MORSE THEORY AND EVASIVENESS

Robin Forman

Department of Mathematics
Rice University
Houston, TX 77251
forman@math.rice.edu

§0. Introduction.

Consider a game played by 2 players, whom we call the hider and the seeker. Let S be a simplex of dimension n, with vertices v_0, $v_1, \ldots, v_n$, and M a subcomplex of S, known to both the hider and the seeker. Let σ be a face of S, known only to the hider. The seeker is permitted to ask questions of the sort "Is vertex v_i in σ?" The seeker's goal is to determine whether σ is in M, using as few questions as possible.

The seeker is permitted to use the answers to the earlier questions when choosing which vertex to ask about next. We assume that the seeker chooses each question, given the answers to the previous questions, according to a deterministic algorithm, which we call a decision tree algorithm. For any decision tree algorithm A, we denote by $Q(\sigma, A, M)$ the number of questions the seeker must ask before it is determined whether or not σ is in M. The *complexity* of M, $c(M)$, is defined by

$$c(M) = \inf_A \sup_\sigma Q(\sigma, A, M).$$

M is said to be *evasive* if $c(M) = n+1$, and *nonevasive* otherwise. That is, M is evasive if for any decision tree algorithm A there is a face σ of S such that the seeker must ask $n+1$ questions before determining if σ is in M.

This decision problem has been the focus of numerous investigations, because it touches upon many subjects, such as complexity theory, graph theory, and topological combinatorics. See, for example, the surveys in section 4.5 of [Bo] and section I.2 of [Bj]. Most of the work has concentrated on the important special case of determining whether or not a graph satisfies a monotone graph property when one is permitted to check one edge at a time. The general problem we have presented above seems to have first appeared in [RV1].

Partially supported by the National Science Foundation and the National Security Agency. The author would like to thank the referees for their very thoughtful comments.

In [KSS] it was shown that there is a fundamental relationship between the evasiveness of the subcomplex M and its topology.

Theorem 1 [KSS]. *If M is nonevasive then M is collapsible.*

See [Gl], page 49, for the definition and basic properties of simplicial collapsing. Theorem 1 has the following immediate corollaries.

Corollary 2. *If M is nonevasive then M is contractible.*

Corollary 3. *If M is nonevasive then for any coefficient field $\mathbf{F}$, $\tilde{H}_*(M, \mathbf{F}) = 0$, where $\tilde{H}_*(M, \mathbf{F}) = 0$, denotes the reduced homology of M.*

The goal of this paper is to make these results quantitative. That is, rather than asking simply "Is M evasive?" we will ask "How evasive is M?" To give an idea of the sort of results we will present, we note that Corollary 3 says that if the reduced homology of M does not vanish, then M is evasive. We will show that more is true. Namely, the more homology M has, the more evasive M is. We will now make this idea precise.

Suppose M is evasive. Then for any decision tree algorithm A there is a simplex σ such that $Q(\sigma, A, M) = n + 1$. We will call such a σ an *evader* of A. Evaders occur in pairs $\{\sigma_1, \sigma_2\}$ where

1) σ_1 is a face of σ_2 (which we indicate by writing $\sigma_1 < \sigma_2$)
2) dimension(σ_2) = dimension$(\sigma_1) + 1$
3) $\sigma_1 \subset M$, $\sigma_2 \not\subset M$.

Namely, $\sigma_2 = \sigma_1 * v_i$ for some vertex v_i, and if $\sigma = \sigma_1$ or $\sigma = \sigma_2$ then the $(n + 1)^{\text{st}}$ question is "Is v_i in σ?" That is, using the algorithm A, the seeker does not distinguish between the possibilities $\sigma = \sigma_1$ (and hence $\sigma \subset M$) and $\sigma = \sigma_2$ (and hence $\sigma \not\subset M$) until $n + 1$ questions have been asked. For any coefficient field $\mathbf{F}$, let $\tilde{H}_p(M, \mathbf{F})$ denote the p-dimensional reduced homology of M, and let $\tilde{b}_p = dim \tilde{H}_p(M, \mathbf{F})$. The following is a precise statement of our above comments.

Theorem 4. *For any decision tree algorithm A*

$$\#\{\textit{pairs of evaders of } A\} \geq \sum_{p=0}^{n} \tilde{b}_p.$$

Note that Theorem 4 does in fact imply Corollary 3. Theorem 4 will be a corollary of some significantly more refined inequalities.

