

Morse decompositions and connection matrices

RICHARD MOECKEL

*School of Mathematics, University of Minnesota–Twin Cities,
Minneapolis, MN 55455, USA*

Abstract. This paper surveys the work of Charles Conley and his students on Morse decompositions for flows on compact metric spaces, as well as the more recent development of the connection matrix formalism for detecting connections between the Morse sets of a Morse decomposition.

0. Introduction

This paper is a survey of some of the ideas of Charles Conley and his students about the qualitative description of flows. It was Charlie's view that a useful qualitative theory should have two key properties. First, it should be applicable to continuous flows on fairly general topological spaces. This is important even in the context of smooth flows on manifolds when it is necessary to restrict attention to a compact invariant set or to form a quotient space. Second, the theory should be stable under perturbation of the flow. This insures that the qualitative predictions are physically meaningful. It also makes the various elements of the qualitative description 'computable' by permitting perturbation of a given flow to a simpler one nearby.

A fundamental result of Conley's is that every flow on a compact metric space can be decomposed as a chain recurrent part and a gradient-like part [3]. The chain recurrent set is a compact invariant set which contains the non-wandering set plus points which, roughly speaking, become recurrent when small errors are introduced into the flow. On each component of the chain recurrent set, the flow is chain transitive. The flow on the rest of the space is gradient-like in the sense that there is a continuous Lyapunov function which is strictly decreasing on orbits which are not chain recurrent. If the components of the chain recurrent set are identified to points, a gradient-like flow is obtained on the quotient space. Thus the problem of providing a qualitative description falls naturally into two parts, the description of the components of the chain recurrent set and the description of how these components are connected to one another.

The simplest instance of this decomposition is the gradient flow of a Morse function ([10], [12]). In this case the components of the chain recurrent set are the critical points of the function. The function itself can be used as a Lyapunov function. Morse's approach to the description of the components was to assign to each restpoint an index which captures the local features of the flow. Connections between restpoints can be found by using the Lyapunov function to construct a filtration of the space by positively invariant sets such that the difference of two successive sets

in the filtration contains a single restpoint. Connections between adjacent restpoints are then reflected in the relative topology of the sets in the filtration.

Conley's approach to the general case is similar. In the light of the decomposition theorem mentioned above, it would be natural to try to assign an index to each component of the chain recurrent set and then use the Lyapunov function to construct a filtration. However, without further assumptions on the flow, this decomposition can be too fine to be stable. In particular, the chain recurrent set may have infinitely many components and these can accumulate on one another. These difficulties led to the idea of a Morse decomposition. A Morse decomposition of a flow is a finite collection of disjoint, compact invariant sets which together contain all the recurrent behaviour of the flow in the sense that if they are identified to points, a gradient-like flow is obtained. The elements of the decomposition are called Morse sets. A given flow may have many such Morse decompositions. The chain recurrent set is precisely the set of points which belongs to every Morse decomposition. In the case where the chain recurrent set has only finitely many components, these components form a Morse decomposition which is the finest possible for the given flow. For example, when the chain recurrent set is hyperbolic, the Morse sets in this finest Morse decomposition are just the basic sets of the spectral decomposition ([1], [6]). In the general case, each Morse set consists of some of the components of the chain recurrent set together with the orbits connecting them. A Morse decomposition is stable in the sense that nearby flows admit similar decompositions.

Once a Morse decomposition has been found, a qualitative description analogous to that of Morse can be given. The individual Morse sets are described by the Conley index which provides a stable description of the local dynamics. Connections between the Morse sets can be found by constructing filtrations by positively invariant sets and studying the topology of these sets. The principal tool for this purpose is the connection matrix. For a Morse decomposition with N Morse sets, the connection matrix is an $N \times N$ matrix whose entries are maps between homology groups associated to the various Morse sets. These entries reflect the existence of connecting orbits in the flow. In particular, if there is no chain of connecting orbits between two given Morse sets, then the corresponding entry of the matrix will be a trivial map. It is sometimes possible to combine the general properties of connection matrices with topological information about a filtration to deduce that a certain entry is non-trivial and so prove the existence of connecting orbits.

Even in the classical case of a gradient flow, the connection matrix can provide more information than the usual analytical methods using the Morse inequalities or spectral sequences. The reason is that the connection matrix describes a more general type of filtration. A filtration constructed from a Lyapunov function consists of a nested sequence of compact positively invariant sets. The Morse sets are between successive sets of the filtration. This imposes a total ordering on the Morse sets which is usually quite unnatural in the sense that a different Lyapunov function could produce a different ordering. The topological analysis of such a filtration can usually find connections only between Morse sets which are adjacent in the ordering. Now there is a natural partial ordering associated to a Morse decomposition; namely,

two Morse sets are related in the order iff there is a sequence of connecting orbits between them. Of course if this partial order were known, there would be no need for a connection matrix. It may happen however that some information about this order is known and it is important to use it to deduce more. The starting point for the connection matrix theory is a Morse decomposition with a partial order known to be compatible with the flow. Then a filtration can be constructed not with nested sets but rather with sets whose intersections and unions reflect the partial order. From this a connection matrix is computed which can detect connections between Morse sets which are adjacent in the partial order.

The goal of this paper is to explain what a connection matrix is, how to construct one for a given flow and how to use one to find connecting orbits. The question of why they exist will remain somewhat mysterious. However, it is true of many algebraic topological machines that they can be used effectively in practice long after the details of their construction have been forgotten. For more information one should read Franzosa's original paper ([7], [8]). Section 1 concerns partially ordered Morse decompositions and the associated filtrations. Section 2 is devoted to connection matrices and includes an axiomatic presentation of their main properties as well as complete computations of the connection matrices for two examples.

Connection matrices have been applied to problems in ecology by Reineck [13] and to travelling waves by Mischaikow [11]. (The author would like to thank them both for their careful lectures on the subject at the University of Minnesota.)

1. Morse decompositions

Let S be a compact Hausdorff space with a flow $\phi(x, t) = x \cdot t$; $x \in S$, $t \in \mathbf{R}$. S could be a compact invariant set in some larger space which need not be compact. If K is any subset of S , we define the forward limit set $\omega(K)$ to be the largest invariant set in the closure of $K \cdot [0, \infty)$ and the backward limit set $\omega^*(K)$ to be the largest invariant set in the closure of $K \cdot (-\infty, 0]$.

Let $\mathbf{M} = \{M(p) : p \in P\}$ be a finite collection of disjoint compact invariant sets. For these sets to form a Morse decomposition it is required that all of the recurrent behaviour of the flow takes place inside them. In other words, orbits cannot cycle back and forth between two of the sets. This can be achieved by insisting that orbits run downhill with respect to some partial order on the collection.

Definition. A Morse decomposition of S is a finite collection $\mathbf{M} = \{M(p) : p \in P\}$ of disjoint, compact invariant subsets of S which can be partially ordered in such a way that for every $x \in S$ either

- (i) $x \in M(p)$ for some p

or

- (ii) there exist $M(p) < M(p^*)$ with $\omega(x) \subset M(p)$ and $\omega^*(x) \subset M(p^*)$.

The collection together with the partial order is called an ordered Morse decomposition of S . A subset of S which belongs to some Morse decomposition is called a Morse set.

It should be pointed out that every Morse decomposition is naturally ordered by the flow. Namely, we can define a relation by setting $M(p) < M(p')$ if there is some

orbit running from $M(p')$ to $M(p)$ and then extend to get transitivity. Because of the no-cycle condition in the definition, this relation is a partial order. It is the smallest admissible partial order for M . However, it is sometimes useful to consider stronger orderings, especially when the natural ordering is unknown.

