Equivalence of Topological and Singular Transition Matrices in the Conley Index Theory

CHRISTOPHER K. MCCORD
& KONSTANTIN MISCHAIKOW*

1. Introduction

Recall that the basic objects of study in the Conley index theory are isolating neighborhoods and their associated isolated invariant sets. To be more precise, let $\phi\colon \mathbb{R}\times X \to X$ be a continuous flow on a locally compact Hausdorff space. A compact set N is an *isolating neighborhood* if

$$\operatorname{Inv}(N,\phi) := \{x\in N \mid \phi(\mathbb{R}R,x)\subset N\}\subset \operatorname{int} N.$$

In this case the maximal invariant set $S := \operatorname{Inv}(N,\phi)$ is called an *isolated invariant set*. The importance of isolating neighborhoods comes from the fact that they are robust under perturbation. In other words, if N is an isolating neighborhood for the flow ϕ, then N is an isolating neighborhood for all nearby flows in the C^0 topology. However, it must be kept in mind that from the point of view of dynamics it is the invariant set S which is of interest. Obviously, if the dynamics on the set S has some complicated internal structure, then we want to be able to decompose that structure into simpler pieces. It is natural to expect that these simpler pieces should also be isolated invariant sets and that the decomposition should, like the isolating neighborhood, possess the feature of being robust with respect to perturbation. This led Conley to propose the following definition. A *Morse decomposition* of S is a collection of mutually disjoint compact invariant subsets of S,

$$\mathcal{M}(S) := \{M(p) \mid p\in P\},$$

indexed by a finite set P, on which it is possible to impose a partial order $<$ such that, if $x\in S\setminus\bigcup_{p\in P}M(p)$, then there exists $p<q$ such that $\alpha(x)\subset M(q)$ and $\omega(x)\subset M(p)$. The individual sets $M(p)$ are called *Morse sets* and the collection of orbits $S\setminus\bigcup_{p\in P}M(p)$ are referred to as the *connecting orbits* of the Morse decomposition. The set of connecting orbits with $\omega(x)\in M(p)$ and $\alpha(x)\in M(q)$ is denoted $C(M(q),M(p))$. Thus, given a Morse

Received August 15, 1994. Revision received December 16, 1994.
*This work was supported in part by the National Science Foundation.
Michigan Math. J. 42 (1995).

decomposition of S, every element of S can be assigned uniquely either to a Morse set $M(p)$ or to a connecting orbit set $C(M(q), M(p))$. Any partial order on $\mathcal{P}$ that satisfies the aforementioned properties is an *admissible order.* Given a Morse decomposition of S, because of the existence of a partial order, it is easy to see that any recurrent dynamics in S must be contained within the Morse sets. Similarly, the dynamics off the Morse sets must be gradient-like. For this reason we refer to the dynamics within the Morse sets as local dynamics (though, of course, it may occur over a large region in phase space) and the dynamics off the Morse sets as global dynamics.

Our interest begins with the question: Given a Morse decomposition, is it possible to re-assemble the dynamics in order to understand the flow on all of S? To obtain even a partial answer requires fairly sophisticated tools, in particular the Conley index and Conley's connection matrix. These topics will be discussed in greater detail in Section 2. For the moment, recall that the (homological) Conley index assigns to each isolated invariant set S a graded module $CH_*(S)$. Since we shall use field coefficients throughout this paper, these modules are, in fact, vector spaces. To each Morse decomposition $\mathfrak{M}(S) = \{M(p) \mid p \in (P, <)\}$, there is associated a family of connection matrices $\mathcal{C}M(\mathfrak{M}, <)$ which are linear maps

$$\Delta\colon \bigoplus_{p\in P} CH_*(M(p)) \to \bigoplus_{p\in P} CH_*(M(p)).$$

Nonzero entries of these matrices imply the existence of connecting orbits, hence connection matrices provide information concerning global dynamics.

The power of the Conley index theory is best realized in the context of continuous families of dynamical systems. Let Λ be a locally contractible, locally arcwise connected space that continuously parameterizes a family of flows ϕ_λ. This is equivalent to the existence of a continuous parameterized flow

$$\phi\colon X\times\Lambda\times\mathbb{R}\to X\times\Lambda$$

defined by $\phi(x,\lambda,t) = (\phi_\lambda(x,t),\lambda)$. We speak of ϕ_λ as the *λ-flow* on X, and denote X with the λ-flow as X_λ.

For any $U\subset\Lambda$, there is a flow ϕ_U on $X\times U$ obtained by restricting ϕ, so we can consider isolated invariant sets in that flow. If S is an isolated invariant set in $X\times U$, then every $S_\lambda = S\cap X_\lambda$ is isolated in X_λ. We then say that the family of isolated invariant sets $\{S_\lambda\}$ is *related by continuation over U* or *continues over U.* If S is an isolated invariant set that continues over $U\subset\Lambda$, then we will write $S_U := \bigcup_{\lambda\in U} S_\lambda$.

The two basic properties of continuation are:

(1) If S_λ is isolated in X_λ, then there is a neighborhood of λ in Λ such that, for every μ in that neighborhood, there is an S_μ related by continuation to S_λ. Of course, that S_μ could be empty.
(2) If S_λ and S_μ are related by continuation, then they have the same Conley index.

Together, these properties imply that the index remains constant for sufficiently small changes in parameters.

Similarly, we say that a Morse decomposition *continues over* U if there is an isolated invariant set S in $X \times U$ with a Morse decomposition $\{M(p) \mid p \in P\}$. This definition of continuation is stronger than the standard definition of continuation of Morse decompositions. Typically, one only assumes that for each $\lambda \in U$, the collection $\{M_\lambda(p) \mid p \in P\}$ is a Morse decomposition. (Observe that this does not imply that there exists a partial order $<$ which is valid over all of U, that is, that there is an admissible order for the decomposition $\{M(p) \mid p \in P\}$ on the level of the parameterized flow.) However, throughout this paper we need to use the stronger concept of continuation and hence, for ease of expression, we shall use our stronger definition.

It is worth making two observations at this point. First, that the Conley index is robust leads to a similar result for connection matrices. Thus, in some sense, connection matrices detect connecting orbits that persist over open sets in parameter space. This, in turn, suggests that they should be reasonably easy to compute; in fact, numerical techniques that rigorously perform these computations are currently being developed [11; 2]. The second observation is that if a Morse decomposition continues over Λ but the sets of connection matrices $\mathcal{C}M(\mathcal{M}, <_{\lambda_0}, \lambda_0)$ and $\mathcal{C}M(\mathcal{M}, <_{\lambda_1}, \lambda_1)$ at different parameter values differ, then some global bifurcation must occur at some set of parameter values separating λ_0 from λ_1.

These comments are meant to suggest that Morse decompositions and connection matrices provide a framework within which one can detect global bifurcations. What is needed is a technique for interpreting the implications of differences in connection matrices. This is the purpose of *transition matrices.* There are currently three different constructions of transition matrices, referred to as singular, topological, and algebraic [7]. The purpose of this paper is to relate the first two. As we shall see, this is not a purely academic quest. Singular and topological transition matrices are constructed in fundamentally different manners, and come equipped with different properties. In particular, questions of dynamics are more easily approached from the framework of the singular theory, while the algebraic properties of the topological matrices are transparent. The results of this paper allow one to conclude, for the first time, that in some instances the algebraic properties of the topological matrices can be applied to singular matrices. To clarify these comments, we shall briefly review these constructions. Again, these issues will be taken up in considerable detail in the later sections of this paper.

