COMBINATORIAL DIFFERENTIAL TOPOLOGY AND GEOMETRY

Robin Forman

Department of Mathematics
Rice University
Houston, TX 77251

Combinatorial Differential Topology and Geometry

0. Introduction

A variety of questions in combinatorics lead one to the task of analyzing a simplicial complex, or a more general cell complex. For example, a standard approach to investigating the structure of a partially ordered set is to instead study the topology of the associated order complex. However, there are few general techniques to aid in this investigation. On the other hand, the subjects of differential topology and differential geometry are devoted to precisely this sort of problem, except that the topological spaces in question are smooth manifolds, rather than combinatorial complexes. These are classical subjects, and numerous very general and powerful techniques have been developed and studied over the recent decades.

A smooth manifold is, loosely speaking, a topological space on which one has a well-defined notion of a derivative. One can then use calculus to study the space. I have recently found ways of adapting some techniques from differential topology and differential geometry to the study of combinatorial spaces. Perhaps surprisingly, many of the standard ingredients of differential topology and differential geometry have combinatorial analogues. The combinatorial theories are often simpler than the corresponding smooth theory, and in some cases imply the smooth theory. In this paper, I will discuss two new general tools to aid in the study of combinatorial spaces. One derived from Morse theory, a standard tool in differential topology, and the other derived from Bochner's method, a standard tool in differential geometry and global analysis. Most of the results in this paper appeared earlier in [Fo3] and [Fo6].

Of course, many others have had the idea of "borrowing" ideas from continuous mathematics to study combinatorial objects. See for example [Bj] and the numerous references therein. As we go along, I will mention those references most closely related to our topic.

This work was partially supported by the National Science Foundation and the National Security Agency. The author would like to thank the staff of the Mathematical Sciences Research Institute, as well as the organizers of the special year in Combinatorics and Topology, for their hospitality while some of this work was completed. The author would also like to thank the referees, whose thoughtful comments have drastically improved this paper.

Typeset by $\mathcal{AMS}$-TEX

1. Morse Theory

Since its introduction in [Mo1], Morse theory has become a fundamental technique for investigating the topology of smooth manifolds. The basic principle is that the topology of a manifold is very closely related to the critical points of a smooth function on the manifold. The simplest example of this relationship is the fact that if the manifold is compact, then any continuous function must have a maximum and a minimum. Morse theory provides a significant refinement of this observation. See [Mi1] for a beautiful exposition of this subject, and [Bo] for a wonderful overview of Morse theory, including some recent developments.

Many others have developed versions of Morse theory for simplicial complexes (e.g. [BK], [Kü]), and other types of non-smooth topological spaces (e.g. [Mo2], [EK],[Kui], [GoM1,2]). When extending Morse theory to such spaces, one must decide what will replace smooth functions, the key ingredient in the classical theory. In the references cited, a number of choices have been made. Attention is focussed on piecewise linear functions ([BK],[Kü]), continuous functions which behave like smooth functions near their critical points ([Mo2],[EK]), and functions which can be extended to a larger domain to be smooth ([GoM1,2]). If one attempts to consider all continuous functions, one can still define the notion of a critical point, but it seems to be impossible to control the topological contribution of each critical point ([Kui],[Gr1]).

In this paper, we take a different approach. We will present a very simple Morse theory for simplicial complexes which is entirely discrete. That is, there is no mention at all of continuous functions. Instead, our functions assign a single real number to each cell of the complex. Because we require so little structure from our functions, we need very little structure in our spaces, so the theory can be applied to very general cell complexes. After defining combinatorial Morse functions and critical points, we will show that the standard theorems of Morse theory, relating the topology of the space to the critical points of the function, are true. We also present discrete analogues of such (seemingly) intrinsically smooth notions as the gradient vector field and the corresponding gradient flow associated to a Morse function. Using these, we define a Morse complex, a differential complex built out of the critical points of our discrete Morse function, which has the same homology as the underlying space. In the smooth setting, the Morse complex has been subject to much recent study, and has played a crucial role in some recent applications (see, for example [Kl],[Fl]). Most of this section will be an informal exposition of the contents of [Fo2] and [Fo3].

This combinatorial theory is not really a new theory, but rather an extraction of the combinatorial essence of the smooth theory. As evidence for this, we will indicate later that it is possible to deduce much of the smooth theory from this combinatorial theory, and, conversely, some of the results we present can be proved by "smoothing" the combinatorial Morse function and applying the corresponding smooth theory (although that is not the approach we will take in this paper).

As we mentioned above, this discrete Morse theory can be defined for any CW complex, but to avoid very minor complications we will restrict attention, in this paper, to simplicial complexes. See [Fo3] for a more general treatment. Let M be any finite simplicial complex. We emphasize that M need not be a triangulated manifold, nor have any other special property. Denote by K the set of (non-empty) simplices of M, and K_p the simplices of dimension p. Write $\alpha^{(p)}$ if α has dimension p, and $\beta > \alpha$ (or $\alpha < \beta$) if α is a face of β.

A discrete Morse function on M will actually be a function on K. That is, we assign a single real number to each simplex in M. Roughly speaking, a function

$$f : K \longrightarrow \mathbf{R}$$

is a Morse function if f usually assigns higher values to higher dimensional simplices, with at most one exception, locally, at each simplex. More precisely,

Definition 1. *A function*

$$f : K \longrightarrow \mathbf{R}$$

is a discrete Morse function if for every $\alpha^{(p)} \in K$

(1) $\#\{\beta^{(p+1)} > \alpha \mid f(\beta) \leq f(\alpha)\} \leq 1$, *and*
(2) $\#\{\gamma^{(p-1)} < \alpha \mid f(\gamma) \geq f(\alpha)\} \leq 1$.

A simple example will serve to illustrate the definition. Consider the two complexes shown in Figure 1. Here we indicate functions by writing next to each cell the value of the function on that cell. The function (i) is not a discrete Morse function as the edge $f^{-1}(0)$ violates rule (2), since it has 2 lower dimensional "neighbors" on which f takes on higher values, and the vertex $f^{-1}(3)$ violates rule (1), since it has 2 higher dimensional "neighbors" on which f takes on lower values. The function (ii) is a Morse function.

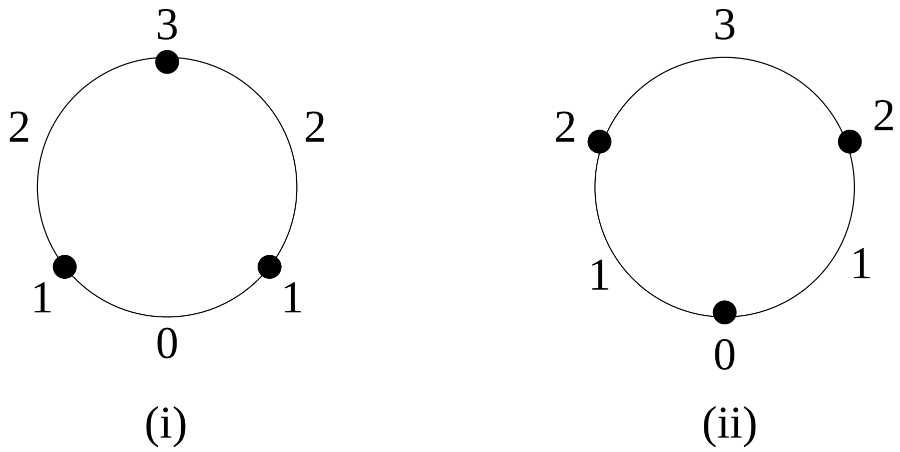

Figure 1

The other main ingredient in Morse theory is the notion of a critical point.

Definition 2. *A simplex* $\alpha^{(p)}$ *is critical (with index* p*) if*

(1) $\#\{\beta^{(p+1)} > \alpha \mid f(\beta) \leq f(\alpha)\} = 0$, *and*
(2) $\#\{\gamma^{(p-1)} < \alpha \mid f(\gamma) \geq f(\alpha)\} = 0$.

For example, Figure 1(ii), the vertex $f^{-1}(0)$ is critical of index 0, the edge $f^{-1}(3)$ is critical of index 1, and there are no other critical simplices. Note that if $\sigma^{(p)}$ is critical, it is necessarily critical of index p.

We will pause here to relate these definitions to the corresponding notions in the smooth category. To illustrate the main idea, we will consider a special case. Suppose x is a critical point of index 1 of a smooth Morse function F on a smooth manifold of dimension n. Then the Morse Lemma (see [Mi1], Lemma 2.2) states that there are coordinates $(t_1, \ldots, t_n)$, with x corresponding to $(0, \ldots, 0)$, such that in these coordinates

$$F(t_1, \ldots, t_n) = F(x) - t_1^2 + \sum_{i=2}^{n} t_i^2.$$

That is, beginning at the point x, F decreases to both sides in the t_1 direction, and increases in the other coordinate directions. Now suppose e is a critical 1-simplex of a discrete Morse function f. Then $f(e)$ is greater than f at either boundary vertex, and less than f at any 2-simplex with e in its boundary (see, for example, Figure 2).

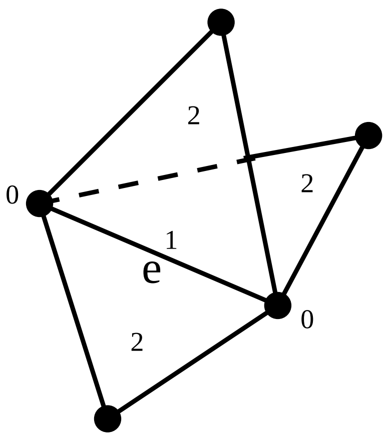

Figure 2

That is, f decreases as one moves from the 1-simplex to either boundary component, and increases in each transverse direction. We can see that this is, in fact, a combinatorial model of the smooth situation. Moreover we see that the critical edge e can be thought of a representing a critical point of index 1 in the smooth case together with a curve pointing in the directions in which the function is decreasing. In the general case, if $\sigma^{(p)}$ is critical then the simplex σ can be thought of as representing the p-dimensional "unstable" directions at a smooth critical point of index p. If our simplicial complex is a smooth triangulation of a smooth manifold, then one can make this discussion more precise. One can think of a discrete Morse function f as asigning a number to the barycenter of each simplex. One can then smoothly interpolate between these values to define a smooth function on all of M.

If this is done carefully, "smoothing" a discrete Morse function f near a critical p-simplex results in a smooth Morse function F with critical point of index p.