Example. Let S be a compact Riemannian manifold and $F: S \rightarrow \mathbf{R}$ a smooth function with finitely many critical points. The vector field $-\text{grad } F$ determines a flow of S . Let M be the set of critical points on F . If the critical values of distinct critical points are distinct, then they provide a total order on M and M is an ordered Morse decomposition of S . The ordering is generally different from the natural one. For example, in figure 1 the partial order induced by the flow is $A < B < D$, $C < D$, $C < E$, $A < E$.




FIGURE 1

Once a Morse decomposition is known, a variety of cruder ones can be obtained by combining some of the Morse sets and their connecting orbits into bigger Morse sets. Some care must be taken to guarantee that the new sets are compact and that the no-cycle condition holds. A subset I of M is called an interval if whenever $M(p), M(p') \in I$ and $M(p) < M(p'') < M(p')$, also $M(p'') \in I$. If no order is specified, the natural one is assumed. Let $C(M(p), M(p'))$ or just $C(p, p')$ denote the set of connecting orbits running from $M(p)$ to $M(p')$; in particular, $C(p, p) = M(p)$. Then the amalgamated Morse set is $M(I) = \cup C(p, p')$, where the union runs over all pairs of Morse sets in I . It is not difficult to show that this is a compact invariant set. If the same construction is attempted with a subcollection which is not an interval, compactness may be lost. For example, figure 2 depicts a flow on a triangle with a Morse decomposition $\{A, B, C\}$. The union of A, C and $C(A, C)$ is not compact because $\{A, C\}$ is not an interval. On the other hand every other subset of $\{A, B, C\}$ is an interval and the corresponding amalgamations are compact. Viewing $M(I)$ as a single Morse set leads to the collection M/I consisting of $M(I)$ together with the Morse sets not contained in I . It can be shown that M/I is a Morse




FIGURE 2


decomposition ([7], [8]). Hence $M(I)$ is a Morse set. If we attempt to combine Morse sets which do not form an interval, then even if the amalgamated set is compact, the no-cycle condition may be violated. For example, in figure 1 the union of A, D and $C(A, D) = \emptyset$ is compact but it is connected to B by orbits running in both directions.

The importance of this construction lies in the fact that cruder Morse decompositions are easier to understand. An analysis of several of the decompositions M/I may lead to an understanding of M itself.

The stability of Morse decompositions derives from their relationship with attractors and repellers which are inherently stable.

Definition. An attractor in S is a compact invariant set which is the ω -limit set of some neighbourhood of itself in S . A repeller in S is a compact invariant set which is the ω^* -limit set of some neighbourhood of itself in S .

If A is an attractor, then the set of all points of S which are not attracted to A is a repeller, A^* , called the dual repeller of A [2]. The collection $\{A, A^*\}$ is called an attractor–repeller pair; it is the simplest instance of a Morse decomposition. Morse decompositions can be amalgamated into attractor–repeller pairs, usually in many different ways. Suppose that I is an attracting interval in M ; that is, an interval such that $M(p) \in I$ and $M(p') < M(p)$ implies $M(p') \in I$. Let $I^* = M \setminus I$. Then I^* is also an interval and $\{M(I), M(I^*)\}$ is an attractor–repeller pair. In the example of figure 1, each of the following is an attracting interval (the natural order is used):



It is evident that the collection of attractors arising from this ordered Morse decomposition is closed under unions and intersections. This holds for every Morse decomposition.

Definition. An attractor filtration in S is a finite collection \mathbf{A} of attractors in S satisfying

- (i) $\emptyset \in \mathbf{A}$ and $S \in \mathbf{A}$,
- (ii) if $A, A' \in \mathbf{A}$, then $A \cap A', A \cup A' \in \mathbf{A}$.

These are just the axioms for a topology on S ; thus the collection of attractors of an ordered Morse decomposition is a (finite) topology.

The ordered Morse decomposition can be recovered from the attractor filtration. First, the Morse sets are just the intersections $A^* \cap A'$, where $A \subset A'$ are attractors in the filtration with no intermediate attractors; i.e. $A \subset A'' \subset A'$ implies $A'' = A$ or $A'' = A'$. Next, a partial order is defined by the rule: $M(p) \leq M(p')$ iff every attractor containing $M(p')$ also contains $M(p)$. One can check that this is the same as the original partial order. For example, from the fact that, in the diagram above, the smallest attractor containing D is $M(\{A, B, C, D\})$, we recover the fact that $A < D$, $B < D$ and $C < D$ but $E \not< D$. In this correspondence, stronger partial orders determine smaller attractor filtrations.

An interesting corollary of this discussion is the characterization of Morse sets in S as intersections of attractors and repellers. We saw that any Morse set can be so represented: conversely, if $M = A^* \cap A'$, where $A \subset A'$ are attractors, then $\mathbf{M} = \{A, M, A'^*\}$ is a Morse decomposition so M is a Morse set.

The one-to-one correspondence between ordered Morse decompositions and attractor filtrations can be used to continue an ordered Morse decomposition to a nearby flow. To see this we will study the continuation of attractors or, more generally, of isolated invariant sets. We will allow perturbations not only of the flow but of the space S itself. Thus we will admit the possibility that S is a compact invariant subset of a flow on a larger space Γ and perturb the larger flow. To keep the theory flexible, it is preferable not to impose stringent hypotheses on Γ , but the constructions to follow require that $S \subset X \subset \Gamma$, where X is a locally compact metrizable space. X need not be invariant under the flow but must possess a kind of local positive invariance, as described in the following definition.

Definition. $X \subset \Gamma$ is a local flow if for every $x \in X$ there is a neighbourhood U relative to X and a $T > 0$ such that

$$U \cdot [0, T) \subset X.$$

As an example, imagine a flow on \mathbf{R}^3 which leaves the (x, y) -plane invariant. Suppose that the flow in the plane has the behaviour indicated in figure 3 on the boundary of a square. Then if Γ is \mathbf{R}^3 , X is the open square and S is the set of all orbits which stay in X for all time, all of the above conditions will be met. Note that X need not have interior in Γ . This approach is useful for studying perturbations; for example, if the flow on \mathbf{R}^3 also preserves the planes parallel to the (x, y) -plane, it can be viewed as a continuous one-parameter family of flows in the plane. In an




FIGURE 3

application to partial differential equations, Γ might be a function space and X a locally compact subset known to contain the solutions of interest.

Consider a perturbation of the local flow X . Clearly some condition on $S \subset X$ is needed to guarantee that there will be a unique compact invariant set of the flow which can be identified as the perturbation of S .

Definition. Let $X \subset \Gamma$ be a local flow. A compact set $N \subset X$ is an isolating neighbourhood in X if the maximal invariant subset of N lies in the interior of N relative to X . $S \subset X$ is an isolated invariant set in X if it is the maximal invariant set in some isolating neighbourhood in X .

The condition for an isolating neighbourhood can be rephrased as follows: boundary points of N eventually leave N in forward or backward time. Since N and its boundary are compact, this condition is stable under perturbations of the local flow which are small in the compact-open topology. Invariant sets in nearby local flows which are isolated by the same isolating neighbourhood are said to be related by continuation. This can be extended to distant flows by using a sequence of isolating neighbourhoods.