The idea of constructing singular transition matrices was due to Conley and carried out by Reineck [14]. It is most easily described in the context of a parameterized family of ordinary differential equations defined on $\mathbb{R}^n$,

$$\dot{x} = f(x, \lambda),$$

where the parameter space Λ is the real line $\mathbb{R}$. Let ϕ_λ denote the flow generated by this equation. Assume that the Morse decomposition $\mathcal{M}(S_\lambda) =$

$\{M_\lambda(p) \mid p \in P\}$ continues over $\mathbb{R}$ and that connection matrices Δ_0 and Δ_1 for the Morse decompositions $\mathcal{M}(S_0)$ and $\mathcal{M}(S_1)$, respectively, are known. Also, for simplicity assume that there exists a set $N \subset \mathbb{R}^n$ that acts as an isolating neighborhood for S_λ for all $\lambda \in \mathbb{R}$.

Now consider the following family of differential equations,

$$\begin{aligned} \dot{x} &= f(x,\lambda), \\ \dot{\lambda} &= \epsilon\lambda(\lambda-1), \end{aligned} \tag{1.1}$$

where now the parameter is $\epsilon \geq 0$, and observe that as $\epsilon \to 0$ the flow ψ_ϵ generated by this system converges to the parameterized flow ϕ for ϕ_λ. Define

$$\begin{aligned} M(p^+) &:= M_1(p), \\ M(p^-) &:= M_0(p). \end{aligned}$$

For $\epsilon > 0$ sufficiently small, $N \times [-2,2]$ is an isolating neighborhood for the flow generated by (1.1). Let $S_\epsilon := \operatorname{Inv}(N \times [-2,2], \psi_\epsilon)$. Now observe that since $\dot{\lambda} < 0$, if $\lambda \in (0,1)$ and $\epsilon > 0$ then

$$\mathcal{M}(S_\epsilon) = \{M(p^\pm) \mid p \in \mathcal{P}\}$$

is a Morse decomposition, and furthermore there is an admissible ordering given by

$$\begin{aligned} &q^- < p^+, \\ &q^- < p^- \Leftrightarrow q <_0 p, \\ &q^+ < p^+ \Leftrightarrow q <_1 p, \end{aligned}$$

where $<_0$ and $<_1$ are admissible orders for $\mathcal{M}(S_0)$ and $\mathcal{M}(S_1)$, respectively. Let Δ_ϵ denote a connection matrix for $\mathcal{M}(S_\epsilon)$. Then, since the dynamics on the subspaces $\mathbb{R}^n_0$ and $\mathbb{R}^n_1$ are given exactly by the flows generated by $\dot{x} = f(x,0)$ and $\dot{x} = f(x,1)$, it is not surprising that

$$\Delta\colon \bigoplus_{p\in\mathcal{P}} CH_*(M(p^-)) \bigoplus_{p\in\mathcal{P}} CH_*(M(p^+)) \to \bigoplus_{p\in\mathcal{P}} CH_*(M(p^-)) \bigoplus_{p\in\mathcal{P}} CH_*(M(p^+))$$

takes the form

$$\Delta_\epsilon = \begin{bmatrix} \Delta_0 & T_\epsilon \\ 0 & \Delta_1 \end{bmatrix}, \tag{1.2}$$

where

$$T_\epsilon\colon \bigoplus_{p\in\mathcal{P}} CH_*(M(p^+)) \to \bigoplus_{p\in\mathcal{P}} CH_*(M(p^-)).$$

Two comments need to be made at this point. The first is that the entries Δ_0 and Δ_1 in Δ_ϵ cannot be the connection matrices Δ_0 and Δ_1 for $\mathcal{M}(S_0)$ and $\mathcal{M}(S_1)$ since they are defined on different spaces. One of the contributions of [14] was to formalize the expression in (1.2). Second, it is possible to make sense of the limit of $T(\epsilon)$ as $\epsilon \to 0$. The resulting matrices are referred to as *singular transition matrices* and the set of singular transition matrices is denoted by $\mathfrak{T}_{0,1}^{\text{sing}}$. In [14] it is shown that a nonzero element of a singular transition matrix implies the existence of connecting orbits between appropriate Morse sets for some parameter value $\lambda \in (0,1)$. Thus, these transition matrices can be used to detect global bifurcations.

This form of transition matrix is by far the most general. In fact, one need only assume the existence of Morse decompositions at different parameter values of isolated invariant sets that are related by continuation. There are two ways to compute these matrices: they can be computed via the dynamics of the slow system (1.1), or as indicated in [12]. However, in such a setting the algebraic properties of T_ϵ for ϵ small are a complete mystery. For example, it seems reasonable to ask whether $T_\epsilon = T_{-\epsilon}^{-1}$. However, one quickly realizes that the spaces on which T_ϵ and $T_{-\epsilon}$ are defined are different, and hence even expressing the question poses some difficulty.

In [10], the authors considered a problem where it was natural to consider the composition of transition matrices. Again, this operation makes no sense in the setting of singular transition matrices, so it was necessary to develop a new matrix,

$$T_{\lambda_0,\lambda_1}^{\text{top}} : \bigoplus_{p \in \mathcal{P}} M_{\lambda_1}(p) \to \bigoplus_{p \in \mathcal{P}} M_{\lambda_0}(p),$$

called a *topological transition matrix*. Using the same setting that was used to discuss the singular transition matrix, the idea behind topological transition matrices can be described as follows. Since each Morse set $M(p)$ continues over $\Lambda = \mathbb{R}$, there exist isomorphisms

$$F_{0,1}(p) : CH_*(M_1(p)) \to CH_*(M_0(p)).$$

Similarly, since S continues over $\mathbb{R}$, there is an isomorphism

$$F_{0,1} : CH_*(S_1) \to CH_*(S_0).$$

An elementary fact from the Conley index theory states that if

$$S_\lambda = \bigcup_{p \in P} M_\lambda(p)$$

(i.e., if the set of connecting orbits is empty), then there exists an isomorphism

$$\Phi_\lambda : \bigoplus_{p \in P} CH_*(M_\lambda(p)) \to CH_*(S_\lambda).$$

Thus, if there are no connections at either λ_0 or λ_1, the following diagram can be constructed:

$$\begin{array}{ccc} \bigoplus_{p \in P} CH_*(M_1(p)) & \xrightarrow{\bigotimes_{p \in P} F_{0,1}(p)} & \bigoplus_{p \in P} CH_*(M_0(p)) \\ \downarrow \Phi_1 & & \downarrow \Phi_0 \\ CH_*(S_1) & \xrightarrow{F_{0,1}} & CH_*(S_0). \end{array}$$

This diagram is not, in general, commutative, and it is precisely its failure to commute that gives information about connecting orbits. For the purpose of applications it is useful to be able to express this last statement in the form of a matrix. Of course, this requires choices of basis on each of the spaces. We do this as follows. Let $\mathcal{B}_1$ be a basis for $\bigoplus_{p\in P} CH_*(M_1(p))$, and let

$$\mathcal{B}_0 = \bigoplus_{p\in P} F_{0,1}(p)(\mathcal{B}_1)$$

be a basis for $\bigoplus_{p\in P} CH_*(M_0(p))$. Using these bases, define the topological transition matrix by

$$T_{0,1}^{\text{top}} = \Phi_0 \circ \bigoplus_{p\in P} F_{0,1}(p) \circ \Phi_1^{-1}.$$

Observe that this construction only makes sense when the Morse decomposition continues over the entire parameter space, and when there are no connecting orbits with respect to the Morse decomposition at the parameter values λ_0 and λ_1.

