A TOPOLOGICAL APPROACH TO THE GAME OF 20 QUESTIONS

ROBIN FORMAN

0. Introduction

Let us begin by clarifying the title of this paper. The game of 20 questions is usually played as follows. There are two players, whom we call the hider and the seeker, and there is a finite set Σ known to both. There is some $\sigma \in \Sigma$ known only to the hider. The seeker is permitted to ask 20 “yes or no” questions about σ, and the goal is to determine σ. The seeker is permitted to use the answers to earlier questions when choosing which question to ask next. Many search problems have this general form. However, in applications one is often limited to a predetermined set of questions, and one is not required to determine the hidden object precisely, but rather only to determine the object up to some sort of equivalence.

For example, suppose one has a communications network, and a storm (or a terrorist attack, or...) knocks out some of the arcs in the network. For each arc of the network, one may test ”Is this arc still working?” (This is the predetermined list of allowable questions.) The goal may not be to determine the surviving network precisely, but instead to determine the answers to a small list of questions of the form ”Is the network still connected?” ”Is there still a direct connection between Point A and the Point B?”, etc.

Suppose we are assured that the set of allowable questions is sufficient to complete the task. The question we ask in this paper is whether the seeker can do any better. Is there any strategy for asking questions that guarantees that the seeker can determine σ up to the desired accuracy before asking all of the questions? The answer depends

This work was partially supported by the National Science Foundation.

on the compatibility of the questions, and it is this phenomenon which can well be studied by topology. The results in this paper generalize the main theorems of [5], which were refinements of a theorem in [6]. In those papers, there was an assumption that guaranteed that the problem was easily stated in terms of topology (namely "monotonicity"). In this paper, we make no such assumption, and we show that a similar topological approach can be applied. The main techniques are slight modifications of some ingredients of the discrete Morse Theory introduced by the author in [3]. In that paper, we presented a discrete Morse Theory that can be applied, in particular, to any simplicial complex. To study the problem at hand, it is necessary to generalize the theory so that it can be applied to more general spaces. In particular, one can begin with any collection of open faces of a given simplex. Perhaps the main new idea is a homology theory for such spaces, different from the usual singular homology, but which agrees with the usual notions in the case of a simplicial complex and somewhat more generally, and which is compatible with the discrete Morse theory. We present here all of the relevant information so that this paper is completely self-contained.

Let us now proceed more precisely. Consider a game played by two players, whom we call the hider and the seeker. Let Σ be a finite set, and $\mathcal{P} : \Sigma = \mathcal{P}_1 \sqcup \mathcal{P}_2 \sqcup \ldots \mathcal{P}_k$ a partition of Σ, known to both the hider and the seeker. Let σ be an element of Σ, known only to the hider. The seeker's task is to determine, as quickly as possible, by asking "yes or no" questions, which element of the partition contains σ. Note that the seeker need not determine σ. For example, Σ might be the list of all professional football players, and one player σ has been chosen. The seeker's job might be to ask "yes or no" questions in order to determine which team σ plays for. We further suppose that the questions the seeker may ask about σ are limited to a set $\mathcal{Q} = \{\mathcal{Q}_1, \mathcal{Q}_2, \ldots, \mathcal{Q}_n\}$. We require that the answer to each $\mathcal{Q}_i$ does not depend on the order in which the questions are asked, so that each question $\mathcal{Q}_i$ can be identified with the set $Q_i \subset \Sigma$ of elements in Σ which result in the answer "yes" to question $\mathcal{Q}_i$.

For each $I \subset \{1, 2, \ldots, n\}$, let $\Sigma_I \subset \Sigma$ denote the elements of Σ that result in the answer "yes" to the questions $\mathcal{Q}_i$ for $i \in I$, and "no" for $i \notin I$. That is,

$$\Sigma_I = (\cap_{\{i \in I\}} Q_i) \cap (\cap_{\{i \notin I\}} Q_i^c),$$

where $Q_i^c \subset \Sigma$ denotes the complement of Q_i. We are unable to distinguish between elements in the same Σ_I using the questions $\mathcal{Q}$. For this reason, we refer to the Σ_I's as *indistinguishable sets.* Note that some indistinguishable sets may be empty.

Definition 0.1. *Say that the questions $\mathcal{Q}$ separate the partition $\mathcal{P}$ if each $\mathcal{P}_i$ is the union of indistinguishable sets.*

The questions $\mathcal{Q}$ contain sufficient information to determine for every possible $\sigma \in \Sigma$ which $\mathcal{P}_i$ contains σ if and only if $\mathcal{Q}$ separates $\mathcal{P}$. For the rest of this paper we will make this assumption. We note that a necessary condition for this is $k \leq 2^n$.

Assuming that $\mathcal{Q}$ separates $\mathcal{P}$, the seeker can certainly complete the desired task by asking all n questions. The main question we address in this paper is: Is it possible to do better?

Say that a partition $\mathcal{P}$ is *evasive* with respect to $\mathcal{Q}$ if $\mathcal{Q}$ separates $\mathcal{P}$, and no search strategy can guarantee a faster solution to the task. That is, $\mathcal{P}$ is evasive if $\mathcal{Q}$ separates $\mathcal{P}$, and no matter which search algorithm one chooses, there is at least one $\sigma \in \Sigma$ which requires that all n questions be asked before one can determine which P_i contains σ.

Conversely, if there is an algorithm which solves the desired task by asking $< n$ questions no matter σ is, then we say that $\mathcal{P}$ is *nonevasive* with respect to $\mathcal{Q}$. Note that if $\mathcal{P}$ is nonevasive with respect to $\mathcal{Q}$, then $\mathcal{Q}$ separates $\mathcal{P}$.

In this paper we will take a topological approach to this problem. Suppose $\mathcal{Q}$ separates $\mathcal{P}$. Let S be an $(n-1)$-dimensional simplex, with vertices labeled $\{1, 2, \ldots, n\}$. In Section 2 we will associate to each $\mathcal{P}_i$ in $\mathcal{P}$ a set P_i of (open) faces of S. These P_i's are pairwise disjoint. The P_i's may not be subcomplexes (i.e. they may not be closed as topological subspaces of S), however, so we refer to them as *simplex spaces* (most of the reasonable terms for such objects have already been assigned other uses in the mathematical literature). In this paper we relate the evasiveness of $\mathcal{P}$ to the topological properties of these simplex spaces.

Definition 0.2. *Say that $\mathcal{Q}$ is full if there are no empty indistinguishable sets.*

Things simplify greatly if $\mathcal{Q}$ is full, and we will make this assumption for much of this introduction. At the end of this introduction, we will state the main results in their full generality.

Theorem 0.3. *Assume that $\mathcal{Q}$ is full and that $\mathcal{P}$ is nonevasive with respect to $\mathcal{Q}$. Then each simplex space P_i can be simplicially collapsed.*

This is a generalization of Proposition 1 of [6]. In Section 6 we will describe precisely what we mean by the simplicial collapse of a simplex space.

Fix a field $\mathbb{F}$. In Section 6, using algebraic ideas presented in Section 3, we will define a homology theory over $\mathbb{F}$ for simplex spaces. Corresponding to this homology theory we will associate a Betti number $B(P_i)$. This Betti number generalizes the usual notion in the following sense: If P_i is a subcomplex of S, then $B(P_i)$ is equal to the dimension of $\tilde{H}_*(P_i, \mathbb{F})$, the usual reduced simplicial homology of P_i. In addition, if $\dot{P}_i = \bar{P}_i - P_i$ is a nonempty subcomplex, where $\bar{P}_i$ denotes the closure of P_i, then $B(P_i)$ is equal to the dimension of $H_*(\bar{P}_i, \dot{P}_i, \mathbb{F})$.

These Betti numbers are obstructions to simplicial collapse.

Theorem 0.4. *If P_i is collapsible, then $B(P_i) = 0$.*

We can use this notion of Betti number to derive a quantitative version of Theorem 0.3. Let $\mathcal{A}$ be a search algorithm. If, for some σ, asking questions in the order determined by $\mathcal{A}$, we must ask all n questions before determining which $\mathcal{P}_i$ contains σ, then we say that σ is *an evader of* A. Every σ' which is in the same indistinguishable set as σ is also an evader of $\mathcal{A}$, since every question results in the same answer for σ and σ'. We refer to the indistinguishable set containing σ as *an evading indistinguishable set of* $\mathcal{A}$. For any deterministic search algorithm $\mathcal{A}$, let $\mathcal{E}(\mathcal{A})$ denote the number of evading indistinguishable sets of $\mathcal{A}$.

Theorem 0.5. *Assume that $\mathcal{Q}$ is full and that $\mathcal{Q}$ separates $\mathcal{P}$. Then, for any deterministic search algorithm $\mathcal{A}$,*

$$\mathcal{E}(\mathcal{A}) \geq \sum_i B(P_i).$$

This is a generalization of Theorem 4 of [5].

We will also present some refinements of this result, based upon a definition of the j'th Betti number of a simplex space which will be equal to $\dim \tilde{H}_j(P_i, \mathbb{F})$ in the case that P_i is a subcomplex. That is, we will define a j'th Betti number $B_j(P_i)$, as well as an even Betti number $B_e(P_i)$ and an odd Betti number $B_o(P_i)$. These satisfy the relations

$$B_e(P_i) \geq \sum_{j \text{ even}} B_j(P_i),$$
$$B_o(P_i) \geq \sum_{j \text{ odd}} B_j(P_i),$$
$$B(P_i) = B_e(P_i) + B_o(P_i).$$

The inequalities are equalities if P_i is a subcomplex, but may be strict inequalities for more general simplex spaces. Define the *index* of an evading indistinguishable set Σ_I to be $|I|-1$, that is, one fewer than the number of questions that are answered "yes" for any $\sigma \in \Sigma_I$. Let $\mathcal{E}_j(\mathcal{A}, \mathcal{P}_i)$ denote the number of evading indistinguishable sets of index j that are contained in $\mathcal{P}_i$, and $\mathcal{E}(\mathcal{A}, \mathcal{P}_i)$ denote the total number of evading indistinguishable sets that are contained in $\mathcal{P}_i$. Let $\mathcal{E}_j(\mathcal{A})$ denote the total number of evading indistinguishable sets of index j, so that $\mathcal{E}(\mathcal{A}) = \sum_j \mathcal{E}_j(\mathcal{A})$.