We can present a simple example to illustrate this result. Suppose G is a graph on $k \geq 3$ nodes which is unknown to us. Our goal is to determine if G is connected, by checking the existence of one edge at a time. Let E be the set of possible edges, so that E has cardinality $\binom{k}{2}$. Let S be a simplex with $\binom{k}{2}$ vertices, with the vertices of S identified with E. Graphs on the fixed set of k nodes can be identified with the subsets of E, and hence with the faces of S. Let $N \subset S$ denote the union of the faces of S which correspond to graphs which are not connected. Then N is a subcomplex of S (that is why we choose to look at the not-connected graphs as opposed to the connected graphs). We can now restate the problem as follows. We must determine whether an unknown face G of S lies in N by inquiring about one vertex of S at a time. This is precisely the sort of decision problem we consider in this paper. Let A be a decision tree algorithm. A graph G is an evader of A if we must examine all $\binom{k}{2}$ possible edges (in the order determined by A) before determining whether or not G is connected. It is known (see, for example, [Va]) that N is homotopy equivalent to a wedge of $(k-1)!$ spheres of dimension $k-3$. We can now apply Theorem 4 to learn

Theorem 5. *Let A be a decision tree algorithm to determine whether or not a graph on k fixed nodes is connected. Then the number of evaders of A is at least $2(k-1)!$.*

J. Jonsson has recently shown that there exist decision tree algorithms which achieve this lower bound [Jo2].

Before going further, we mention that the conclusion of the next theorem depends, in a minor way, on whether or not the emptyset $\emptyset$ is an evader. For many purposes, one may restrict to the situation in which M contains all of the vertices of S. In this case, $\emptyset$ is not an evader of any decision tree algorithm. The following theorem is a refinement of Theorem 1.

Theorem 6. *Let A be a decision tree algorithm, and let k denote the number of pairs of evaders of A. Suppose that the emptyset is not an evader of A. Then there is*
(i) an ordering $\sigma_1^1, \sigma_1^2, \ldots \sigma_1^k$ of the evaders of A which lie in M, and an ordering $\sigma_2^1, \sigma_2^2, \ldots \sigma_2^k$ of the evaders of A which do not lie in M
and
(ii) a nested sequence of subcomplexes of S

$$S \supset S_k \supset S_{k-1} \supset \cdots \supset S_1 \supset S_0 = M \supset M_k \supset M_{k-1} \supset \cdots \supset M_1 \supset M_0 = v,$$

where v is a vertex of M which is not an evader of A,
with the following properties.

1) S collapses to S_k, and M collapses to M_k.
2) For each $i = 1, 2, \ldots, k$, σ_2^i is a maximal simplex of S_i and is not contained in S_{i-1}. Similarly, σ_1^i is a maximal simplex in M_i and is not contained in M_{i-1}.
3) For each $i = 1, 2, \ldots, k$, $S_i - \sigma_2^i$ collapses onto S_{i-1} and $M_i - \sigma_1^i$ collapses onto M_{i-1}.
If $\emptyset$ is an evader of A the theorem remains true if we let $\sigma_1^1 = \emptyset$, and set $M_0 = M_1 = \emptyset$. This requires that $M_2 = \sigma_1^2$ is a vertex.

That is, if $\emptyset$ is not an evader of A, one can build S as follows. Beginning with the vertex v of M, which is not an evader of A, one can uncollapse to the subcomplex $M_1 - \sigma_1^1$ (which contains the entire boundary of σ_1^1). Now we add σ_1^1 to obtain M_1. We can then uncollapse and add σ_1^2 to obtain M_2. Continuing in this manner, uncollapsing and then adding an evader of S which lies in M, we eventually obtain the subcomplex M_k, which uncollapses to M. Now we begin uncollapsing and adding an evader not in M, and we eventually obtain the subcomplex S_k, which uncollapses to the original simplex S. If $\emptyset$ is an evader of A, the story is the same, except that we begin the process with a vertex of M that is an evader of A.

We pause here to present a simple example. Let S be a 2-dimensional simplex, with vertices v_0, v_1, v_2, and M the subcomplex consisting of vertices v_0 and v_1 along with the edge (v_0v_1) between them. Then M is collapsible, so that, in particular, the reduced homology is 0. Suppose that A_1 is the decision tree algorithm which consists of asking three questions in the fixed order "Is v_0 in σ?", "Is v_2 in σ?" and "Is v_1 in σ?", independent of the answers received. This algorithm A_1 has no evaders. A chain of subcomplexes satisfying the conditions of Theorem 6 is simply

$$S \supset M \supset v_1.$$

Another chain is possible, but this is the chain produced in the proof of Theorem 6.

We now consider the decision tree algorithm A_2 shown in Figure 1 (where we indicate which vertex to ask about, given the answers to the previous questions). In this case there are 4 evaders consisting of the 2 pairs $\{\emptyset, v_2\}$ and $\{v_1, (v_1v_2)\}$. In particular, $\emptyset$ is an evader of A_2. A chain of subcomplexes satisfying the conditions of Theorem 6 is given by

$$S \supset S_2 = M \cup (v_0v_2) \supset S_1 = M \cup v_2 \supset S_0 = M \supset M_2 = v_1 \supset M_1 = M_0 = \emptyset.$$

In this case, $\sigma_1^1 = \emptyset, \sigma_1^2 = v_1, \sigma_2^1 = v_2$, and $\sigma_2^2 = (v_0v_2)$.