In this setting it is perfectly clear what is meant by the composition of two transition matrices. It is equally obvious that

$$T_{0,1}^{\text{top}} = (T_{1,0}^{\text{top}})^{-1}.$$

Furthermore, $T_{0,1}^{\text{top}}$ shares many properties with elements of $\mathfrak{J}_{0,1}^{\text{sing}}$. In particular, nonzero off-diagonal entries imply the existence of connecting orbits for some parameter values between 0 and 1.

The content of this paper is the proof that when both singular and topological transition matrices are defined, they are "equal". Equal needs to be put in quotation marks because the maps are defined on different vector spaces. However, in Section 7 we shall exhibit a canonical isomorphism between these spaces that justifies our claim.

Having presented this sketch of these two transition matrices, we are now in a position to justify our interest in showing that they are equal. In [8], H. Kokubu, H. Oka, and the second author considered a set of equations that can be related to singular boundary-value problems of the form

$$\dot{x} = f(x,\lambda),$$
$$\dot{\lambda} = \epsilon g(\lambda),$$

where $g(-1) = g(1) = 0$ and $g(\lambda) > 0$ for $\lambda \in (-1,1)$. The problem was to prove the existence of connecting orbits from a critical point $(x_0,-1)$ to a critical point $(x_0,1)$ for small $\epsilon > 0$. In this particular problem, the singular transition matrix T_-^{sing} arising from $\epsilon < 0$ is easy to compute. However, it does not appear possible to compute directly T_+^{sing}, the singular transition matrix arising from $\epsilon > 0$. Fortunately, in this problem the topological transition matrix is defined. Thus we have the following sequence of identities that determines T_+^{sing} and allows one to conclude the existence of the desired connecting orbits:

$$\begin{aligned} T_+^{\text{sing}} &= T_{1,-1}^{\text{top}} \\ &= (T_{-1,1}^{\text{top}})^{-1} \\ &= (T_-^{\text{sing}})^{-1}. \end{aligned}$$

The reader is referred to [8] for the details.

This paper is organized as follows. We begin in Section 2 with a review of the relevant aspects of the Conley index theory. While it is assumed that the reader is familiar with the index, our intention is to provide sufficient background to make the paper accessible to readers unfamiliar with the technical aspects of Conley's connection matrix. In Section 3 we define singular connection matrices and in Section 4 show how they apply in the context of parameterized families of flows. Finally, in Section 5 we obtain singular transition matrices. In Section 6 topological transition matrices are reviewed. The mathematics in Sections 3–5 is technically not new; what is different is the framework in which the continuation theory for connection matrices is presented. We construct an abstract framework for Conley's ideas and Reineck's results. However, this is not *abstractus gratia abstractum.* The constructions of transition matrices contain some ambiguities, which we needed to remove before we could understand the relationship between singular and topological transition matrices. The framework developed in Sections 3–5 gives us the required precision. We then outline the construction of topological transition matrices in Section 6, and conclude by proving the equivalence between topological and singular transition matrices in Section 7.

2. Connection Matrices and Continuation

We assume that the reader is familiar with the basics of the Conley index theory for flows, as developed in [1; 15; 16], and with Conley's connection matrix (see [3; 4; 5; 6]). However, to fix notation and terminology, we briefly review the relevant aspects of the theory: Morse decompositions, connection matrices, and continuation.

2.1. Morse Decompositions

The simplest example of a Morse decomposition is an *attractor–repeller decomposition.* This is simply a Morse decomposition with two Morse sets, an attractor A and a repeller R. The sets A and R are isolated invariant sets in their own right, and S decomposes as the disjoint union $S = R \sqcup C(R, A) \sqcup A$. Attractor–repeller decompositions are fairly easy to generate. If the flow on S is not chain-recurrent, then there will be a proper nonempty subset A that is an attractor in S (in the sense of Conley). Corresponding to every attractor in A is its dual repeller A^*. The pairs (A, A^*) are attractor–repeller pairs for S.

In an attractor–repeller decomposition, the index set P is simply a two-element set (without loss of generality, $P = \{0, 1\}$) with partial order $0 < 1$. However, when a more general Morse decomposition is considered, some added complications surrounding the partial order are introduced. With a larger index set, the number of different partial orders that can be assigned to it grows. The two problems are to decide which partial orders on P are acceptable for a given Morse decomposition, and to find the appropriate mechanism for collating the information generated by the partial order.

Given a Morse decomposition $\mathcal{M}$ indexed by P, there is an intrinsic partial order, called the *flow-defined order,* that can be associated with the decomposition. This order is generated by first declaring $p <_f q$ if $C(M(q), M(p)) \neq \emptyset$, then taking the transitive closure of this relation. In this order, $p <_f q$ if and only if there are elements $p = p_0, p_1, \ldots, p_n = q$ in P such that

$$C(M(p_i), M(p_{i-1})) \neq \emptyset \quad \text{for all } i = 1, \ldots, n.$$

While very natural to define, this partial order has one obvious disadvantage: to identify completely this partial order, we must know exactly which connecting orbit sets are nonempty. But this sort of dynamic information is usually the goal, not the starting point, of our study of Morse decomposition. In other words, if our goal is to determine the structure of the connecting orbit set, we cannot start with tools that require knowledge of the flow-defined order.

Although all of the connection matrix machinery depends on the existence of a partial order, it has two features we can use to our advantage:

(1) Much of the machinery requires only that a partial order *exist;* it does not require the partial order to be explicitly given.
(2) The machinery can be applied equally well with an *admissible* order.

An admissible order is any refinement of the flow-defined order: an order $<$ is admissible if $p <_f q$ implies $p < q$. That is, to establish that an order is admissible, we need only to establish that if $p \not< q$ then there exist no connecting orbits (or chains of connecting orbits) from $M(q)$ to $M(p)$. In practice, this can often be established by means of Lyapunov functions. Once an admissible order is found, the goal is to extract the flow-defined order from it. Much of the connection matrix theory can be viewed as machinery for doing so.

Once an admissible partial order $<$ has been fixed, a great deal of structure for the Morse decomposition ensues. Given a partially ordered set $(P, <)$, we define an *interval* in P to be a subset $I \subset P$ that is closed under the partial order: if $p, r \in I$ and $p < q < r$, then $q \in I$. An interval is *attracting* if $q \in I$ with $p < q$ implies $p \in I$; *repelling* if $p \in I$ with $p < q$ implies $q \in I$. Disjoint intervals I, J are *adjacent* if $IJ := I \cup J$ is also an interval.