Theorem 0.6. *Assume that $\mathcal{Q}$ is full and that $\mathcal{Q}$ separates $\mathcal{P}$. Let $\mathcal{A}$ be any deterministic search algorithm. Then for each j and i,*

$$\mathcal{E}_j(\mathcal{A}, \mathcal{P}_i) \geq B_j(P_i),$$
$$\sum_{j \textit{ even}} \mathcal{E}_j(\mathcal{A}, \mathcal{P}_i) \geq B_e(P_i),$$
$$\sum_{j \textit{ odd}} \mathcal{E}_j(\mathcal{A}, \mathcal{P}_i) \geq B_o(P_i),$$

and

$$\mathcal{E}(\mathcal{A}, \mathcal{P}_i) \geq B(P_i). \tag{0.1}$$

Summing (0.1) over all i yields Theorem 0.5.

The results in this paper generalize the main results presented in [5], which are refinements of a theorem in [6]. These papers were concerned with the case in which the partition contained only two sets: $\mathcal{P} : \Sigma = \mathcal{P}_1 \sqcup \mathcal{P}_2$, and the simplex space P_1 corresponding to $\mathcal{P}_1$ is a subcomplex of S. In this case the statement that P_2 simplicially collapses is equivalent to the statement that S collapses onto P_1. Thus Theorem 0.3 implies the following statement.

Theorem 0.7. *Suppose that $\mathcal{Q}$ is full and that the partition $\mathcal{Q}$ separates $\mathcal{P}$. Suppose further that $\mathcal{P}$ contains two sets $\Sigma = \mathcal{P}_1 \sqcup \mathcal{P}_2$, and that the simplex space P_1 corresponding to $\mathcal{P}_1$ is a subcomplex of S. Then if $\mathcal{P}$ is nonevasive with respect to $\mathcal{Q}$, S simplicially collapses to P_1, and P_1 simplicially collapses.*

This is essentially Proposition 1 of [6] (See the comments after Corollary 10 of [5]).

In this special case, if $\mathcal{P}_1$ is a non-empty proper subset of Σ, we have $\bar{P}_2 = S$ and $\dot{P}_2 = P_1$, so that Alexander duality implies

$$\dim H_*(\bar{P}_2, \dot{P}_2, \mathbb{F}) = \dim \tilde{H}_*(S, P_1, \mathbb{F}) = \dim \tilde{H}_*(P_1, \mathbb{F})$$

so that in this special case Theorem 0.5 implies the following result.

Theorem 0.8. *With all hypotheses as in Theorem 0.7, for any decision tree algorithm A, the number of evading indistinguishable sets of A is at least $2 \dim \tilde{H}_*(P_1, \mathbb{F})$.*

This is Theorem 4 of [5]

In the remainder of this section, we examine the general situation, in which we do not assume that $\mathcal{Q}$ is full. In this case, the partition $\mathcal{P} : \Sigma = \mathcal{P}_1 \sqcup \mathcal{P}_2 \sqcup \ldots \sqcup \mathcal{P}_k$ induces pairwise disjoint sets $P_1, P_2, \ldots, P_k$ of open faces of S. However, the P_i's do not contain all of the faces of S. Let $P^\emptyset$ denote the set of faces that are not contained in any P_i, that is, the faces that correspond to the empty indistinguishable sets. In particular, the simplex spaces $P_1, P_2, \ldots, P_k$ along with the simplex space $P^\emptyset$ form a partition of the set of faces of S. We introduce the following notation. For any pairwise disjoint simplex spaces P and P', we say that $P \to P'$ *is collapsible* if there is a collapsible simplex space P^* with $P \subseteq P^* \subseteq P \cup P'$. We also define generalized Betti numbers

$$B_j(P \to P') := \min_{P \subseteq P^* \subseteq P \cup P'} B_j(P^*).$$

That is, one is permitted to add to P any subset of P', and $B_j(P \to P')$ measures how small the j^{th} Betti number can be. We similarly define

$$B_e(P \to P') := \min_{P \subseteq P^* \subseteq P \cup P'} B_e(P^*),$$

$$B_o(P \to P') := \min_{P \subseteq P^* \subseteq P \cup P'} B_o(P^*),$$

and

$$B(P \to P') := \min_{P \subseteq P^* \subseteq P \cup P'} B(P^*).$$

For these Betti numbers we have the usual inequalities

$$B_e(P \to P') \geq \sum_{j \text{ even}} B_j(P \to P'), \tag{0.2}$$

and

$$B_o(P \to P') \geq \sum_{j \text{ odd}} B_j(P \to P'), \tag{0.3}$$

along with a new inequality (which is equality if $P' = \emptyset$)

$$B(P \to P') \geq B_e(P \to P') + B_o(P \to P'). \tag{0.4}$$

More generally, suppose $P_1, P_2, \ldots, P_k, P'$ are pairwise disjoint simplex spaces, and we set $P = \{P_1, P_2, \ldots, P_k\}$. We say that $P \to P'$ is collapsible if there are pairwise disjoint collapsible simplex spaces $P_1^*, P_2^*, \ldots, P_k^*$ satisfying

$$P_i \subseteq P_i^* \subseteq P_i \cup P' \text{ for each } i. \tag{0.5}$$

We also define

$$B_j(P \to P') := \min \sum_i B_j(P_i^*)$$

where now the minimum is taken over all sets $\{P_1^*, P_2^*, \ldots, P_k^*\}$ of pairwise disjoint simplex spaces satisfying (0.5). We then set

$$B_e(P \to P') := \min \sum_i B_e(P_i^*),$$

$$B_o(P \to P') := \min \sum_i B_o(P_i^*),$$

and

$$B(P \to P') := \min \sum_i B(P_i^*).$$

The inequalities (0.2),(0.3) and (0.4) also hold for these Betti numbers. We can now state the main theorem.

Theorem 0.9. *Assume that $\mathcal{Q}$ separates $\mathcal{P}$. We do not require that $\mathcal{Q}$ be full.*
(i) If $\mathcal{P}$ is nonevasive with respect to $\mathcal{Q}$ then $P \to P^{\emptyset}$ is collapsible, where $P = \{P_1, P_2, \ldots, P_k\}$.
(ii) Let $\mathcal{A}$ be any deterministic search algorithm. For each i, if there are no evading indistinguishable sets in $\mathcal{P}_i$, then $P_i \to P^{\emptyset}$ is collapsible. Moreover, for each j and i,

$$\mathcal{E}_j(\mathcal{A}, \mathcal{P}_i) \geq B_j(P_i \to P^{\emptyset}),$$

$$\sum_{j \ \mathit{even}} \mathcal{E}_j(\mathcal{A}, \mathcal{P}_i) \geq B_e(P_i \to P^{\emptyset}),$$

$$\sum_{j \ \mathit{odd}} \mathcal{E}_j(\mathcal{A}, \mathcal{P}_i) \geq B_o(P_i \to P^{\emptyset}),$$

and

$$\mathcal{E}(\mathcal{A}, \mathcal{P}_i) \geq B(P_i \to P^{\emptyset}).$$

Let $P = \{P_1, P_2, \ldots, P_k\}$. For each j

$$\mathcal{E}_j(\mathcal{A}) \geq B_j(P \to P^c),$$

$$\sum_{j \ \mathit{even}} \mathcal{E}_j(\mathcal{A}) \geq B_e(P \to P^{\emptyset}),$$

$$\sum_{j \ \mathit{odd}} \mathcal{E}_j(\mathcal{A}) \geq B_o(P \to P^{\emptyset}).$$

and

$$\mathcal{E}(\mathcal{A}) \geq B(P \to P^{\emptyset}).$$

The reader should consult the paper [1] for a topological aproach to a similar sort of search problem, in which the Betti numbers of the search space again provide a lower bound for the complexity.

1. Some Examples

In this section we present some simple examples. We will be stating some conclusions without proofs, and the proofs will be supplied throughout the rest of this paper. In particular, we revisit these examples in Section 3 in order to restate them in a more topological language. We note that these examples are rather trivial, and can easily be analyzed without using the results of this paper. They are included to illustrate the meaning of the theorems.

Let Σ denote the set $\{0, 1, 2, 3, 4, 5, 6, 7\}$, and the $\mathcal{Q}$ the set consisting of the following 3 questions.

$\mathcal{Q}_1$: Is σ even?

$\mathcal{Q}_2$: Is σ congruent to 0 or 1 $(mod\ 4)$?

$\mathcal{Q}_3$: Is $\sigma \leq 3$?

Since the indistinguishable sets for $\mathcal{Q}$ are all the subsets of Σ consisting of a single element, $\mathcal{Q}$ separates all partitions of Σ.

Example 1: Let $\mathcal{P} = \mathcal{P}_1 \sqcup \mathcal{P}_2$ denote the partition given by

$$\mathcal{P}_1 = \{0, 1, 2, 3\},\ \mathcal{P}_2 = \{4, 5, 6, 7\}.$$

Then $\mathcal{P}$ is certainly nonevasive with respect to $\mathcal{Q}$, since it is sufficient to ask only question $\mathcal{Q}_3$.

Example 2: Let $\mathcal{P}' = \mathcal{P}'_1 \sqcup \mathcal{P}'_2$ denote the partition given by

$$\mathcal{P}'_1 = \{0, 1, 2, 4, 5, 6\},\ \mathcal{P}'_2 = \{3, 7\}.$$

Then $\mathcal{P}'$ is also nonevasive with respect to $\mathcal{Q}$. To see this, consider the following question algorithm: Ask question $\mathcal{Q}_1$. If the answer is "yes", then σ is in $\mathcal{P}'_1$. If the

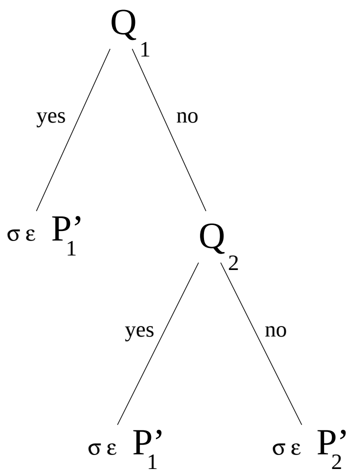

Figure 0.1. The search tree of example 2.

answer is "no", then ask question $\mathcal{Q}_2$. If the answer is "yes", then $\sigma \in \{1,5\} \subset \mathcal{P}'_1$. If the answer is "no", then $\sigma \in \{3,7\} \subset \mathcal{P}'_2$. Thus, at most two questions are needed.

We illustrate this guessing algorithm, using the standard "tree" format, in Figure 1.