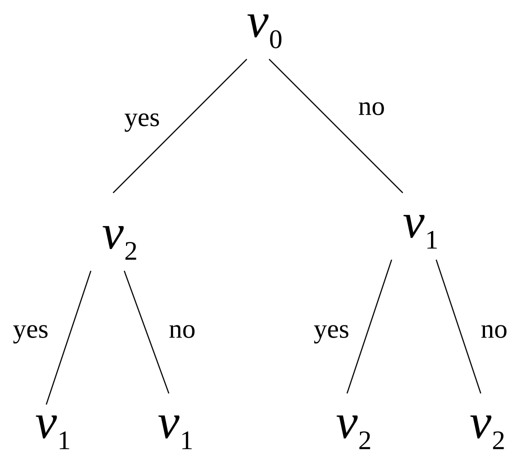

Figure 1

In this paper we will assume that all decision trees give instructions for the order in which one is to ask all $n+1$ questions, even after one has sufficient information to determine whether of not σ is in M. This is merely for convenience. There is no topological information contained in the questions one asks after the problem has been solved. More explicitly, if one has been given a decision tree which stops when one has sufficient information to determine if $\sigma \in M$, then one can ask the remaining $n+1$ questions in any order. The order chosen does not affect which faces are evaders, and so does not affect the conclusions of any of the results in this paper. However, the order chosen may affect the precise chain of complexes which arises in the proof of Theorem 6.

Before presenting some corollaries to Theorem 6, we must introduce some notation. If $\{\sigma_1 < \sigma_2\}$ is a pair of evaders of a decision tree A, define the index of the pair to be the dimension of σ_1 ,where we consider the dimension of the empty set $\emptyset$ to be -1. Let $e_p(A)$ denote the number of pairs of evaders of A which have index p. Note that $e_{-1}(A) \in \{0, 1\}$, and $e_{-1}(A) = 1$ if and only if $\emptyset$ is an evader of A. The following corollary of Theorem 6 is a refinement of Corollary 2.

Corollary 7. *Let A be a decision tree algorithm. Then M is homotopy equivalent to a CW complex with exactly $e_p(A)$ cells of dimension p, for $p \geq 1$, and $e_0(A) - e_{-1}(A) + 1$ cells of dimension* 0.

This implies the following inequalities which can be viewed as quantitative versions of Corollary 3.

Corollary 8. *Let A be a decision tree algorithm. For each $p = 0,\ 1,\ 2, \ldots, n+1$*

$$e_p(A) - e_{p-1}(A) + \cdots \pm e_0(A) \mp e_{-1}(A) \geq \tilde{b}_p - \tilde{b}_{p-1} + \cdots \pm \tilde{b}_0.$$

Subtracting consecutive inequalities from Corollary 8 yields the following weaker, but perhaps more intuitive relations.

Corollary 9. *Let A be a decision tree algorithm.*
(i) For each $p = 0,\ 1,\ 2, \ldots, n$

$$e_p(A) \geq \tilde{b}_p.$$

(ii) $-e_{-1}(A) + e_0(A) - e_1(A) + \cdots + (-1)^n e_n(A) = \chi(M) := \tilde{b}_0 - \tilde{b}_1 + \cdots + (-1)^n \tilde{b}_n.$

Theorem 4 easily follows from Corollary 9 (i).

To prove Theorem 6, will apply the discrete Morse theory developed in [Fo2] (see also [Fo1,3]). This theory is presented in a more combinatorial language in [Ch]. See [BBLSW], [Sha] and [Jo1] for other combinatorial applications. The main idea of the proof of Theorem 6 is that every decision tree algorithm A gives rise to a natural Morse function on S (or, more precisely, the gradient vector field of a Morse function - all of this is reviewed in §1). The critical points of this function are precisely the evaders of A (if $\emptyset$ is an evader of A, otherwise there is an additional critical vertex).

It is interesting to observe that specializing Theorem 6 to the case of a nonevasive complex yields the following strengthening of Theorem 1.

Corollary 10. *Let M be a nonempty subcomplex of S. If M is nonevasive, then S collapses to M, and M is collapsible.*

In fact, Corollary 10 is essentially proved in [KSS] in the discussion following the proof of Theorem 1 (which appears as Proposition 1 in [KSS]). In the language of this discussion in [KSS] a proof of the remaining part of Corollary 10 can be given as follows. If M is nonevasive, then so is M^*, the dual of M. Therefore, M^* collapses to a vertex v, which implies that v^* collapses to M. Since S collapses to v^*, the result follows.

Before proceeding, we note that most of the previous work on this subject begins with a hypothesis that M has a large symmetry group. In this paper, we do not make any such assumptions. See the discussion in §3 for further thoughts on this topic.