Given an interval $I \subset P$, define the Morse interval

$$M(I) = \left(\bigcup_{p \in I} M(p) \right) \cup \left(\bigcup_{p,q \in I} C(M(q), M(p)) \right).$$

That is, $M(I)$ is the union of all of the Morse sets in I and all of the connecting orbits between Morse sets in I. The significance of this construction is that it nicely captures the relationship between the structure of the partial order and dynamics on S as follows.

PROPOSITION 2.1. *If $\mathcal{M}$ is a $(P, <)$-Morse decomposition of S and if I, J are an adjacent pair of intervals in P, then:*

(1) *$M(I)$ is an isolated invariant set; and*
(2) *if $J \not< I$ (i.e. if $j \not< i$ for every $i \in I$ and $j \in J$), then $(M(I), M(J))$ is an attractor–repeller pair in $M(IJ)$.*

Thus, knowing an admissible partial order tells us all of the possible attractor–repeller decompositions in S. To further develop that knowledge, we turn to connection matrices.

2.2. Connection Matrices

Once we have a Morse decomposition $\mathcal{M}$ of an isolated invariant set S, we have a Conley index $CH_*(M(I))$ defined for every interval $I \subset P$. We now look at how those indices are interrelated, and what those relations reveal about the flow on S.

In the simplest case, an attractor–repeller decomposition, there are three isolated invariant sets: S, A, and R. The Conley indices of these three sets are related by a long exact sequence, called the *attractor–repeller sequence:*

$$\cdots \longrightarrow CH_p(A) \longrightarrow CH_p(S) \longrightarrow CH_p(R) \xrightarrow{\partial} CH_{p-1}(A) \longrightarrow \cdots.$$

The map ∂ is called the *connection homomorphism.* Its basic property is that if $\partial \neq 0$ then there exist connecting orbits from R to A in S. In some cases, it can give more refined information about the set of connecting orbits. For example, if A and R are hyperbolic fixed points of indices p and $p-1$ respectively, then the only nontrivial portion of the attractor–repeller sequence is

$$0 \longrightarrow CH_{p+1}(S) \longrightarrow \mathbb{Z} \xrightarrow{\partial} \mathbb{Z} \longrightarrow CH_p(S) \longrightarrow 0.$$

Then ∂ can be thought of as an integer (without loss, a nonnegative integer). If the flow has the additional property that $W^s(A)$ and $W^u(R)$ intersect transversely, then the connecting orbit set consists of a disjoint set of orbits. In this setting, there are at least ∂ connecting orbits, and the number of connecting orbits is equal to $\partial \bmod 2$ [9].

In general, if $\mathcal{M}$ is a Morse decomposition with an admissible order $(P, <)$, then there is an attractor–repeller sequence for every adjacent pair of intervals in P. In [3; 4; 5], Franzosa introduced *connection matrices* as devices for simultaneously encoding the information expressed in all of these sequences. In brief, connection matrices are matrices defined on the sum of the homology indices of the Morse sets; when treated as boundary maps, these matrices allow all of the attractor–repeller sequences to be reconstructed.

More precisely, for every interval $I \subset P$, let $C_*\Delta(I) = \bigoplus_{p \in I} CH_*(M(p))$. Suppose that $\Delta(P): C_*\Delta(P) \to C_*\Delta(P)$ is a degree-1 endomorphism such that

(1) $\Delta(P)^2 = 0$; and
(2) if $p \not< q$, then $\Delta(p,q): CH_*(M(q)) \to CH_*(M(p))$ is zero.

Such a matrix is said to be an *upper triangular boundary map.* Given any two intervals $I, J \subset P$, define $\Delta(I, J): C_*\Delta(J) \to C_*\Delta(I)$ to be the obvious restriction of $\Delta(P)$ and denote $\Delta(I, I)$ by $\Delta(I)$. Then the two conditions on $\Delta(P)$ are inherited by $\Delta(I)$. In particular, given an adjacent pair of intervals I, J in P, there is a commutative diagram

$$\begin{array}{ccccccccc} 0 & \longrightarrow & C_*\Delta(I) & \xrightarrow{i} & C_*\Delta(IJ) & \xrightarrow{p} & C_*\Delta(J) & \longrightarrow & 0 \\ & & \downarrow{\scriptstyle \Delta(I)} & & \downarrow{\scriptstyle \Delta(IJ)} & & \downarrow{\scriptstyle \Delta(J)} & & \\ 0 & \longrightarrow & C_*\Delta(I) & \xrightarrow{i} & C_*\Delta(IJ) & \xrightarrow{p} & C_*\Delta(J) & \longrightarrow & 0, \end{array}$$

where i and p are (respectively) the inclusion and projection homomorphisms. This can be interpreted as a short exact sequence of chain complexes, with the matrices Δ acting as boundary homomorphisms. If the homology of the complex $\{C_*\Delta(I), \Delta(I)\}$ is denoted $H_*\Delta(I)$, then the preceding diagram produces a long exact sequence

$$\longrightarrow H_k\Delta(I) \xrightarrow{i_*} H_k\Delta(IJ) \xrightarrow{p_*} H_k\Delta(J) \xrightarrow{[\Delta(J,I)]} H_{k-1}(I) \longrightarrow .$$

Hence, an upper triangular boundary map produces a long exact sequence for every adjacent pair of intervals; $\Delta(P)$ is a connection matrix if all of these sequences are canonically isomorphic to the attractor-repeller sequences. That is, we require for every interval I that there be an isomorphism $\phi(I)$: $H_*\Delta(I) \to CH_*(M(I))$ such that $\phi(p) = \text{id}$ for every $p \in P$, and for every adjacent pair of intervals I, J that there be a commutative diagram

$$\begin{array}{ccccccccccc} \longrightarrow & H_k\Delta(I) & \xrightarrow{i_*} & H_k\Delta(IJ) & \xrightarrow{p_*} & H_k\Delta(J) & \xrightarrow{\partial} & H_{k-1}(I) & \longrightarrow \\ & \downarrow & & \downarrow & & \downarrow & & \downarrow & \\ \longrightarrow & CH_k(M(I)) & \xrightarrow{i_*} & CH_k(M(IJ)) & \xrightarrow{p_*} & H_k(M(J)) & \xrightarrow{\partial} & H_{k-1}(M(I)) & \longrightarrow . \end{array} \tag{2.1}$$

Connection matrices exist for all Morse decompositions, but are not necessarily unique. In general, the set of connection matrices depends on the indices of the Morse sets and the partial order on P. We therefore use the notation $\mathcal{C}M(\mathcal{M}, <)$ to denote the set of connection matrices for a given partial order. If the partial order is the flow-defined order, we write simply $\mathcal{C}M(\mathcal{M})$.