Example 3: Now suppose that our goal is to solve both of our previous tasks. That is, we must ask questions until we determine whether or not $\sigma \in \mathcal{P}_1$, and whether or not $\sigma \in \mathcal{P}'_1$. Equivalently, we must determine which set of the partition

$$\mathcal{P}^* = (\mathcal{P}_1 \cap \mathcal{P}'_1) \sqcup (\mathcal{P}_1 \cap \mathcal{P}'_2) \sqcup (\mathcal{P}_2 \cap \mathcal{P}'_1) \sqcup (\mathcal{P}_2 \cap \mathcal{P}'_2)$$

contains σ. We will see in Section 3 that $\mathcal{P}^*$ is evasive with respect to $\mathcal{Q}$, and that for any guessing algorithm, there are at least 4 values of σ that require one to ask all 3 questions. That is, every guessing algorithm has at least 4 evaders. In fact, we can make the more precise statement that for every seach algorithm, each of the four blocks of the partition contains at least one evader.

Example 4: Now consider the partition $\mathcal{P} = \mathcal{P}_1 \sqcup \mathcal{P}_2$, where

$$\mathcal{P}_1 = \{0,1,3\}, \;\; \mathcal{P}_2 = \{2,4,5,6,7\}.$$

In section 3, we will see that this partition is evasive with respect to $\mathcal{Q}$. The simplex space P_1 corresponding to $\mathcal{P}_1$ is not a subcomplex, and neither is $\dot{P}_1$, so the results of [5] and [6] do not apply to this example.

Example 5: With $\mathcal{Q}$ as in Examples 1,2,3, and 4, let Σ' be the finite set $\{1, 2, 3, 4, 5, 6, 7\} = \Sigma - \{0\}$. In this example, $\mathcal{Q}$ is not full, since there is no element in Σ which results in the answer "yes" to each of the three questions. Consider the trivial partition $\mathcal{P} : \Sigma = \mathcal{P}_1$. This is certainly nonevasive, since we can complete the task of determining which block of the partition contains σ by asking 0 questions. The point of including this example will become clear when we examine this example from a topological point of view in Section 3.

Example 6: With $\mathcal{Q}$ and Σ' as in Example 5, we consider the following partition: $\mathcal{P} : \Sigma' = \{1\} \cup \{2\} \cup \{3, 4, 5, 6, 7\}$. As we will see in sections 3 and 10, this partition is evasive. Unlike in the previous examples, we cannot state in advance which block of the partition will contain an evader. There are search algorithms such that 1 is not an evader, and search algorithms such that 2 is not an evader, but, as we will see later, for every search algorithm, either 1 or 2 is an evader.

2. A Topological Reformulation

Let us begin by stating a search problem in a more topological setting. Let S denote the simplex of dimension $n - 1$ with vertices labeled $\{1, 2, , \dots, n\}$. By *a face of S* we mean a subset of the vertices of S. Let $P = \{P_1, P_2, \dots, P_k\}$ denote a collection of pairwise disjoint sets of faces of S. We suppose that S and P are known to both the hider and the seeker. Suppose that s is a face of S which is unknown to the seeker. The seeker is given the information that s belongs to one of the P_i's, and the seeker's task is to determine precisely which P_i contains s. It is permitted to ask questions of the form " Is vertex i in s?", Of course, the desired task can be completed by asking all n possible questions, since that is sufficient to determine s uniquely. Say that P is *evasive* if there is no question algorithm which guarantees that one can do better. That is, P is evasive if no matter what search algorithm one uses, there is at least one face s which requires that one ask all n questions. We say that P is *nonevasive* otherwise.

The search problem described in the introduction can be easily translated into an equivalent problem in this topological language. Let Σ be a finite set, $\mathcal{Q} = \{\mathcal{Q}_1, \mathcal{Q}_2, \ldots, \mathcal{Q}_n\}$ a set of permitted questions, and $\mathcal{P} : \Sigma = \mathcal{P}_1 \sqcup \mathcal{P}_2 \sqcup \ldots \sqcup \mathcal{P}_k$ a partition which is separated by $\mathcal{Q}$. Let S denote the simplex of dimension $n-1$ with vertices labeled $\{1, 2, , \ldots, n\}$. There is a natural identification of the indistinguishable sets of of Σ with the faces of S. Namely, identify Σ_I, $I \subset \{1, 2, \ldots, n\}$, with the subset I.

Since $\mathcal{Q}$ separates $\mathcal{P}$, we may express each $\mathcal{P}_i$ uniquely as a union of nonempty indistinguishable sets. Using the identification of indistinguishable sets with faces of S, we may associate each $\mathcal{P}_i$ with a set of faces P_i of S.

For any element σ of Σ, let s_σ denote the face of S which corresponds to the indistinguishable set which contains σ. We can now see that asking question $\mathcal{Q}_i$ about $\sigma \in \Sigma$ is equivalent to asking the question "Is vertex i in s" about the face s_σ of S. Thus there is a 1-1 correspondence between search algorithms $\mathcal{A}$ for the original problem and search algorithms A for the topological reformulation. Determining that σ is an element of $\mathcal{P}_i$ is equivalent to determining that the corresponding face s_σ belongs to P_i. Thus, if $\mathcal{A}$ is a guessing algorithm for the original problem, and A is the corresponding guessing algorithm for the topological reformulation, then σ is an evading indistinguishable set of $\mathcal{A}$ if and only if s_σ is an evader of A. Thus, the number of evaders of A is equal to the number of evading indistinguishable sets of $\mathcal{A}$. In particular, $\mathcal{P}$ is evasive with respect to $\mathcal{Q}$ if and only if P is evasive.

3. Another Look at the Examples of Section 1

We now return to the examples considered in Section 1. We let $\Sigma = \{0, 1, 2, 3, 4, 5, 6, 7\}$ and let $\mathcal{Q}$ be the set of three questions listed before the discussion of Example 1 in Section 1. In this case, the relevant simplex S is two-dimensional, with its vertices labeled 1,2 and 3. Each face s of S can be identified with the element σ of Σ which results in an affirmative answer precisely for those questions corresponding to the vertices of s.

For example, the empty face corresponds to the $\sigma \in S$ which results in each question being answered "no", i.e. $\sigma = 7$. The edge $[1, 2]$ corresponds to the $\sigma \in \Sigma$ which results in the answer "yes" to questions 1 and 2, and the answer "no" to question 3, i.e. $\sigma = 4$. In Figure 3.1 we have written the corresponding value of σ next to each face of S.

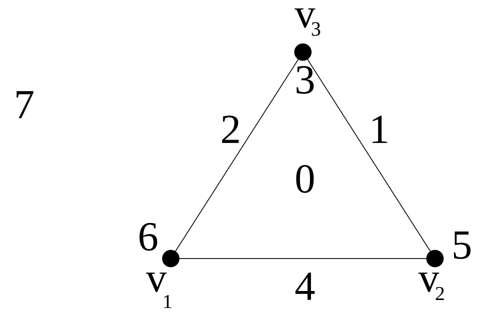

Figure 3.1. The correspondence between elements of Σ and faces of S.

Example 1: The set $\mathcal{P}_1 = \{0, 1, 2, 3\}$ corresponds to the simplex space P_1 consisting of all simplices which contain the vertex 3. That is, $P_1 = \{[3], [1, 3], [2, 3], [1, 2, 3]\}$. The simplex space P_1 is not a subcomplex, but $\dot{P}_1$ is, so, as we will see in Theorem 6.1, $B(P_1) = \dim H_*(\bar{P}_1, \dot{P}_1) = 0$.

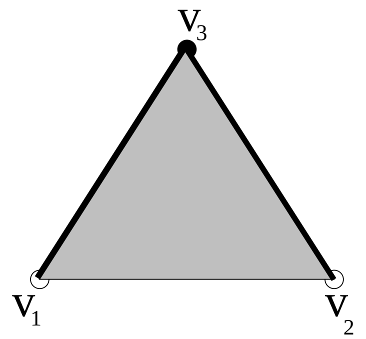

Figure 3.2. The simplex space P_1.

The set $\mathcal{P}_2$ corresponds to the simplex space $P_2 = \{\emptyset, [1], [2], [1,2]\}$. This set is a subcomplex, so that, by Theorem 6.1, we have $B(P_2) = \dim \tilde{H}_*(P_2) = 0$. Thus, Theorem 0.5 does not provide an obstruction to being nonevasive, and, as we have seen, $\mathcal{P}$ is nonevasive.

Example 2: The set $\mathcal{P}'_1 = \{0,1,2,4,5,6\}$ corresponds to the simplex space $P'_1 = \{[1],[2],[1,2],[1,3],[2,3],[1,2,3]\}$. The simplex space P'_1 is not a subcomplex, but $\dot{P}'_1$ is, so by Theorem 6.1 $B(P'_1) = \dim H_*(\bar{P}'_1, \dot{P}'_1) = 0$.

The set $\mathcal{P}'_2$ corresponds to the simplex space $P'_2 = \{\emptyset, [1], [3]\}$. This set is a subcomplex, so that by Theorem 6.1 we have $B(P'_2) = \dim \tilde{H}_*(P'_2) = 0$. Thus, just as in Example 1, Theorem 0.5 does not provide an obstruction to $\mathcal{P}'$ being nonevasive, and, in fact, $\mathcal{P}'$ is nonevasive.

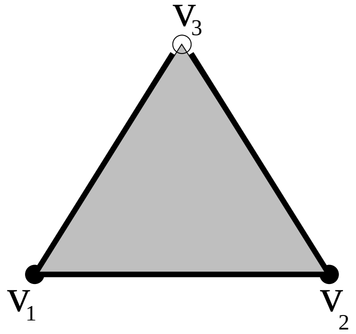

Figure 3.3. The simplex space P'_1.

