}d_{p-1}^* + d_p^* d_p : C^p(X) \longrightarrow C^p(X).$$

There is much that can be said about this fascinating operator but for us, the main point is the following.

Theorem 76. *For each p*

$$\mathrm{Ker}(\Box^p) \cong H^p(X).$$

Proof. Using basic linear algebra, and the fact that $\mathrm{Ker}(d_p) \supset \mathrm{Im}(d_{p-1})$, we have that

$$\begin{aligned} H^p(X) &\cong \frac{\mathrm{Ker}(d_p)}{\mathrm{Im}(d_{p-1})} \\ &\cong \mathrm{Ker}\, d_p \cap (\mathrm{Im}\, d_{p-1})^\perp \\ &\cong \mathrm{Ker}\, d_p \cap \mathrm{Ker}\, d_{p-1}^* \\ &= \mathrm{Ker}(\Box^p). \end{aligned}$$

The last equality follows from the observation that if $\alpha \in \mathrm{Ker}(\Box^p)$, then

$$0 = \langle \Box^p\alpha, \alpha\rangle = |d_{p-1}^*\alpha|^2 + |d_p\alpha|^2,$$

so that

$$\alpha \in \mathrm{Ker}\, d_p \cap \mathrm{Ker}\, d_{p-1}^*.$$

□

Cochains in the kernel of $\Box^p$ are called *harmonic*. We will denote the space of harmonic p-cochains in X by $\mathcal{H}^p(X)$.

So far, in this section, we have been considering the case of a finite CW complex. How do things change in the case of an infinite complex? Of particular interest to us is the case of the universal cover of a finite complex. With that in mind, let Y denote an infinite CW complex that is the covering space of some finite complex. Let us take a look at Hodge theory on Y. Let

$$C^*(Y) : 0 \longrightarrow C^0(Y) \xrightarrow{d_0} C^1(Y) \xrightarrow{d_1} C^2(Y) \longrightarrow \cdots$$

denote the cochain complex of Y. Hodge theory requires inner products. We quickly realize that the standard formula (11) does not yield a well-defined inner product in the infinite setting. There are various possible ways to proceed. However, if one desires to work with Hilbert spaces, there is a natural choice. Let $C_2^p(Y)$ denote the L^2 p -cochains on Y. That is, if $S_p(Y)$ denotes the set of p-cells in Y, each endowed with an orientation, then

$$C_2^p(Y) = \{\alpha \in C^p(Y) \text{ s.t. } \sum_{y\in S_p(Y)} (\alpha(y))^2 < \infty\}.$$

Then $C_2^p(Y)$ is a Hilbert space with respect to the inner product

$$\langle \alpha, \beta\rangle = \sum_{y\in S_p(y)} \alpha(y)\beta(y).$$

The next step is to replace the standard cochain complex on Y by the complex of L^2 cochains. To do this, one requires the following lemma.

Lemma 77. $d_p(C_2^p(Y)) \subset C_2^{p+1}(Y)$.

The proof is left as an exercise. (See Exercise 1 at the end of this lecture.)

Now consider the L^2 cochain complex

$$C_2^*(Y) : 0 \longrightarrow C_2^0(Y) \xrightarrow{d_0} C_2^1(Y) \xrightarrow{d_1} C_2^2(Y) \longrightarrow \cdots .$$

One might wish to proceed by defining the L^2-cohomology of Y by the usual formula

$$\frac{\mathrm{Ker}(d_p : C_2^p(Y) \to C_2^{p+1}(Y))}{\mathrm{Im}(d_{p-1} : C_2^{p-1}(Y) \to C_2^p(Y))}$$

This is certainly possible (this is called the unreduced L^2 cohomology). However, it does lead to certain difficulties, since $\mathrm{Im}\, d_{p-1}$ need not be a closed subspace of $\mathrm{Ker}\, d_p$, and hence the quotient need not inherit the structure of a Hilbert space. With that in mind, we define the L^2 cohomology of Y to be

$$H_2^p(Y) := \frac{\mathrm{Ker}(d_p)}{\overline{\mathrm{Im}}(d_{p-1})}$$

where $\overline{\mathrm{Im}}(d_{p-1})$ indicates that we take the closure of $\mathrm{Im}(d_{p-1})$ in $\mathrm{Ker}(d_p)$. (This quotient is sometimes called the *reduced* L^2-cohomology.)

Following the same procedure as before, we can construct the adjoint operator $d^* C_2^p(Y) \to C_2^{p-1}(Y)$ (see exercise 1 at the end of this section), and the Laplace operator

$$\Box_2^p : C_2^p(Y) \longrightarrow C_2^p(Y).$$

Let $\mathcal{H}_2^p(Y)$ denote the kernel of the operator $\Box_2^p$. The proof of Theorem 76, applied in this setting, yields the following result.

Theorem 78. $\mathcal{H}_2^p(Y) \cong H_2^p(Y)$.