Let S' denote a continuation of S to a nearby flow. If A is an attractor in S , we will show that A continues to an attractor A' in S' . Let U be a neighbourhood of A in S such that $\omega(U) = A$ and let V be an isolating neighbourhood of S and S' in X . Now U is of the form $W \cap S$, where W is a neighbourhood of A in X ; then $V \cap W$ is an isolating neighbourhood of A in X . Thus A continues to some isolated invariant set $A' \subset S'$. Since A is an attractor, there will be some $T > 0$ with $U \cdot T \subset V \cap W$ and this will still hold if U is fattened to a neighbourhood W' of A

in X . If the perturbation is small enough, we will have $A' \subset W'$ and $(W' \cap S')T \subset V \cap W \cap S'$. This implies that $\omega(W' \cap S') = A'$, so A' is an attractor in S .

Now we can continue an ordered Morse decomposition of S to an ordered Morse decomposition of S' by making use of the equivalent attractor filtration. Each attractor in the filtration continues to every sufficiently near flow. Since there are only finitely many attractors in the filtration, we can continue them all at once. Moreover, continuation of attractors commutes with unions, intersections and inclusions, so we obtain an attractor filtration of S' with the same order structure. From this we can construct an ordered Morse decomposition by the method described above.

Consider figure 1 again. If we take a perturbation which eliminates the restpoint B , then the Morse set corresponding to B in the perturbed Morse decomposition will be empty. The ordering will no longer be the natural one since the new B is not connected to anything. A different perturbation produces a Morse set consisting of three restpoints; the ordering is the natural one since the connections from D to B and from B to A persist.

In the next section we take up the qualitative analysis of a flow with a Morse decomposition. However, it is often very difficult to find even a crude Morse decomposition for a given flow. Most of the applications of the theory have been to flows where some kind of Lyapunov function is available to limit the recurrent behaviour. We will now describe a flow of this type which will be as an example throughout the rest of the paper. It is modelled on a problem in celestial mechanics: the study of the collision manifold of the planar three-body problem.

Consider a flow on $\Gamma = \mathbf{S}^2 \times \mathbf{S}^2$ which is gradient-like with respect to a function $f: \Gamma \rightarrow \mathbf{R}$, i.e. f is strictly decreasing on orbits which are not restpoints. We will study the restpoint connections in the positively invariant region $X = \{f \leq 0\}$. X is homeomorphic to $\mathbf{S}^2 \times \mathbf{D}^2$. More generally, for $c < 0$, let $X_c = \{f \leq c\}$. Let $\pi: X \rightarrow \mathbf{S}^2$ be the projection to the first factor. It is given that the projections $\pi(X_c)$ are the sets $K_c = \{g \leq c\}$, where g is a function on \mathbf{S}^2 whose level sets are shown in figure 4. Moreover $\pi: X_c \rightarrow K_c$ is a homotopy equivalence for every c ; all of the topology of the level sets of f is reflected in the level sets of g . The local flow on X has eight restpoints and their projections are the eight critical points of g visible in figure 4. All of the restpoints are hyperbolic and their unstable manifolds have the same dimensions as those of the projected restpoints do in the flow of $-\text{grad } g$. The restpoint connections in the latter flow are clear from the figure: each saddle is connected to the two adjacent sinks and each source is connected to all six saddles and sinks. The problem is to show that the corresponding restpoint connections occur in the local flow X . It is relatively easy to study the connections originating at the restpoints C , since the unstable manifolds are one-dimensional. Suppose it is given that these connections are the same as those in \mathbf{S}^2 ; we will see below that restpoint connections originating at E' and E'' can be found very easily by using the connection matrix. For now, we will just choose convenient Morse decompositions. The main Morse decomposition of interest is simply the set of restpoints, but to simplify computations later we will also consider the coarser decomposition





FIGURE 4

obtained by amalgamating E' and E'' into a single Morse set E . The natural orderings are as yet unknown, but the assumptions above imply that the natural partial ordering of the corresponding restpoints in S^2 is admissible. We will use this partial ordering in what follows rather than the ordering by f -values, since it is preferable to use the weakest order known to be admissible. The partial order is represented by the graph below, where the existence of a directed path from A to B represents the relation $B < A$.



2. Connection matrices

Let Γ be a topological space with a flow, X a locally compact metric local flow in Γ , and S an isolated invariant set in X . Suppose an ordered Morse decomposition of S , $M = \{M(p) : p \in P\}$, is known. In this section we will associate to M an algebraic structure which reflects the topological structure of the Morse sets of M , the nature of the local flow nearby and the pattern of connecting orbits between them. This algebraic model of M will be stable with respect to perturbations of the local flow X .

The Conley index provides a stable description of an isolated invariant set in a local flow. The basic properties of this index will be given briefly below (see [2], [3], [5] and [14] for details). First we will show that not only S , but every Morse set in S , is isolated. Let I be an interval in M and let $M(I)$ be the corresponding

Morse set. We saw in § 1 that $M(I)$ is the intersection of an attractor and a repeller in S . Since attractors and repellers in S are themselves isolated invariant sets in X , $M(I)$ is an isolated invariant set in X .

We will use the Conley index of $M(I)$ to represent $M(I)$ in the algebraic model. The index of an isolated invariant set S is actually more a topological structure than an algebraic one. From the geometrical information contained in the family of isolating neighbourhoods of S , it condenses the stable part. The construction depends on the existence of special neighbourhoods of S in X .

Definition. Let $N_0 \subset N_1$ be compact subsets of X . N_0 is positively invariant in N_1 if whenever $x \in N_0$, $t > 0$ and $x \cdot t \in N_1$, then $x \cdot t \in N_0$. N_0 is an exit set for N_1 if it is positively invariant in N_1 and if for any $x \in N_1$ with $x \cdot [0, \infty) \not\subset N_1$ there is a $T > 0$ with $x \cdot [0, T) \subset N_1$ and $x \cdot T \in N_0$.

A pair of compact sets (N_1, N_0) in X is an index pair for S in X if N_0 is an exit set for N_1 and if the closure of $N_1 \setminus N_0$ is an isolating neighbourhood for S .

In figure 3 the compact pair consisting of the square and the vertical line segments in its boundary is an index pair for whatever invariant set is inside the square. Another index pair for the same isolated invariant set is shown in figure 5 (assuming




FIGURE 5

nothing stays in the smaller square). It can be shown that any isolated invariant set has index pairs [2]. In practice, they are often found with Lyapunov functions. In the example described at the end of § 1, compact pairs of the form (X_c, X_d) with $d < c$ are index pairs provided c and d are not 'critical values'.

The concept of an index pair evolved from that of an isolating block [4]. From our point of view, an isolating block is just a particularly nice index pair which exists in the setting of smooth flows on manifolds. A fundamental property of isolating blocks is that for points in N_1 which eventually leave, the time required to hit the exit set depends continuously on the initial conditions. An index pair with this property is called regular, and it is known that regular pairs always exist [14]. The pairs that one constructs in examples are invariably regular.

A comparison of figures 3 and 5 shows that a given isolated invariant set may have topologically distinct index pairs. However, the relative topology of any two index pairs is the same up to homotopy. From an index pair one can construct a quotient space N_1/N_0 with a distinguished point (the equivalence class of N_0). It is a remarkable fact that the quotient spaces associated to two different index pairs are homotopy equivalent ([2], [14]).

Definition. The Conley homotopy index of S , $h(S)$, is the common homotopy type of the quotiented index pairs of S .

In figures 3 and 5, if the exit sets are collapsed to a point, the resulting spaces can be deformed into circles. Thus the homotopy index of the isolated invariant set in the square is the homotopy type of a pointed circle, i.e. a circle with a distinguished point. It is possible (but in no way necessary) that the invariant set in the square is a saddle point. Thus the index of a saddle point in the plane is a pointed circle. More generally, it is easy to show that the homotopy index of a hyperbolic restpoint in a flow on a manifold is the homotopy type of a pointed sphere of dimension equal to that of the unstable manifold of the restpoint; the dimension of this sphere is just the usual Morse index.