Example 3: Here, the relevant partition is $\mathcal{P}^* = \mathcal{P}^*_{11} \sqcup \mathcal{P}^*_{12} \sqcup \mathcal{P}^*_{21} \sqcup \mathcal{P}^*_{22}$, where $\mathcal{P}^*_{ij} = \mathcal{P}_i \cap \mathcal{P}'_j$. The set $\mathcal{P}^*_{11} = \{0,1,2\}$ corresponds to the simplex space $P^*_{11} = \{[13],[23],[123]\}$. The set $\mathcal{P}^*_{12} = \{3\}$ corresponds to the simplex space $P^*_{11} = \{[3]\}$. The set $\mathcal{P}^*_{21} = \{4,5,6\}$ corresponds to the simplex space $P^*_{21} = \{[1],[2],[12]\}$. The set $\mathcal{P}^*_{22} = \{7\}$ corresponds to the simplex space $P^*_{22} = \{\emptyset\}$. The simplex space P^*_{22} is a subcomplex, so $B(P^*_{22}) = \dim \tilde{H}_*(P^*_{22}) = 1$. Since $\dot{P}^*_{11}$ is a subcomplex, we have that $B(P^*_{11}) = \dim H_*(\bar{P}^*_{11}, \dot{P}^*_{11}) = 1$. The simplex spaces P^*_{12} and P^*_{21} are not subcomplexes, but each would be a subcomplex if the empty simplex were to be added. As we will see in Theorem 6.1, in this case the Betti number is equal to the dimension of the unreduced homology, so that $B(P^*_{12}) = B(P^*_{21}) = 1$. From Theorem 0.5 is follows that any question algorithm has at least $B(P^*_{11})+B(P^*_{12})+B(P^*_{21})+B(P^*_{22}) = 4$ evaders. In fact, for any search algorithm $\mathcal{A}$, there must be at least one evader of $\mathcal{A}$ in each of the simplex spaces $\mathcal{P}^*_{11}, \mathcal{P}^*_{12}, \mathcal{P}^*_{21}, \mathcal{P}^*_{22}$.

Example 4: The simplex space P_1 corresponding to the set $\mathcal{P}_1$ of example 4 consists of the faces $[3], [2,3]$ and $[1,2,3]$.

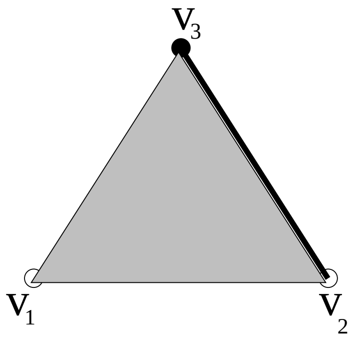

Figure 3.4. The simplex space P_1 of example 4.

This space is not a subcomplex, and $\dot{P}_1$ is not a subcomplex. We will see in Example 6.2 that for the homology theory defined in this paper, $B(P_1) = 1$. Thus, for any search algorithm $\mathcal{A}$, there will be at least one evading basic set in $\mathcal{P}_1$. The same holds for $\mathcal{P}_2$.

Example 5: Here we consider the case in which $\Sigma' = \{1,2,3,4,5,6,7\}$ and $\mathcal{P}$ is the trivial partition $\mathcal{P}' : \Sigma' = \mathcal{P}_1$. In this case, $\mathcal{Q}$ is not full, and $P^{\emptyset}$ consists of the single two-dimensional face $[1,2,3]$, and the simplex space P_1 consists of all of the remaining faces of Σ. The partition $\mathcal{P}$ is trivially nonevasive, and yet the simplex space P_1 has some homology. This illustrates that in the case that $\mathcal{Q}$ is not full, one cannot ignore $P^{\emptyset}$. In fact, $P_1 \to P^{\emptyset}$ has no homology and in fact is collapsible, since adding $P^{\emptyset} \to P_1$ yields the entire closed simplex S, as Theorem 0.9 predicts.

Example 6: Here we begin with the questions $\mathcal{Q}$, and the set $\Sigma' = \{1,2,3,4,5,6,7\}$ of Example 5. We consider the partition $\mathcal{P} : \Sigma' = \mathcal{P}_1 \cup \mathcal{P}_2 \cup \mathcal{P}_3$, where $\mathcal{P}_1 = \{1\}$, $\mathcal{P}_2 = \{2\}$ and $\mathcal{P}_3 = \{3,4,5,6,7\}$. The simplex space P_1 corresponding to $\mathcal{P}_1$ consists of just the edge $[1,3]$. The simplex space P_2 corresponding to $\mathcal{P}_2$ consists of just the edge $[2,3]$. The simplex space P_3 corresponding to $\mathcal{P}_3$ is $\{\emptyset,[1],[2],[3],[1,2]\}$, and $P^{\emptyset}=\{[1,2,3]\}$. The simplex space P_3 is a simplicial complex, with $B_0(P_3)=1$ and all other Betti numbers 0. Adding $P^{\emptyset}$ to P_3 only increases the Betti numbers, so $B_0(P_3\to P^{\emptyset})=1$, and $B_i(P_3\to P^{\emptyset})=0$ for $i\neq 0$. By Theorem 0.9, every search algorithm is guaranteed to have an evader in P_3. We see that $B_1(P_1)=B_1(P_2)=1$ while all other Betti numbers of P_1 and P_2 vanish. On the other hand, both $P_1\cup P^{\emptyset}$ and $P_2\cup P^{\emptyset}$ have vanishing homology. (These facts follow from the fact that for $P=P_1,P_2,P_1\cup P^{\emptyset}$ or $P_2\cup P^{\emptyset}$, $\dot{P}=\bar{P}-P$ is a nonempty simplicial complex, so the homology we are calculating is $H_*(\bar{P},\dot{P})$. See Theorem 6.1.) Therefore, $B(P_1\to P^{\emptyset})=B(P_2\to P^{\emptyset})=0$. On the other hand, we cannot add $P^{\emptyset}$ to both P_1 and P_2, so $B(\{P_1,P_2\}\to P^{\emptyset})=1$. Thus, Theorem 0.9 does not guarantee that there will be an evader in $\mathcal{P}_1$ nor does it guarantee that there will be an evader in $\mathcal{P}_2$, there must be at least one evader in $\mathcal{P}_1\cup\mathcal{P}_2\}$.

4. The Homology of Nondifferential complexes

Let $\mathbb{F}$ be any field. Let

$$\mathcal{V}: 0 \longrightarrow V_n \xrightarrow{\partial_n} V_{n-1} \xrightarrow{\partial_{n-1}} V_{n-2} \xrightarrow{\partial_{n-2}} \dots \xrightarrow{\partial_1} V_0 \xrightarrow{\partial_0} V_{-1} \longrightarrow 0$$

be a complex of finite dimensional vector spaces over the field $\mathbb{F}$. The ∂_i's are assumed to be linear maps, but we are not assuming that $\partial_d\circ\partial_{d+1}=0$. Our goal is to define the "homology" of this complex. If $\partial_d\circ\partial_{d+1}=0$ for each d, then we say that $\mathcal{V}$ is a *differential complex* and that the ∂_i's form a *differential*. We recall that one defines the homology of a differential complex by the formula

$$H_p(\mathcal{V}) := \operatorname{Ker}\partial_d/\operatorname{Im}\partial_{d+1}. \tag{4.1}$$

In this section we will define a homology of the complex $\mathcal{V}$ that reduces to the standard definition when ∂ is a differential.

Definition 4.1. *A homological decomposition $\mathcal{D}$ of the complex $\mathcal{V}$ is a decomposition*

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

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

(By the notation $V_i = X_i \oplus Y_i \oplus Z_i$ we mean that X_1, Y_2 and Z_i are linear subspaces ao V_i, such that their pairwise intersections are $\{0\}$, and they sum to give all of V_i.) Homological decompositions mimic the usual Hodge decompositions of differential complexes. Homological compositions exists, since one can take $X_i = 0, Y_i = 0$, and $Z_i = V_i$, for each i.

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

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

the odd Betti number of $\mathcal{D}$

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

and the total Betti number of $\mathcal{D}$

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

We now define the Betti numbers of $\mathcal{V}$ by

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

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

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

and

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

We observe the following facts.

Theorem 4.2. *Let $\mathcal{V}$ be any finite complex of finite dimensional vector spaces.*

(1) $B_e(\mathcal{V}) \geq \sum_{i \text{ even}} B_i(\mathcal{V})$ and $B_o(\mathcal{V}) \geq \sum_{i \text{ odd}} B_i(\mathcal{V})$.

Let $\chi(V) = -\dim(V_{-1}) + \dim(V_0) - \dim(V_1) + \dim(V_2) - \cdots + (-1)^n \dim(V_n)$.

(2) For any homological decomposition $\mathcal{D}$ of $\mathcal{V}$,

$$B_e(\mathcal{V}, \mathcal{D}) - B_o(\mathcal{V}, \mathcal{D}) = \chi(V).$$

(3) For any homological decomposition $\mathcal{D}$ of $\mathcal{V}$, $B_e(\mathcal{V},\mathcal{D}) = B_e(\mathcal{V})$ if and only if $B_o(\mathcal{V},\mathcal{D}) = B_o(\mathcal{V})$.

(4) $B_e(\mathcal{V}) - B_o(\mathcal{V}) = \chi(V)$.

(5) $B(\mathcal{V}) \geq |\chi(V)|$.

Example 4.3. A simple example will serve to show that the inequalities in part (1) of the theorem can be strict when $\mathcal{V}$ is not a differential complex. Consider the complex $\mathcal{V}$ with $V_0 = V_1 = V_2 = \mathbb{F}$, and $V_i = 0$ for $i = -1$ and $i > 1$. Suppose that ∂_1 and ∂_0 are both the identity map.

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

Let $\mathcal{D}_1$ denote the homological decomposition

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

We have that $H_1(\mathcal{V},\mathcal{D}_1) \cong H_2(\mathcal{V},\mathcal{D}_1) \cong 0$, while $H_0(\mathcal{V},\mathcal{D}_1) \cong \mathbb{F}$, which implies that $B_1(\mathcal{V}) = B_2(\mathcal{V}) = 0$, and $B_0(V) \leq 1$.

Let $\mathcal{D}_2$ denote the homological decomposition

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

In this case we see that $H_0(\mathcal{V},\mathcal{D}_1) \cong H_1(\mathcal{V},\mathcal{D}_1) \cong 0$, while $H_2(\mathcal{V},\mathcal{D}_1) \cong \mathbb{F}$, which implies that $B_0(\mathcal{V}) = B_1(\mathcal{V}) = 0$, and $B_2(V) \leq 1$.

Thus we learn that $B_i(\mathcal{V}) = 0$ for every i. On the other hand $B_o(\mathcal{V},\mathcal{D}_1) = B_o(\mathcal{V},\mathcal{D}_2) = 0$, which implies that $B_e(\mathcal{V}) = 0$. We note that $B_e(\mathcal{V},\mathcal{D}_1) = B_e(\mathcal{V},\mathcal{D}_2) = 1$, and, in fact, it follows from part (4) of Theorem 4.2 that $B_e(\mathcal{V}) = 1$.

Summarizing, in this little example it can easily be seen that we can choose a homological decomposition $\mathcal{D}$ so that either $B_0(\mathcal{V},\mathcal{D}) = 0$ or so that $B_2(\mathcal{V},\mathcal{D}) = 0$, but there is no homological decomposition which satisfies both conditions. This phenomenon is precisely captured by the inequality

$$B_e(\mathcal{V}) > B_0(\mathcal{V}) + B_2(\mathcal{V}).$$

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

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

A remnant of this equation holds for general complexes.