Now let X be a finite CW complex. We define the p^{th} Betti number of X to be the dimension of $H^p(X)$, and denote this number by $b^p(X)$. From Theorem 76 we know that

$$b^p(X) = \dim \mathcal{H}^p(X).$$

Let $\pi^p : C^p(x) \to \mathcal{H}^p(x)$ denote the orthogonal projection. Then we can also write

$$b^p(X) = \mathrm{trace}(\pi^p).$$

A useful way to calculate the trace of an operator is to express the operator as a matrix with respect to some basis, and then to take the sum of the diagonal elements of the matrix. Let us carry out this procedure here, and represent π^p as a matrix with respect to the standard basis for the cochain space. Let $\alpha_1, \ldots, \alpha_k$ denote an orthogonal basis for $\mathcal{H}^p(X)$ (so that $k = b^p(X)$). Then we can write

$$\pi^p = \sum_{i=1}^k \alpha_i \otimes \alpha_i^*$$

where $\alpha_i^* : C^p(X) \to \mathbb{R}$ is the map that takes β to $\langle \beta, \alpha_i \rangle$. The function

$$K^p : S_p(X) \times S_p(X) \longrightarrow \mathbb{R}$$

given by

$$K^p(x,y) = \sum_{i=1}^{k} \alpha_i(x)\alpha_i(y)$$

is a matrix for the operator π^p, in the sense that for any $\beta \in C^p(X)$, and any $x \in S_p(X)$,

$$[\pi^p(\beta)](x) = \sum_{i=1}^{k} \alpha_i(x)\langle\beta,\alpha_i\rangle = \sum_{i=1}^{k} \alpha_i(x)\left(\sum_{y\in S_p(X)} \alpha_i(y)\beta(y)\right)$$
$$= \sum_{y}\left(\sum_{i=1}^{k} \alpha_i(x)\alpha_i(y)\right)\beta(y) = \sum_{y} K^p(x,y)\beta(y).$$

It follows that

$$b^p(X) = \text{trace}(\pi^p) = \sum_{x\in S_p(X)} K^p(x,x)$$

(This identity can easily be proved directly, without the preceeding discussion.) The function

$$K^p(x,x) = \sum_{i=1}^{k} \alpha_i^2(x)$$

is sometimes called the p^{th} *local Betti number* (as its integral gives the p^{th} Betti number). We chose an orthonormal basis for the space of harmonic cochains on order to define this function, but one can easily check that it is independent of the choices.

Now let us consider again the case of an infinite CW complex Y which has an action by a group G, such that Y/G is finite. In this case, if $\mathcal{H}_2^p(Y) \neq 0$, then it is necessarily infinite-dimensional. Still, much of the previous discussion makes sense in this setting. That is, one can define the orthogonal projection

$$\pi^p : C_2^p(Y) \to \mathcal{H}_2^p(Y).$$

While the trace of this operator is not defined, we can still construct a kernel $K^p(x,y)$, $x, y \in S_p(Y)$, given by

$$K^p(x,y) = \sum_{i} \alpha_i(x)\alpha_i(y)$$

where $\{\alpha_i\}$ is an orthonormal basis for $\mathcal{H}^p(Y)$. Just as for the finite complex, we can restrict this operator to the diagonal and consider the p^{th} local Betti number $K^p(x,x) = \sum_i \alpha_i^2(x)$, $x \in S_p(Y)$. Summing these entries, however, yields an infinite result. At this point, we use the extra information we have, namely the fact that everything is invariant under the action of the group G. Let $S_p^*(Y) \subset S_p(Y)$ denote a set of p-cells of Y containing exactly one p-cell from each G-orbit in $S_p(Y)$. Then $S_p^*(Y)$ is a finite set, and the values $K^p(x,x)$ for $x \in S_p^*(Y)$ determine $K^p(x,x)$ for all x.

With that in mind, define the *G-trace* of π^p, denoted by $\tau_G(\pi^p)$ to be the result of summing $K^p(x,x)$ over $x \in S_p^*(Y)$. That is

$$\tau_G(\pi^p) = \sum_{x\in S_p^*(Y)} K^p(x,x) \in [0,\infty).$$

(The G-trace, denoted by τ_G, of any G-equivariant operator on $C_2^p(Y)$ can be defined by the same procedure.) We will denote the number $\tau_G(\pi^p)$ by $b_G^p(Y)$, and call it the p^{th} *L^2-Betti number of Y*. One simple, but essential observation, is that $b_G^p(Y) = 0$ if and only if $\mathcal{H}_2^p(Y) = 0$. The definition of $b_G^p(Y)$ may seem a bit *ad hoc*. However, the reader can consult [**4**] to see how this definition results from natural notions in the study of operator algebras. In that context, $b_G^p(Y)$ is the von Neumann trace of the operator $\pi^p : C_2^p(Y) \to \mathcal{H}_2^p(Y)$.

Now let us restrict attention to the case in which $\widetilde{X}$ is the universal cover of a finite CW complex X, and we take the group G to be the fundamental group of X, acting freely on $\widetilde{X}$ in the usual way. In this setting, Dodziuk proved that the L^2-Betti numbers of $\widetilde{X}$ computed combinatorially from a cell decomposition are equal to those calculated from the Riemannian Laplacian, and that these numbers are homotopy invariants of X [**23**]. For our purposes, the main property of the L^2-Betti numbers of $\widetilde{X}$ is the following result of Atiyah [**4**].

Theorem 79. *Let $\widetilde{X}$ be the universal cover of a finite CW complex X, and take the group G to be the fundamental group of X. Then*

$$\chi(X) = \sum_i (-1)^i b_G^p(\widetilde{X}). \tag{12}$$

Thus, L^2-Betti numbers are another tool at our disposal for investigating the Euler characteristic. It may not be clear how one could use this new information to investigate the Hopf-Charney-Davis conjectures, but a link is provided by the following beautiful conjecture of Singer.