The homotopy equivalences between the different quotiented index pairs are constructed by using the local flow. These homotopy equivalences can be included in the index object to provide additional structure when necessary.

Definition. The Conley index of S , $I(S)$, is the collection of all quotiented index pairs of S in X with the additional structure of a homotopy class of flow-induced homotopy equivalences between any two of them.

This index is a so-called connected simple system of spaces and homotopy classes of maps. It squeezes virtually all of the useful information out of the local dynamics near S in X .

We will not make use of the full structure of the Conley index in what follows. In fact, it is useful to replace the homotopy index with something even weaker but more algebraic. Most of the algebraic objects associated to spaces in algebraic topology are homotopy invariant. For example, the singular homology groups of a space depend only on the homotopy type. Any such object associated to N_1/N_0 can therefore be viewed as an algebraic index of S . The homology index of S , $H(S)$, is the graded group $\{H_n(h(S)): n = 0, 1, \dots\}$, where $H_n(h(S))$ stands for the n th singular homology group of any of the pairs $(N_1/N_0, *)$, where $*$ is the distinguished point, with coefficients in some Abelian group G which is suppressed in the notation. The homology index of a hyperbolic restpoint with d -dimensional unstable manifold is: $H_d(S) = G$, $H_n(S) = 0$ for $n \neq d$. It is often simpler to compute with the homology index than with the homotopy index, although some information is lost in the transition.

To study all of the Morse sets $M(I)$ in an ordered Morse decomposition M simultaneously, it is important to choose index pairs systematically. The main problem is to find index pairs for the attractors in the corresponding attractor filtration.

Definition. Let \mathbf{A} be an attractor filtration in S . An index filtration for \mathbf{A} in X is a collection of compact subsets of X , $\{N(A) : A \in \mathbf{A}\}$, such that

(i) $N(A \cap B) = N(A) \cap N(B)$ and $N(A \cup B) = N(A) \cup N(B)$ for every $A, B \in \mathbf{A}$.

(ii) $(N(A), N(\emptyset))$ is an index pair for A in X .

If (N_1, N_0) is an index pair for S in X , then there exists such an index filtration for any given attractor filtration with $N(S) = N_1$ and $N(\emptyset) = N_0$ ([7], [8]). For the special case of a totally ordered decomposition, this fact was known earlier. The difficulty in the partially ordered case is in trying to satisfy both the intersection condition and the union condition. It is instructive to experiment with the flow in figure 1.

Once an index filtration has been constructed, we have index pairs for any Morse set in the decomposition. Let \mathbf{I} be any interval in \mathbf{M} . Then $M(\mathbf{I})$ can be represented as the intersection of an attractor and a repeller, say $A \cap B^*$, where $A, B \in \mathbf{A}$ and $B \subset A$. It is then easy to see that $(N(A), N(B))$ is an index pair for $M(\mathbf{I})$.

We will choose an index filtration for the model problem from § 1. For the sake of simplicity we will use the coarser Morse decomposition. First we choose $N(S) = X = X_0$ and $N(\emptyset) = \emptyset$. This is an index pair for S . Note that for c large and negative the set K_c in S^2 is a union of three discs, one around each of the projected restpoints B_j . The corresponding components of X_c will serve as the neighbourhoods $N(B_j)$; they are four-dimensional balls. For $\{i, j, k\} = \{1, 2, 3\}$ there is an attractor in S consisting of B_i, B_j , and the unstable manifold of C_k . Call this attractor A_k . For $N(A_k)$ we choose a small ball around the respoint C_k and append forward orbits to make a positively invariant set. To this we add the previously chosen neighbourhoods $N(B_i)$ and $N(B_j)$. Thus $N(A_k)$ is a four-dimensional tube connecting two four-dimensional balls. Clearly this construction assures that $N(A_i) \cap N(A_j) = N(B_k)$ as required. Aside from S itself, the only other attractors in the attractor filtration are unions of those already discussed, and so the corresponding N are determined. Of particular interest later will be the attractor $A_1 \cup A_2 \cup A_3$ which we call simply A . A and E form an attractor–repeller pair in S . $N(A)$ consists of three balls cyclically connected by three tubes. A more natural choice for $N(A)$ might be a set X_c with c chosen just below the level of the restpoints which comprise E . We will use such a set later, but this choice would violate the union condition and so would not yield an index filtration.

When an index filtration has been constructed, all of the indices of the Morse sets $M(\mathbf{I})$ can be found. But we are also in a position to study connecting orbits between the Morse sets. We begin with connections between adjacent Morse sets.

Definition. A pair (\mathbf{I}, \mathbf{J}) of disjoint intervals of \mathbf{M} is adjacent if $\mathbf{I} \cup \mathbf{J}$ is an interval and if for any $M(p) \in \mathbf{I}$ and $M(q) \in \mathbf{J}$, $M(p) \prec M(q)$. The corresponding Morse sets are then an adjacent pair of Morse sets.

Since by definition the Morse set $M(\mathbf{I})$ lies above $M(\mathbf{J})$ in the ordering (or at least no lower), we will be looking for connecting orbits running from $M(\mathbf{I})$ to $M(\mathbf{J})$. One easily finds attractors $A \supset B \supset C$ in the attractor filtration such that $M(\mathbf{I}) = A \cap B^*$ and $M(\mathbf{J}) = B \cap C^*$. From the index filtration we obtain a triple of

compact sets $(N(A), N(B), N(C))$ such that the first two sets form an index pair for $M(I)$, the second pair form an index pair for $M(J)$, and the first and third form an index pair for $M(I \cup J)$. At this point we will retreat to the homology index, although it is possible to continue the discussion at the geometrical level by introducing exact sequences of topological spaces ([9], [14]). If we suppose that the index pairs above are all regular, then the homology of the quotiented index pairs is the same as the relative homology of the pairs themselves: for example $H(M(I)) = H(N(A), N(B))$. Therefore the exact homology sequence of the triple becomes a sequence relating the indices of the invariant sets:

$$\rightarrow H_n(J) \rightarrow H_n(IJ) \rightarrow H_n(I) \xrightarrow{\delta} H_{n-1}(J) \rightarrow .$$

Here we have written $H(J)$ instead of $H(M(J))$ and IJ instead of $I \cup J$. The connecting homomorphism $\delta : H_n(I) \rightarrow H_{n-1}(J)$ is called the connection map of the adjacent pair; it is independent of the choice of the triple used to compute it. If there were no connecting orbits from $M(I)$ to $M(J)$, then the Morse set $M(IJ)$ would just be their union. In this case, by choosing a triple which splits into parts near $M(I)$ and parts near $M(J)$, it is easy to show that the exact sequence above represents $H_n(IJ)$ as the direct sum of $H_n(I)$ and $H_n(J)$. It follows that the connection map is the trivial homomorphism. Therefore a non-trivial connection map implies the existence of a connecting orbit.