Lemma 4.4. *For any complex $\mathcal{V}$,*

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

Proof: For any homological decomposition $\mathcal{D} : V_i = X_i \oplus Y_i \oplus Z_i$, we know that $\partial_i|_{X_i}$ is injective, so $\dim(\operatorname{Ker}(\partial_i)) \leq \dim V_i - \dim X_i = \dim Y_i + \dim Z_i$. Therefore, $\dim(Z_i) \geq \dim(\operatorname{Ker}(\partial_i)) - \dim(Y_i)$. Since $Y_i \subset \operatorname{Im}\partial_{i+1}$ it follows that $\dim Y_i \leq \dim(\operatorname{Im}\partial_{i+1})$ so that $\dim Z_i \geq \dim(\operatorname{Ker}\partial_i) - \dim(\operatorname{Im}\partial_{i+1})$. Clearly $\dim Z_i \geq 0$. Thus, for any homological decomposition , we have the inequality

$$\dim Z_i \geq max\{\dim(\operatorname{Ker}\partial_i) - \dim(\operatorname{Im}\partial_{i+1}), 0\},$$

which implies that

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

To prove the converse, we will construct a homological decomposition $\mathcal{D}$ with the property that $\dim Z_i = max\{\dim(\operatorname{Ker}\partial_i) - \dim(\operatorname{Im}\partial_{i+1}), 0\}$. Let Y_i be the image of ∂_{i+1}. Let X_i be a subspace of V_i which is maximal with respect to being transverse both to Y_i and $\operatorname{Ker}\partial_i$. Then

$$\dim X_i = \dim V_i - max\{\dim(\operatorname{Ker}\partial_i), \dim Y_i\}.$$

Let Z_i be a subspace of V_i which is complementary to $X_i \oplus Y_i$. Then

$$\dim Z_i = \dim V_i - (\dim X_i + \dim Y_i) = max\{\dim(\operatorname{Ker}\partial_i) - \dim(\operatorname{Im}\partial_{i+1}), 0\}.$$

Choose a subspace X_{i+1} of V_{i+1} with the property that ∂_{i+1} maps X_{i+1} isomorphically onto Y_i. Let Z_{i+1} be any complementary subspace to X_{i+1} in V_{i+1}, and let $Y_{i+1} = 0$. Let $Y_{i-1} = \operatorname{Im}\partial_i|_{X_i}$. Note that $\partial_i|_{X_i}$ is injective. Let Z_{i-1} be any complementary subspace to Y_{i-1} in V_{i-1} and let $X_{i-1} = 0$. For $j \neq i+1, i, i-1$ set $X_j = Y_j = 0, Z_j = V_j$. It is easy to check that these decompositions form a homological decomposition with the desired property. □

We note that these definitions do, in fact, generalize the standard definition those for a differential complex.

Proposition 4.5. *Suppose that $\mathcal{V}$ is a differential complex, and let $H_i(\mathcal{V})$ denote the homology of $\mathcal{V}$ as defined by the standard formula (4.1). Then for each i,*

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

Moreover,

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

and

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

Proof: It follows from Lemma 4.4 that $B_i(\mathcal{V}) = \dim H_i(\mathcal{V})$. It follows from part (i) of Proposition 4.2 that to prove the remaining statements, it is sufficient to construct a homological decomposition $\mathcal{D}$ with the property that for each i, $\dim Z_i = \dim H_i(\mathcal{V})$. For each i, let Y_i denote the image of ∂_{i+1}, and let Z_i be a complementary subspace to Y_i in $\operatorname{Ker} \partial_i$. Let X_i be any complementary subspace to $\operatorname{Ker}(\partial_i) = Y_i \oplus Z_i$ in V_i. It can be easily seen that the resulting decomposition is a homological decomposition with the desired property. □

5. Some Properties of the Homology of Nondifferential Complexes

The main goal of this section is to see how the homology of $\mathcal{V}$ changes under certain algebraic operations. Some topological implications of these results will be presented in the following section. In what follows, we will begin with the complex

$$\mathcal{V} : 0 \longrightarrow V_n \xrightarrow{\partial_n} V_{n-1} \xrightarrow{\partial_{n-1}} V_{n-2} \xrightarrow{\partial_{n-2}} \dots \xrightarrow{\partial_1} V_0 \xrightarrow{\partial_0} V_{-1} \longrightarrow 0.$$

The complex

$$\mathcal{V}' : 0 \longrightarrow V'_n \xrightarrow{\partial'_n} V'_{n-1} \xrightarrow{\partial'_{n-1}} V'_{n-2} \xrightarrow{\partial'_{n-2}} \dots \xrightarrow{\partial'_1} V'_0 \xrightarrow{\partial'_0} V_{-1} \longrightarrow 0 \tag{5.1}$$

will denote a "perturbation" of the original complex.

Proposition 5.1. *Let d be an element of $\{-1,0,1,2,\ldots,n\}$. In the comparison complex 5.1, we assume that $V_i' = V_i$ for all $i \neq d$ and that V_d is a codimension-one subspace of V_d'. We require that the maps be compatible, so that $\partial_i = \partial_i'$ for $i \neq d$ (which implies that $\mathrm{Im}(\partial_{d+1}') \subseteq V_d \subset V_d'$) and that ∂_d is the restriction of ∂_d' to V_d. Then the following are true.*

(i) $B_i(V) = B_i(V')$ for $i \neq d, d-1$.

(ii) $B_d(\mathcal{V}) = B_d(\mathcal{V}')$ or $B_d(\mathcal{V}') - 1$, and $B_{d-1}(\mathcal{V}) = B_{d-1}(\mathcal{V}')$ or $B_{d-1}(\mathcal{V}') + 1$.

(iii) If d is even, then either $B_e(\mathcal{V}) = B_e(\mathcal{V}') + 1$ and $B_o(\mathcal{V}) = B_o(\mathcal{V}')$ or $B_e(\mathcal{V}) = B_e(\mathcal{V}')$ and $B_o(\mathcal{V}) = B_o(\mathcal{V}') - 1$. If d is odd, then either $B_e(\mathcal{V}) = B_e(\mathcal{V}')$ and $B_o(\mathcal{V}) = B_o(\mathcal{V}') + 1$ or $B_e(\mathcal{V}) = B_e(\mathcal{V}') - 1$ and $B_o(\mathcal{V}) = B_o(\mathcal{V}')$,

(iv) $B(\mathcal{V}) = B(\mathcal{V}') \pm 1$.

Proof:

Choose a complementary one-dimensional space U_d to V_d in V_d', so that $V_d' = V_d \oplus U_d$. Let $\mathcal{D} : V_i = X_i \oplus Y_i \oplus Z_i$ denote any homological decomposition of $\mathcal{V}$. Then one can construct a homological decomposition $\mathcal{D}'$ of $\mathcal{V}'$ by setting and $V_d' = X_d \oplus Y_d \oplus (Z_d \oplus U_d)$, and $V_i' = X_i \oplus Y_i \oplus Z_i$ for each $i \neq d$. Thus $B_i(\mathcal{D}') \leq B_i(\mathcal{D}) for i \neq d$, $B_d(\mathcal{D}') \leq B_d(\mathcal{D}) + 1$, and $B(\mathcal{D}') \leq B(\mathcal{D}) + 1$. If d is even, $B_e(\mathcal{D}') \leq B_e(\mathcal{D}) + 1$ and $B_o(\mathcal{D}') \leq B_o(\mathcal{D}) + 1$. If d is odd, $B_o(\mathcal{D}') \leq B_o(\mathcal{D}) + 1$ and $B_e(\mathcal{D}') \leq B_e(\mathcal{D})$.

This implies that

$$B_i(\mathcal{V}') \leq B_i(\mathcal{V}) \text{ for } i \neq d, \tag{5.2}$$

$$B_d(\mathcal{V}') \leq B_d(\mathcal{V}) + 1, \tag{5.3}$$

and

$$B(\mathcal{V}') \leq B(\mathcal{V}) + 1. \tag{5.4}$$

If d is even,

$$B_e(\mathcal{V}') \leq B_e(\mathcal{V}) + 1, \;\; B_o(\mathcal{V}') \leq B_o(\mathcal{V}) + 1.$$

and if d is odd,

$$B_o(\mathcal{V}') \leq B_o(\mathcal{V}) + 1, \;\; B_e(\mathcal{V}') \leq B_e(\mathcal{V}).$$

Now let $\mathcal{D}' : V_i' = X_i' \oplus Y_i' \oplus Z_i'$ denote a homological decomposition of $\mathcal{V}'$. Define a homological decomposition $\mathcal{D}$ of $\mathcal{V}$ by setting $V_i = X_i' \oplus Y_i' \oplus Z_i'$ for $i \neq d, d-1$. Set $V_d = (X_d \cap V_d) \oplus Y_d' \oplus Z_d$, where $X_d = X_d' \cap V_d$, and Z_d is any complement of $X_d \oplus Y_d'$ in V_d. (Note that since Y_d' is in the image of $\partial_{d+1}' = \partial_{d+1}$, we must have $Y_d' \subset V_d$.) Now set $V_{d-1} = X_{d-1}' \oplus (\partial_d(X_d)) \oplus Z_{d-1}$, where Z_{d-1} is any complement of $X_{d-1}' \oplus (\partial_d(X_d))$ in V_{d-1}. These decompositions form a homological decomposition.