Our next step is to show that the topology of a general simplicial complex M is intimately related to the critical points of a discrete Morse function on M. Before stating the main theorems, we need one more definition. Suppose f is a discrete Morse function on a simplicial complex M. For any $c \in \mathbf{R}$ define the level subcomplex $M(c)$ by

$$M(c) = \bigcup_{f(\beta) \leq c} \bigcup_{\alpha \leq \beta} \alpha.$$

That is, $M(c)$ is the subcomplex of M consisting of all simplices β with $f(\beta) \leq c$, as well as all of their faces.

Theorem 1. *Suppose the interval $(a, b]$ contains no critical values of f. Then $M(a)$ is a deformation retract of $M(b)$. Moreover, $M(b)$ simplicially collapses onto $M(a)$.*

To define simplicial collapse, suppose M is a simplicial complex, $\alpha^{(p)}$ is a p-simplex of M which is not the face of any other simplex, and $\gamma^{(p-1)} < \alpha$ a $(p-1)$-face of α which is not the face of any other simplex. Then we say that M simplicially collapses onto $M - (\alpha \cup \gamma)$ (Figure 3). More generally, a simplicial collapse is any sequence of such operations. See Figure 4 for the simplicial collapse of a 2-simplex onto a vertex.

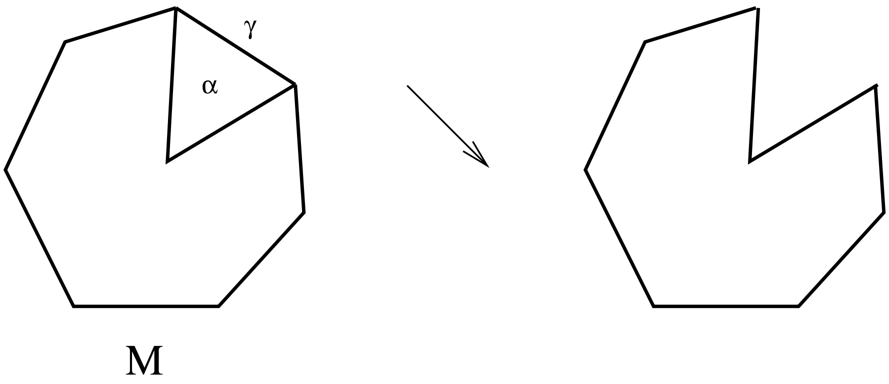

Figure 3

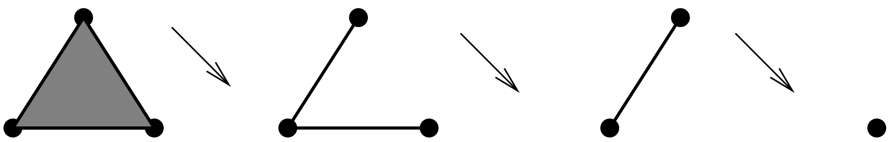

Figure 4

The equivalence relation generated by simplicial collapse is called *simple-homotopy equivalence*. This indicates that our Morse theory is particularly well suited to handling questions in this category.

Theorem 2. *Suppose $\alpha^{(p)}$ is a critical simplex with $f(\alpha) \in (a, b]$, and there are no other critical simplices with values in $(a, b]$. Then $M(b)$ is (simple-)homotopy equivalent to*

$$M(a) \bigcup_{\dot{e}^{(p)}} e^{(p)}$$

where $e^{(p)}$ is a p-cell, and it is glued to $M(a)$ along its entire boundary $\dot{e}^{(p)}$.

Before discussing the proofs of these theorems, we pause to consider some important corollaries.

Corollary 1. *Suppose M is a simplicial complex with a discrete Morse function. Then M is (simple-)homotopy equivalent to a CW complex with exactly one cell of dimension p for each critical simplex of dimension p.*

This corollary is often most usefully applied via inequalities, known collectively as the Morse Inequalities, between numerical invariants. Let M be a simplicial complex with a discrete Morse function. Let m_p denote the number of critical simplices of dimension p. Let $\mathbf{F}$ be any field, and $b_p = dim H_p(M, \mathbf{F})$ the p^{th} Betti number with respect to $\mathbf{F}$. Then we have the following inequalities.

Corollary 2. I. The Weak Morse Inequalities.
(i)For each $p = 0, 1, 2, \ldots, n$ (where n is the dimension of M)

$$m_p \geq b_p.$$

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

II. The Strong Morse Inequalities.
For each $p = 0, 1, 2, \ldots, n, n+1$,

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

The Morse inequalities follow immediately from the previous corollary (see [Mi1]). Note that the Strong Morse Inequalities imply the Weak Morse Inequalities.

At this point it is worthwhile to observe that every simplicial complex supports a discrete Morse function. Namely, for any complex M and any simplex α of M, set $f(\alpha) = \dim(\alpha)$. Then f is a discrete Morse function, and every simplex of M is critical for f.

The results we have stated above apply to any simplicial complex. In the case that M is a combinatorial manifold ([Gl],p. 19) one can often say more. For example, one can state a combinatorial analogue of the standard Sphere theorem of smooth Morse theory (Theorem 4.1 of [Mi1]).

Corollary 3. *Suppose M is a combinatorial n-manifold without boundary with a Morse function with exactly two critical points. Then M is a combinatorial n-sphere.*

The proof of this corollary rests heavily on Whitehead's wonderful theorem on the uniqueness of regular neighborhoods [Wh] which implies, as a special case, that a combinatorial n-manifold with boundary which simplicially collapses onto a vertex is a combinatorial n-ball. Suppose M is a combinatorial n-manifold with a discrete Morse function f with exactly two critical points. Then the maximum of f must occur at an n-simplex $\alpha^{(n)}$ which is a critical simplex of index n. Thus, f restricts on $M - \alpha$ to a Morse function with exactly one critical simplex. The minimum of f must occur at a vertex which is a critical simplex of index 0, and this is the only critical simplex of f restricted to $M - \alpha$. Theorem 1 implies that $M - \alpha$ simplicially collapses to that vertex. Now we apply Whitehead's theorem to conclude that $M - \alpha$ is a combinatorial n-ball. It follows that $M = (M - \alpha) \cup \alpha$ is a combinatorial n-sphere.

Rather than prove the main theorems here (rigorous proofs appear in [Fo3]), we will illustrate the main ideas with an example (see Figure 5).

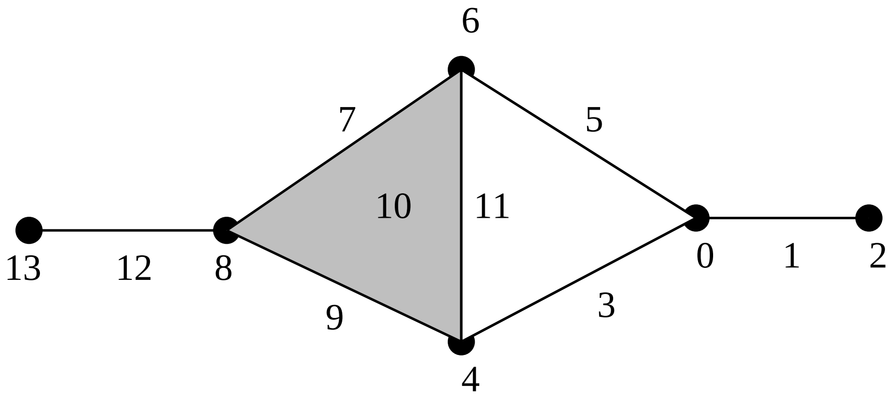

Figure 5

Here $f^{-1}(0)$ is a critical simplex of index 0, $f^{-1}(9)$ is a critical simplex of index 1, and there are no other critical simplices. Thus, it follows from the above corollary that M is homotopy equivalent to a CW complex built from a single 0-cell, and a single 1-cell. The only CW complex which can be built from two such cells is a circle. Therefore, we can conclude that M is homotopy equivalent to a circle, as is evident from the picture.

Consider the sequence of level subcomplexes (Figure 6). We can see why the theorems are true. If α is not critical, then one of two cases must occur. The first possibility is that α has a codimension one face γ with $f(\gamma) \geq f(\alpha)$ (for example $\alpha = f^{-1}(1), \gamma = f^{-1}(2)$). In this case, when α is added to M, γ is added at the same time. The new subcomplex collapses onto the previous complex by pushing in α along the free face γ. (There are a few things to check. For example, we must make sure that γ is, in fact, a free face of α. That is, we must check that γ is not a face of any other simplex β with $f(\beta) \leq f(\alpha)$, but this is

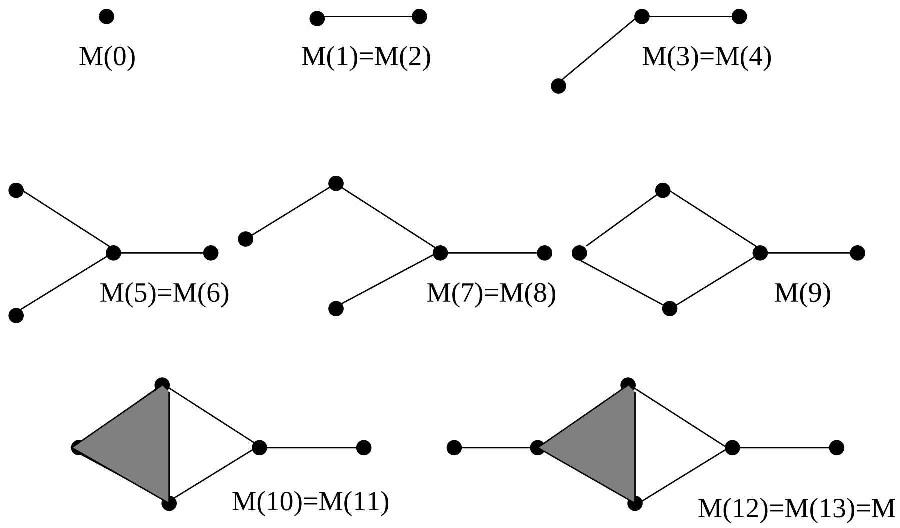

Figure 6

true.) The other possibility is that there is a simplex β with α as a codimension-one face and with $f(\beta) \leq f(\alpha)$ (for example $\alpha = f^{-1}(2), \beta = f^{-1}(1)$). In this case, α is added to M when β is added to M, and again the new complex collapses onto the previous complex by pushing in β along the free face α.