As an example, consider the model problem. We will take $I = \{E', E''\}$ and $J = M \setminus I = \{C_i, B_i : i = 1, 2, 3\}$. Then $IJ = M$. We have $M(I) = \{E', E''\}$, $M(J) = A$ and $M(IJ) = S$. Thus we are looking for connections from the restpoints E' or E'' to the attractor A . These obviously exist since the unstable manifolds of E' and E'' must go somewhere, but we will compute the connection map δ anyway. The triple of attractors $S \supset A \supset \emptyset$ can be used to represent these Morse sets as intersections of attractors and repellers. The triple (S, X_c, \emptyset) , with c just below the level of E' and E'' , can be used to compute the connection map. We have seen that homologically this triple can be replaced by the triple (S^2, K_c, \emptyset) in the projection. The relevant portion of the exact homology sequence of this triple is:

$$0 \rightarrow H_2(S^2) \rightarrow H_2(S^2, K_c) \rightarrow H_1(K_c) \rightarrow 0.$$

The left and right groups are isomorphic to G , the coefficient group, while the middle group is isomorphic to $G \oplus G$. The natural way to decompose the middle group is to use the generators α' and α'' depicted in figure 6. These two-dimensional discs are relative two-cycles in (S^2, K_c) . Also shown in figure 6 is a one-cycle β in K_c generating its first homology. With this basis the boundary map $\delta : G \oplus G \rightarrow G$ is given by the matrix $(1, -1)$, which means that $\delta(\alpha') = -\delta(\alpha'') = \beta$ and $\delta(g\alpha' + h\alpha'') = (g - h)\beta$ for $g, h \in G$.

The number of adjacent pairs of Morse sets arising from a given Morse decomposition can be enormous, and each adjacent pair determines a connection map. It is clear, however, that these maps cannot be completely independent of one another. In the model problem, for example, connections from E to A are also connections from E' or E'' to some of the restpoints in A . The connection maps from E to A and from E' and E'' to C_j must be related. Thus there is some hope of packing all of the connection map data into a reasonably simple structure.





FIGURE 6

Consider attractors $A \supset B \supset C \supset D$ and intervals $I = A \cap B^*$, $J = B \cap C^*$, $K = C \cap D^*$, $IJ = A \cap C^*$, $JK = B \cap D^*$ and $IJK = A \cap D^*$. The homology exact sequences of the triples $(N(A), N(B), N(C))$, $(N(A), N(C), N(D))$, $(N(A), N(B), N(D))$ and $(N(B), N(C), N(D))$ form a braided diagram which shows how the corresponding connection maps are related:



Using such braided diagrams, Franzosa was able to condense all of the information contained in the connection maps of an ordered Morse decomposition into a single matrix of maps, the connection matrix.

Before presenting the axioms for connection matrices, we will show how to use the braids to obtain a matrix in the relatively simple case of a totally ordered Morse decomposition with three Morse sets. These can play the role of I, J and K in the braid above if the ordering is $K < J < I$. The attractors are $A = S = IJK, B = JK, C = K$ and $D = \emptyset$. Then the six Morse sets which occur in the braid diagram are the only Morse sets of the decomposition, and the four connection maps in the diagram are the only connection maps of the decomposition. Thus the complete algebraic description of the decomposition consists of six homology indices and four connecting maps. The connection matrix approach will have the effect of storing the same information in three homology indices and three connection maps.

Assume that the coefficient group is the field $G = Z_2$, so that all homology groups are (graded) vector spaces. The three homology indices in the connection matrix description are those of I, J and K . Let $v = H(I) \oplus H(J) \oplus H(K)$. An element of V is a triple $(\alpha, \alpha', \alpha'')$, where the components are homology classes of the same degree. V is again a graded vector space. We will define a degree -1 mapping $\Delta: V \rightarrow V$ which summarizes all of the connection data. Such a mapping can be represented by a 3×3 matrix, each entry of which is a degree -1 mapping between two of the summands of V . To be specific, we associate the first row and column with I , the second with J and the third with K . So, for example, the entry in the first row and second column of Δ is a degree -1 map $\Delta(1, 2): H(I) \rightarrow H(J)$. Since all of the vector spaces are graded, the matrix should be visualized as having some 'depth'; alternatively, one can think of a sequence of matrices Δ_n taking degree n homology to degree $n - 1$ homology.

Some of the entries can be filled in on the basis of general principles which apply to all connection matrices. There can be no connecting orbits from a Morse set to another one which is higher with respect to the ordering, nor are there connecting orbits from a Morse set to itself. So it is natural to take all of the entries on or below the diagonal to be zero (the trivial map of graded vector spaces). Because (I, J) and (J, K) are adjacent pairs of Morse sets, we already have degree -1 connection maps between their homologies. We can set $\Delta(1, 2) = \delta(I, J)$ and $\Delta(2, 3) = \delta(J, K)$. From just this much information, we can recover the homology indices $H(IJ)$ and $H(JK)$ up to isomorphism by a simple procedure. Let the upper 2×2 diagonal block of Δ act on $V(IJ) = H(I) \oplus H(J)$. Call this block $\Delta(IJ)$:

$$\Delta(IJ) = \begin{bmatrix} 0 & \delta(I, J) \\ 0 & 0 \end{bmatrix}.$$

Since $\Delta(IJ)^2 = 0$, the graded vector space $V(IJ)$ together with the degree -1 map $\Delta(IJ)$ forms a chain complex and has a graded homology vector space: $H\Delta(IJ) = \text{image}(\Delta IJ) / \text{ker}(\Delta IJ)$. It turns out that $H\Delta(IJ)$ and $H(IJ)$ are isomorphic. Similarly $H(JK)$ can be recovered using the lower 2×2 diagonal block of Δ . It should be pointed out that the isomorphism between $H\Delta$ and H is not natural or 'flow-induced';

in fact it need not be unique. This artificiality is the price to be paid for storing $H(\mathbf{IJ})$ in the connection matrix. For the construction of such isomorphisms it is required that the coefficient group be a field.

It remains to recover $H(\mathbf{IJK})$ and the other two connection maps $\delta(\mathbf{IJ}, \mathbf{K})$ and $\delta(\mathbf{I}, \mathbf{JK})$. We will define the entry $\Delta(1, 3)$ so as to make this possible. We can try to find $H(\mathbf{IJK})$ by viewing V as a chain complex with boundary map Δ . First we should check that $\Delta^2 = 0$. The only entry of Δ^2 which is not obviously 0 is $\Delta^2(1, 3) = \delta(\mathbf{I}, \mathbf{J})\delta(\mathbf{J}, \mathbf{K})$ (where the maps are to be composed from left to right). From the braid diagram we find that $\Delta^2(1, 3)$ can be written as a composition of three maps, two of which are successive maps in an exact sequence. It follows that this entry of Δ^2 is indeed 0. Thus we can use Δ itself as a boundary map on V and so we obtain a homology group $H\Delta(\mathbf{IJK}) = \text{image}(\Delta)/\ker(\Delta)$. We will see that an isomorphism of $H\Delta(\mathbf{IJK})$ and $H(\mathbf{IJK})$ can be constructed if $\Delta(1, 3)$ is properly chosen.