As noted in the introduction, the significance of connection matrices is that nonzero entries detect connecting orbits: if $\Delta(p, q) \neq 0$ for some $\Delta \in \mathcal{C}M(\mathcal{M}, <)$, then $p < q$. In particular, if the partial order is the flow-defined order $<_f$, then $p <_f q$ and there is a sequence of connecting orbits connecting M_q to M_p. Similarly, if $\Delta(p, r)\Delta(r, q) \neq 0$ then $p < r < q$. If $< = <_f$, then there is a sequence of connecting orbits from M_q to M_p whose closure intersects M_r.

2.3. Parameterized Families of Flows

We want to make more precise the continuation invariance of the index referred to in Section 1. Suppose S_0 and S_1 are isolated invariant sets in X_{λ_0} and X_{λ_1} respectively, and are related by continuation. Then there is a path ω in Λ from λ_0 to λ_1; that is, there exist both a continuous function $\omega\colon [0, 1] \to \Lambda$ such that $\omega(0)=\lambda_0$ and $\omega(1)=\lambda_1$, and an isolated invariant set S over $\omega(I)$ such that $S_{\lambda_i}=S_i$. Then the inclusion $f_i: X_{\lambda_i}\to X\times\omega(I)$ induces an isomorphism $CH_*(S_i)\xrightarrow{f_{i*}} CH_*(S)$, where $CH_*(S_i)$ indicates the index of S_i in X_{λ_i}, and $CH_*(S)$ indicates the index of S in $X\times\omega(I)$. Thus there is an isomorphism

$$F_\omega: CH_*(S_0)\xrightarrow{f_{1*}^{-1}\circ f_{0*}} CH_*(S_1).$$

This isomorphism depends on the endpoint-preserving homotopy class of ω. However, if $\pi_1(\Lambda)=0$ (which we will usually assume), then F_ω is independent of the path ω. In this case we can speak of the unique continuation isomorphism $F_{\lambda_1\lambda_0}$, though we will still write F_ω when we wish to emphasize the path chosen to carry out continuation.

The continuation isomorphism is well-behaved with respect to composition of paths: $F_{\lambda,\lambda}=\mathrm{id}$ and $F_{\lambda,\mu}\circ F_{\mu,\nu}=F_{\lambda,\nu}$. Of course, it follows then that $F_{\lambda,\mu}=F_{\mu,\lambda}^{-1}$.

Similarly, we say that a Morse decomposition continues over U if there is an isolated invariant set S in $X\times U$ with a Morse decomposition $\{M(p)\}_{p\in P}$ such that each $M(p)$ continues over U. It follows then that every Morse interval $M(I)$ also continues over U. If $V\subset U$ and I is an interval in P, then $M_V(I)$ will denote the Morse interval of I over V:

$$M_V(I)=\bigcup_{\lambda\in V} M_\lambda(I).$$

We will denote the continuation isomorphism for a Morse interval $M(I)$ by $F_{\lambda,\mu}(I)$, and will denote the direct sum $\bigoplus_{p\in I}F_{\lambda,\mu}(p): C_*\Delta_\lambda(I)\to C_*\Delta_\mu(I)$ by $F\Delta_{\lambda,\mu}(I)$. (Note: this map was denoted by $E_{\lambda,\mu}(I)$ in [10].)

We now turn to the question of how the flow-defined partial order on P behaves with respect to continuation. If $\mathcal{M}$ is a Morse decomposition that continues over Λ, we want to understand the relation between subsets of Λ and partial orders on the index set P. First, if $U\subset\Lambda$ then $\{M_U(p)\}$ is a Morse decomposition over U. Let $<_U$ denote the flow-defined partial order of this decomposition. That is, $p<_U q$ if and only if there is a sequence $p_0,\dots,p_n$ in P and a sequence $\lambda_1,\dots,\lambda_n$ in U such that $p=p_0$, $q=p_n$ and $p_{i-1}<_{\lambda_i}p_i$ (without loss of generality, there is a connecting orbit from $M_{\lambda_i}(p_i)$ and $M_{\lambda_i}(p_{i-1})$). Thus we can associate to each subset of Λ a partial order on P.

Proposition 2.2. *For every $K\subset\Lambda$, there is an open neighborhood U of K such that $<_U$ equals $<_K$.*

Proof. First note that the statement is true when K is a point [4]. Then, for general K, every point $k\in K$ has a neighborhood U_k with $<_{U_k}$ equal to $<_k$. The set $U=\bigcup_{k\in K}U_k$ is an open neighborhood of K, with $<_U$ the common refinement of the $<_{U_k}$s and $<_K$ the common refinement of the $<_k$s. □

This property is usually referred to as *lower semicontinuity.*

Conversely, we can attempt to associate to every partial order on P a subset of Λ. This will only be meaningful for partial orders that are refined by $<_\Lambda$.

Given such a partial order $<$ on P, let $\Lambda(<)$ be the maximal subset of Λ such that $<$ is the flow-defined order on $\{M_{\Lambda(<)}(p)\}$. Of course, $\Lambda(<)$ could be empty. Another way of formulating lower semicontinuity is as follows.

PROPOSITION 2.3. *For every partial order $<$ that is refined by $<_\Lambda$, $\Lambda(<)$ is an open set.*

We have associated to every subset of Λ a partial order, and to every partial order an open subset. These two processes are related, as the next proposition shows.

PROPOSITION 2.4. *Supose that $\mathfrak{M}$ is a Morse decomposition which continues across Λ, and that $<_\Lambda$ is the minimal partial order on P which continues across Λ. Suppose K, L are subsets of Λ and $<_1$, $<_2$ are partial orders on P which are refined by $<_\Lambda$.*

(1) *If $K \subset L$, then $<_L$ refines $<_K$.*
(2) *If $<_1$ refines $<_2$, then $\Lambda(<_1) \subset \Lambda(<_2)$.*
(3) *$\Lambda(<_K)$ is an open neighborhood of K.*
(4) $<_{\Lambda(<_1)} = <_1$.

We need this precision of language because we are interested in the continuation properties of connection matrices across Λ, and connection matrices are tied to partial orders on P. That is, it only makes sense to ask about continuation of $\mathcal{C}M(\mathfrak{M}, <)$ over $\Lambda(<)$. And, as we have just seen, all matrices in $\mathcal{C}M(\mathfrak{M}, <)$ do continue across $\Lambda(<)$. Now, however, we want to express this continuation property more explicitly.

PROPOSITION 2.5. *If $\mathfrak{M}$ is a Morse decomposition that continues over Λ, and if $<$ is an admissible ordering that continues over $U \subset \Lambda$, then the connection matrices over different points in U are conjugate via the continuation isomorphisms. That is, if $\{\Delta_\mu(P), \phi_\mu(I)\}$ is a $(P, <)$-connection matrix over μ then*

$$\{F\Delta_{\lambda,\mu}(P)\circ\Delta_\mu(P)\circ F\Delta_{\mu,\lambda}(P), F_{\lambda,\mu}(I)\circ\phi_\mu(I)\circ F_{\mu,\lambda}(I)\}$$

is a $(P, <)$-connection matrix over λ.