If α is critical (for example $f^{-1}(9)$), then for every simplex β with $\beta > \alpha$ we have $f(\beta) > f(\alpha)$ (this requires a bit of proof). Thus α appears first in $M(f(\alpha))$. On the other hand, for every face γ of α, $f(\gamma) < f(\alpha)$, and is added earlier. Therefore, when α is added, it is simply glued to $M(f(\alpha) - \epsilon)$ along its boundary. This completes our "proof" of the main theorems.

While the Morse inequalities are sufficient for many applications, it is sometimes desirable to replace the Morse inequalities with equalities. That is, to understand the precise nature of any "extra" critical points which arise. This can be done, and requires the introduction of the gradient vector field. In fact, as we will see later, some of the results we have stated above are best understood in terms of the gradient vector field, rather than the original Morse function. Returning to the example in Figure 5, let us now indicate pictorially the simplicial collapse referred to in the theorem. Suppose $\alpha^{(p)}$ is a non-critical simplex with $\beta^{(p+1)} > \alpha$ satisfying $f(\beta) \leq f(\alpha)$. We then draw an arrow from α to β. The resulting diagram can be seen in Figure 7. A simplex is critical if and only if it is neither the tail nor the head of an arrow. These arrows can be viewed as the discrete analogue of the gradient vector field of the Morse function. (To be precise, when we say "gradient vector field" we are really refering to the negative of the gradient vector field.)

It is better to think of this gradient vector field, which we now call V, as a map of *oriented* simplices. That is, if v is a boundary vertex of an edge e with $f(e) \leq f(v)$, we want to think of $V(v)$ as a discrete tangent vector *leaving* v i.e., with e given the orientation indicated by the arrow in Figure 8.

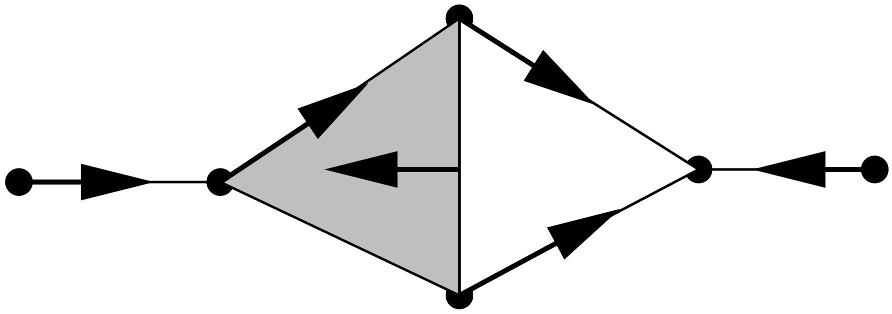

Figure 7

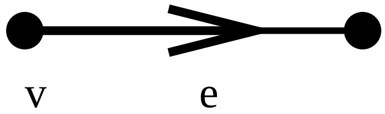

Figure 8

More generally, if $\beta^{(p+1)} > \alpha^{(p)}$ are oriented simplices and satisfy $f(\beta) \leq f(\alpha)$ then we set $V(\alpha) = \pm\beta$ with the sign chosen so that

$$< \alpha, \partial V(\alpha) > = \ -1$$

where $<,>$ is the canonical inner product on oriented chains with respect to which the simplices are orthonormal (in other words, $< \alpha, \partial V(\alpha) >$ is the incidence number of α with respect to $V(\alpha)$). Define $V(\alpha)$ to be 0 for all simplices α for which there is no such β. Now V can be extended linearly to a map

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

where, for each p, $C_p(M, \mathbf{Z})$ is the space of integer p-chains on M.

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 [Sm1]. To proceed further in our combinatorial setting, we now define a (discrete-time) flow along the gradient vector field V. Define a map

$$\Phi : C_p(M, \mathbf{Z}) \longrightarrow C_p(M, \mathbf{Z}),$$

the discrete-time flow, by

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

We illustrate by an example. Consider the complex hown in Figure 9, with the indicated gradient vector field V. Let e be the top edge oriented from left to right. We calculate

$$\Phi(e) \ = \ e + \partial V(e) + V\partial(e)$$

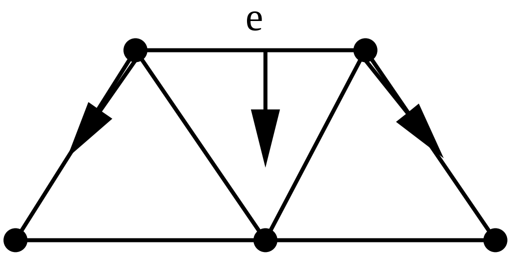

Figure 9

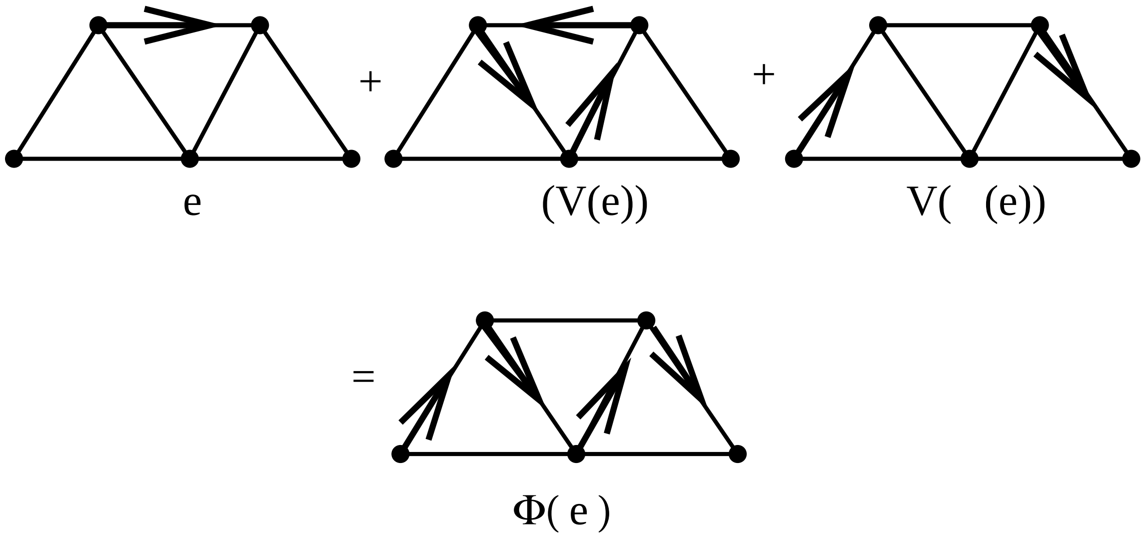

Figure 10

in Figure 10.

The map Φ has many of the properties that one expects from the gradient flow. For example, if $\alpha^{(p)}$ is not a critical simplex, then $\Phi(\alpha)$ is a p-chain consisting entirely of oriented $p-$simplices on which f is less than $f(\alpha)$, so, loosely speaking, Φ decreases f. The discrete gradient flow differs from the gradient flow in the smooth case primarily in that it stabilizes in finite time, i.e. there is an N such that

$$\Phi^N = \Phi^{N+1} = \ldots \Phi^{\infty}.$$

This observation can be quite useful (for example, see the proof of the following theorem). See [Fo3], section 6, for the proofs of the main properties of Φ

Let $C_p^{\Phi} \subseteq C_p(M, \mathbf{Z})$ denote the Φ-invariant p-chains, i.e., those p-chains c such that $\Phi(c) = c$. Since

$$\partial\Phi \;=\; \Phi\partial,$$

the space of Φ-invariant chains is preserved by the boundary operator ∂, so we can consider the Morse complex

$$\mathcal{C}^{\Phi} : 0 \longrightarrow C_n^{\Phi} \xrightarrow{\partial} C_{n-1}^{\Phi} \xrightarrow{\partial} C_{n-2}^{\Phi} \xrightarrow{\partial} \cdots .$$

The crucial point is the following theorem.

Theorem 3. $H_*(\mathcal{C}^{\Phi}) \;=\; H_*(M, \mathbf{Z})$.

That is, the Morse complex has the same homology as the underlying space. The proof of this theorem is not difficult. To prove that the stabilization map

$$\Phi^{\infty} : C_*(M, \mathbf{Z}) \to C_*^{\Phi}$$

induces an isomorphism on homology (with the inverse isomorphism induced by the natural injection), it is sufficient to find an algebraic homotopy operator, i.e. an operator $L : C_*(M, \mathbf{Z}) \to C_{*+1}(M, \mathbf{Z})$ such that $\Phi^{\infty} - 1 = \partial L + L\partial$. We now use the fact that $\Phi^{\infty} = \Phi^N$ for some N. If $N = 1$, then $L = V$ works. The general case is not much harder.

Our goal in introducing this complex was to further our understanding of the critical simplices. With this in mind, we now show that the Morse complex can be alternately defined in terms of the critical simplices. For each p, let $\mathcal{M}_p \subseteq C_p(M, \mathbf{Z})$ denote the span of the critical p-simplices. We then obviously have

(1) $$\dim \mathcal{M}_p = \#\{\textit{critical simplices of dimension } p\}$$

For each p, we can apply the stablization map Φ^{∞} to get a map from $\mathcal{M}_p$ to C_p^{Φ}.

Theorem 4. *For each p, the map*

$$\Phi^{\infty} : \mathcal{M}_p \to C_p^{\Phi}.$$

is an isomorphism.