This paper is organized as follows. In §1 we review the basics of discrete Morse theory. In §2 we show that any decision tree algorithm induces a gradient vector field on M, and deduce Theorem 6. In §3 we discuss some areas for future research.

§1. A Review of Discrete Morse Theory.

In this section we review the essentials of discrete Morse theory. A more complete treatment appears in [Fo2]. An informal introduction appears in [Fo1].

Let N be a finite simplicial complex. That is, we are given a finite set V of vertices, and N is a set of subsets of V with the property that if $\beta \in N$ and α is a subset of β, then $\alpha \in N$. The elements of N are called simplices, and if $\beta \in N$ and α is a subset of β, we write $\alpha < \beta$ (or $\beta > \alpha$) and say that α is a face of β. If $\alpha \in N$ contains $p+1$ vertices, then we say that α has dimension p, and we indicate this with a superscipt $\alpha^{(p)}$. In what follows, we will also use the symbol N to indicate the topological realization of this simplicial complex, and we hope that this will not cause any confusion.

The emptyset $\emptyset$ is necessarily an element of N. Let N' denote the set of nonempty simplices of N. A function

$$f : N' \longrightarrow \mathbf{R}$$

is a *(discrete) Morse function* if, for each $\alpha^{(p)}$

(i) $$\#\{\beta^{(p+1)} > \alpha \mid f(\beta) \leq f(\alpha)\} \leq 1$$

and

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

Say $\alpha^{(p)}$ is a critical simplex if

(i) $$\#\{\beta^{(p+1)} > \alpha \mid f(\beta) \leq f(\alpha)\} = 0$$

and

(ii) $$\#\{\gamma^{(p-1)} < \alpha \mid f(\gamma) \geq f(\alpha)\} = 0.$$

For any $c \in \mathbf{R}$, let

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

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

We are now ready to state the main results. For proofs, see [Fo2]

Theorem 1.1. *Suppose $f^{-1}((a, b])$ contains no critical simplices. Then $N(b)$ collapses onto $N(a)$.*

Theorem 1.2. *Suppose $f^{-1}((a, b])$ contains a single simplex $\alpha^{(p)}$, and α is critical. Then $\alpha \not\subset N(a)$, $\dot{\alpha} \subset N(a)$ (where $\dot{\alpha}$ denotes the boundary of α) and $N(b) = N(a) \cup_{\dot{\alpha}} \alpha$.*

In the discussion that follows, it will be useful to assume that f is 1-1 on N'. It is easy to see that any Morse function can be perturbed slightly to become 1-1 without affecting which simplices are critical, so we will assume this from now on. It follows easily from the definitions that if f is a Morse function on N then the minimum of f must occur at a vertex, and this vertex is critical. This corresponds to the fact that every nonempty CW complex contains at least one vertex. The role of the other critical simplices is clarified in the following theorem which follows from Theorems 1.1 and 1.2.

Theorem 1.3. *Let f be a Morse function on N which is 1-1. Let $\sigma_1, \sigma_2, \ldots, \sigma_k$ be the critical simplices of f, where $f(\sigma_1) < f(\sigma_2) < \cdots < f(\sigma_k)$. Let $N_i = N(f(\sigma_i))$. Then N collapses to N_k and, for each $i = 2, 3 \ldots k$, σ_i is a maximal simplex of N_i, and $N_i - \sigma_i$ collapses to N_{i-1}. Lastly, $N_1 = \sigma_1$ consists of a single vertex.*

As an immediate corollary, we learn

Corollary 1.4. *N is homotopy equivalent to a CW complex with exactly 1 cell of dimension p for each critical simplex of dimension p.*

This result is sometimes most conveniently applied via some numerical inequalities. Fix a coefficient field $\mathbf{F}$, and let $b_p = \dim H_p(N, \mathbf{F})$ denote the p^{th} Betti number of N. Let m_p denote the number of critical simplices of dimension p. Corollary 1.4 implies the following inequalities.

Corollary 1.5. (The Strong Morse Inequalities) *For each $p = 0, 1, 2, \ldots$*

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

Subtracting consecutive strong Morse inequalities yields the following results.

Corollary 1.6. (The Weak Morse Inequalities)

(i) *For $p = 0, 1, 2, \ldots$*

$$m_p \geq b_p.$$

(ii)

$$m_0 - m_1 + m_2 - \cdots = b_0 - b_1 + b_2 - \ldots .$$

The next lemma follows easily from the definitions.