Conjecture 80. *Let X be a compact aspherical n-manifold, and $\widetilde{X}$ the universal cover of X. Then for all $p \neq n/2$, $\mathcal{H}_2^p(\widetilde{X}) = 0$.*

Applying Theorem 79, we see that Singer's conjecture immediately implies the Hopf conjecture (as generalized by Thurston, Conjecture 58). While the Hopf conjecture is trivial for odd dimensional manifolds, Singer's conjecture is not. Singer's conjecture can quite easily be shown to be true for surfaces (it follows from the fact that there are no L^2 harmonic functions on a complete Riemannian manifold of infinite volume). It has also been shown to hold for 3-manifolds for which Thurston's geometrization conjecture is true [**66**], locally symmetric spaces [**24**], negatively curved Kähler manifolds [**49**], manifolds with sufficiently pinched negative curvature [**26**], aspherical manifolds whose fundamental group contains an infinite amenable normal subgroup [**16**], and manifolds which fiber over S^1 [**67**].

It is quite natural to guess that Singer's conjecture also holds for suitable piecewise Euclidean manifolds. The following conjecture, along with a series of related conjectures, appears in Section 8 of [**22**].

Conjecture 81. *Let X be a compact nonpositively curved PE manifold of dimension n, and let $\widetilde{X}$ be the universal cover of X. Then for all $p \neq n/2$, $\mathcal{H}_2^p(\widetilde{X}) = 0$.*

This conjecture implies the Charney-Davis conjecture 64. In [**22**], Davis and Okun used this circle of ideas to establish Conjecture 72 for 3-dimensional flag simplicial spheres (and somewhat more generally).

Theorem 82. *Let L be any flag simplicial 3-sphere, then $\widetilde{k}(L) \geq 0$, where $\widetilde{k}$ is as in (10).*

Very roughly speaking, for any flag simplicial 3-sphere, a special nonpositively curved 4-dimensional cubical PE manifold M is constructed (using the structure of right-angled Coxeter groups) which has the properties that the link of every vertex is identified with L, and Singer's conjecture can be shown to hold for M. This is a wonderful result, requiring a lot of hard work, and Davis and Okun introduce some powerful new ideas into the subject. The reader is strongly encouraged to consult their paper. Following the discussion in Section 1, Theorem 82 implies the following general result.

Theorem 83. *If X is any finite nonpositively curved cubical PE manifold of dimension 4, then $\chi(X) \geq 0$.*

2. The Charney-Davis Conjecture and the h-vector

In this section we focus attention on Conjectures 72 and 73, and show that they reside quite naturally in the well-developed circle of ideas surrounding the investigation of f-vectors of simplicial complexes. In (10) the formula for the relevant function $\widetilde{k}(L)$ is given in terms of the f-numbers of L. In a number of settings, especially those related to commutative algebra and toric varieties, it has proved very useful to study certain special linear combinations of the f-numbers, called the h-numbers. In [**15**] Charney and Davis observe that Conjectures 72 and 73 can be restated in a very nice way in terms of the h-numbers. For any finite n-dimensional simplicial complex K, define the f-polynomial of K to be the generating function of the f-numbers. More explicitly, set $f(K,t) = \sum_{i=0}^{n+1} f_{i-1}(K)t^i$,where we define f_{-1} to be 1. Define the h-polynomial of K, $h(K,t) = \sum_{i=0}^{n+1} h_i(K)t^i$, by the formula

$$h(K,t) = (1-t)^{n+1} f(K, \frac{t}{t-1}). \tag{13}$$

(We will often write $h(t)$ for $h(K,t)$ if it will not cause any confusion.) It follows immediately from (10) and (13) that for any n-dimensional simplicial complex K,

$$h(K,-1) = 2^{n+1}\widetilde{k}(K).$$

Hence, we can now restate Conjecture 73 as

Conjecture 84. *If K is any simplicial Gorenstein* $(2n-1)$-complex which is flag, then $(-1)^n h(K,-1) \geq 0$.*

In this form, the conjecture can more easily be compared to other conjectures and results concerning the h-vectors of simplicial spheres and related spaces. One advantage of the h-polynomial is that it is quite easy to state the Dehn-Somerville relations. Say that K is *Eulerian* if the link of every i-simplex, $i \geq -1$, has the same Euler characteristic as a sphere of dimension $n-i-1$, so that, in particular, every Gorenstein* complex, and thus every triangulated sphere, is Eulerian. If K is Eulerian, then

$$h(t) = t^{n+1} h(t^{-1}), \tag{14}$$

(equivalently, $h_i = h_{n+1-i}$ for each i).

An n-dimensional simplicial complex is said to be *Cohen-Macaulay* if for every i the link of every i-simplex has nonzero reduced homology only in dimension $n-i-1$. So, for example, every Gorenstein* complex, and thus every triangulated sphere, is Cohen-Macaulay. It can be shown using algebraic methods that if K is Cohen-Macaulay, then all of the coefficients of the h-polynomial are nonnegative (see Corollary II.3.2 of [**90**]).

To every rational polytope P that is simple (i.e. its boundary complex is dual to a simplicial complex – and every such polytope can be perturbed slightly to be rational), one can associate a toric variety X_P in a natural way (e.g. see [**42**] [**19**]). Danilov [**19**] showed that for each i, $b_{2i}(X_P)$, the $2i^{th}$ Betti number of X_P is equal to $h_i(P)$. This, and related indentities, has proved to be an immensely powerful source of information about the h-polynomial, as well as a new inspiration for the study of toric varieties.

For example, in [**61**] Reiner and Leung proved that if K is the boundary complex of a simple $2n$-polytope, then $h(K, -1)$ is equal to the signature of an associated toric variety. They were able to then show, using the Hirzebruch signature formula that if such a complex K satisfies a certain local convexity property (which is stronger than flag) then $(-1)^n h(K, -1) \geq 0$. Probably the most striking application of toric techniques is Stanley's proof of Theorem 87, stated below.

From another direction, special tools are available for simplicial complexes which arise as the order complex of a poset. For example, the following result is due to Babson.