The connection map $\delta(\mathbf{IJ}, \mathbf{K})$ comes from an exact sequence relating the indices $H(\mathbf{IJ})$, $H(\mathbf{IJK})$ and $H(\mathbf{K})$ (see the braid diagram). There is an analogous exact sequence involving the artificial homology indices $H\Delta$. There is a short exact sequence of vector spaces

$$0 \rightarrow V(\mathbf{IJ}) \rightarrow V \rightarrow V(\mathbf{K}) \rightarrow 0$$

where the maps are simply inclusion and projection of direct sums. Each of these vector spaces was made into a chain complex above by using an appropriate submatrix of Δ for the boundary map (the boundary map for $V(\mathbf{K}) = H(\mathbf{K})$ is the lower 1×1 diagonal submatrix, i.e. 0). From homology theory we obtain a long exact sequence which covers the one containing $\delta(\mathbf{IJ}, \mathbf{K})$:

$$\begin{array}{ccccccc}
 \rightarrow & H\Delta(\mathbf{K}) & \rightarrow & H\Delta(\mathbf{IJK}) & \rightarrow & H\Delta(\mathbf{IJ}) & \xrightarrow{\Delta(1,3)} & H\Delta(\mathbf{K}) & \rightarrow \\
 & \downarrow \text{id} & & \downarrow & & \downarrow \Phi(\mathbf{IJ}) & & \downarrow \text{id} & \\
 \rightarrow & H(\mathbf{K}) & \rightarrow & H(\mathbf{IJK}) & \rightarrow & H(\mathbf{IJ}) & \xrightarrow{\delta(\mathbf{IJ}, \mathbf{K})} & H(\mathbf{K}) & \rightarrow
 \end{array}$$

Here Φ is the isomorphism mentioned above. The connecting homomorphism $\Delta(1, 3)$ turns out to be induced by the map $(\Delta(1, 3), \Delta(2, 3)): V(\mathbf{IJ}) \rightarrow V(\mathbf{K})$. By choosing $\Delta(1, 3)$ properly, the right square can be made to commute. Now it follows that $H\Delta(\mathbf{IJK})$ and $H(\mathbf{IJK})$ have the same dimension and so are isomorphic; moreover, there is an isomorphism making the diagram commute. Thus by choosing $\Delta(1, 3)$ properly, we can recover both $H(\mathbf{IJK})$ and $\delta(\mathbf{IJ}, \mathbf{K})$. Similarly, some choice of $\Delta(1, 3)$ will allow us to recover $\delta(\mathbf{I}, \mathbf{JK})$. By using the braid diagram, one can show that there is some way to choose $\Delta(1, 3)$ which is compatible with both exact sequences. Still $\Delta(1, 3)$ need not be uniquely determined by the construction.

This construction can be carried out for any ordered Morse decomposition. To make this precise, we will present axioms that a connection matrix for the given decomposition should satisfy. Then the theorem will be that at least one such matrix exists.

Let $\mathbf{M} = \{M(p) : p \in P\}$ be the given decomposition and let V be the direct sum of the homology indices of the Morse sets: $V = \bigoplus_{p \in P} H(M(p))$. We always assume that the homology has field coefficients. Then V is a graded vector space over the

field. A connection matrix will be a degree -1 linear map $\Delta: V \rightarrow V$ with the properties described below. Such a map can be viewed as a matrix in the following way: Δ is determined by the linear maps $\Delta(p, q): H(M(p)) \rightarrow H(M(q))$, for $p, q \in P$, which are the components of the restrictions of Δ to the summands. If we assign a column to each summand, then elements of V can be viewed as row vectors with entries in the summands. Then the map Δ is represented by a matrix with entries $\Delta(p, q)$ which acts on the row vectors from the right. The entries $\Delta(p, q)$ are themselves linear maps of graded vector spaces, and so if bases were to be chosen for these spaces, they could be represented by matrices as well. Often the vector spaces $H(M(p))$ are one-dimensional in some degree and 0 in the other degrees. In this case the maps $\Delta(p, q)$ can be viewed as elements of the field and Δ becomes an ordinary matrix.

The first two properties of a connection matrix are:

CM1. Δ is a boundary map; that is, a degree -1 linear map with $\Delta^2 = 0$.

CM2. Δ is upper triangular with respect to the partial order; that is, $\Delta(p, q) = 0$ unless $p > q$.

These axioms are motivated by the need to recover the indices of the Morse sets $M(I)$, for intervals I , by a procedure like the one described above. Namely, let $V(I) = \bigoplus_{p \in I} H(M(p))$ and let $\Delta(I): V(I) \rightarrow V(I)$ be the ‘diagonal block’ of Δ with entries $\Delta(p, q)$, $p, q \in I$. Then one can show that the two axioms hold for $\Delta(I)$ as well. The graded vector space $V(I)$ together with the boundary map $\Delta(I)$ can be viewed as a chain complex, and so there is a graded homology vector space: $H\Delta(I) = \text{image}(\Delta(I))/\text{ker}(\Delta(I))$. We want this artificial homology vector space to be isomorphic to the real homology index $H(I)$.

CM3. For each interval I there is an isomorphism $\Phi(I): H\Delta(I) \rightarrow H(I)$. Moreover, if $I = \{p\}$ for some $p \in P$, then $\Phi(I) = \text{id}$.

We do not require these isomorphisms to be unique or in any way natural. We do want them to be compatible with the exact sequences containing the connection maps. If I and J are adjacent intervals and IJ the union interval, then there is an obvious inclusion map $V(J) \rightarrow V(IJ)$ and projection map $V(IJ) \rightarrow V(I)$. Since they commute with the restrictions of Δ , they can be viewed as chain maps of chain complexes. This leads to an exact sequence to be isomorphic to the real exact sequence relating the homology indices.

CM4. For each adjacent pair of intervals (I, J) the following diagram commutes:

$$\begin{array}{ccccccc}
 \rightarrow & H\Delta(J) & \rightarrow & H\Delta(IJ) & \rightarrow & H\Delta(I) & \xrightarrow{\bar{\Delta}(I,J)} & H\Delta(J) & \rightarrow \\
 & \downarrow \phi_{(I,J)} & & \downarrow \phi_{(IJ)} & & \downarrow \phi_{(I)} & & \downarrow \phi_{(J)} & \\
 \rightarrow & H(J) & \rightarrow & H(IJ) & \rightarrow & H(I) & \xrightarrow{\bar{\Delta}(I,J)} & H(J) & \rightarrow
 \end{array}$$

where $\Delta(I, J)$ is the submatrix of Δ consisting of all the $\Delta(p, q)$ for $p \in I$ and $q \in J$, and $\bar{\Delta}(I, J)$ is the induced map on homology.

To say that $\bar{\Delta}$ is induced by Δ means that $\bar{\Delta}[v] = [\Delta(v)]$, where the brackets denote homology classes; it is easy to check that this is well defined. When $I = \{p\}$ and

$\mathbf{J} = \{q\}$ for adjacent Morse sets $M(p)$ and $M(q)$, then the vertical maps $\Phi(\mathbf{I})$ and $\Phi(\mathbf{J})$ are identity maps by CM3. Thus $\Delta(\mathbf{I}, \mathbf{J}) = \Delta(p, q) = \delta(p, q)$. This observation is a double-edged sword. If a connection map $\delta(p, q)$ is known, then the corresponding entry in the connection matrix can be filled in. On the other hand, if such an entry in the connection matrix is deduced, we get information about the connection map $\delta(p, q)$ and hence about the orbits connecting the Morse sets $M(p)$ and $M(q)$. This is the approach we will use below in the examples.

An ordered Morse decomposition may have a great number of intervals and adjacent pairs of intervals, especially if the partial order is weak. Thus CM3 requires that a great deal of information be stored in the connection matrix, and CM4 involves many compatibility conditions on the isomorphisms. Nevertheless, Franzosa was able to show that every ordered Morse decomposition has at least one connection matrix ([7], [8]). The proof uses braid diagrams to construct Δ and the isomorphisms $\Phi(\mathbf{I})$ in such a way that all the compatibility conditions hold. In applications, the knowledge that at least one such matrix exists is often enough to deduce what some of the entries must be and so to obtain information about connecting orbits. For this purpose, no information about the construction of the rest of Δ or about the isomorphisms Φ is required.