COROLLARY 2.6. *If (A, R) is an attractor-repeller decomposition that continues over Λ, then the attractor-repeller sequence commutes with the continuation isomorphisms. That is, for any $\lambda, \mu \in \Lambda$, there is a commutative diagram*

$$\begin{array}{ccccccccc}
\longrightarrow & CH_p(A_\lambda) & \longrightarrow & CH_p(S_\lambda) & \longrightarrow & CH_p(R_\lambda) & \xrightarrow{\partial} & H_{p-1}(A_\lambda) & \longrightarrow \\
 & \downarrow F_{\mu\lambda} & & \downarrow F_{\mu\lambda} & & \downarrow F_{\mu\lambda} & & \downarrow F_{\mu\lambda} & \\
\longrightarrow & CH_p(A_\mu) & \longrightarrow & CH_p(S_\mu) & \longrightarrow & CH_p(R_\mu) & \xrightarrow{\partial} & H_{p-1}(A_\mu) & \longrightarrow .
\end{array}$$

This defines in a natural way a function $\mathfrak{F}_{\lambda,\mu}\colon \mathcal{C}M(\mathfrak{M}_\mu, <) \to \mathcal{C}M(\mathfrak{M}_\lambda, <)$, which has the following properties.

PROPOSITION 2.7. *As a function on $(P, <)$-connection matrices:*

(1) $\mathfrak{F}_{\lambda,\lambda} = \mathrm{id}$ *for every* $\lambda \in U$;

(2) $\mathfrak{F}_{\lambda,\nu} = \mathfrak{F}_{\lambda,\mu} \circ \mathfrak{F}_{\mu,\nu}$ *for every* $\lambda, \mu, \nu \in U$; *and*

(3) *for every* $\lambda, \mu, \nu \in U$, $\mathfrak{F}_{\lambda,\mu}$ *is a bijection with inverse* $\mathfrak{F}_{\mu,\lambda}$.

Proof. These properties of $\mathfrak{F}_{\lambda,\mu}$ follow immediately from the analogous properties of $F_{\lambda,\mu}$. □

We will say that $(P, <)$-connection matrices $\{\Delta_\mu(P), \phi_\mu(I)\}$ and $\{\Delta_\lambda(P), \phi_\lambda(I)\}$ are *related by continuation* if

$$\mathfrak{F}_{\lambda,\mu}(\{\Delta_\mu(P), \phi_\mu(I)\}) = \{\Delta_\lambda(P), \phi_\lambda(I)\}$$

or, equivalently, if

$$\Delta_\lambda(P) \circ F\Delta_{\lambda,\mu}(P) = F\Delta_{\lambda,\mu}(P) \circ \Delta_\mu(P)$$

and

$$\phi_\lambda(I) \circ F_{\lambda,\mu}(I) = F_{\lambda,\mu}(I) \circ \phi_\mu(I)$$

for every $I \in \mathfrak{I}(P, <)$. Clearly *relation by continuation* defines an equivalence relation on the collection of $(P, <)$-connection matrices $\bigcup_{\lambda \in U} \mathcal{C}M(\mathfrak{M}_\lambda, <)$. We will denote the set of equivalence classes by $\mathcal{C}M(\mathfrak{M}, <, U)$. If the partial order is the flow-defined order $<_U$, then the corresponding set of connection matrices will be denoted $\mathcal{C}M(\mathfrak{M}, U)$. Note that, since $\pi_i(\Lambda)$ acts trivially on $\mathcal{S}(\Lambda)$, every $F_{\lambda,\lambda} = \mathrm{id}$ and no two elements of $\mathcal{C}M(\mathfrak{M}_\lambda, <)$ are identified in $\mathcal{C}M(\mathfrak{M}, <, U)$. That is, $\mathcal{C}M(\mathfrak{M}, <, U)$ is bijective with each $\mathcal{C}M(\mathfrak{M}_\lambda, <)$.

While in some sense this is just bookkeeping, the notation is suggestive of the point of view we wish to adopt. We want to think of a *single* set of connection matrices which continues over U, and which yields $\mathcal{C}M(\mathfrak{M}_\lambda, <)$ when restricted to $\lambda \in U$. That is, each element of $\mathcal{C}M(\mathfrak{M}, <, U)$ is a connection matrix which, via a canonical change of coordinates (i.e., the continuation isomorphisms), defines a $(P, <)$-connection matrix at every point in U.

We now want to consider comparisons between sets of connection matrices. We consider a fixed Morse decomposition $\mathfrak{M}$, and consider changes in the parameter set U and the partial order $<$. That is, given $U_1, U_2 \subset \Lambda$ and partial orders $<_1, <_2$ such that $<_i$ continues over U_i, how (if at all) are $\mathcal{C}M(\mathfrak{M}, <_1, U_1)$ and $\mathcal{C}M(\mathfrak{M}, <_2, U_2)$ related? Note that both the power set 2^Λ and the set of all possible partial orders on P are themselves partially ordered sets (ordered by inclusion and refinement, respectively).

PROPOSITION 2.8. *If* $U_1 \subset U_2$ *and* $<_2$ *refines* $<_1$, *then there is an injection* $\iota(U_1, <_1; U_2, <_2)\colon \mathcal{C}M(\mathfrak{M}, <_1, U_1) \to \mathcal{C}M(\mathfrak{M}, <_2, U_2)$. *This map is a bijection if* $<_1$ *and* $<_2$ *coincide.*

If $<_i = <_{U_i}$, the flow-defined order on U_i, and if $U_1 \subset U_2$, then $<_2$ refines $<_1$; hence there is an injection $\iota(U_1, <_1; U_2, <_2)$, denoted $\iota(U_1; U_2)$. The collection $\{\mathcal{C}M(\mathfrak{M}, U), \iota(U_1; U_2)\}$ is then a directed system, indexed by 2^Λ.

PROPOSITION 2.9. *For every $K \subset \Lambda$, there is an open neighborhood U such that $\mathcal{C}M(\mathcal{M}, K) = \mathcal{C}M(\mathcal{M}, U)$.*

The neighborhood U is of course $\Lambda(<_K)$.

COROLLARY 2.10. *For every $K \subset \Lambda$, $\mathcal{C}M(\mathcal{M}, K) = \varprojlim \mathcal{C}M(\mathcal{M}, U)$, as U ranges over supersets of K.*

These results are just a reformulation of the well-known lower semicontinuity property of connection matrices. This reformulation was carried out not for its own sake, but in preparation for an extension of the existing results. Up until now, we have considered Morse decompositions that continue across the entire parameter space. We now turn to Morse decompositions that continue across some subset Λ' of a parameter space, and ask what information about the flow at a parameter value in the closure of Λ' can be extracted from the connection matrices defined on Λ'.

3. Singular Connection Matrices

Assume S is an isolated invariant set that continues over Λ, and $\mathcal{M}$ is a Morse decomposition of S that continues over $\Lambda' \subset \Lambda$. In this setting, the most we can hope for is information on the flow over parameter values in $\mathrm{cl}_\Lambda(\Lambda')$. Consequently, we will assume that Λ' is dense in Λ. Since a Morse decomposition defined at any parameter value continues to an open neighborhood of parameter values, we can assume without loss that Λ' is open.