Thus, the Morse complex can be defined equivalently as the complex

$$\mathcal{M} : 0 \longrightarrow \mathcal{M}_n \xrightarrow{\widetilde{\partial}} \mathcal{M}_{n-1} \xrightarrow{\widetilde{\partial}} \mathcal{M}_{n-2} \xrightarrow{\widetilde{\partial}} \cdots$$

where $\widetilde{\partial}$ is the differential induced by the above isomorphism (that is, $\widetilde{\partial} = (\Phi^{\infty})^{-1}\partial\Phi^{\infty}$). In particular

(2) $$H_*(\mathcal{M}) \;=\; H_*(M, \mathbf{Z}).$$

It is interesting to note that together (1) and (2) imply the Morse inequalities, so this gives an alternate proof of Corollary 2. From (2), we see that the Morse complex $\mathcal{M}$ contains a more complete description of the relationship between the critical simplices of f, and the

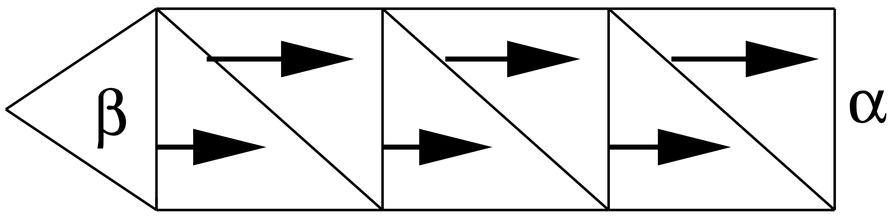

Figure 11

homology of M. All that remains is to get a explicit description of the differential $\widetilde{\partial}$. This requires the introduction of the notion of a combinatorial gradient path. Let α and $\tilde{\alpha}$ be p-simplices. A *gradient path* from $\tilde{\alpha}$ to α is a sequence of simplices

$$\tilde{\alpha} = \alpha_0^{(p)}, \beta_0^{(p+1)}, \alpha_1^{(p)}, \beta_1^{(p+1)}, \alpha_2^{(p)}, \ldots, \beta_r^{(p+1)}, \alpha_{r+1}^{(p)} = \alpha$$

such that for each $i = 1, \ldots r$, $f(\alpha_i) \geq f(\beta_i) > f(\alpha_{i+1})$. Equivalently, we require $V(\alpha_i) = \pm\beta_i$, and $\beta_i > \alpha_{i+1} \neq \alpha_i$. In Figure 11 we show a single gradient path from the boundary of a critical 2-simplex β to a critical edge α, where the arrows indicate the gradient vector field V.

We are now ready to state the desired formula.

Theorem 5. *Choose an orientation for each simplex. Then for any critical* $(p+1)$*-simplex* β

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

where

$$c_{\alpha,\beta} = \sum_{\tilde{\alpha}^{(p)} < \beta} < \partial\beta, \tilde{\alpha} > \sum_{\gamma \in \Gamma(\tilde{\alpha},\alpha)} m(\gamma)$$

where $\Gamma(\tilde{\alpha}, \alpha)$ *is the set of gradient paths from* $\tilde{\alpha}$ *to* α.

For any gradient path from $\tilde{\alpha} < \beta$ to α, the orientation on β induces an orientation on α. The coefficient $m(\gamma)$ is equal to ± 1, depending on whether the orientation on β induces the chosen orientation on α, or the opposite orientation.

As stated earlier, it is often easier to work with the gradient vector field, rather than the original Morse function. In fact, one need not refer to the Morse function at all, as long as one knows that the vector field under consideration is the gradient of a Morse function. Thus, it is important to characterize those vector fields which are gradient vector fields of Morse functions. In the smooth case, Smale provided such a characterization in [Sm1]. We can provide an analogous characterization of discrete gradient vector fields. First, we define an object we call a discrete vector field. Such objects have been previously studied, under a different name, in [Du] and [St] and some of the references therein.

Definition 3. *A discrete vector field is any one-to-one map*

$$U : K \to K \cup \{0\}$$

satisfying for each $\alpha^{(p)}$

(1) $U(\alpha) = 0$ *or* α *is a codimension-one face of* $U(\alpha)$,
(2) *If* $\alpha^{(p)} \in$ *Image*(U) *then* $U(\alpha) = 0$.

A discrete gradient vector field is not quite such an object, since it was defined for oriented simplices, but for now we will ignore such subtleties. We have presented the definition in this form so that the relationship to the notion of a gradient vector field is apparent. However, there is a simpler, equivalent, definition. Namely, one can define a combinatorial vector field to be a disjoint collection $\mathcal{P}$ of pairs $\{\alpha^{(p)} < \beta^{(p+1)}\}$ of simplices of M. To see that these definitions are equivalent, note that given a combinatorial vector field U (as defined in Definition 3), the pairs $\{\alpha < U(\alpha)\}$, for all α with $U(\alpha) \neq 0$, are disjoint. Conversely, given any set $\mathcal{P}$ of disjoint pairs $\{\alpha^{(p)} < \beta^{(p+1)}\}$, we can define a combinatorial vector field U by setting

$$U(\alpha) = \beta \iff \{\alpha < \beta\} \in \mathcal{P}.$$

Such collections of disjoint pairs of simplices were studied in [St] and [Du] as a tool for investigating the possible f-vectors for a simplicial compex. In [Fo5] we investigated the dynamical properties of the discrete-time flow corresponding to a general discrete vector field.

If U is a combinatorial vector field, we define a $U-$path to be any path as in Figure 10 (where now the arrows indicate the map U). More precisely, a $U-$path is a sequence of simplices

$$\alpha_0^{(p)}, \beta_0^{(p+1)}, \alpha_1^{(p)}, \beta_1^{(p+1)}, \alpha_2^{(p)}, \ldots, \beta_r^{(p+1)}, \alpha_{r+1}^{(p)}$$

such that for each $i = 1, \ldots r$, $\beta_i = U(\alpha_i)$ and $\beta_i > \alpha_{i+1} \neq \alpha_i$. We say such a path is a non-trivial closed path if $r \geq 0$ and $\alpha_0 = \alpha_{r+1}$. There is a simple characterization of those discrete vector fields which are the gradient of a discrete Morse function.

Theorem 6. *A discrete vector field U is the gradient vector field of a discrete Morse function if and only if there are no non-trivial closed U-paths.*

This theorem has a very nice purely combinatorial description, using which we can recast the Morse theory in an appealing form. We begin with the Hasse diagram of M, that is, the partially ordered set of simplices of M ordered by the face relation. Consider the Hasse diagram as a directed graph. The vertices of the graph are in 1-1 correspondence with the simplices of M, and there is a directed edge from β to α if and only if α is a codimension-one face of β. Now let U be a combinatorial vector field. We modify the directed graph as follows. If $U(\alpha) = \beta$ then reverse the orientation of the edge between α and β (so that it now goes from α to β). A U-path is just a directed path in this modified graph. From Theorem 6, U is a gradient vector field if and only if there are no closed directed paths. That is, 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 above sense.

We can now restate some of our earlier theorems in this language. There is a very minor complication in that one usually includes the empty set as an element of the Hasse diagram (considered as a simplex of dimension -1) while we have not considered the empty set previously.

Theorem 7. *Let U be an acyclic partial matching of the Hasse diagram of M (of the sort described above). Assume that the empty set is matched with a vertex. Let u_p denote the number of unpaired p-simplices. Then M is homotopy equivalent to a CW-complex with $u_0 + 1$ vertices, and, for each $p \geq 1$, u_p cells of dimension p.*

An important special case of Theorem 7 is when U is a complete matching, that is, every simplex is paired with another (possible empty) simplex.

Theorem 8. *Let U be a complete acyclic matching of the Hasse diagram of M, then M collapses onto a vertex, so that, in particular, M is contractible.*

This language was introduced in [Cha] and [Sh], and the reader should consult these sources for a more complete presentation of the Morse theory from this point of view. Theorem 8 was used to good effect in [BBLSW] to investigate the homology of the simplicial complex of "not k-connected graphs" on a fixed set of vertices. This homology plays an important role in Vassiliev's theory of invariants for links [Va]. Let us be a bit more precise. Let V be a fixed set of n vertices. A graph on V is determined by specifying a subset of the $\binom{n}{2}$ possible edges. Let S be a simplex with $\binom{n}{2}$ vertices (i.e. $dim(S) = \binom{n}{2} - 1$). A graph on V can now be determined by specifying a face of S. We say a graph is connected if every pair of vertices can be connected by a path consisting of edges. We say a graph is k-connected if it is connected after the removal of an arbitrary set of ℓ vertices (and the edges connected to them) for any $\ell < k$. If a face α of S corresponds to a graph which is not k-connected, then every face β of α also corresponds to a graph which is not k-connected. Therefore, the set of "not k-connected" graphs forms a subcomplex of S. This subcomplex, which we denote by Δ_n^k, is the object of study in [BBLSW],[Tu] and [Sh]. (Actually, of interest is the homology of S/Δ_n^k, but since S is contractible, these questions are essentially equivalent.) We state their result in the special case $k = 2$.

Theorem 9 ([BBLSW],[Tu],[Sh]). *The complex Δ_n^2 has the homotopy type of a wedge of $(n-2)!$ spheres of dimension $2n - 5$.*

For the applications to the investigations in [Va] it was necessary to find explicit generators of the homology of Δ_n^k. This was done in [Sh] in which the more general Theorem 7 was featured. In [Sh], Shareshian contructs, by hand, an explicit acyclic partial matching on the simplicial complex of not 2-connected graphs. From the Morse theory he can then deduce the homotopy type of the complex. Moreover, the unmatched (i.e. critical) simplices provide a set of explicit generators for the homology.

In [Fo7] we apply these methods to a problem in complexity theory. We will describe a special case of this problem. Suppose an oracle chooses a graph Γ on the set of n vertices V. The graph Γ is unknown to us, but we are permitted to ask questions to the oracle, one at a time, of the sort "Is the edge (v_i, v_j) in Γ?", where $v_i, v_j \in V$. We may use the answers to our previous questions when choosing our next question. The only restriction is that we choose each question according to a deterministic algorithm (which takes as its input the previous questions and answers). Our task is to decide if Γ is k-connected, using as few questions as possible. In particular, we must try to accomplish our task before asking $\binom{n}{2}$ questions. (Of course, after asking $\binom{n}{2}$ questions we have completely determined Γ.) See §4.5 of [Bol] for a survey of the work on this and related problems.