Some ordered Morse decompositions have several connection matrices. If two partial orderings of the same Morse sets are given, then the weaker one will admit fewer connection matrices because it has more intervals and adjacent pairs of intervals and so CM3 and CM4 are harder to satisfy. Thus non-uniqueness may arise from using too strong a partial order, i.e. from lack of knowledge of the true flow-induced order. But even if the flow-induced order is used, there may be several connection matrices. This can be a reflection of the occurrence of unstable connections between the Morse sets. For the case of a flow-ordered Morse decomposition consisting entirely of hyperbolic restpoints with transverse stable and unstable manifolds, Reineck has shown that the connection matrix is unique [13].

Connection matrices are stable with respect to continuation of ordered Morse decompositions. If a Morse decomposition \mathbf{M} of S is continued to a Morse decomposition \mathbf{M}' of S' , then for any interval \mathbf{I} the indices of the corresponding Morse sets are isomorphic ([9], [14]). Moreover, these isomorphisms are natural, so that if \mathbf{I} and \mathbf{J} are an adjacent pair, the exact sequences containing the connection maps $\delta(M(\mathbf{I}), M(\mathbf{J}))$ and $\delta(M'(\mathbf{I}), M'(\mathbf{J}))$ are equivalent. From this it follows that a connection matrix for \mathbf{M} can be conjugated to a connection matrix for \mathbf{M}' with respect to the given order. We have seen in examples above that the flow-induced order on \mathbf{M} may not be the flow-induced order on \mathbf{M}' ; the new flow-induced order may be weaker because some connections are broken. Thus the continued connection matrices may not all work for the flow-induced order. However, the connection matrices for the flow-induced order on \mathbf{M}' will be a subset of the continued matrices. In particular, if the connection matrix for \mathbf{M} is unique, it will continue to the unique connection matrix for \mathbf{M}' .

We will bring this survey to a close by computing connection matrices for two examples. First consider the flow on the circle shown in figure 1 with the Morse

decomposition $\{A, B, C, D, E\}$. Restpoints D and E have index pairs consisting of an interval and its endpoints. The quotient spaces are pointed circles and so the graded homology indices with coefficients in $G = \mathbb{Z}_2$ are isomorphic to G in degree 1 and to 0 in all other degrees. Restpoints A and C have index pairs consisting of intervals with empty exit sets. The quotient spaces are intervals together with a disjoint distinguished point (the equivalence class of the empty set). The graded homology indices are isomorphic to G in degree 0 and to 0 elsewhere. Finally B has an index pair consisting of an interval and one endpoint. The quotient space is the same and so the graded homology index is 0 in every degree. This reflects the fact that B can be perturbed away.

The graded vector space V is the direct sum of these homology indices and is therefore trivial except in dimensions 0 and 1, so the degree -1 map $\Delta: V \rightarrow V$ is trivial except in dimension 1. Thus Δ reduces to a single 5×5 matrix of maps between degree 1 homology indices and degree 0 homology indices. Since these indices are all either trivial or one-dimensional, the entries of this matrix are either trivial or isomorphisms. With respect to any bases, the trivial map is represented by the 1×1 matrix 0 and an isomorphism by the 1×1 matrix 1. Thus Δ is represented by a 5×5 matrix of 0s and 1s.

We will assign the first row and column to E , the second to D and so on. Since the indices of A, B and C are trivial in degree 1, all of the entries in the corresponding rows are 0. Similarly all entries in the B, D and E columns are 0. This leaves the entries $\Delta(D, A)$, $\Delta(D, C)$, $\Delta(E, A)$ and $\Delta(E, C)$ undetermined:

$$\Delta_1 = \begin{bmatrix} 0 & 0 & \Delta(E, C) & 0 & \Delta(E, A) \\ 0 & 0 & \Delta(D, C) & 0 & \Delta(D, A) \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix}.$$

Now (D, C) , (E, A) and (E, C) are adjacent pairs of Morse sets, so by the observation following CM4, the corresponding entries in the matrix are connection maps δ . These could be computed by choosing an index filtration and constructing the exact sequences of appropriately chosen triples as described above in the definition of a connection map. But the main point of the connection matrix is to avoid such computations. To illustrate this we will compute the matrix from the axioms.

First of all we should be able to recover the index of the whole circle as the homology of the chain complex (V, Δ) . An index pair for the circle is the circle itself with empty exit set. The quotient space is a circle with a disjoint distinguished point and the homology index is G in degrees 0 and 1 and 0 in all other degrees. We have seen that V has dimension 2 in degrees 0 and 1 and dimension 0 in all other degrees. If Δ_n denotes Δ in degree n , then the homology of the complex (V, Δ) is

$$H\Delta_1 = \ker \Delta_1 / \text{image } \Delta_2 = \ker \Delta_1,$$

$$H\Delta_0 = \ker \Delta_0 / \text{image } \Delta_1 = G \oplus G / \text{image } \Delta_1.$$

Now Δ_1 is just the 5×5 matrix above. To reproduce the homology index of the circle, it must have one-dimensional kernel and image, i.e. rank 1. To get more information, note that since D is a repeller and C an attractor, $\{A, B, E\}$ is an interval I in the flow-induced order. An index pair for $M(I)$ is obtained by deleting neighbourhoods of C and D from S^1 . This pair is two disjoint intervals with one endpoint of each in the exit set. From this we see that $H(I)$ is trivial. Now the artificial homology $H\Delta(I)$ is obtained from the chain complex $(V(I), \Delta(I))$, where $V(I) = H(A) \oplus H(B) \oplus H(E) = H(A) \oplus H(E)$ and

$$\Delta(I) = \begin{bmatrix} 0 & 0 & \Delta(E, A) \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}.$$

For this complex to have trivial homology we must have $\Delta(E, A) = 1$. In a similar way we can show that the other three unknown entries of Δ are 1s, and so

$$\Delta = \begin{bmatrix} 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix}.$$

It is interesting to compare the amount of information necessary to derive this connection matrix with the amount of information which can be extracted from it. We started with the flow in figure 1 about which everything was already known, but in computing Δ we used only the following facts: the existence of a Morse decomposition with five Morse sets; the homology indices of these Morse sets and of the whole space; the fact that certain subsets of the decomposition form adjacent pairs of intervals; and the homology indices of these intervals. We did not explicitly use any information about connecting orbits other than the obvious lack of connecting orbits leading into D and E or out of C . To put it another way, we did not use our knowledge of the flow-induced partial order, except for the facts that in this order D and E are maximal while A and C are minimal. Nevertheless, we uniquely determined Δ . It follows that this Δ is a connection matrix for the flow-induced order. By CM2 we conclude that in the flow-induced order, $E > A$, $E > C$, $D > A$ and $D > C$. This means that each of D, E is connected to each of A, C by a connecting orbit or by a chain of connecting orbits. While this is obvious from the figure, it is not obvious how it follows from the minimal information we used to compute the matrix.

Now we turn to the model problem in $S^2 \times S^2$. Recall that we have a Morse decomposition $\{E', E'', C_j, B_j; j = 1, 2, 3\}$ consisting of eight restpoints. At the end of § 1 we gave a partial order which is admissible. Since all restpoints are hyperbolic, the homology indices are the homologies of pointed spheres of dimensions given by the unstable manifolds. Thus the indices of E' and E'' are G in degree 2 and 0 otherwise; the index of C_j is G in degree 1 and 0 otherwise; and the index of B_j is G in degree 0 and 0 otherwise. Thus Δ will be trivial except in degrees 1 and 2.

Many of the entries of the matrices Δ_1 and Δ_2 are also trivial because one of the corresponding vector spaces is zero-dimensional.