Theorem 85. *If K is the the boundary complex of a simplicial $2n$-polytope, and is the order complex of a poset, then $(-1)^n h(K, -1) \geq 0$.*

Note that (i) the order complex of a poset is always flag and (ii) the barycentric subdivision of any simplicial complex is the order complex of a poset. Therefore this theorem implies the Charney-Davis conjecture for the barycentric subdivision of the boundary complex of any $2n$-dimensional simplicial polytope.

We can give an outline of the proof. For Eulerian n-dimensional order complexes, it has proved very useful to introduce a refinement of the h-vector, known as the cd-index, a homogeneous polynomial of degree n, $\Phi_K(c, d)$ in two noncommuting variables (where d is considered to have degree 2). For any $(2n-1)$-dimensional Eulerian complex K which is the order complex of a poset, Babson observed that we have the relationship

$$\Phi_K(0, -2) = h(K, -1),$$

and $\Phi_K(0, -2)$ is equal to $(-2)^n$ times the coefficient of $d^n$of a polytope, then all coefficients of the cd-index are nonnegative [**89**], so the result follows.

Stanley has made the following conjecture.

Conjecture 86. *If K is a Gorenstein* order complex, then every coefficient of the cd-index is nonegative.*

By Babson's argument, this would imply Conjecture 73 for any Gorenstein* complex which is the order complex of a poset. We note that the barycentric subdivision of any Gorenstein* complex is Gorenstein* and is the order complex of a poset.[1]

The Charney-Davis conjecture is related to some of the central conjectures in the subject. For example, consider the Generalized Lower Bound Theorem (GLBT) of Stanley (proved with toric methods [**88**], and then later reproved by McMullen using only ingredients from convex geometry [**69**][**70**]) .

[1]After this lecture was delivered, the preprint by K. Karu, *The cd-index of fans and lattices*, math.AG/0410513, appeared in which Karu claims to prove this conjecture. See also his follow-up preprint *Lefschetz decomposition and the cd-index of fans*, math.AG/0509220.

Theorem 87. *If K is an n-dimensional simplicial complex which is the boundary of an $(n+1)$-polytope, then the h-numbers of K are unimodal. That is*

$$h_0(K) \leq h_1(K) \leq \cdots \leq h_{\lfloor \frac{n+1}{2} \rfloor}(K).$$

These are the most general linear inequalities satisfied by the f-vectors of simplicial polytopes. One of the central open problems in the study of f-vectors is to determine precisely for which simplicial complexes the conclusion of the GLBT holds. For example, does it hold for all Gorenstein* complexes (this is sometimes called the Generalized Lower Bound Conjecture), or all triangulated spheres, or all PL spheres? Independently, Kalai and Stanley have shown that the conclusion holds for the boundary complex of any triangulated $(n+1)$-ball which appears as a subcomplex of a simplicial $(n+1)$-polytope, but it is not clear which spheres arise in this way. This unimodality property has some relation to the Charney-Davis conjecture.

Theorem 88. *Let K be a Gorenstein* complex of dimension $2n-1$. Suppose that $h(K,t)$ only real roots. Then the following two conclusions hold.*

(1) the h-numbers of K are unimodal, i.e. the conclusion of the GLBT holds;
(2) $(-1)^n h(K,-1) \geq 0$.

While we have stated this result in terms of h-polynomials of simplicial complexes, this theorem is really just a statement about polynomials with real coefficients satisfying a symmetry relation as in (14). Part two of this theorem is due to Charney and Davis (see Lemma 7.5 of [**15**]). In fact, they prove the stronger statement that the conclusion holds as long as $h(K,t)$ has no nonreal roots of modulus 1. The first part of this theorem is due to Isaac Newton!

With this theorem in mind, it is natural to make the following *Real Root Conjecture* (apparently due originally to Januzkiewicz, see [**20**]).

Conjecture 89. *For any Gorenstein* complex K which is flag, $h(K,t)$ has only real roots.*

In [**75**], Reiner and Welker consider these questions for the order complex of a graded poset P. This special case of the real root conjecture was formulated earlier, and is known as the Neggers-Stanley conjecture. Without proving the Neggers-Stanley conjecture, they are able to prove the implications of this conjecture. More precisely, they construct a simplicial polytope with the same h-polynomial as the order complex of P, and thus the unimodality of the h numbers of the order complex follows from Theorem 87. By other means (using the results of [**61**]) they establish the Charney-Davis conjecture for K_P for graded posets of width 2.

More recently, Brändén [**12**] has proved the Charney-Davis conjecture for K_P, as well as the unimodality of the h-numbers, for *any* graded poset P. That is, he establishes both conclusions of Theorem 88 for such complexes. He does not do this by proving that the h-polynomial has real roots, however. Let us take a moment to discuss Brändén's approach, an approach that was also presented, independently, in the recent work of Gal [**43**]. We know from the Dehn-Sommerville relations (14) that the h-polynomial of any Eulerian complex of dimension $2n-1$ satisfies the symmetry

$$h(t) = t^{2n} h(t^{-1}).$$

The polynomials

$$p_i(t) = t^i(1+t)^{2n-2i}, \quad i = 0, 1, 2, \ldots, n$$

form a basis for the vector space of polynomials of degree $2n$ with this symmetry. Thus, for any Eulerian complex K of dimension $2n$ we can write

$$h(K,t) = \sum_{i=0}^{n} a_i(K)p_i(t). \tag{15}$$

and the $a_i(K)$'s are uniquely determined by this identity. We can make two simple observations. First, we see that

$$1 = h(K,0) = a_0(K).$$

Second, we see that

$$h(K,-1) = (-1)^n a_n(K). \tag{16}$$

Both Brändén and Gal make the following observation.