For any $K \subset \Lambda$, we want to define *singular connection matrices* on K by taking neighborhoods U of K in Λ, taking the set of connection matrices $\mathcal{C}M(\mathcal{M}, U')$ defined for $K' = K \cap \Lambda'$, and then taking the inverse limit of this collection as U nests down onto K. While this process is straightforward enough, the important issue is its interpretation. If Δ is a singular connection matrix on K, what does that tell us about the dynamics on K? The interpretation is basically the same as in [14]: the Morse sets limit to (not necessarily isolated) invariant sets, and nonzero entries in singular connection matrices indicate a chain of invariant sets connecting these "limiting Morse sets".

DEFINITION 3.1. Suppose that Λ' is an open dense set in Λ, and that $\mathcal{M}$ is a Morse decomposition which continues over Λ'. Then, for every $K \subset \Lambda$, the set of *singular connection matrices* for $\mathcal{M}$ on K is

$$\mathcal{C}M(\mathcal{M}, K, \Lambda) = \varprojlim \mathcal{C}M(\mathcal{M}, U'),$$

the inverse limit of the directed system $\{\mathcal{C}M(\mathcal{M}, U'), \iota(U'; V') \mid U' = U \cap \Lambda'$ for some neighborhood U of K in $\Lambda\}$.

As the above discussion shows, all that really changes as we take a sequence $\cdots \supset U_n \supset U_{n+1} \supset \cdots$ nesting onto K is the partial order $<_n$ associated to each U_n'. We can then define a partial order $<_K$ to be the maximal partial order on P which is refined by $<_{U'}$ for all neighborhoods U of K.

PROPOSITION 3.2. *If $K' = K \cap \Lambda' \neq \emptyset$, then $\mathcal{C}M(\mathcal{M}, K, \Lambda) = \mathcal{C}M(\mathcal{M}, K')$.*

PROPOSITION 3.3. *For every $K \subset \Lambda$, there exists an open neighborhood U of K in Λ such that $\mathcal{C}M(\mathcal{M}, K, \Lambda) = \mathcal{C}M(\mathcal{M}, U')$, where $U' = U \cap \Lambda'$.*

Clearly, the neighborhood U in the proposition is $U(<_K)$, where $<_K$ is the limiting partial order.

We now turn to the question of interpretation. In the usual (nonsingular) setting, connection matrices provide information about connections between Morse sets. In the singular setting, the Morse sets may not persist as isolated invariant sets. They will, however, limit to closed invariant sets. That is, for $k \in K$, take a sequence λ_n' in Λ' that converges to k. Consider the corresponding sequence of compact subsets $\{M_n(p)\}$ in X. In the Hausdorff topology, this sequence converges to some compact subset $M_k(p)$ in X. Further, by the continuity of the flow, $M_k(p)$ is invariant under the k-flow. However, $M_k(p)$ is not necessarily isolated, and not necessarily nonempty. Further, if we carry out the same process for $M(q)$, the resulting $M_k(q)$ may not be disjoint from $M_k(p)$. However, once the sets $M_k(p)$ are so understood, the interpretation of connection matrix entries remains essentially unchanged.

PROPOSITION 3.4. *If Δ is a singular connection matrix for $\mathcal{M}$ on K, and if $\Delta_{p,q}$ is a nonzero entry between adjacent indices, then for every $\lambda \in K$ there is a nonempty closed connected invariant set $c_\lambda(p,q)$ such that $c_\lambda(p,q) \cap M_\lambda(p)$ and $c_\lambda(p,q) \cap M_\lambda(q)$ are nonempty.*

PROPOSITION 3.5. *If Δ is a singular connection matrix for $\mathcal{M}$ on K, and if $\Delta_{p,q}$ is a nonzero entry, then there exists a sequence $p = p_0, \dots, p_n = q$ in P, a sequence $\lambda_1, \dots, \lambda_n$ in K, and a nonempty closed connected set c_i which is invariant under the λ_i-flow and which has nonempty intersection with $M_{\lambda_i}(p_{i-1})$ and $M_{\lambda_i}(p_i)$.*

4. One-Parameter Families of Flows

The development of singular connection matrices presented in the last section is essentially a generalization of Reineck's construction of transition matrices [14]. We now return to his original setting of one-parameter flows. Here, too, we will generalize Reineck's development. He considered a single one-parameter family of flows—a flow on $X \times [0,1]$. Instead, we will continue to work over an arbitrary parameter space Λ, and obtain one-parameter families by choosing paths in Λ. Most of the results stated in this section are simply the obvious generalization to this new setting of results in [14].

One of the basic ingredients in Reineck's development is the introduction of a "drift flow" on the parameter space. The original one-parameter family is then recovered in the limit as the drift goes to zero. To capture this as a limit in a parameter space, we must create a new parameter space that incorporates the drift flows and the one-parameter families in Λ. To do so, let $\mathcal{P}(\Lambda) = \{\alpha : [0,1] \to \Lambda\}$ be the set of paths in Λ, and let

$$\mathcal{G} = \{g : [-1,2] \to \mathbb{R} \mid g(-1,0) \geq 0, g(0,1) \leq 0, g(1,2) \geq 0, \\ g(i) = 0 \text{ for } i = -1,0,1,2\}.$$

Let $\mathfrak{D} = \mathcal{P}(\Lambda) \times \mathcal{G}$.

We will use $\mathfrak{D}$ as a parameter space for flows on $X \times [-1,2]$. At first glance, there are two things about the construction of this parameter space that might seem unnatural: the paths are parameterized from -1 to 2, rather than from 0 to 1; and there are sign restrictions on the functions in $\mathcal{G}$. Both of these arise for the following reason: we want to examine flows on $X \times [0,1]$ for which the drift flow $\dot{s} = g(s)$ on $[0,1]$ has 0 as a hyperbolic attractor and 1 as a hyperbolic repeller. To do so, we need to extend the flow to a neighborhood of $[0,1]$, and for convenience we choose $[-1,2]$ as the neighborhood. The sign conventions are likewise chosen so that $g = 0$ is included in $\mathcal{G}$, and so that drift flows with $\{0,1\}$ a hyperbolic attractor–repeller pair for $[0,1]$ are dense in $\mathcal{G}$.

Having justified our choices for the parameter space, we now construct the family of dynamical systems on $X \times [-1,2]$ which they parameterize. First, choose a function $\tau : [-1,2] \times [0,1]$ with $\tau_{[0,1]} = \text{id}$. To be concrete, say,

$$\tau(s) = \begin{cases} -s & \text{for } -1 \leq s \leq 0, \\ s & \text{for } 0 \leq s \leq 1, \\ 2-s & \text{for } 1 \leq s \leq 2. \end{cases}$$

Then, if $(\alpha, g) \in \mathfrak{D}$, construct a flow on $X \times [-1,2]$ over (α, g) by

$$\dot{x} = f(x, \alpha\tau(s)),$$
$$\dot{s} = g(s).$$

Note that if α is a constant path $\overline{\lambda_0}$, then we have a product flow

$$\dot{x} = f(x, \lambda_0),$$
$$\dot{s} = g(s),$$

while if $g \equiv 0$, we have the original parameterized family of flows

$$\dot{x} = f(x, \alpha\tau(s))$$

restricted to the image of α.