We will assign rows and columns to the restpoints in the order E', E'', C_1, \dots, B_3 . Then Δ_1 takes the form

$$\Delta_1 = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & * & * \\ 0 & 0 & 0 & 0 & 0 & * & 0 & * \\ 0 & 0 & 0 & 0 & 0 & * & * & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}.$$

Here * denotes an entry which is as yet unknown. All of the 0s are clear except for the ones embedded in the 3×3 block containing the *s. These are obtained from CM2 and the partial order; for example, $\Delta(C_1, B_1) = 0$ because in the partial order the relation $B_1 < C_1$ does not hold. We can determine the unknown entries by the following simple argument. Recall the attractor A consisting of the last six restpoints and their unstable manifolds. We found an attractor neighbourhood for A in the form of four-dimensional tubes cyclically connected to four-dimensional balls. Thus the homology index of A is that of a circle with a disjoint distinguished point. By CM3 we must be able to recover this homology from the chain complex $(V(A), \Delta(A))$, where $V(A)$ is the direct sum of the indices of the six restpoints in A and $\Delta(A)$ is the lower right 6×6 submatrix of Δ . By arguments very similar to those used in the first example it is found that Δ_1 must have rank 2 and that all of the *s are actually 1s. This could have been anticipated because of the existence of the connecting orbits between the restpoints. In fact, it is possible to compute these entries directly from the exact sequences of appropriate triples.

We now turn to Δ_2 which contains the entries relevant to the study of connections from E' and E'' to the C_j . It is clear that

$$\Delta_2 = \begin{bmatrix} 0 & 0 & * & * & * & 0 & 0 & 0 \\ 0 & 0 & * & * & * & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}.$$

CM1 requires that the matrix product $\Delta_2 \Delta_1 = 0$. Using the information about Δ_1 obtained above shows that the entries in each row of Δ_2 sum pairwise to 0 mod 2 and so we have $\Delta(E, C_1) = \Delta(E, C_2) = \Delta(E, C_3)$ and $\Delta(D, C_1) = \Delta(D, C_2) = \Delta(D, C_3)$. By CM3 the homology of the chain complex (V, Δ) is isomorphic to that of the isolated invariant set S consisting of all eight restpoints and their unstable manifolds.

But S has an attractor neighbourhood X_0 (see the end of § 1) homotopy equivalent to S^2 . Thus the homology index of S is G in degrees 0 and 2 and 0 otherwise. On the other hand, the artificial homology group $H\Delta_2 = \ker \Delta_2$. Therefore Δ_2 must have rank 1. This means that one or both of the unknown rows consist of 1s. We will see that all of the unknown entries are 1s.

Recall the attractor-repeller pair (E, A) , where $E = \{E', E''\}$ and A is the attractor described above. We computed the connection map $\delta(E, A)$ as an example earlier. We had $H_2(E) = G \oplus G$ and $H_1(A) = G$. With respect to certain bases for these spaces, $\delta(E, A)$ was represented by the matrix $(1, -1)$. With $G = Z_2$ this becomes $(1, 1)$. Now in the direct sum representation of $H_2(E)$ the factors can be taken as $H_2(E')$ and $H_2(E'')$. Since E' and E'' are an adjacent pair of intervals, CM4 provides a commutative diagram

$$\begin{array}{ccccccc}
 \rightarrow & H\Delta(E') & \rightarrow & H\Delta(E) & \rightarrow & H\Delta(E'') & \rightarrow \\
 & \downarrow \text{id} & & \downarrow \Phi(E) & & \downarrow \text{id} & \\
 \rightarrow & H(E') & \rightarrow & H(E) & \rightarrow & H(E'') & \rightarrow
 \end{array}$$

The bottom row is the direct sum representation of $H(E)$ and it follows that the top row is a similar representation for $H\Delta(E)$. A portion of the analogous diagram for the adjacent pair (E, A) is

$$\begin{array}{ccccccc}
 \rightarrow & 0 & \rightarrow & H\Delta_2(S) & \rightarrow & H\Delta_2(E) & \xrightarrow{\bar{\Delta}(E,A)} & H\Delta_1(A) & \rightarrow \\
 & \downarrow & & \downarrow \Phi(S) & & \downarrow \Phi(E) & & \downarrow \Phi(A) & \\
 \rightarrow & 0 & \rightarrow & H_2(S) & \rightarrow & H_2(E) & \xrightarrow{\delta(E,A)} & H_1(A) & \rightarrow
 \end{array}$$

We know that $\delta(E, A)$ restricts to an isomorphism on each summand. Since $\Phi(E)$ respects the direct sum decomposition, the same must be true for $\bar{\Delta}(E, A)$. This latter map is the map on homology induced by the upper right 2×6 submatrix of Δ_2 . Even without looking into the details of how $\bar{\Delta}(E, A)$ is induced, it is clear that its restrictions to the summands of $H\Delta_2(E)$ are determined by the individual rows of this submatrix. If these are to be non-trivial, neither row can vanish. Therefore all of the *s in Δ_2 are indeed 1s.

Since E' and E'' are adjacent to the C_j in the partial order, the corresponding entries Δ are connection maps. We have seen that a non-trivial connection map implies the existence of a connecting orbit, so we can conclude that there are connecting orbits from each of the restpoints E' and E'' to each of the restpoints C_j ; $j = 1, 2, 3$. Because of the stability of connection matrices, we can even conclude that these connecting orbits are in some sense transverse. One cannot find these connections using Morse theory in the usual way. Once E' and E'' have each been connected to one of the restpoints C_j , they are cancelled homologically and no further connections can be detected.

These examples illustrate the way in which the connection matrix exploits the partial order structure to reveal the rich network of connections between the Morse sets. The connection matrix thrives on complexity.

REFERENCES

- [1] R. Bowen. *On Axiom A Diffeomorphisms*. CBMS Regional Conference Series 35. AMS, Providence, RI (1978).
- [2] C. Conley. *Isolated Invariant Sets and the Morse Index*. CBMS Regional Conference Series 38. AMS, Providence, RI (1978).
- [3] C. Conley. The gradient structure of a flow: I. *Ergod. Th. & Dynam. Sys.* 8* (1988), 11–26.
- [4] C. Conley & R. Easton. Isolated invariant sets and isolating blocks. *Trans. AMS* 158 (1971), 35–61.
- [5] C. Conley & E. Zehnder. Morse type index theory for flows and periodic solutions of Hamiltonian systems. *Commun. Pure Appl. Math.* XXXVII (1984), 207–253.
- [6] J. Franks. *Homology and Dynamical Systems*. CBMS Regional Conference Series 49. AMS, Providence, RI (1982).
- [7] R. Franzosa. Index filtrations and the homology index braid for partially ordered Morse decompositions. *Trans. AMS* 298 (1986), 193–213.
- [8] R. Franzosa. The connection matrix theory for Morse decompositions. *Preprint*. University of Maine (1986).
- [9] H. Kurland. Homotopy invariants of repeller–attractor pairs, I. *J. Differential Equations* 46 (1982), 1–31; II. *J. Differential Equations* 49 (1983), 281–329.
- [10] J. Milnor. *Morse Theory*. Annals of Mathematics Studies 51. Princeton University Press, Princeton, NJ.
- [11] K. Mischaikov. Classification of traveling wave solutions of reaction–diffusion equations. *Report* 86-5. Lefschetz Center for Dynamical Systems (1985).
- [12] M. Morse. *The Calculus of Variations in the Large*. AMS, New York (1934).
- [13] J. Reineck. The connection matrix and the classification of flows arising from ecological models. *Thesis*. University of Wisconsin–Madison (1985).
- [14] D. Salamon. Connected simple systems and the Conley index of isolated invariant sets. *Trans. AMS* 291 (1985), 1–41.