Either $\dim X_d = \dim X_d'$ or $\dim X_d = \dim X_d' - 1$. In either case, we have $B_i(\mathcal{D}) = B_i(\mathcal{D}')$ for $i \neq d, d-1$, which implies that $B_i(\mathcal{V}) \leq B_i(\mathcal{V}')$ for $i \neq d, d-1$. Comparing with (5.2) yields part (i) of the proposition. In the first case we have that $B_d(\mathcal{D}) = B_d(\mathcal{D}') - 1$ and $B_{d-1}(\mathcal{D}) = B_{d-1}(\mathcal{D}')$. In the second case we have $B_d(\mathcal{D}) = B_d(\mathcal{D}')$ and $B_{d-1}(\mathcal{D}) = B_{d-1}(\mathcal{D}') + 1$. Thus $B_d(\mathcal{V}) \leq B_d(\mathcal{V}')$, $B_{d-1}(\mathcal{V}) \leq B_{d-1}(\mathcal{V}') + 1$, and $B(\mathcal{V}) \leq B(\mathcal{V}') + 1$. A comparison with (5.3) and (5.4) yields the remainder of the proposition. □

Proposition 5.2. *Fix a $d \in \{0, 1, 2, \ldots, n\}$. In the comparison complex (5.1), we assume that $V_i' = V_i$ for $i \neq d, d-1$, V_d is a codimension-one subspace of V_d', and V_{d-1} is a codimension-one subspace of V_{d-1}'. We require that the maps be compatible, so that $\partial_i = \partial_i'$ for $i \neq d, d-1$ and that ∂_d and ∂_{d-1}' are the restrictions of ∂_d' and ∂_{d-1}', resp. Suppose there exist a linear complement U_d of V_d in V_d' (so that $V_d' \cong V_d \oplus U_d$) and a linear complement U_{d-1} of V_{d-1} in V_{d-1}' so that $\pi_U \circ \partial_d$ maps U_d isomorphically onto U_{d-1}. Then*

(i) $B_i(\mathcal{V}) = B_i(\mathcal{V}')$ for all $i \neq d-1, d-2$,

(ii) $B_{d-1}(\mathcal{V}) = B_{d-1}(\mathcal{V}')$ or $B_{d-1}(\mathcal{V}') + 1$ and $B_{d-2}(\mathcal{V}) = B_{d-2}(\mathcal{V}')$ or $B_{d-2}(\mathcal{V}') + 1$,

(iii) Either $B_e(\mathcal{V}) = B_e(\mathcal{V}')$ and $B_o(\mathcal{V}) = B_o(\mathcal{V}')$, so that $B(\mathcal{V}) = B(\mathcal{V}')$, or $B_e(\mathcal{V}) = B_e(\mathcal{V}') + 1$ and $B_o(\mathcal{V}) = B_o(\mathcal{V}') + 1$, so that $B(\mathcal{V}) = B(\mathcal{V}') + 2$.

Proof: We will first show that for any homological decomposition $\mathcal{D} : V_i = X_i \oplus Y_i \oplus Z_i$ of $\mathcal{V}$, there is a homological decomposition $\mathcal{D}' : V_i' = X_i' \oplus Y_i' \oplus Z_i'$ of $\mathcal{V}'$ such that

$$Z_i = Z_i' \quad \text{for every } i,$$

so that $B_i(\mathcal{V}', \mathcal{D}') = B_i(\mathcal{V}, \mathcal{D})$ for every i, $B_e(\mathcal{V}', \mathcal{D}') = B_o(\mathcal{V}, \mathcal{D})$, $B_o(\mathcal{V}', \mathcal{D}') = B_o(\mathcal{V}, \mathcal{D})$, and $B(\mathcal{V}', \mathcal{D}') = B(\mathcal{V}, \mathcal{D})$. This is sufficient to prove that

$$B_i(\mathcal{V}') \leq B_i(\mathcal{V}') \text{ for every } i$$

and

$$B_e(\mathcal{V}') \leq B_e(\mathcal{V}),\ B_o(\mathcal{V}') \leq B_o(\mathcal{V}),\ B(\mathcal{V}') \leq B(\mathcal{V}).$$

Let $\mathcal{D} : V_i = X_i \oplus Y_i \oplus Z_i$ be a homological decomposition of $\mathcal{V}$. We begin our construction of $\mathcal{D}'$ by setting $V_i' = X_i \oplus Y_i \oplus Z_i$ for $i \neq d, d-1$. Let $X_d' = X_d \oplus U_d$ and decompose V_d' as $V_d' = X_d' \oplus Y_d \oplus Z_d$. The hypotheses imply that $\partial_d'|_{X_d'}$ is injective. Let Y_{d-1}' denote $\partial_d'(X_d)$. Then Y_{d-1} is a codimension-one subspace of Y_{d-1}', and $Y_{d-1}' \cap V_{i-1} = Y_{d-1}$. Thus $V_{d-1}' = X_{d-1} \oplus Y_{d-1}' \oplus Z_i$. It is easy to check that these decompositions form a homological decomposition of $\mathcal{V}'$ with the desired property.

We will now show that for any homological decomposition $\mathcal{D}'$ of $\mathcal{V}'$, there is a homological decomposition $\mathcal{D}'$ of $\mathcal{V}$ such that $B_i(\mathcal{V}, \mathcal{D}) = B_i(\mathcal{V}', \mathcal{D}')$ for $i \neq d-1, d-2$, $B_{d-1}(\mathcal{V}, \mathcal{D}) \leq B(\mathcal{V}', \mathcal{D}') + 1$, and $B_{d-2}(\mathcal{V}, \mathcal{D}) \leq B(\mathcal{V}', \mathcal{D}') + 1$.

One can then apply Proposition 4.2 (iii) to deduce the desired result. Let $\mathcal{D}' : V_i' = X_i' \oplus Y_i' \oplus Z_i'$ be any homological decomposition of $\mathcal{V}'$. We begin our construction of $\mathcal{D}$ by setting $V_i = X_i' \oplus Y_i' \oplus Z_i'$ for all $i \neq d-1, d-2$. Let $X_d = X_d' \cap V_d$, and let Z_d be any subspace of V_d which is complementary to $X_d \oplus Y_d'$. Thus $V_d = X_d \oplus Y_d' \oplus Z_d$. Moreover, either $\dim X_d = \dim X_d'$, so that $\dim Z_d = \dim Z_d'$, or $\dim X_d = \dim X_d' - 1$, so that $\dim Z_d = \dim Z_d'$. Let $Y_{d-1} = \partial_d(X_d)$, $X_{d-1} = X_{d-1}' \cap V_d$ and let Z_{d-1} denote any subspace of V_{d-1} which is complementary to $X_{d-1} \oplus Y_{d-1}$. Then $\dim Z_{d-1} \leq \dim Z_{d-1}' + 1$. Let $Y_{d-2} = \partial_{d-1}(X_{d-1})$, $X_{d-2} = X_{d-2}'$ and let Z_{d-2} denote any subspace of V_{d-2} which is complementary to $X_{d-2} \oplus Y_{d-2}$. Then $\dim Z_{d-2} \leq \dim Z_{d-2}' + 1$. It is easy to check that these decompositions form a homological decomposition of $\mathcal{V}$ with the desired property. $\square$

In the case of differential complexes, we can strengthen the conclusion of Theorem 5.2. First note that $\mathcal{V}$ is a differential complex if and only if $\mathcal{V}'$ is. In that case it can be shown that all Betti numbers are the same for $\mathcal{V}$ and $\mathcal{V}'$. (As we will see in the next section, this corresponds to the fact that in the category of simplicial complexes, a simplicial collapse is a homotopy equivalence, and hence does not chanhe the homology of the complex.) This is definitely not the case for a general complex. Consider the complex

$$\mathcal{V}' : 0 \longrightarrow \mathbb{F} \xrightarrow{\partial_1'} \mathbb{F}^2 \xrightarrow{\partial_0'} \mathbb{F} \longrightarrow 0,$$

where $\partial_1'(x) = (x, x)$, and $\partial_0'(x, y) = y$ for every $x, y \in \mathbb{F}$. All Betti numbers are zero for this complex, as can be seen from the following homological decomposition:

$$\mathcal{V}' : 0 \longrightarrow \mathbb{F} \oplus 0 \oplus 0 \xrightarrow{\partial_1'} X_0 \oplus Y_0 \oplus 0 \xrightarrow{\partial_0'} 0 \oplus \mathbb{F} \oplus 0 \longrightarrow 0$$

where $X_0 = \{(x, x)\}$ and $Y_0 = \{(0, x)\}$. Now consider the complex

$$\mathcal{V} : 0 \longrightarrow 0 \longrightarrow V_0 \xrightarrow{\partial_0} \mathbb{F} \longrightarrow 0,$$

where $V_0 = \{(x, 0)\}$ and ∂_0 is the zero map. One can easily check that $B_{-1}(\mathcal{V}) = B_0(\mathcal{V}) = B_e(\mathcal{V}) = B_o(\mathcal{V}) = 1$. These complexes satisfy the hypotheses of Theorem 5.2. Thus, in the case of differential noncomplexes, an "algebraic collapse" form $\mathcal{V}'$ to $\mathcal{V}$, as defined in the hypotheses to Theorem 5.2 can result in an increase in the Betti numbers. For a more topological point of view on this phenomenon, the reader is refered to the example at the end of the next section.

6. The Homology of Simplex Spaces

Let S be a simplex. A simplex space K is a set of faces of S (we consider the empty set to be a face of S). There is (at least) one potentially confusing point in all of this. We let $\emptyset$ denote the face of S containing no vertices. Let the *empty simplex space* denote the simplex space consisting of no faces of S. In particular, the simplex space consisting of just $\emptyset$ is not the empty simplical set, it is a simplex space with one element.

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

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

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

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

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

$$0 \longrightarrow C_n(K,\mathbb{F}) \xrightarrow{\partial_n} C_{n-1}(K,\mathbb{F}) \xrightarrow{\partial_{n-1}} \ldots \xrightarrow{\partial_2} \ldots C_1(K,\mathbb{F}) \xrightarrow{\partial_1} C_0(K,\mathbb{F}) \xrightarrow{\partial_0} C_{-1}(K,\mathbb{F}) \longrightarrow 0.$$

Using the notation of the previous section, we denote by

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

the numbers

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

We observe that if K is a simplicial complex, and contains the empty simplex, then $C(K,\mathbb{F})$ is the usual reduced chain complex. If K is a simplicial complex except that K does not contain the empty simplex, then $C(K,\mathbb{F})$ is the usual (nonreduced) chain complex. If $\dot{K} = \bar{K} - K$ is a nonempty subcomplex (where $\bar{K}$ denotes the closure of K, i.e. K along with all of the faces of the elements in K, and $\ldots K$ denotes the elements in $\bar{K}$ which are not in K), then $C(K, \mathbb{F})$ is isomorphic to the relative chain complex $C(\bar{K}, \dot{K}, \mathbb{F})$. In particular, in these cases, $C(K, \mathbb{F})$ is a differential complex. We summarize these observations in the following theorem.

Theorem 6.1. *(i) If K is a simplicial complex (containing the empty simplex), then $B_i(K, \mathbb{F}), B_e(K, \mathbb{F}), B_o(K, \mathbb{F})$ and $B(K, \mathbb{F})$ are equal to the usual reduced i^{th} Betti number, even Betti number, odd betti number and total Betti number of K.*

(ii) If K is a simplicial complex except that K does not contain the empty simplex, then $B_i(K, \mathbb{F}), B_e(K, \mathbb{F}), B_o(K, \mathbb{F})$ and $B(K, \mathbb{F})$ are equal to the usual nonreduced i^{th} Betti number, even Betti number, odd Betti number and total Betti number of K.