Lemma 1.7. *For any Morse function f and any $\alpha^{(p)}$, either*

$$\#\{\beta^{(p+1)} > \alpha \mid f(\beta) \leq f(\alpha)\} = 0$$

or

$$\#\{\gamma^{(p-1)} < \alpha \mid f(\gamma) \geq f(\alpha)\} = 0.$$

This is Theorem 2.1 of [Fo2]. Any Morse function f defines a disjoint collection X of pairs of simplices of N. Namely, X is the set of all pairs $\{\alpha^{(p)}, \beta^{(p+1)}\}$ with $\alpha < \beta$ and $f(\beta) \leq f(\alpha)$. It follows from Lemma 1.7 that these pairs are, in fact, disjoint. In [Fo2], X is called the (discrete) gradient vector field of f. In [Ch] and [Sha] X is called a Morse matching. Note that $\alpha \in N$ is a critical simplex of f if and only if α is not contained in any pair in X. We will also call such a simplex α a critical simplex of X.

We define an *X-path of dimension p* to be any sequence of simplices

$$\alpha_0^{(p)}, \beta_0^{(p+1)}, \alpha_1^{(p)}, \beta_1^{(p+1)}, \ldots, \beta_{k-1}^{(p+1)}, \alpha_k^{(p)} \tag{1.1}$$

such that for each $i = 0, 1, \ldots, k$

$$\{\alpha_i, \beta_i\} \in X$$

and

$$\beta_i > \alpha_{i+1} \neq \alpha_i.$$

It follows from the definition of a Morse function that for any such X-path

$$f(\alpha_0) \geq f(\beta_0) > f(\alpha_1) \geq f(\beta_1) > \ldots .$$

In particular, there are no closed X-paths, where we say the X-path in (1.1) is closed if $\alpha_k = \alpha_0$.

Expanding our view a bit, we define a discrete vector field on N to be any collection of disjoint pairs $\{\alpha^{(p)}, \beta^{(p+1)}\}$ where $\alpha < \beta$. General discrete vector fields were studied from a topological point of view in [Fo4]. For any discrete vector field X, we define an X-path as in (1.1). It is important to have a characterization of those discrete vector fields which are the gradient vector field of some Morse function. This is provided by the following theorem.

Theorem 1.8. *A discrete vector field X is the gradient vector field of a Morse function if and only if there are no closed X-paths.*

From now on, if X is the gradient of a Morse function, we will refer to X simply as a gradient vector field.

We end this section with some simple observations concerning the restriction of a Morse function to a subcomplex. Suppose N is a simplicial complex, and $\tilde{N}$ is a subcomplex of N. Let X denote a discrete vector field on N. We define the restriction of X to $\tilde{N}$, which we denote by $\tilde{X}$ as follows. For two cells $\alpha^{(p)} < \beta^{(p+1)}$ of $\tilde{N}$, we set $\{\alpha^{(p)}, \beta^{(p+1)}\} \in \tilde{X}$ if and only if $\{\alpha^{(p)}, \beta^{(p+1)}\} \in X$. That is, $\tilde{X}$ consists precisely of those pairs in X such that both simplices lie in $\tilde{N}$. We first note that it follows from Theorem 1.8 that if X is a gradient vector field, then so is $\tilde{X}$. Now let f be a Morse function on N, and let $\tilde{f}$ denote the restriction of f to $\tilde{N}$. That is, for any simplex α of N, we set $\tilde{f}(\alpha) = f(\alpha)$. Theorem 1.9 follows immediately from the definitions.

Theorem 1.9. *Let N be a simplicial complex with a Morse function f, and corresponding gradient vector field X. Let $\tilde{N}$ be a subcomplex of N. Then $\tilde{f}$ is a Morse function on $\tilde{N}$, with gradient vector field $\tilde{X}$.*

§2. From Decision Tree Algorithms to Morse Functions.

In this section we prove that every decision tree algorithm A gives rise to a gradient vector field on S, and hence, by restriction, a gradient vector field on M. For any decision tree algorithm A, the induced gradient vector field has only one critical simplex in S, which is a vertex v^*, reflecting the fact that S is collapsible. The critical simplices of the gradient vector field restricted to M are the the evaders of A which lie in M (along with the vertex v^*, if it lies in M).

Let A be a decision tree algorithm. Then A induces a partition $\mathcal{P}$ of all of the faces of S into disjoint pairs $\{\alpha^{(p)}, \beta^{(p+1)}\}$ where α and β are not distinguished by A until all $n+1$ questions have been asked. More explicitly, suppose $\sigma = \alpha$. Asking questions according to A, the $(n+1)^{\text{st}}$ question will be "Is v_i in σ?" for some i. Suppose v_i is not in α. Let $\beta = \alpha * v_i$. Then α and β are not distinguished until the $(n+1)^{\text{st}}$ question, so $\{\alpha, \beta\}$ is a pair in $\mathcal{P}$. If v_i is in α, then let $\gamma^{(p-1)} = \alpha - v_i$ so that $\alpha = \gamma * v_i$. Then γ and α are not distinguished until the $(n+1)^{\text{st}}$ question, so $\{\gamma, \alpha\}$ is a pair in $\mathcal{P}$. Note that the evaders are precisely the pairs $\{\alpha^{(p)}, \beta^{(p+1)}\}$ with $\alpha \subset M$ and $\beta \not\subset M$.