Let A denote our algorithm for choosing questions. Say a graph Γ is an *evader* of A if, asking questions according to A, we do not determine whether or not Γ is k-connected until we ask $\binom{n}{2}$ questions. We observe that A induces a complete pairing of the set of graphs on V. Namely, pair Γ_1 and Γ_2 if we can not distinguish between them until all $\binom{n}{2}$ questions have been asked. By definition, a not-k-connected graph is an evader of A if and only if it is paired with a k-connected graph, and vice versa. The main result of [Fo7] is that this is an acyclic pairing, and hence, by Theorem 7, the gradient vector field of a Morse function. We can restrict this vector field to Δ_n^k, to get the gradient of a Morse function on Δ_n^k. The critical simplices of this restricted gradient vector field are the simplices of Δ_n^k which were paired with simplices not in Δ_n^k. These are precisely the not-k-connected graphs which were paired with a k-connected graph, i.e. the evaders of A. We can now use the weak Morse inequalities (Corollary 2) along with Theorem 9 to conclude

Theorem 10. *Suppose our goal is to determine whether or not an unknown graph is 2-connected. For any question algorithm A, there are at least $(n-2)!$ evaders of A which are not 2-connected, and $(n-2)!$ evaders of A which are 2-connected.*

Before ending this section, we make some final remarks about the combinatorial Morse theory, and its relation to the smooth theory. One of the main problems in Morse theory is finding a Morse function, for a given space, with the fewest possible critical points. 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. (See [Mi2] for a completely Morse theoretic presentation of Smale's proof [Sm2] of the higher dimensional Poincaré conjecture.) As an application of Theorem 6, we easily prove a "cancellation" theorem, which enables one, under certain conditions, to simplify a Morse function (this result is a discrete analogue of Theorem 5.4, the "First Cancellation Theorem", in [Mi2]). That is, if $\alpha^{(p)}$ and $\beta^{(p+1)}$ are 2 critical simplices, and if there is exactly 1 gradient path from $\partial\beta$ to α, then α and β can be cancelled. More precisely,

Theorem 11. *Suppose f is a discrete Morse function on M such that $\beta^{(p+1)}$ and $\alpha^{(p)}$ are critical, and there is exactly one gradient path from $\partial\beta$ to α. Then there is another Morse function g on M with the same critical simplices except that α and β are no longer critical. Moreover, the gradient vector field associated to g is equal to the gradient vector field associated to f except along the unique gradient path from $\partial\beta$ to α.*

In the smooth case, the proof, either as presented originally by Morse in [Mo3] or as presented in [Mi2], is rather technical. In our discrete case the proof is simple. If, in Figure 11, the

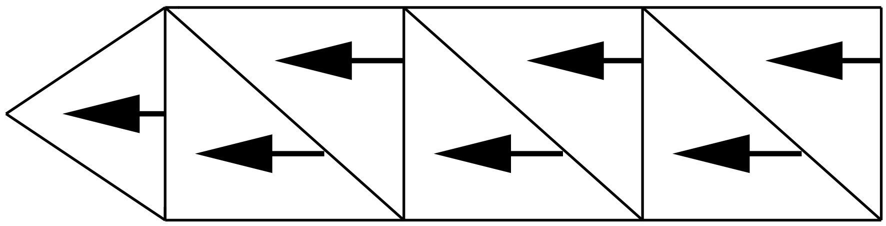

Figure 12

indicated gradient path is the only gradient path from $\partial\beta$ to α, then we can reverse the gradient vector field along this path, replacing the figure by the vector field in Figure 12. Theorem 6 immediately implies that this new vector field is the gradient vector field associated to some Morse function, and α and β are no longer critical. 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 $\partial\beta$ to α, one must also adjust all nearby vectors so that the resulting vector field is smooth. Moreover, one must check that the new vector field is the gradient of a function, 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.

We end this section by mentioning that much of the discrete Morse theory implies the corresponding smooth theory.

Theorem 12 ([Fo8]). *Let M be a smooth manifold with a smooth Morse function. Then there is a C^1 triangulation of M and a discrete Morse function on the resulting simplicial complex which has the same Morse data as the original function (that is, the same number of critical points, and isomorphic Morse complexes).*

2. Differential Geometry

In this section, we borrow ideas from Riemannian geometry, and define local "curvature" invariants for cell complexes from which one can deduce global topological properties. In particular, given a cell complex M, we define a combinatorial Ricci curvature, which is a function on the set of edges of M. The Ricci curvature of an edge is completely determined by those cells of M which are neighbors of e. If the Ricci curvature is always positive, or nonnegative, then one can deduce strong statements about the topology of M.

Much fascinating and beautiful mathematics has resulted from earlier efforts to extend the basic notions of curvature to combinatorial spaces. In the mathematics literature, one can see, for example, [AW] [Ba1,2], [Ch], and [Sto1,2]. In the mathematical physics literature such ideas appear, most notably, in the voluminous work on the so-called Regge calculus (see [CMS],[Re],[Wi] for an introduction). Most (but not all) of this work begins with a rigid presentation of the cell complex, which we think of as a piecewise flat Riemannian manifold, so that the curvature is localized to the lower dimensional cells. The curvature is then defined in terms of how the neighborhood of a cell differs from what one would see in a flat complex. This section is also related to efforts to derive local combinatorial formulas for the characteristic classes of combinatorial manifolds (see [BM],[GGL],[GM],[Le],[LR],[Mac], [Sto],[Yu]). In the case of smooth manifolds, such local formulas are expressed in terms of curvature, so, presumably, local combinatorial formulas for characteristic classes can be thought of as combinatorial formulas for curvature (this connection is made explicit in [Ch]). In the mathematical physics literature, similar questions have been studied under the name of lattice gauge theory (e.g. [PS]). The main result of [LS] is also similar in spirit to the results we will present, in that the combinatorics of the edges and faces of a combinatorial 3-manifold are related to the topology of the underlying manifold. One should also compare with [Al], [Gr2], and the references therein, in which one defines notions of curvature on combinatorial spaces by comparing distances between points with the distances one would find on smooth Riemannian manifolds with constant curvature.

As will soon become clear, our approach is completely different from those followed in these works, and is based instead on an analysis of the combinatorial Laplace operator (this operator also appears in [Yu], but plays a very different role there). The combinatorial Laplace operator has a long history. In the case of a 1-dimensional simplicial complex, this operator appeared, implicitly, in Kirchoff's work on electrical networks [Ki]. See [Chu] (and[Fo1]) for more recent work on this case. The combinatorial Laplace operator on general complexes was introduced by Eckmann in [Ec]. The close relationship between the combinatorial and Riemannian Laplace operators has been explored in a number of papers, most notably by Dodziuk and his collaborators (e.g. [Do1,2],[DK],[DP]). The work by Garland in [Ga1,2,3], is very closely related to ours, in that he also uses the combinatorial Laplace operator to define combinatorial curvatures, which are then used to deduce global topological properties of the complex.

Let us begin with some examples of the theory we develop. Although the theory can be applied to very general cell complexes, in this introduction we will restrict attention to CW complexes satisfying a combinatorial condition modelled on the notion of convexity. The reader may think only of simplicial complexes, but it is insightful to allow more general convex cells.

Definition 4. *A regular CW complex M is quasi-convex if for each pair of p-cells α_1, α_2, if either $\overline{\alpha}_1 \cap \overline{\alpha}_2 = \emptyset$, or $\overline{\alpha}_1 \cap \overline{\alpha}_2 = \overline{\gamma}$ for some cell γ. For example, all simplicial complexes and all polyhedra are quasi-convex. From now on, by a quasi-convex complex we will always mean a compact regular CW complex which is quasi-convex.*

Throughout this section, a primary role will be played by the notion of a *neighbor*.

Definition 5. *If α_1 and α_2 are p-cells of M, say α_1 and α_2 are neighbors if*

(1) *α_1 and α_2 share a $(p+1)$-cell, that is, there is a $(p+1)$-cell β with $\beta > \alpha_1$ and $\beta > \alpha_2$,*
or
(2) *α_1 and α_2 share a $(p-1)$-cell, that is, there is a $(p-1)$-cell γ with $\gamma < \alpha_1$ and $\gamma < \alpha_2$.*

We say that α_1 and α_2 are parallel neighbors if either (1) or (2) is true but not both. If both (1) and (2) are true, we say α_1 and α_2 are transverse neighbors.

The example shown in Figure 13 illustrates the origin of these terms. This is an illustration of the case $p = 1$. The edge e_0 has 8 neighbors (e_1 through e_8). Of these, 4 are parallel neighbors (e_2, e_4, e_6 and e_8). The edges e_2 and e_6 share a vertex with e_0 but not a face, and the edges e_4 and e_8 share a face with e_0 but not a vertex.

We are now ready to state the main new definition of this section.

Definition 6. *For each p, define the p^{th} curvature function*

$$\mathcal{F}_p : K_p \to \mathbf{R}$$

as follows. For any p-cell α, set

$$\mathcal{F}(\alpha) = \#\{(p+1)\text{-}cells\ \beta > a\} + \#\{(p-1)\text{-}cells\ \gamma < a\} - \#\{parallel\ neighbors\ of\ \alpha\}.$$

The significance of this definition is made apparent in the following combinatorial analogue of "Bochner's Theorem" ([Bo1,2,3]).

Theorem 13. *Let M be a quasi-convex complex. Suppose, for some p, that $\mathcal{F}(\alpha) > 0$ for each p-cell α. Then*

$$H_p(M, \mathbf{R}) = 0.$$

Of particular interest is the case $p = 1$. In this case we give the curvature function a special name. For any edge (i.e. 1-cell) e of M, define the *Ricci curvature* of e by

$$\mathrm{Ric}(e) = \mathcal{F}_1(e).$$

More explicitly,

$$\mathrm{Ric}(e) = \#\{\text{2-cells } f > e\} + 2 - \#\{\text{parallel neighbors of } e\}.$$

For example, in Figure 13

$$\mathrm{Ric}(e_0) = 2 + 2 - 4 = 0.$$

Thus, Figure 13 is a "Ricci flat" cell complex. In Figure 14(i) we show a 2-dimensional sphere with "cubic" cell decomposition. In this case the edge e_0 has 2 parallel neighbors (the edges e_1 and e_2) so that

$$\mathrm{Ric}(e_0) = 2 + 2 - 2 = 2.$$

One purpose of Figure 14(ii) is to demonstrate that the definitions and theorems discussed in this paper can, with a few exceptions, be applied to very general complexes. The complex need not be a manifold, or have any other special structure. In Figure 14(ii) we have indicated the Ricci curvature of each edge.