Theorem 90. *Let K be an Eulerian complex of dimension $2n-1$. Suppose that $a_i(K) \geq 0$ for each $i = 0, 1, \ldots, n$. Then*
(i) the h-numbers of K are unimodal (i.e. the conclusion of the GLBT holds), and
(ii) $(-1)^n h(K,-1) \geq 0$.

Note that (i) follows immediately from (15) and the observation that the coefficients of each of the p_i's is unimodal, and (ii) from (16).

In [**12**], Brändén establishes the nonnegativity of the $a_i(K)$'s for the order complex of any graded poset, and hence he established the unimodality of the h-polynomial and the Charney-Davis conjecture for such complexes. In [**43**] Gal filled in more of the picture. In particular, he proved

Theorem 91. *(i) The real root conjecture is true for spheres up to dimension 4. (This follows from the results of Davis and Okun.)*

(ii) The real root conjecture is false in all dimensions ≥ 5, and counterexamples are found among boundary complexes of simplicial polytopes.

This certainly seems to put an end to the idea of using Theorem 88 to prove both the Charney-Davis conjecture, and the Generalized Lower Bound Conjectures. In [**43**], Gal turns his attention to Theorem 90 and (in a slightly different notation) makes the following conjecture.

Conjecture 92. *If K is a Gorenstein* complex which is flag, then for all i, $a_i(K) \geq 0$.*

From Theorem 90 we see that this conjecture implies Conjecture 73. Moreover, Gal shows that his conjecture is weaker than the real root conjecture.

Theorem 93. *Let K be a Gorenstein* complex of dimension $2n-1$. Suppose that $h(K,t)$ has only real roots. Then for all i, $a_i(K) \geq 0$.*

Proof. Let $\gamma(K,t)$ denote the generating function of the a_i's. That is

$$\gamma(K,t) = \sum_{i=0}^{n} a_i(K)t^i.$$

One can easily check the following relation.

$$(1+t)^{2n}\gamma(K, \frac{t}{(1+t)^2}) = h(K,t). \tag{17}$$

Assume that all of the roots of $h(K,t)$ are real. Since the coefficients of $h(K,t)$ are all nonnegative and $h(K,0) = 1$, the roots must all be negative. It follows from (17) that the roots of $\gamma(K,t)$ occur either when $t = -1$ or at places when $t/(1+t)^2$ is real and negative, which implies that t is real and negative. Since the roots are real, we can apply Isaac Newton's result Theorem 88 (1) to deduce that the $a_i(K)$'s are unimodal. We also know that $\gamma(K,0) = a_0(K) = 1$. Since all of the roots are negative, it follows that the coefficient $a_n(K)$ of the highest order term is ≥ 0, and now, by the unimodality, it follows that all of the coefficients are ≥ 0. □

As evidence for Conjecture 92, Gal shows that for a flag Gorenstein* complex, the coefficient of t in γ is nonnegative. This is equivalent to the statement that a flag complex simplicial sphere of dimension n has at least $2(n+1)$ vertices (with equality only for the boundary complex of a cross polytope). The other coefficients remain quite mysterious. Much work remains to be done, to discover the geometric/topological meaning behind these numbers, and to begin to assess the truth of this new conjecture. At present, based on the evidence presented here, there is every reason to believe that with Conjecture 92 we have found the right formulation of the problem.

3. Exercises for Lecture 5

Exercises for Section 1 of Lecture 5.

(1) Let Y be an infinite CW complex which is a covering space of a finite CW complex.
   (a) Show that $d(C_2^p(Y)) \subset C_2^{p+1}(Y)$, and, moreover, that d is a bounded operator. That is, there is a constant c so that for any $\alpha \in C_2^p(Y)$, $|d\alpha| \leq c|d\alpha|$.
   (b) Prove that one can define an adjoint operator $d^* : C_2^{p+1}(Y) \to C_2^p(Y)$, and that this operator is also bounded.

(2) Let X be a finite polyhedron. Let c_d denote the number of d-cells for each d. The goal in this problem is to prove that

$$\sum_p (-1)^p \dim \mathcal{H}^p(M) = \sum_p (-1)^p c_p.$$

Of course, using Theorem 76 this is just the standard formula for the Euler characteristic, but we have something else in mind. For $t \in [0,\infty)$, let

$$I(t) = \sum_p (-1)^p \text{trace}(e^{-t\,\square^p}),$$

where $\square^p : C^p(X) \to C^p(X)$ is the Laplace operator, and $e^{-t\,\square^p}$ is defined by a power series expansion.

[The operator $e^{-t\,\Box^p}$ is the unique solution to the differential equation on the space of operators on $C^p(X)$

$$\frac{\partial}{\partial t}L(t) = -\Box^p L(t), \;\; L(0) = I,$$

where I denotes the identity map on $C^p(X)$, and this characterization, can be used to give an alternate definition for $e^{-t\,\Box^p}$.]

(a) Show that $I(0) = \sum_p (-1)^p c_p$.

(b) Show that $\lim_{t\to\infty} I(t) = \sum_p (-1)^p \dim \mathcal{H}^p(M)$.

(c) Show that $d/dt\; I(t) = 0$ for all $t \in [0, \infty)$.

(3) Let us use the approach from exercise 2 to prove Theorem 79. In this case, let X be a finite polyhedron, Y the universal cover, and G the fundamental group of X. Define

$$I(t) = \sum_p (-1)^p \tau_G(e^{-t\,\Box^p}),$$

where $\Box^p : C^p_2(Y) \to C^p_2(Y)$ is the Laplace operator on Y. Show that

(a) $I(0)$ is the left hand side of (12) and $\lim_{t\to\infty} I(t)$ is the right hand side of (12).