For example, the decision tree illustrated in Figure 1 induces the following partition of the faces of the two-dimensional simplex.

$$\{\{\emptyset, (v_2)\}, \{(v_0), (v_0, v_1)\}, \{(v_1), (v_1, v_2)\}, \{(v_0, v_2), (v_0, v_1, v_2)\}\}.$$

The matching $\mathcal{P}$ is not quite a vector field on S as defined in §1, since a vector field on S is a partial matching of the nonempty faces of S, while $\mathcal{P}$ is a matching of all faces of S, including $\emptyset$. The empty set $\emptyset$ is paired in $\mathcal{P}$ with some vertex, i.e. there is a vertex v^* such that $p = \{\emptyset, v^*\}$ is a pair in $\mathcal{P}$. Let $X = \mathcal{P} - p$. Then X is a vector field on S.

We can now state the main theorem of this section.

Theorem 2.1. *The vector field X is a gradient vector field.*

Before proving this theorem, we indicate how Theorem 2.1 implies the main results of this paper.

Proof of Theorem 6 (assuming Theorem 2.1): The first step is to modify the vector field X. Let $Y \subset X$ denote the set of pairs $\{\alpha, \beta\} \in X$ such that either α and β both lie in M, or neither lies in M. It follows from Theorems 1.8 and 2.1 that Y is a gradient vector field. We will now describe the critical simplices of Y, that is the simplices of S which are not paired in Y with any other simplex. If $\{\alpha, \beta\} \in X$ satisfies $\alpha \subset M$ and $\beta \not\subset M$ then α and β are critical for Y. Moreover, the vertex v^* (the vertex which had been paired with $\emptyset$ in $\mathcal{P}$) is critical. These are the only critical simplices. That is, the critical simplices of Y are precisely v^* and the evaders of A. Note that v^* is an evader of A if and only if $\emptyset$ is an evader of A, and this occurs if and only if v^* is not a vertex of M. The proof depends, in a minor way, on whether or not v^* is an evader of A, and this is the reason for the two different cases which arise in the statement of Theorem 6.

Let $f : S' \to \mathbf{R}$ be a Morse function on S with gradient vector field Y. It follows from the definition of Y that if $\alpha^{(p)} \subset M$ and $\gamma^{(p+1)} \not\subset M$ satisfy $\gamma > \alpha$, then $\{\alpha, \gamma\} \notin Y$, so

$$f(\gamma) > f(\alpha). \tag{2.1}$$

Let

$$\begin{aligned} a &= \sup_{\alpha \subset M} f(\alpha) \\ b &= \inf_{\alpha \not\subset M} f(\alpha) \\ c &= 1 + a - b \\ d &= \inf_{\alpha \subset S} f(\alpha). \end{aligned}$$

Define a new function

$$g : S' \longrightarrow \mathbf{R}$$

by setting

$$g(\alpha) = \begin{cases} f(\alpha) & \text{if } \ \alpha \subset M \\ f(\alpha) + c & \text{if } \ \alpha \not\subset M. \end{cases}$$

If $v^* \in M$, so that $\emptyset$ is not an evader of A, modify f further by setting $g(v^*) = d - 1$ (so that v^* minimizes g). Then for every $\alpha \subset M$ and $\beta \not\subset M$

$$g(\beta) \geq c + 1 > c \geq g(\alpha).$$

Moreover, for every pair of simplices $\alpha^{(p)} < \beta^{(p+1)}$

$$g(\beta) > g(\alpha) \Longleftrightarrow f(\beta) > f(\alpha).$$

This implies that g is a Morse function on S with the same critical simplices as f. In particular, the critical simplices of g are precisely $\{v^*\} \cup \{$ the evaders of A $\}$. Theorem 6 now follows from Theorem 1.3 $\square$

It should be noted that if all one cares about is Corollary 7 (and its implications Corollaries 8 and 9) then one can give a direct proof avoiding many of the details of our proof of Theorem 6. Namely, Let X_M denote the restriction of the vector field X to the subcomplex M. The critical simplices of X_M are the evaders of A which lie in M, and the vertex v^* (if v^* is in M, i.e. if $\emptyset$ is not an evader of A). Corollary 7 now follows from Corollary 1.4.

We now return to the proof of Theorem 2.1.

Proof of Theorem 2.1: We will show that there are no closed X-paths. The theorem then follows from Theorem 1.8. More precisely, we will define an ordering $\succ$ on the simplices of S with the property that if

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

is any X-path then

$$\alpha_0 \succ \beta_0 \succ \alpha_1 \succ \beta_1, \ldots .$$

This obviously implies that there are no closed X-paths. (In fact, one can get by with less. It is sufficient to define a partial order with the property that in any X-path as above we have $\alpha_0 \succ \alpha_1 \succ \alpha_2 \ldots$ This is a slightly easier task.)