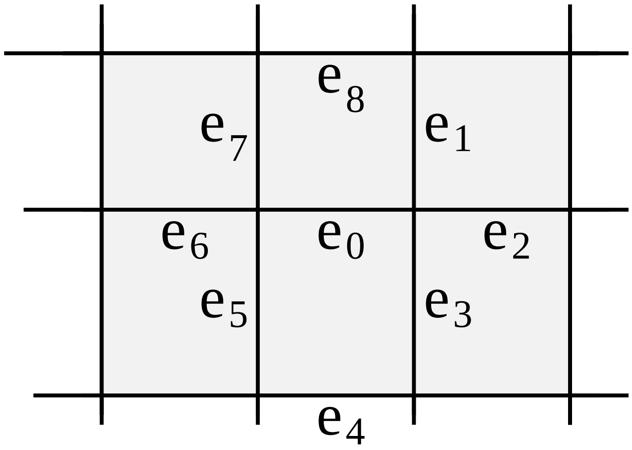

Figure 13

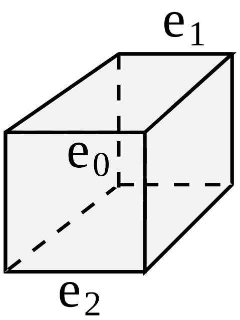

(i)

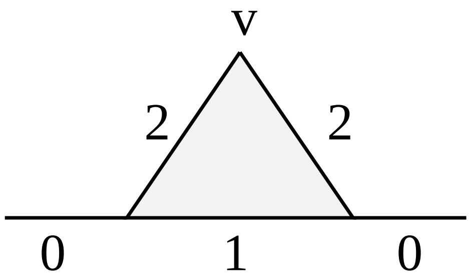

(ii)

Figure 14

In the Riemannian case Bochner showed that one does not need the p^{th} curvature function to be strictly positive, as we required in Theorem 13, in order to draw some topological conclusions. In fact, non-negativity also has strong implications. In the combinatorial setting, we can prove a corresponding result in the case p=1.

Theorem 14. *Let M be a connected quasi-convex complex.*
(i) Suppose that $\mathrm{Ric}(e) \geq 0$ *for all edges e of M. If there is a vertex v such that* $\mathrm{Ric}(e) > 0$ *for all $e > v$ then*

$$H_1(M, \mathbf{R}) = 0.$$

(ii) Suppose that M is a combinatorial n-manifold, $n \leq 3$, and $\operatorname{Ric}(e) \geq 0$ *for all edges of M, then*

$$\dim H_1(M, \mathbf{R}) \leq n.$$

(This theorem is true for all $n \geq 0$ with an additional technical hypothesis.)

We will later indicate the proof of these results. We conjecture that, just as in the Riemannian setting, Theorem 14 holds for general p (with the right hand side of the inequality in (ii) replaced by $\binom{n}{p}$). However, we have run into interesting obstacles in trying to prove Theorem 14 in the desired generality. We will later give an idea as to how $p = 1$ differs from the general case. For a more complete discussion of this point, see §4 of [Fo6].

It is interesting to see how these theorems apply to the very simple complexes shown in Figures 13 and 14. In the case of the 2-sphere S^2 given the cell decomposition illustrated in Figure 14(i), each edge has Ricci curvature 2. Thus we may apply Theorem 13 to conclude

$$H_1(S^2, \mathbf{R}) = 0.$$

Each edge of the complex M in Figure 14(ii), has Ricci curvature ≥ 0. Moreover, every edge $e > v$ has positive Ricci curvature. In this case, Theorem 14(i) implies

$$H_1(M, \mathbf{R}) = 0.$$

The 2-dimensional torus T^2 can be given a cell decomposition which looks everywhere like the cell complex shown in Figure 13. With this cell structure T^2 is a connected combinatorial 2-manifold. Since every edge has Ricci curvature $= 0$, we may apply Theorem 14(ii) to learn

$$\dim H_1(T^2, \mathbf{R}) \leq 2.$$

More generally, for every n, the n-dimensional torus T^2 can be given such a cubic cell decomposition decomposition. In this case, T^n will be a connected combinatorial n-manifold, with every edge having Ricci curvature 0. Moreover, T^n satisfies the (unstated) technical hypothesis of the above theorem, so we may conclude

$$\dim H_1(T^n, \mathbf{R}) \leq n.$$

Since, in fact, $\dim H_1(T^n, \mathbf{R}) = n$, we see that the conclusion of this combinatorial Bochner theorem, just as the conclusion of the classical Bochner theorem, is sharp.

In [My], a precursor to Bochner's work, Myers proves a theorem (in the Riemannian setting) which is a strengthening of Theorem 13 in the case $p = 1$. The analogous result holds in the combinatorial setting.

Theorem 15. *Let M be a quasi-convex complex. Suppose that for every edge e,* $\operatorname{Ric}(e) > 0$. *Then $\pi_1(M)$ is finite.*

In [Sto1,2] a different notion of curvature is defined, in a somewhat more restricted setting, from which a version of Myers' theorem is deduced. Even though the proof of Theorem 15 is very similar to the proof in [Sto1,2] (and the proof in [My]), the relationship between the two notions of combinatorial curvature is not clear to this author at present.

These theorems show that there are strong topological restrictions for a topological space to have a quasi-convex cell decomposition with $\mathrm{Ric}(e) > 0$ (or ≥ 0) for each edge e. It is natural to consider the other extreme, and to investigate the possibility of decompositions with $\mathrm{Ric}(e) < 0$ for every edge e. In [Gao],[GY] and [Lo1,2] it was shown that every smooth manifold of dimension ≥ 3 has a Riemannian metric with everywhere negative Ricci curvature. The same, and more, is true in the combinatorial setting.

Theorem 16. *Let M be a (not necessarily compact) combinatorial n-manifold with $n \geq 2$. Then there is a finite subdivision M^* of M such that $\mathrm{Ric}(e) < 0$ for all edges e of M^*.*

It is interesting to note that the corresponding Riemannian theorem is only true for manifolds of dimension ≥ 3. The Gauss-Bonnet Theorem provides an obstruction for a surface to have a metric with negative Ricci ($=$ Gauss if $n = 2$) curvature. Theorem 16 implies that there is no simple Gauss-Bonnet theorem for combinatorial Ricci curvature. See §5 of [Fo6] for a way to recover a version of Gauss-Bonnet in this setting.

We will now give an idea of the proofs of Theorems 13 and 14. The readers familar with Bochner's original proof will notice that we are simply following his proof (now known as "Bochner's method") in a combinatorial setting. Let M be a finite CW complex. Consider the cellular chain complex

$$0 \longrightarrow C_n(M, \mathbf{R}) \xrightarrow{\partial} C_{n-1}(M, \mathbf{R}) \xrightarrow{\partial} \cdots \xrightarrow{\partial} C_0(M, \mathbf{R}) \longrightarrow 0$$

Endow each $C_p(M, \mathbf{R})$ with a positive definite inner product such that the cells of M are orthogonal. This involves choosing a positive "weight" ω_α for each cell α and setting

$$\langle \alpha, \alpha \rangle = \omega_\alpha.$$

(The formulas presented in this paper result from setting $\omega_\alpha = 1$ for every cell α.) Let $\partial^* : C_p(M, \mathbf{R}) \longrightarrow C_{p+1}(M, \mathbf{R})$ denote the adjoint of ∂ and

$$\Box_p = \partial\partial^* + \partial^*\partial : C_p(M, \mathbf{R}) \longrightarrow C_p(M, \mathbf{R})$$

the corresponding Laplacian. The combinatorial Hodge theorem

$$\mathrm{Ker}\,\Box_p \cong H_p(M, \mathbf{R})$$

follows from basic linear algebra.

We now make a brief aside. The next step in Bochner's original proof is the Bochner-Weitzenbock formula

$$\Box_p = (\nabla_p)^* \nabla_p + F_p$$

where now $\Box_p$ is the Riemannian Laplace operator on p-forms on a compact Riemannian manifold, ∇_p is the (Levi-Civita) covariant derivative operator on $\Omega^p(M)$, and F_p is a 0^{th} order operator whose value at $m \in M$ depends only on derivatives of the Riemannian metric at m. The operator $(\nabla_p)^*\nabla_p$ is called the Bochner Laplacian. The main point is that this is a non-negative operator, in the sense that for every p-form ω on M

$$\langle(\nabla_p)^*\nabla_p\omega, \omega\rangle = \langle\nabla_p\omega, \nabla_p\omega\rangle \geq 0$$

where $\langle\, ,\, \rangle$ is the L^2- inner product on p-forms induced by the Riemannian metric. It immediately follows that if F_p is a positive operator, in the sense that $m \in M$ and all $0 \neq \omega \in \Lambda^p T^*_m M$

$$\mathcal{F}_p(\omega) \equiv (F_p(m)\omega, \omega) > 0,$$

then $\Box_p$ is a positive operator, and hence the kernel is trivial. This implies that $H_p(M, \mathbf{R}) = 0$, the desired conclusion of Theorem 13. For $p > 1$, the function $\mathcal{F}_p$ is rather mysterious. On the other hand, we can recognize $\mathcal{F}_1$ as Ricci curvature. That is, for all $\omega \in T^*M$

$$\mathrm{Ric}(\omega) = \mathcal{F}_1(\omega). \tag{3}$$

Now we go back to the combinatorial setting. Following Bochner's proof, our next step is to develop a Bochner-Weitzenbock formula for $\Box_p$, the combinatorial Laplace operator. This is rather mysterious, since we begin by knowing neither a combinatorial analogue for $(\nabla_p^*)\nabla_p$ nor for F_p. However, we show that there is, in fact, a canonical decomposition

$$\Box_p = B_p + F_p$$

where B_p is a non-negative operator and, when expressed as a matrix with respect to a basis of $C_p(M, \mathbf{R})$ consisting of the p-cells of M, F_p is diagonal.

To describe this decomposition, we begin by defining a decomposition of any $n \times n$ symmetric matrix. Rather than writing a general definition, we will illustrate the decomposition in the case of a symmetric 3×3 matrix.

$$\begin{aligned} \begin{pmatrix} a & b & c \\ b & d & e \\ c & e & f \end{pmatrix} &= \begin{pmatrix} |b|+|c| & b & c \\ b & |b|+|e| & e \\ c & e & |c|+|e| \end{pmatrix} \\ &+ \begin{pmatrix} a-(|b|+|c|) & 0 & 0 \\ 0 & d-(|b|+|e|) & 0 \\ 0 & 0 & f-(|c|+|e|) \end{pmatrix} \\ &= B + F. \end{aligned} \tag{4}$$

The linear map defined by the matrix B is non-negative (with respect to the standard inner product on R^3). This is a special case of the fact that any symmetric matrix $(a_{i,j})$ such that $a_{i,i} \geq \sum_{j \neq i} |a_{i,j}|$ for each i is non-negative. This is one of those very useful facts in matrix theory which is constantly being rediscovered. For the history of this result (not stated in quite the same way) and extensions see [Ta]. The desired conclusion is also implied by the classical eigenvalue estimate known as "the Gersgorin circle method" [Ge], which says that the eigenvalues lie in circles centered at $a_{i,i}$ with radius $\sum_{j \neq i} |a_{i,j}|$.