(iii) If $\dot{K} = \bar{K} - K$ is a nonempty subcomplex, then $B_i(K, \mathbb{F}), B_e(K, \mathbb{F}), B_o(K, \mathbb{F})$ and $B(K, \mathbb{F})$ are equal to the usual i^{th} Betti number, even Betti number, odd Betti number and total Betti number of $\bar{K}$ relative to $\dot{K}$.

Example 6.2. Let S denote the two-dimensional simplex with vertices labelled 1,2,3. Let K denote the simplex space consiting of the faces $[3], [2, 3]$ and $[1, 2, 3]$. The chain complex for K can easily seen to be isomorphic to the complex examined in Example 4.3. From the calculations done there, we see that $B_i(K) = 0$ for each $i = -1, 0, 2, \ldots$, and $B_o(K) = 0$, while $B(K) = B_e(K) = 1$.

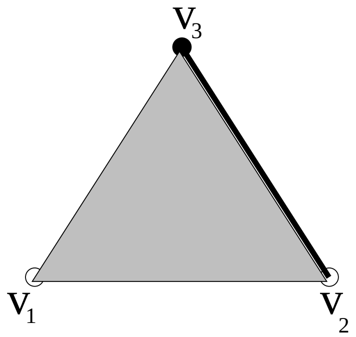

Figure 6.1. The simplex space K of Example 6.2.

7. Collapsing Simplex Spaces

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

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

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

Given K, our goal is to understand whether K is collapsible. If K is not collapsible, an additional goal is to find a lower bound for the number of simplical removals which must appear in any complete reduction of K. In the case that K is a subcomplex of S then standard algebraic topology provides us with an answer. In this case, thinking of K as a topological space, a simplical collapse results in a topological space which is homotopy equivalent to K. Thus if K is collapsible, it must be contractible. Moreover, a simplicial removal changes the sum of the reduced Betti numbers by at most one, while a simplicial collapse does not change the sum of the reduced Betti mnumbers. Thus, the sum of the reduced Betti numbers of K gives a lower bound on the number of removals in any total reduction of K.

Our immediate goal is to generalize these comments to all simplex spaces.

Note that a simplicial removal changes the complex $C(K, \mathbb{F})$ in precisely the manner investigated in Theorem 5.1. We can restate Theorem 5.1 in the following language.

Theorem 7.1. *Let K be a simplex space, and $K' = K - \alpha$, for some maximal d-simplex α. Then*

(i) $B_i(K') = B_i(K)$ for $i \neq d, d-1$.

(ii) $B_d(K') = B_d(K)$ or $B_d(K) - 1$, and $B_{d-1}(K') = B_{d-1}(K)$ or $B_{d-1}(K) + 1$.

(iii) If d is even, then either $B_e(K') = B_e(K)+1$ and $B_o(K') = B_o(K)$ or $B_e(K') = B_e(K)$ and $B_o(K') = B_o(K) - 1$. If d is odd, then either $B_e(K') = B_e(K)$ and $B_o(K') = B_o(K) + 1$ or $B_e(K') = B_e(K) - 1$ and $B_o(K') = B_o(K)$,

(iv) $B(K') = B(K) \pm 1$.

A simplicial collapse changes the complex $C(K, \mathbb{F})$ in precisely the manner investigated in Theorem 5.2. We can restate Theorem 5.2 in the following language.

Theorem 7.2. *Let K be a simplex space, and $K' = K - (\alpha \cup \beta)$ the result of a simplicial collapse, where α is a maximal d-simplex, and β is a free face of α. Then*

(i) $B_i(K') = B_i(K)$ for all $i \neq d-1, d-2$,

(ii) $B_{d-1}(K') = B_{d-1}(K)$ or $B_{d-1}(K)+1$ and $B_{d-2}(K') = B_{d-2}(K)$ or $B_{d-2}(K)+1$,

(iii) Either $B_e(K') = B_e(K)$ and $B_o(K') = B_o(K)$, so that $B(K') = B(K)$, or $B_e(K') = B_e(K) + 1$ and $B_o(K') = B_o(K) + 1$, so that $B(K') = B(K) + 2$.

The main point is that a simplicial collapse cannot result in a decrease in B_i, B_e, B_o, or B. Simplicially removing a d-simplex cannot result in a decrease in B_i for $i \neq d$, and can decrease B_d by at most 1. If d is even, then B_o will not decrease, and B_e can decrease by at most 1. The corresponding statement holds if d is odd. This implies that B can decrease by at most 1. Thus we can deduce the following corollary.

Corollary 7.3. *Let K be a simplex space. In any complete reduction of K, the number of d-simplices which are taken out by a simplicial removal is at least $B_d(K, \mathbb{F})$, the number of even and odd dimensional simplices which are taken out by simplicial removal is at least $B_e(K, \mathbb{F})$ and $B_o(K, \mathbb{F})$, resp., and the total number of simplices which are taken out by simplicial removal is at least $B(K, \mathbb{F})$.*

We end this section with an example to illustrate that, unlike in the case of a simplcial complex, a simplicial collapse can increase the Betti numbers of a simplex space. Let S denote the two-dimensional simplex with vertices labelled 1,2,3. Let K denote the simplex space consisting of the four faces $[1], [1,2], [2,3], [1,2,3]$. Then K is collapsible, since one can remove $[2,3]$ and $[1,2,3]$ by one simplicial collapse, and the remaining two faces with a second simplicial collapse. Thus, all Betti numbers of K are zero. One the other hand, beginning with K, one can also remove the faces $[1,2]$ and $[1,2,3]$ by a simplicial collapse, resulting in the simplex space K' consisting of the faces $[1]$ and $[2,3]$. One can easily check that $B_0(K') = B_1(K') = B_e(K') = B_o(K') = 1$. This example is just a topological realization of the algebraic example presented at the end of Section 5.

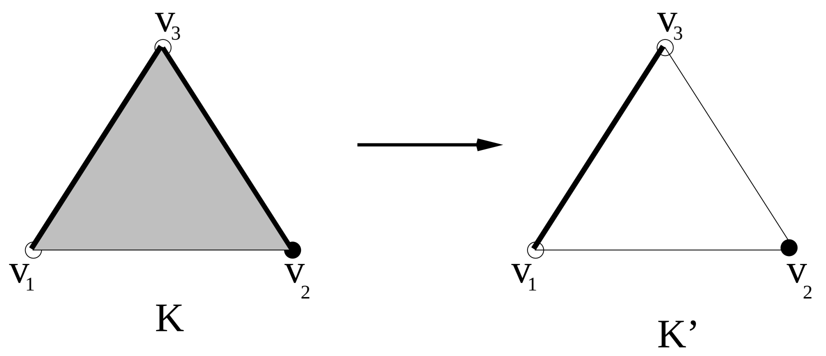

Figure 7.1. The simplex space K collapses to the simplex space K'.

8. Discrete Morse Theory

In this section we give a very brief look at some of the relevant aspects of discrete Morse Theory. The reader should consult [2], [3] or [4] for a more complete overview of the theory. Let K be a simplex space. If α is a simplex in K, we use the notation $\alpha^{(p)}$ to indicate that α is p-dimensional. If α and β are simplices in K, we use the notation $\alpha \prec \beta$, or $\beta \succ \alpha$ to indicate that α is a facet of β, i.e. a codimension-one face of β.

A *discrete vector field* on K is a collection V of pairs of simplices of the form $\{\alpha \prec \beta\}$, such that no simplex is in more than one pair. A complete reduction of K gives rise to a discrete vector field. Namely, a complete reduction of K is a sequence of simplicial removals, and simplicial collapses. Every simplicial collapse corresponds to removing from K a pair of simplices of the form $\{\alpha \prec \beta\}$. Given the simplical reduction of K, let V denote the set of all pairs $\{\alpha \prec \beta\}$ which are removed from K by a simplicial collapse. We call a discrete vector field V on K a *(discrete) gradient vector field* if V arises in this fashion from some complete reduction of K. Note that in this case, the simplices which are removed from K by a simplicial removal are precisely the simplices which do not appear in any pair in V. We refer to these faces as the *critical faces of* V. Corollary 7.3 can now be translated into the language of this section as follows.

Theorem 8.1. *Let K be a simplex space and let V be any gradient vector field on K. Then the number of critical d-dimensional faces of V is at least $B_d(K,\mathbb{F})$, the number of even and odd dimensional critical faces of V is at least $B_e(K,\mathbb{F})$ and $B_o(K,\mathbb{F})$, resp., and the total number of critical faces of V is at least $B(K,\mathbb{F})$.*

We warn the reader that the definition of gradient vector field whcih appears here is not the same definition that appears in the references sited above. Instead, we have chosen to present an definition that is much better suited for our current purposes. This definition is equivalent to that presented earlier except for the very minor point that in this paper the empty simplex is permitted to occur in some pair in V, while this possibility is not permitted in the earlier references.

Our next goal is to determine when a discrete vector field is a gradient vector field. Let V be a discrete vector field on K. We define a *V-path of dimension p* to be any sequence of simplices of K

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

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

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

and

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

We say that such a V-path is *closed* if $k > 0$ and $\alpha_k^{(p)} = \alpha_0$.

The main theorem of this section is the following.

Theorem 8.2. *A discrete vector field V on a simplex space K is a gradient vector field if and only if there are no closed V-paths.*

This theorem is essentially Theorem 9.3 of [3]. However, it seems to us that the context is sufficiently different to make it worthwhile presenting a self-contained proof.

Proof: First we will assume that V is a gradient vector field corresponding to some complete reduction of K, and prove that there are no closed V-paths. Suppose that $\alpha_0^{(p)}, \beta_0^{(p+1)}, \alpha_1^{(p)}, \beta_1^{(p+1)}$ is a V-path. That is, α_0 and β_0 are paired in V, as are α_1 and β_1,and $\beta_0 > \alpha_1 \neq \alpha_0$. This implies that in the complete reduction of K, α_1 and β_1 are taken out by a simplicial collapse. This requires that at some point in the complete reduction of K, there appears a simplex space K' with the property that β_1 is a maximal simplex of K', and α_1 is a free face of β_1 in K'. Since α_0 is also a face of β_0 in K, it must be the case that in the complete reduction of K, β_0 is taken out of K before β_1. It easily follows by induction that there can be no closed V-paths.