(b) $d/dt\; I(t) = 0$ for all $t \in [0, \infty)$.

(4) Show that if Y is an infinite polyhedron, then $\mathcal{H}^0_2(\widetilde{Y}) = 0$.

(5) Let Y denote the real line given the structure of an infinite polyhedron by placing a vertex at each integer point. What is the (reduced) L^2-cohomology of Y? What is the unreduced L^2-cohomology of Y?

Exercises for Section 2 of Lecture 5.

(1) The best exercise is to calculate $f(t)$, $h(t)$, and $\gamma(t)$ for your favorite Eulerian complexes. Start with simple complexes, and then keep going.

(2) If you have never done this before: Prove the Dehn-Somerville relations (14) for any Eulerian complex.

(3) Prove identity (17).

(4) Find explicit formulas for the first few coefficients of $\gamma(K, t)$ in terms of the f-vector of K.

(5) Show that the coefficient of t in $\gamma(K, t)$ is always ≥ 0 for a Gorenstein* complex that is flag.

BIBLIOGRAPHY

1. C. Allendoerfer, *The Euler number of a Riemannian manifold*, Amer. J. Math., **62** (1940), pp. 243–248.
2. C. Allendoerfer and A. Weil, *The Gauss-Bonnet theorem for Riemannian polyhedra*, Trans. Amer. Math. Soc., **53** (1943), pp. 101–129.
3. A. Alexandrov, *A theorem on triangles in a metric space and some of its applications*, Trudy Math. Inst. Stekl., **38** (1951), pp. 5–23.
4. M. Atiyah, *Elliptic operators, discrete groups and von Neumann algebras*, Astérisque, **32-33** (1976), pp. 43–72.
5. E. Babson, L. Billera and C. Chan, *Neighborly cubical spheres and a cubical lower bound conjecture,* Israel J. Math., **102** (1997), pp. 297–315.
6. E. Babson, A. Björner, S. Linusson, J. Shareshian and V. Welker, *Complexes of not i-connected graphs*, Topology, **38** (1999), pp. 271–299.
7. E. Babson, P. Hersh, *Discrete Morse functions from lexicographic orders*, Trans. Amer. Math. Soc., **357** (2005), pp. 509–534.
8. T. Banchoff, *Critical points and curvature for embedded polyhedra,* J. Differential Geom., **1** (1967), pp. 245–356.
9. E. Batzies and V. Welker, *Discrete Morse theory for cellular resolutions*, J. Reine Angew. Math., **543** (2002), pp. 147–168.
10. L. Billera, S. Holmes and K. Vogtmann, *Geometry of the space of phylogenic trees*, Adv. in Appl. Math., **29** (2002), pp. 733–767.
11. A. Björner, *Topological methods*, in Handbook of combinatorics, vol. 2, R. Graham, M. Grötschel, and L. Lovász eds., North-Holland/Elsevier, Amsterdam, 1995, pp. 1819–1872.
12. P. Brändén, *Sign-graded posets, unimodality of W-polynomials and the Charney-Davis conjecture*, Electron. J. Combin., **110** (2005), pp. 127–145.
13. M. Bridson and A. Haefliger, *Metric spaces of nonpositive curvature*, Grundlehren der Math. Wiss. 319, Springer-Verlag, Berlin, 1999.
14. M. Chari, *On discrete Morse functions and combinatorial decompositions*, Discrete Math, **217** (2000), pp. 101–113.
15. R. Charney and M. Davis, *The Euler characteristic of a nonpositively curved, piecewise Euclidean manifold*, Pacific J. Math., **171** (1995), pp. 117–137.
16. J. Cheeger and M. Gromov, *L_2-cohomology and group cohomology*, Topology, **25** (1986), pp. 189–215.

17. S. Chern, *A simple intrinsic proof of the Gauss-Bonnet formula for closed Riemannian manifolds,* Ann. of Math. (2), **45**(1944), pp. 747–752.
18. S. Chern, *On curvature and characteristic classes of a Riemannian manifold*, Abh. Math. Sem. Univ. Hamburg, **20** (1956), pp. 117–126.
19. V. Danilov, *The geometry of toric varieties*, Russian Math. Surveys, **33** (1978), pp. 97–154.
20. M. Davis, J. Dyrnara, T. Januszkiewicz and B. Okun, *Decompositions of Hecke-von Neumann modules and the L^2-cohomology of buildings*, preprint (2004).
21. M. Davis, *Nonpositive curvature and reflection groups*, in Handbook of Geometric Topology, R. Daverman, R. Sher (eds.), Elsevier, Amsterdam, 2002, pp. 373–422.
22. M. Davis and B. Okun, *Vanishing theorems and conjectures for the ℓ^2-homology of right-angled Coxeter groups*, Geometry and Topology, **5** (2001), pp. 7–74.
23. J. Dodziuk, *de Rham-Hodge theory for L^2-cohomology of infinite covers*, Topology, **16** (1977), pp. 157–165.
24. J. Dodziuk, *L^2-harmonic forms on rotationally symmetric Riemannian manifolds*, Proc. Amer. Math. Soc., **77** (1979), pp. 395–400.
25. X. Dong, *Topology of bounded degree graph complexes*, J. Algebra, **262** (2003), pp. 287–312.
26. H. Donnelly and F. Xavier, *On the differential form spectrum of negatively curved Riemannian manifolds*, Amer. J. of Math., **106** (1984), pp. 169–185.
27. A. Duval, *A combinatorial decomposition of simplicial complexes*, Israel J. of Math., **87** (1994), pp. 77–87.
28. B. Eckmann, *Introduction to l^2-methods in topology: reduced l^2-homology, harmonic chains, l^2-Betti numbers*, (notes prepared by Guido Mislin), Israel J. Math., **117** (2000), pp.183–219.
29. W. Fenchel, *On total curvature of Riemannian manifolds: I,* J. London Math. Soc., **15** (1940), pp. 15–22.
30. S. Fomin and N. Reading, *Root systems and generalized associahedra*, this volume.
31. R. Forman, *A discrete Morse theory for cell complexes*, in Geometry, topology & physics, S.T. Yau (ed.), Conf. Proc. Lecture Notes Geom. Topology, International Press, Cambridge, MA, 1995, pp. 112–125.
32. R. Forman, *Morse theory for cell complexes*, Adv. in Math., **134** (1998), pp. 90–145.
33. R. Forman, *Witten-Morse theory for cell complexes*, Topology, **37** (1998), pp. 945–979.
34. R. Forman, *Combinatorial vector fields and dynamical systems*, Math. Zeit., **228** (1998), pp. 629–681.
35. R. Forman, *Combinatorial differential topology and geometry*, in New Perspectives in Algebraic Combinatorics, Math. Sci. Res. Inst. Publ. 38, Cambridge Univ. Press, Cambridge, 1999, pp. 177–206.
36. R. Forman, *Morse theory and evasiveness*, Combinatorica, **20** (2000), pp. 498–504.
37. R. Forman, *Novikov-Morse theory for cell complexes*, Int. J. of Math., **4** (2002), pp. 333–368.