The non-negativity of B implies that in this simple example we can already make the Bochner-like statement that if each of the 3 numbers $a-(|b|+|c|)$, $d-(|b|+|e|)$, $f-(|c|+|e|)$ is positive then the original 3×3 matrix has a trivial kernel.

To apply this decomposition to the combinatorial Laplace operator, we must first represent the Laplace operator as a symmetric matrix, and this requires choosing an ordered orthonormal basis for the space of chains. We can define such a basis by choosing an orientation for each cell of M, and then choosing an ordering of the cells. Making a different choice for the orientation of a cell has the effect of multiplying the corresponding row and column of the symmetric matrix by -1. Permuting the ordering of the cells has the effect of applying the same permutation to both the rows and the columns of the matrix. Therefore, for a decomposition of symmetric matrices to induce a well-defined decomposition of the combinatorial Laplace operator, the decomposition must be equivariant under multiplying a row and column by -1, and applying a single permutation to both the rows and the columns. It is easy to see that the decomposition shown in (4) has this property, and hence induces a decomposition $\Box_p = B_p + F_p$. In analogy with the Bochner-Weitzenbock formula, we call B_p the combinatorial Bochner Laplacian, and F_p the p^{th} combinatorial curvature function.

It follows immediately that if $F_p > 0$ then $\Box_p > 0$ so that $H_p(M, \mathbf{R}) = 0$. For any p-cell α, we define

$$\mathcal{F}_p(\alpha) = < F_p(\alpha), \alpha > .$$

In analogy with formula (3) we define Ricci curvature by setting, for each edge e of M

$$\mathrm{Ric}(e) = \mathcal{F}_1(e).$$

We note that we have no independent definition of combinatorial Ricci curvature, so that (3) is a theorem in the Riemannian setting, but a definition in the combinatorial case. Proving Theorem 13 is now reduced to explicitly calculating $\mathcal{F}_p$ as defined above, and showing that this is equivalent to the formula given in Definition 6. See [Fo6] for the details of this straight-forward calculation.

In the Riemannian case (I hope the reader is not getting dizzy from all of this bouncing back and forth between categories), Theorem 14 follows from the Bochner-Weitzenbock formula with just a bit more work. Namely, if F_p is a non-negative operator, then since $(\nabla_p)^*\nabla_p$ is also a non-negative operator, we have

$$\mathrm{Ker}(\Box_p) = \mathrm{Ker}((\nabla_p)^*\nabla_p) \cap \mathrm{Ker}(F_p).$$

We note that

$$\mathrm{Ker}((\nabla_p)^*\nabla_p) = \mathrm{Ker}(\nabla_p).$$

Any $\omega \in \mathrm{Ker}(\nabla_p)$ is said to be *parallel*, and is completely determined by its value at any one point. This is essentially all that is required to complete the proof.

The proofs of Theorem 14(i) and (ii) in the combinatorial case, just as in the Riemannian case, would follow if we knew that a 1-chain $c \in \mathrm{Ker}\, B_1$ is completely determined by its values in some neighborhood in M, that is, if there was a unique continuation theorem for 1-chains in $\mathrm{Ker}\, B_1$. This is not quite true. However, if $c \in \mathrm{Ker}\,\Box_1 \cap \mathrm{Ker}\, B_1$, then c is, in fact, determined by its value in a neighborhood of any vertex. This is sufficient for the proof of Theorem 14(i) and (ii) . It can also be shown that for $p > 1$, p-chains $c \in \operatorname{Ker} \Box_p \cap \operatorname{Ker} B_p$, in drastic contradiction to the Riemannian case, are not determined by local information. This explains the difficulty in extending Theorem 14 to the case of general p (see §4 of [Fo6]).

We briefly indicate the proof of the unique continuation theorem for $c \in \operatorname{Ker} \Box_1 \cap \operatorname{Ker} B_1$. Consider, for example, a vector $x = (x_1, x_2, x_3)$ in the kernel of the matrix B in (4). It is not too hard to see that if $b \neq 0$ then x_1 determines x_2. Namely,

$$x_2 = -(sign(b))x_1.$$

More generally, if $x \in \operatorname{Ker} B$, and the $(i, j)^{th}$ entry of B is non-zero, then x_i determines x_j (so that x_i is zero if and only if x_j is zero).

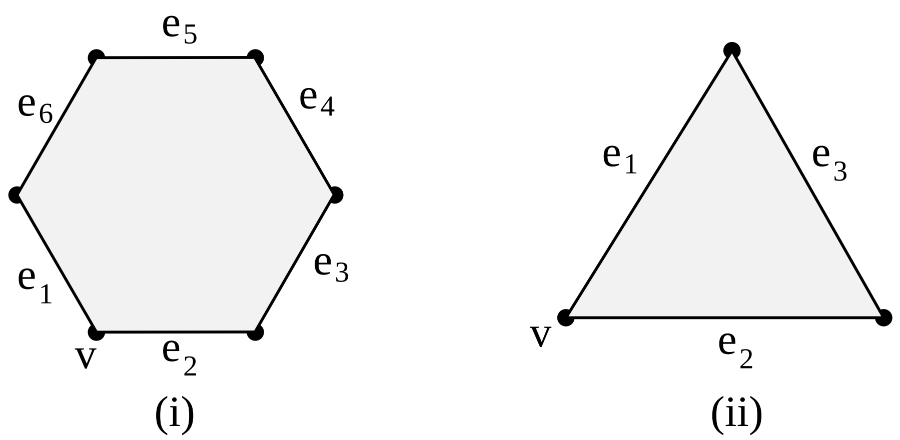

Figure 15

In the case of the Laplace operator, if α_1 and α_2 are p-cells, then the entry of B_p corresponding to α_1 and α_2 is non-zero if α_1 and α_2 are parallel neighbors. Therefore, if $c = \sum_{\alpha^{(p)}} c_\alpha \alpha \in \operatorname{Ker} B_p$, and α_1 and α_2 are parallel neighbors, then c_{α_1} determines c_{α_2}. Consider the polygon shown in Figure 15 (i). Every side of the polygon is a parallel neighbor of either e_1 of e_2. This is the case whenever the polygon has at least 4 sides. Therefore, if $c = \sum_e c_e e \in \operatorname{Ker} B_1$, then the value of c on the boundary of the polygon is completely determined by c_{e_1} and c_{e_2}. This argument doesn't quite work if the polygon has 3 sides, as in Figure 15 (ii). Here, e_3 is a transverse neighbor of both e_1 and e_2. However, if we knew that $c \in \operatorname{Ker} \partial^*$, then c_{e_3} would be determined by c_{e_1} and c_{e_2}. Putting these ideas together, it is not too hard to get to the general result that if M is connected, and $c \in \operatorname{Ker} \partial^* \cap \operatorname{Ker} B_1$, then for any vertex v, c is completely determined by c_e for $e > v$. Since

$$\operatorname{Ker} \Box_p = \operatorname{Ker} \partial \cap \operatorname{Ker} \partial^* \subset \operatorname{Ker} \partial^*$$

the desired theorem follows.

To see how Theorem 14(i) follows, suppose $\mathrm{Ric}(e) \geq 0$ for each edge e. Then, since B_1 is also a non-negative operator,

$$\mathrm{Ker}(\Box_1) = \mathrm{Ker}(B_1) \cap \mathrm{Ker}(\mathrm{Ric}) = (\mathrm{Ker}(B_1) \cap \mathrm{Ker}(\Box_1)) \cap \mathrm{Ker}(\mathrm{Ric})$$

$$= (\mathrm{Ker}(B_1) \cap \mathrm{Ker}(\partial^*)) \cap \mathrm{Ker}(\mathrm{Ric}).$$

Now suppose $c = \sum_e c_e e \in \mathrm{Ker}(\Box_p)$. Then $c \in \mathrm{Ker}(\mathrm{Ric})$. Since $\mathrm{Ric}(e) > 0$ for each edge $e > v$, we must have $c_e = 0$ for each edge $e > v$. Now we observe that since $c \in \mathrm{Ker}(B_1) \cap \mathrm{Ker}(\partial^*)$, c is completely determined by c_e for $e > v$, and hence we must have $c = 0$. This demonstrates that $H_1(M, \mathbf{R}) = 0$.

The results we have presented here, as well as those in the references cited earlier, indicate that there should exist more complete theories of combinatorial differential topology and geometry. However, at this point the combinatorial theories capture only a small portion of these beautiful and far-reaching classical subjects. There remains much to be done.

References

[Al] A.D. Alexandrov, *A Theorem on Triangles in a Metric Space and some of its Applications*, Trudy Math. Inst. Steks **38** (1951), 5–23.

[AW] C. Allendoerfer and A. Weil, *The Gauss-Bonnet Theorem for Riemannian Polyhedra*, Trans. A.M.S. **53,** (1943,), 101–129.

[BBLSW] E. Babson, A. Björner, S. Linusson, J. Shareshian and V. Welker, *Complexes of not i-Connected Graphs*, to appear in Topology.

[Ba1] T. Banchoff, *Critical Points and Curvature for Embedded Polyhedra*, J. Diff. Geo. **1** (1967), 245–256.

[Ba2] ______, *Critical Points and Curvature for Embedded Polyhedra II*, in Differential Geometry(COllege Park, Md, 1981/1982), Progress in Mathematics, vol. 32. Birkäuser Boston, Boston, Mass. (1983), 34–55.

[BM] T. Banchoff and C. McCrory, *A Combinatorial Formula for Normal Stiefel-Whitney Classes*, Proc. Am. Math. Soc. **76** (1979), 171–177.

[Bj] A. Björner, *Topological Methods*, in Handbook of Combinatorics, Vol. 1,2, Elsevier, Amsterdam (1995), 1819–1872.

[Bo1] S. Bochner, *Vector Fields and Ricci Curvature*, Bull. Am. Math. Soc. **52** (1946), 776–797.

[Bo2] ______, *Curvature and Betti Numbers*, Ann. Math. **49** (1948), 379–390.

[Bo3] ______, *Curvature and Betti Numbers II*, Ann. Math. **50** (1949), 77–93.

[Bol] B. Bollobás, *Extremal Graph Theory*, in Handbook of Combinatorics, Vol.1,2, 1231–1292, Elsevier, Amsterdam, (1995)..

[Bo] R. Bott, *Morse Theory Indomitable*, Publ. Math. I.H.E.S. **68** (1988), 99–117.

[BK] U. Brehm and W. Kühnel, *Combinatorial Maifolds with Few Vertices*, Topology **26** (1987), 465–473.