Let α be any p-simplex of S. We associate to α a set of $p+1$ nonnegative integers $n_0(\alpha)$, $n_1(\alpha), \ldots, n_p(\alpha)$ such that

$$1 \leq n_0(\alpha) < n_1(\alpha) < n_2(\alpha) < \cdots < n_p(\alpha) \leq n+1.$$

Namely, suppose $\sigma = \alpha$. Then the seeker asks, according to the decision tree algorithm A, $n+1$ questions of the form "Is v_i in σ?" We number these questions consecutively from 1 to $n+1$. Let $n_0(\alpha) < \cdots < n_p(\alpha)$ be the numbers of the questions which receive the answer "yes." We let $n(\alpha)$ denote the string $n_0(\alpha), \ldots, n_p(\alpha)$.

We now define the ordering. Let $\alpha^{(p)}$ and $\beta^{(q)}$ be 2 simplices of S with corresponding numerical strings $n(\alpha)$, $n(\beta)$.

Suppose first that $n(\beta)$ is an extension of $n(\alpha)$. That is, $q > p$ and $n_i(\alpha) = n_i(\beta)$ for $i \leq p$. Then we set $\alpha \succ \beta$. If neither $n(\alpha)$ nor $n(\beta)$ is an extension of the other, we order them lexicographically and set $\alpha \succ \beta$ if there is a $k \leq \min\{p, q\}$ with

$$n_i(\alpha) = n_i(\beta) \quad \text{for} \quad i < k$$

and

$$n_k(\alpha) < n_k(\beta).$$

It is easy to check that this defines a transitive ordering.

Suppose $\alpha_0^{(p)}$, $\beta_0^{(p+1)}$, $\alpha_1^{(p)}$ is a segment of an X-path, so that we have $\{\alpha^{(p)}, \beta^{(p+1)}\} \in X$, and $\alpha_1^{(p)} \neq \alpha_0^{(p)}$ is another face of β_0. Since $\{\alpha_0, \beta_0\} \in X$, we have by definition

$$n_i(\beta) = n_i(\alpha) \quad \text{for} \quad i \leq p$$

and

$$n_{p+1}(\beta) = n+1.$$

That is, $n(\beta)$ is an extension of $n(\alpha)$, so $\alpha \succ \beta$.

Number the vertices of β_0 u_0, $u_1, \ldots, u_p$, u_{p+1} so that if $\sigma = \alpha_0$ or β_0 then question number $n_i(\beta_0)$ is "Is u_i in σ?" Then the vertices of $\alpha_1^{(p)}$ are u_0, $u_1, \ldots, u_{k-1}$, $u_{k+1}, \ldots, u_{p+1}$ for some $k \leq p$. If $\sigma = \beta_0$, the first $n_k(\beta_0) - 1$ questions involve only $u_0, \ldots, u_{k-1}$ and vertices which are not in β_0 (and hence not in α_0 or α_1). Therefore, the first $n_k(\beta_0) - 1$ answers are the same whether $\sigma = \alpha_0$, β_0 or α_1. This implies that if $\sigma = \alpha_0$, β_0 or α_1 then question number $n_k(\beta_0)$ is "Is u_k in σ?" If $\sigma = \beta_0$ this question is answered "yes", while if $\sigma = \alpha_1$ this question is answered "no". Thus

$$n_i(\alpha_1) = n_i(\beta_0) \quad \text{for} \quad i < k$$

and

$$n_k(\alpha_1) > n_k(\beta_0),$$

which implies that $\beta_0 \succ \alpha_1$. $\square$

Before moving on, it is perhaps worthwhile to remark upon the linear ordering constructed in this proof. The purpose of finding such an ordering was to show that the vector field X is the gradient of a Morse function. Of course, any Morse function with gradient vector field X provides such an ordering. The ordering constructed in the proof of Theorem 2.1 is not derived from a Morse function. (For many $\alpha^{(p)}$ there will be more than one $\beta^{(p+1)} > \alpha$ with $\alpha \succ \beta$.) However, we can modify the ordering slightly as follows. If neither $n(\alpha)$ nor $n(\beta)$ is an extension of the other, then define $\alpha \succ \beta$ as above. If $n(\beta^{(p+1)})$ is an extension of $n(\alpha^{(p)})$ then set $\alpha \succ \beta$ if $n(\beta) = n_0(\alpha), n_1(\alpha), \ldots, n_p(\alpha), n+1$, and $\beta \succ \alpha$ otherwise. Then any map $f : S' \to \mathbf{R}$ which preserves this order will be a Morse function with gradient vector field X.

§3. Concluding Remarks.

We will conclude with mention of two directions (of many possible) for much further investigation.