38. R. Forman, *The cohomology ring and discrete Morse theory*, Trans. of the Amer. Math. Soc., **354** (2002), pp. 5063–5085.
39. R. Forman, *A user's guide to discrete Morse theory*, Seminaire Lotharingien de Combinatoire (electronic), **48** (2002), article B48c.
40. R. Forman, *Some applications of combinatorial differential topology*, Graphs and patterns in mathematics and theoretical physics, Proc. Sympos. Pure Math., 73, Amer. Math. Soc., Providence, RI, 2005, pp. 281–313.
41. R. Forman, *A topological approach to the game of "20 Questions"*, preprint.
42. W. Fulton, *Introduction to toric varieties*, Annals of Mathematics Studies 131, Princeton University Press, 1993.
43. S. Gal, *Real Root Conjecture fails for five and higher dimensional spheres*, Discrete Comput. Geom., **34** (2005), pp. 269–284.
44. R. Geroch, *Positive sectional curvature does not imply positive Gauss–Bonnet integrand*, Proc. Amer. Math. Soc., **54** (1976), pp. 267–270.
45. L. Glaser, *Geometrical combinatorial topology*, Vol. 1, Van Nostrand Reinhold Mathematical Studies # 27, 1970.
46. M. Goresky and R. MacPherson, *Stratified Morse theory*, in Singularities, Part I (Arcata, CA, 1981), Proc. Sympos. Pure Math., 40, Amer. Math. Soc., R.I., 1983, pp. 517–533.
47. M. Goresky and R. MacPherson, *Stratified Morse theory*, Ergebnisse der Mathematik und ihrer Grenzgebiete (3), 14, Springer Verlag, Berlin-New York, 1988.
48. M. Gromov, *Hyperbolic groups*, in Essays in group theory, (S.M. Gersten (ed.)), Math. Sci. Res. Inst. Publ. 8, Springer-Verlag, New York, 1987, pp. 75–263.
49. M. Gromov, *Kähler hyperbolicity and L^2-Hodge theory*, J. Diff. Geom., **33** (1991), pp. 263–292.
50. P. Hersh, *On optimizing discrete Morse functions*, Adv. in Appl. Math., **35** (2005), pp. 294–322.
51. P. Hersh and V. Welker, *Gröbner basis degree bounds on $Tor^{k[\Lambda]}_{\cdot}(k,k)_{\cdot}$ and discrete Morse theory for posets*, in Integer points in polyhedra - geometry, number theory, algebra, optimization, Contemp. Math., 374, Amer. Math. Soc., Providence, RI, (2005), pp, 101–138.
52. J. Jonsson, *On the topology of simplicial complexes related to 3-connected and hamiltonian graphs*, J. Combin. Theory, Ser. A., **104** (2003), pp. 169–199.
53. J. Jonsson, *Optimal decision trees for simplicial complexes*, Electron. J. Combin., **12** (2005), pp. 211–239.
54. M. Joswig and M. Pfetsch, *Computing optimal discrete Morse functions*, in Workshop on Graphs and Combinatorial Optimization (electronic), Electron. Notes Discrete Math., 17. Elsevier, Amsterdam, (2004), pp. 191–195.
55. J. Kahn, M. Saks and D. Sturtevant, *A topological approach to evasiveness*, Combinatorica, **4** (1984), pp. 297–306.
56. D. Kleitman and D. Kwiatkowski, *Further results on the Aanderaa-Rosenberg conjecture*, J. Combin. Theory, Ser. B, **28** (1980), pp. 85–95.
57. D. Kozlov, *Collapsibility of $\Delta(\pi_n)/S_n$ and some related CW complexes*, Proc. of the Amer. Math. Soc., **128** (2000), pp. 2253–2259.
58. D. Kozlov, *Topology of spaces of hyperbolic polynomials and combinatorics of resonances*, Israel J. of Math., **132** (2002), 189–206.