[Cha] M. Chari, *On Discrete Morse Functions and Combinatorial Decompositions*, preprint.

[Ch] J. Cheeger, *Spectral Geometry of Singular Riemannian Spaces*, J. Diff. Geo. **18** (1983), 575–657.

[CMS] J. Cheeger, M. Muller, and R. Schrader, *On the Curvature of Piecewise Flat Spaces*, Comm. Math. Phys. **92** (1984), 405–454.

[Chu] F. Chung, *Spectral Graph Theory*, CBMS regional Conference Series in Mathematics # 92, AMS, 1997.

[Do1] J. Dodziuk, *Combinatorial and continuous Hodge theories*, Bull. A.M.S. **80** (1974,), 1014–1016.

[Do2] ______, *Sobolev spaces of differential forms and de Rham Hodge isomorphism*, J. Diff. Geo. **16** (1981,), 63–73.

[DK] J. Dodziuk and L. Karp, *Spectral and function theory for combinatorial Laplacians*, in Geometry of Random Motion (Ithaca, N.Y., 1987), Contemp. Math. 73, Pitman Res. Notes Math. Ser. 150., Longman Sci. Tech., 1986, 25–40.

[DP] J. Dodziuk and V. Patodi, *Riemannian structures and triangulations of manifolds*, J. Indian Math. Soc. **40** (1976,), 1–52.

[Du] A. Duval, *A Combinatorial Decomposition of Simplicial Complexes*, Israel J. of Math. **87** (1994), 77–87.

[Ec] B. Eckmann, *Harmonische Funktionen und Randwertaufgaben in einem Komplex*, Comm. Math. Helvetici **17** (1945), 240–255.

[EK] J. Eells, Jr. and N. Kuiper, *Manifolds Which are Like Projective Planes*, Inst. Hautes Études Sci. Publ. Math. No. 14, (1962), 5–46.

[Fl] A. Floer, *Morse Theory for Lagrangian Intersections*, J. Diff. Geo. **28** (1988), 513-547.

[Fo1] R. Forman, *Determinants of Laplacians on Graphs*, Topology **32** (1993), 35-46.

[Fo2] ______, *A Discrete Morse Theory for Cell Complexes*, in Geometry, Topology & Physics for Raoul Bott, S.T. Yau (ed.), International Press, 1995.

[Fo3] ______, *Morse Theory for Cell Complexes*, Adv. in Math **. 134** (1998), 90–145.

[Fo4] ______, *Witten-Morse Theory for Cell Complexes*, Topology **. 37** (1998), 945–979.

[Fo5] ______, *Combinatorial Vector Fields and Dynamical Systems*, to appear in Math. Zeit.

[Fo6] ______, *Combinatorial Ricci Curvature*, preprint.
[Fo7] ______, *Morse Theory and Evasiveness*, preprint.
[Fo8] ______, *On Triangulations of Vector Fields*, in preparation.
[GGL] A.M. Gabrielov, I.M. Gelfand, and M.V. Losik, *The Combinatorial Computation of Charactersitic Classes*, Func. Anal. and Appl. **10** (1976), 12–15.
[Gao] L.Z. Gao, *The Construction of Negatively Ricci Curved Manifolds*, Math. Ann. **271** (1985), 185–208.
[GY] L.Z. Gao and S.-T. Yau, *The Existence of Negatively Ricci Curved Metrics on Three-Manifolds*, Invent. Math. **85** (1986), no. 3, 637–652.
[Ga1] H. Garland, *p-adic Curvature and a Conjecture of Serre.*, Bull. Am. Math. Soc. **78** (1972), 259-261.
[Ga2] ______, *p-adic Curvature and the Cohomology of Discrete Subgroups of p-adic Groups*, Ann. of Math. **97** (1973), 375–423.
[Ga3] ______, *On the cohomology of discrete subgroup of semi-simple Lie groups*, in Discrete Subgroups of Lie Groups and Applications to Moduli (Internat. Colloq., Bombay, 1973), 9pp. 21-30. Oxford Uniov. Press, Bombay, 1975.
[GM] I.M. Gelfand and R.D. MacPherson, *A Combinatorial Formula for the Pontrjagin Classes*, Bull. Am. Math. Soc. **26** (1992), 304–309.
[Ge] S. Gersgorin, *Über die Abgrenzung der Eigenverte einer Matrix*, Izv. Akad. Nauk., USSR, **7** (1931), 749–754.
[Gl] L.C. Glaser, *Geometrical Combinatorial Topology* I, Van Nostrand Reinhold Mathematical Studies #28, Van Nostrand Reinhold, 1972.
[GoM1] 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), 517-533.
[GoM2] ______, *Stratified Morse Theory*, Ergebnisse der Mathematik und ihrer Grenzgebeite (3), 14, Springer Verlag, Berlin-New York, 1988.
[Gr1] M. Gromov, *Cyrvature, Diameter and Betti Numbers*, Comment. Math. Helv. **56** (1981), 179–195.
[Gr2] ______, *Hyperbolic Groups*, in "Essays in Group Theory," S.M. Gersten (ed.), M.S.R.I. Publ. 8 (1987), 75-263..
[HT] S. Halperin and D. Toledo, *Stiefel-Whitney Homology Classes*, Ann. of Math.(2) **96** (1972), 511–525.
[Ki] G. Kirchoff, *Über die Auflösung der Gleichungen auf Welche Man beider Untersuchen der Linearen Vertheilung Galvanischer Ströme Gefüht Wird*, Ann. der Phys. und Chem. **72**, 495–508.
[Kl] W. Klingenberg, *The Morse Complex*, Symposia Mathematica, Vol. XXVI (Rome, 1980), Academic Press (1982), 117-122.
[Kü] W. Kühnel, *Triangulations of Manifolds with few Vertices*, in Advances in Differential Geometry and Topology, World Sci. Publishing, N.J., (1990), 59–114.
[Kui] N. Kuiper, *Morse Relations for Curvature and Tightness*, in Proceedings of the Liverpool Singularities Symposium (1969/1970), Lecture Notes in Math **209** (1971), 77–89.
[Le] N. Levitt, *Grassmannians and Gauss Maps in Piecewise-Linear Topology*, Lecture Notes in Mathematics #1366, Springer-Verlag, 1989.
[LR] N. Levitt and C. Rourke, *The Existence of Combinatorial Formulae for Characteristic Classes*, Trans. Am. Math. Soc. **239** (1978), 391–397.
[Lo1] J. Lohkamp, *Metrics of Negative Ricci Curvature*, Ann. of Math.(2) **140** (1994), no. 3, 655–683.
[Lo2] ______, *Negative Bending of Open Manifolds*, J. of Diff. Geo. **40** (1994), 461–474.
[LS] F. Luo and R. Stong, *COmbinatorics of Triangulations of 3-Manifolds*, Trans. Amer. Math. SOc. **337** (1993), 891–906.
[Mac] R.P. MacPherson, *Combinatorial Differentiable Manifolds*, in Topological Methods in Modern Mathematics (Stony Brook, NY, 1991), Publish or Perish, Houston, TX, (1993), 203–221.
[Mi1] J. Milnor, *Morse Theory*, Annals of Mathematics Study No. 51, Princeton University Press, 1962.
[Mi2] ______, *Lectures on the h-Cobordism Theorem*, Princeton Mathematical Notes, Princeton University Press, 1965.

[Mo1] M. Morse, *Relations Between the Critical Points of a Real Function of n Independent Variables*, Trans. Amer. Math. Soc. **27** (1925), 345–396.

[Mo2] ———, *The Calculus of Variations in the Large*, Amer. Math. Soc., Colloquium Publications **18** (1934).

[Mo3] ———, *Bowls of a Non-Degenerate Function on a Compact Differentiable Manifold, in*, Differential and Combinatorial Topology (A Symposium in Honor of M. Morse), Princeton University Press (1965), 81–104.

[My] S.B. Myers, *Riemannian Manifolds with Positive Mean Curvature*, Duke Math J. **8** (1941), 401–404.

[PS] A. Phillips and D. Stone, *The Computation of Charactersitic Classes of Lattice Gauge Fields*, Comm. Math. Phys. **131** (1990), 255–282.

[Re] T. Regge, *General Relativity Without Coordinates*, Nuovo Cimento **19** (1961), 558–571.

[Sc] M. Schwartz, *Morse Homology*, Progress in Mathematics, 111, Birkhäuser verlag, Basel (1993).

[Sh] J. Shareshian, *Discrete Morse Theory for Complexes of 2-Connected Graphs*, preprint.

[Sm1] S. Smale, *On Gradient Dynamical Systems*, Annals of Math. **74** (1961), 199–206.

[Sm2] ———, *The Generalized Poincaré Conjecture in Dimensions Greater than Four*, Annals of Math. **74** (1961), 391–406.

[St] R. Stanley, *A Combinatorial Decomposition of Acyclic Simplicial Complexes*, Discrete Math. **118** (1993).

[Sto1] D. Stone, *A Combinatorial Analogue of a Theorem of Myers*, Ill. J. of Math. **20** (1976), 12–21.

[Sto2] ———, *Correction to my Paper: "A Combinatorial Analogue of a Theorem of Myers"*, Ill. J. of Math. **20** (1976), 551–554.

[Sto3] ———, *On the Combinatorial Gauss Map for C^1 Submanifolds of Euclidean Space*, Topology **20** (1981), 247–272.

[Ta] O. Taussky, *A recurring Theorem on Determinants*, Am. Math. Monthly **56** (1949), 672–676.

[Tu] V. Turchun, *Homology of Complexes of Biconnected Graphs*, Uspekhi Mat. Nauk. **52** (1997), 189–190.

[Va] V. Vassiliev, *Complexes of Connected Graphs*, in The Gelfand Mathematical Smeinars, 1990–1992, pp. 223–235. Birkhäuser Boston, Boston, (1993)..

[Wh] J.H.C. Whitehead, *Simplicial Spaces, Nuclei, and m-Groups*, Proc. London Math. Soc..

[Wi] R.M. Williams, *Discrete Quantum Gravity: The Regge Calculus Approach*, Topological and Quantum Group Methods in Field Theory and Condensed Matter Physics, Int. J. of Mod. Phys. B **6** (1992), 2097–2108.

[Yu] Y.L. Yu, *Combinatorial Gauss-Bonnet-Chern Formula*, Topology **22** (1983), 153–163.

[Zi] G. Ziegler, *Lectures on Polytopes*, Graduate Texts in Mathematics, 152. Springer-Verlag, New York, 1995.