1) In this paper we have not at all examined the role of symmetry. Subcomplexes arising in applications often have large symmetry groups. Central to understanding the role of symmetry is the following conjecture of Rivest and Vuillemin [RV1–2]

Conjecture: Let M be a proper subcomplex of a simplex S, and suppose the isomorphism group of M acts transitively on the vertices of M. Then M is evasive.

Rivest and Vuillemin proved this conjecture in the case that the number of vertices of S is a prime power (see also [KK]). The Rivest-Vuillemin conjecture is a generalization of a conjecture, attributed to Karp, which states that any nontrivial monotone graph property is evasive. In [KSS] the Karp conjecture was verified in the special case of a graph property of graphs on a prime power number of vertices. These proofs use only the conclusion of Corollary 3. The proof of the complete conjecture will probably require the use of more of the information contained in Corollary 10. However, neither decision tree algorithms nor collapsing need preserve symmetry, so it is not clear how one can best make use of the hypotheses. This is certainly fertile ground for future important investigations.

2) For a subcomplex M of S, and a decision tree algorithm A, let $e(A)$ denote the number of pairs of evaders of A. Let

$$e^*(M) = \min_A e(A)$$

so that M is nonevasive if and only if $e^*(M) = 0$. The determination of $e^*(M)$ seems to be a very difficult problem. Note that *a priori* $e^*(M)$ may depend on the dimension of S, but a little thought shows that that is not the case.

Let $m^*(M)$ denote the minimal number of critical simplices for a Morse function on M. The ideas explored in this paper show that

$$e^*(M) \geq m^*(M) - 1.$$

The determination of $m^*(M)$ has a long history, at least in the case where M is a smooth compact manifold and we consider smooth Morse functions, e.g. [Mi],[Sh], but many of the ideas apply in the combinatorial setting ([Fo2]). We emphasize that this inequality is not, in general, equality. For example, if M is collapsible, so that $m^*(M) = 1$. Then $e^*(M) = m^*(M) - 1 = 0$ if and only if M is nonevasive. As discussed earlier, $e^*(M) + 1$ is the number of critical points of a Morse function f on M, but f must be the restriction to M of a Morse function on S with only a single critical point. The main question is whether it is possible to express $e^*(M)$, or at least to find other bounds for $e^*(M)$, in terms of more classical geometric and topological invariants.

References

[BBLSW] E. Babson, A. Björner, S. Linusson, J. Shareshian and V. Welker, *Complexes of not i-connected graphs*, Topology **38**, 271–299.

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

[Bo] B. Bollobás, *Extremal graph theory*, in Handbook of Combinatorics, Vol.1,2, 1231–1292, Elsevier, Amsterdam, (1995).

[Ch] M. Chari, *On discrete Morse functions and combinatorial decompositions*, preprint.

[Fo1] R. Forman, *A discrete Morse theory for cell complexes*, in Geometry, Topology & Physics for Raoul Bott, S.T. Yau (ed.), International Press, 1995.

[Fo2] ———, *Morse theory for cell complexes*, Adv. in Math **134** (1998), 90–145.

[Fo3] ———, *Witten-Morse theory for cell complexes*, Topology **37** (1998), 945–979.

[Fo4] ———, *Combinatorial vector fields and dynamical systems*, Math. Zeit **228** (1998), 629-681.

[Gl] L.C. Glaser, *Geometric Combinatorial Topology, Vol. I*, Van Nostrand Reinhold Mathematical Studies # 27, Van Nostrand Reinhold, 1970.

[Jo1] J. Jonsson, *On the homology of some complexes of graphs*, preprint, 1998.

[Jo2] ———, *The decision tree method*, preprint, 1999.

[KSS] J. Kahn, M. Saks and D. Sturtevant, *A topological approach to evasiveness*, Combinatorica **4** (1984), 297–306.

[KK] D.J. Kleitman and D.J. Kwiatkowski, *Further results on the Aanderaa-Rosenberg conjecture*, J. of Comb. Theory Series B **28** (1980), 85–95.

[Mi] J. Milnor, *Lectures on the h-cobordism theorem. Notes by L. Siebenmann and J. Sondow*, Princeton University Press, Princeton, N.J. 1965.

[RV1] R.L. Rivest and J. Vuillemin, *On recognizing graph properties from adjacency matrices*, Theor. Comp. Sci, **3** (1976), 371–384.

[RV2] ———, *A generalization and proof of the Aanderaa-Rosenberg conjecture*, in Proceedings of the 7th SIGACT Conference, Albuquerque (1975), ACM, 1976.

[Sh] V. Sharko, *Functions on Manifolds*, Translations of Mathematical Monographs, 131. American Mathematical Society, Providence, RI, 1993.

[Sha] J. Shareshian, *Discrete Morse theory for complexes of 2-connected graphs*, preprint.

[Va] V. Vassiliev, *Complexes of connected graphs*, in The Gelfand Mathematical Seminars, 1990–1992, Birkhäuser Boston, Boston, MA, (1993), 223–235.