59. W. Kühnel, *Triangulations of manifolds with few vertices*, in Advances in Differential Geometry and Topology, World Sci. Publishing, N.J., 1990, pp. 59–114.
60. C. Lee, *The associahedron and triangulations of the n-gon*, Europ. J. of Comb., **10** (1989), pp. 551–560.
61. N. Leung and V. Reiner, *The signature of a toric variety*, Duke Math. J., **111** (2002), pp. 253–286.
62. T. Lewiner, H. Lopes, G. Tavares, *Visualizing Forman's discrete vector field*, in Mathematical Visualization III, H.-C. Hege, K. Polthier (eds.), Springer-Verlag, Heidelberg, 2002, pp. 95–112.
63. T. Lewiner, H. Lopes, G. Tavares, *Optimal discrete Morse functions for 2-manifolds*, Comput. Geom., **26** (2003), pp. 221–233.
64. T. Lewiner, H. Lopes, G. Tavares, *Towards optimality in discrete Morse theory*, Experiment. Math., **12** (2003), pp. 271–285.
65. S. Linusson and J. Shareshian, *Complexes of t-colorable graphs*, SIAM J. Discrete Math., **16** (2003), pp. 371–389.
66. J. Lott and W. Lück, *L^2-topological invariants of 3-manifolds*, Invent. Math, **120** (1995), pp. 15–60.
67. W. Lück, *L^2-Betti numbers of mapping tori and groups*, Topology, **33** (1994), pp. 203–214.
68. A. Lundell and S. Weingram, *The topology of CW complexes*, Van Nostrand Reinhold Company, New York, 1969.
69. P. McMullen, *On simple polytopes,* Invent. Math, **113** (1993), pp. 419–444.
70. P. McMullen, *Weights on polytopes*, Discrete Comput. Geom., **15** (1996), pp. 363–388.
71. J. Milnor, *Morse theory*, Annals of Mathematics Study No. 51, Princeton University Press, 1962.
72. J. Milnor, *Lectures on the h-cobordism theorem*, Princeton Mathematical Notes, Princeton University Press, 1965.
73. M. Morse, *The calculus of variations in the large*, Amer. Math. Soc. Colloquium Pub. 18, Amer. Math. Soc., Providence R.I., 1934.
74. M. Morse, *Bowls of a non-degenerate function on a compact differentiable manifold*, in Differential and Combinatorial Topology (A Symposium in Honor of Marston Morse), Princeton University Press (1965), pp. 81–104.
75. V. Reiner and V. Welker, *On the Charney-Davis and the Neggers-Stanley conjectures*, J. Combin. Theory Ser. A, **109** (2005), pp. 247–280.
76. R. Rivest and J. Vuillemin, *A generalization and proof of the Aanderaa-Rosenberg conjecture*, 7th Annual ACM Symposium on Theory of Computing, Albuquerque, New Mexico, 1975, pp. 6–12.
77. R. Rivest and J. Vuillemin, *On recognizing graph properties from adjacency matrices*, Theor. Comp. Sci, **3** (1976), pp. 371–384.
78. M. Schwartz, *Morse homology*, Progress in Mathematics 111, Birkhäuser Verlag, Basel, 1993.
79. J. Shareshian, *Discrete Morse theory for complexes of 2-connected graphs*, Topology, **40** (2001), pp. 681–701.
80. V. Sharko, *Functions on manifolds, algebraic and topological aspects*, Translations of Math. Monographs 131, Amer. Math. Soc., Providence, R.I., 1993.

81. E. Sköldberg, *Combinatorial discrete Morse theory from an algebraic viewpoint*, Trans. Amer. Math. Soc., **358** (2006), pp. 115–129.
82. S. Smale, *On gradient dynamical systems*, Annals of Math., **74** (1961), pp. 199–206.
83. S. Smale, *The generalized Poincaré conjecture in dimensions greater than four*, Annals of Math., **74** (1961), pp. 391–406.
84. D. Soll, *Evasive Simpliziale Komplexe und Diskrete Morse Theorie*, Diploma Thesis, Philipps-Universität Marburg, 2002.
85. R. Stanley, *A combinatorial decomposition of acyclic simplicial complexes*, Discrete Math., **118** (1993), pp. 175–182.
86. J. Stasheff, *Homotopy associativity of H-spaces*, Trans. of the Amer. Math. Soc., **108** (1963), pp. 275–292.
87. J. Stasheff, *The prehistory of operads*, in Operads: proceedings of the Renaissance conferences (Hartford, CT/Luminy, 1995), Cotemp. Math. 2002, pp. 9–14.
88. R. Stanley, *The number of faces of a simplicial convex polytope,* Adv. in Math., **35** (1980), pp. 236–238.
89. R. Stanley, *Flag f-vectors and the cd-index*, Math. Zeitschrift, 216 (1994), pp. 483–499.
90. R. Stanley, *Combinatorics and commutative algebra*, 2nd ed., Birkäuser Verlag, Basel, 1996.
91. R. Stanley, *Positivity problems and conjectures in algebraic combinatorics*, in Mathematics: frontiers and perspectives (V. Arnold, M. Atiyah, P. Lax, and B. Mazur, eds.), AMS, Providence, RI, 2000, pp. 295–319.
92. V. Turchun, *Homology of complexes of biconnected graphs*, Uspekhi Mat. Nauk., **52** (1997), pp. 189–190.
93. V. Vassiliev, *Complexes of connected graphs*, in The Gelfand mathematical seminars, 1990–1992, Birkhäuser Verlag, Boston, 1993, pp. 223–235.
94. V. Welker, *Constructions preserving evasiveness and collapsibility*, Discrete Math., **207** (1999), pp. 243–355.
95. J. Whitehead, *Simplicial spaces, nuclei, and m-groups*, Proc. London Math. Soc., **45** (1939), pp. 243–327.