Now suppose that V is a discrete vector field with the property that there are no closed V-paths. We must show that V arises as the gradient vector field corresponding to some complete reduction of K. Actually, we must only show that there exists at least one elementary reduction that is consistent with V, that is, a maximal simplex α which is not paired in V aand hence can be taken out of K by simplicial removal, or a pair $\{\alpha \prec \beta\}$ in V such that β is maximal and α is a free face of β, which can be taken out of K by a simplicial collapse. The general result then follows by induction. Let n denote the dimension of K. Define a partial order on the n-simplices of K by declaring that $\beta_0^{(n)} \overset{*}{>} \beta_1^{(n)}$ if and only if there is some V-path in which β_0 appears before β_1. This is clearly a transitive order. Moreover, the nonexistence of closed V-paths implies that if $\beta_0 \overset{*}{>} \beta_1$ then it is not true that $\beta_1 \overset{*}{>} \beta_0$, as well as the fact that it cannot happen that $\beta \overset{*}{>} \beta$ for some β.

Let $\alpha\beta$ be an n-simplex of K which is maximal with respect to this partial order. If β is not paired in V, then β can be taken out by a simplicial removal (β must be maximal since $n = \dim K$). Suppose that β is paired in V with some $(n-1)$-simplex

α. Since β is maximal, if α is a free face of β then $\alpha \prec \beta$ can be taken out of K by a simplicial collapse. If α is not a free face of β, then there must be another n-simplex β' which has α as a face. If β' is not paired in V, then β' can be taken out by a simplicial removal. Otherwise, β' is paired in V with some $(n-1)$-simplex α'. However, in this case $\alpha', \beta', \alpha, \beta$ is a V-path, which contradicts the maximality of β. □

9. From Deterministic Search Algorithms to Gradient Vector Fields

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.

Let A be a decision tree algorithm. Here we require A to be a complete decision tree algorithm, i.e. A gives instructions for the order in which each of the n questions is to be asked, even if we complete the task by asking fewer questions. Then A induces a vector field V consisting of disjoint pairs $\{\alpha^{(p)} \prec \beta^{(p+1)}\}$ where α and β are not distinguished by A until all n questions have been asked. More precisely, suppose $\sigma = \alpha$. Asking questions according to A, the n^{th} question will be "Is x_i in σ?" for some i. Suppose x_i is not in α. Let $\beta = \alpha * x_i$. Then α and β are not distinguished until the n^{th} question, so $\{\alpha, \beta\}$ is a pair in V. If x_i is in α, then let $\gamma^{(p-1)} = \alpha - x_i$ so that $\alpha = \gamma * x_i$. Then γ and α are not distinguished until the n^{th} question, so $\{\gamma, \alpha\}$ is a pair in V.

For example, if $n = 3$, the decision tree shown in Figure 9.1 induces the following pairing of the faces of the two-dimensional simplex.

$$V = \{\{\phi \prec [v_3]\}, \{[v_2] \prec [v_2, v_3]\}, \{[v_1] \prec [v_1, v_2]\}, \{[v_1, v_3] \prec \{v_1, v_2, v_3\}\}\}.$$

Note that if $\mathcal{P} : \Sigma = P_1 \sqcup P_2 \ldots P_n$ is a partition then the evaders of A are precisely the pairs $\{\alpha^{(p)}, \beta^{(p+1)}\}$ in V such that α and β are in different P_i's.

We can now state the main theorem of this section.

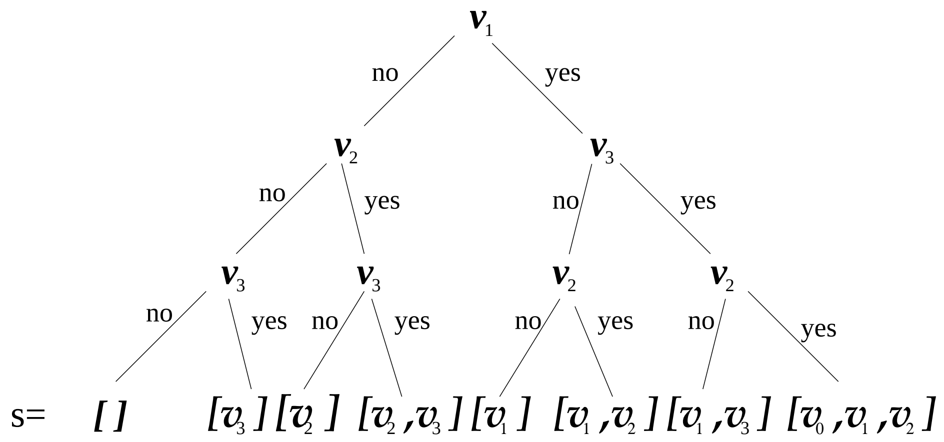

Figure 9.1. A search algorithm.

Theorem 9.1. *For any search algorithm A, the vector field V is a gradient vector field.*

Proof: We will show that there are no closed V-paths. The theorem then follows from Theorem 8.2. More precisely, we will define an ordering $\rightarrow$ on the set of 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 V-path then

$$\alpha_0 \rightarrow \alpha_1 \rightarrow \alpha_2 \rightarrow \ldots \ .$$

This obviously implies that there are no closed V-paths.

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

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

Namely, suppose $\sigma = \alpha$. Then the seeker asks, according to the decision tree algorithm A, n questions of the form "Is x_i in σ?" We number these questions consecutively from 1 to n. 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 order the simplices lexicographically according to their n-strings. That is, if $\alpha^{(p)}$ and $\beta^{(q)}$ are faces of S, we set $\alpha \to \beta$ if there is a $k \le \min\{p, q\}$ 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 a V-path, so that we have $\{\alpha^{(p)}, \beta^{(p+1)}\} \in V$, and $\alpha_1^{(p)} \ne \alpha_0^{(p)}$ is another face of β_0. Since $\{\alpha_0, \beta_0\} \in V$, we have by definition

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

and

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

Number the vertices of β_0 $v_0, v_1, \ldots, v_p, v_{p+1}$ so that if $\sigma = \alpha_0$ or β_0 then the $n_i(\beta_0)$ question is "Is v_i in σ?" Then the vertices of $\alpha_1^{(p)}$ are $v_0, v_1, \ldots, v_{k-1}, v_{k+1}, \ldots, v_{p+1}$ for some k. If $\sigma = \beta_0$, the first $n_k(\beta_0) - 1$ questions involve only $v_0, \ldots, v_{k-1}$ and vertices which are not in β_0 (and hence not in α_0 or α_1). Therefore, the first $n_k(\beta_0)$ questions are the same whether $\sigma = \alpha_0$, β_0 or α_1. If $\sigma = \alpha_1$, questions $n_0(\beta_0), \ldots, n_{k-1}(\beta_0)$ are answered "yes", and all other questions before $n_k(\beta_0)$ are answered "no", since they concern vertices not in β_0. Moreover, the $n_k(\beta_0)$ question is answered "no" because $v_k \notin \alpha_1$. Thus

$$n_i(\alpha_1) = n_i(\alpha_0) \quad \text{for} \quad i < k$$

and

$$n_k(\alpha_1) > n_k(\alpha_0),$$

which implies that $\alpha_0 \succ \alpha_1$. □

10. The Proof of Theorem 0.9

In this section we piece together the ingredients of the previous sections to show how they are sufficient to prove Theorem 0.5. Let Σ be a finite set, $\mathcal{Q} = \{\mathcal{Q}_1, \mathcal{Q}_2, \ldots, \mathcal{Q}_n\}$ be a set of n "yes-no" questions about Σ, and $\mathcal{P} : \Sigma = \mathcal{P}_1 \sqcup \mathcal{P}_2 \sqcup \ldots \sqcup \mathcal{P}_k$ a partition of Σ which is separated by $\mathcal{Q}$. Let S be an $(n-1)$ dimensional simplex with vertices labelled $\{1, 2, \ldots, n\}$ and $P_1, P_2, \ldots, P_k$ the simplex spaces corresponding to the sets of the partition $\mathcal{P}$. Let $P^{\emptyset}$ denote the simplex space of all faces of S which are not in any P_i. Let $\mathcal{A}$ be a deterministic search algorithm. As in Section 9, even though we may stop asking questions once we determine which $\mathcal{P}_i$ contains the unknown $\sigma \in \Sigma$, we assume that $\mathcal{A}$ is a complete algorithm, determining the order in which all n questions are to be asked, given the answers to the preceeding questions.

As shown in Section 9, a deterministic search algorithm A induces a gradient vector field V on Σ, in which each face s of Σ is paired with the face s' with the property that s and s' result in the same answers to each of the first $n-1$ questions. In particular, every face of S is paired with another face of Σ, so V correponds to a simplicial collapse of Σ.

Now for each i, let P_i^* denote the union of P_i and all elements of $P^{\emptyset}$ which are paired in V with a face of P_i. Now let V_i denote the restriction of V to P_i^*. That is, V_i denotes the set of pairs in V such that both elements of the pair belong to P_i^*. A face s in P_i^* is not paired in V_i if and only if s lies in P_i, and s is paired in V with a face s' which is in P_j, for some $j \neq i$. This is true if and only if one must ask all n questions before determining whether or not s is P_i (given that s is not in $P^{\emptyset}$), i.e. if and only if s is a face of P_i which is an evader of the algorithm A.

Thus the evaders of A which lie in P_i correspond to the critical faces of a gradient vector field V_i on P_i^*. The P_i^*'s are disjoint, and satisfy $P_i \subseteq P_i^* \subseteq P_i \cup P^{\emptyset}$. Theorem 0.9 (ii), and hence Theorem 0.5, now follows from Theorem 8.1 which provides lower bounds for the number of critical faces of any gradient vector field in terms of the Betti numbers of P_i^*. Theorem 0.9 (i), and hence Theorem 0.3, follows from the observation that if $\mathcal{P}$ is nonevasive, then there exists a search algorithm A with no evader, which, in turn, induces a gradient vector field on each P_i^* with no critical faces. Such a gradient vector field corresponds to a complete simplicial collapse of each P_i^*. $\square$

References

[1] A Björner and L. Lovász, *Linear Decision Trees, Subspace Arrangements and Möbius Functions*, J. of the Amer. Math. Soc., **7** (1994), pp. 677-706.

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

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

[4] ______, *Combinatorial differential topology and geometry*, in New Perspectives in algebraic Combinatorics (Berkeley, CA. 1996-97), Math. Sci. Res. Inst. Publ. 38, Cambridge Univ. Press, Cambridge, 91999), pp. 177-206.

[5] ______, *Morse Theory and Evasiveness*, Combinatorica, **20** (2000), pp.498-504.

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

Department of Mathematics, Rice University, Houston, TX 77251