We need to understand the relationship between continuation over $\mathfrak{D}$ and continuation over Λ. First, note that for every $s \in [0,1]$ the projection ρ_s: $\mathfrak{D} \to \Lambda$, defined by $\rho_s(\alpha, g) = \alpha(s)$, is a homotopy equivalence. In particular, if Λ is simply-connected, so is $\mathfrak{D}$. More generally, if continuation across Λ is path-independent, then the same is true for continuation over $\mathfrak{D}$. Before making further comparisons between continuations over $\mathfrak{D}$ and Λ, we must understand what are the possible isolated invariant sets over a point (α, g) in $\mathfrak{D}$. Even in the comparatively simple structure we are dealing with, this is not an easy question to answer. For a fixed (α, g), the flow on $X\times[-1,2]$ fibers over the g-flow on $[-1,2]$. Thus, invariant sets in $X\times[-1,2]$ must project to invariant sets in $[-1,2]$. However, simple examples show that other, equally desirable correspondences do not hold:

(1) The projection $\pi: X\times[-1,2]\to[-1,2]$ does not preserve isolation; if S is an isolated invariant set in $X\times[-1,2]$, then $\pi(S)$ is invariant but not necessarily isolated.
(2) If Z is an isolated invariant set for the g-flow on $[-1,2]$, there may be no isolated invariant set in $\pi^{-1}(Z)$.
(3) If S is an isolated invariant set that continues over $\alpha(I)$ and Z is an isolated invariant set for the g-flow on $[-1,2]$, then $S(Z)=\bigcup_{z\in Z}S_z$ need not be an invariant set.

However, there is one important special case that is valid.

Proposition 4.1. *Supose S is an isolated invariant set that continues across $\alpha(I)$ for some $\alpha\in\mathcal{P}(\Lambda)$. Further, suppose that $C\subset[-1,2]$ is a connected, isolated set of zeros for $g\in\mathcal{G}$. Then $S_\alpha(C)=\{(x,s)\mid s\in C,\ x\in S_{\alpha(s)}\}$ is an isolated invariant set for the flow over (α, g).*

Proof. As a connected, isolated set of zeros for g, C is an isolated invariant set for the g-flow on $[-1,2]$ with isolating neighborhood M.

For every $c\in C$, there is a neighborhood U_c of $\alpha(c)$ in Λ and a compact set $N_c\subset X$ such that N_c is an isolating neighborhood of S_λ for every $\lambda\in U_c$. Choose $c_1,\dots,c_n\in C$ and compact sets $C_i\subset M$ such that $c_i\in C_i$, $C\subset\bigcup\operatorname{int}C_i$, and $\alpha(C_i)\subset U_{c_i}$. Let $N=\bigcup_i N_i\times C_i$. Then $\bigcup C_i$ is an isolating neighborhood for C in $[-1,2]$ and N is an isolating neighborhood for $S(C)$ in $X\times[-1,2]$.

To see this, take a point $(x,s)\in N$. If $s\notin C$ then $s\in M\setminus C$, and the drift flow carries s out of M in either forward or backward time. Then (x,s) is similarly carried out of $\pi^{-1}(M)$, and in particular out of N. If $s\in C$ then $\dot g(x)=0$, and the solution through (x,s) is simply the solution to $\dot x=f(x,\alpha(s))$. But then N_s is an isolating neighborhood for S_s, so either $x\in S_s$ (and $(x,s)\in S_\alpha(C)$) or the orbit through x leaves N_s, and hence N. □

This will suffice to establish the continuation results we require. For these, we introduce some subsets of $\mathcal{G}$ and $\mathfrak{D}$. Suppose $C\subset[-1,2]$ is a closed subset with a finite number of components. Let $\mathcal{G}(C)$ denote the set of functions in $\mathcal{G}$ such that C is an isolated set of zeros of g. Since C has a finite number of components, it follows that C is also an isolated invariant set for the flow $\dot s=g(s)$. For later use, we also define $\mathcal{G}^+\subset\mathcal{G}$ to be the set of functions with $g^{-1}(0)=\{-1,0,1,2\}$, and let $\mathfrak{D}(C)$ and $\mathfrak{D}^+$ denote the products of these sets with $\mathcal{P}(\Lambda)$. Note that, in general, $\mathcal{G}(C)$ (and hence $\mathfrak{D}(C)$) is disconnected with a finite number of components, while $\mathcal{G}^+$ is a contractible open dense set in $\mathcal{G}$. Thus $\mathfrak{D}^+$ is a contractible open dense set in $\mathfrak{D}$.

PROPOSITION 4.2. *Suppose S is an isolated invariant set which continues across Λ, and suppose that C is a closed subset of $[-1,2]$ with a finite number of components. Then, for every component $\mathcal{C}$ of $\mathcal{G}(C)$, there is a family of isolated invariant sets $\mathcal{S}$ in $X\times[-1,2]$ which are related by continuation over a path-connected open neighborhood $\mathcal{U}$ of $\mathcal{P}(\Lambda)\times\mathcal{C}$ in $\mathcal{D}$. Over $\mathcal{P}(\Lambda)\times\mathcal{C}$, this family is $\mathcal{S}_{(\alpha,g)}=S_\alpha(C)$.*

Proof. Without loss, we may assume that C is connected. Given $(\alpha,g)\in\mathcal{C}$, we must construct an isolating neighborhood N for $S_\alpha(C)$ with the property that, for α' near α and g' near g in $\mathcal{C}$, N isolates $S_{\alpha'}(C)$. This will give continuation over $\mathcal{C}$. This continuation will then automatically extend to an open neighborhood $\mathcal{U}$ of $\mathcal{C}$ in $\mathcal{D}$.

Fix $(\alpha,g)\in\mathcal{C}$, and construct an isolating neighborhood N for $S_\alpha(C)$ as in the previous proposition. Now consider the neighborhood $\mathcal{U}=\bigcap_i S(C_i,U_{s_i})$ of α in $\mathcal{P}(\Lambda)$. For every $(\alpha',g')\in\mathcal{U}\times\mathcal{C}$, the maximal invariant set in N is contained in $X\times C$, and consists of the union of the maximal invariant sets in each $X\times s$. But since $\alpha'(C_i)\subset U_{s_i}$, if $s\in C\cap C_i$ then the maximal invariant set in N over s is the maximal invariant set in N_{s_i}, or $S_{\alpha(s)}$. □

Having established the necessary continuation properties, we now turn to the question of computing Conley indices for isolated invariant sets over $\mathcal{D}$.

PROPOSITION 4.3. *If C is a connected, isolated set of zeros of g, then $h(C)$, the Conley index of C in $[-1,2]$, is either Σ^0, Σ^1 or $\bar{0}$. Let $\mathcal{C}_g$ denote the component of $\mathcal{G}(C)$ that contains g. If S is an isolated invariant set that continues over Λ then, over $\mathcal{P}(\Lambda)\times\mathcal{C}_g$, the index of $S_\alpha(C)$ in $X\times[-1,2]$ is $h(S)\wedge h(C)$.*

Proof. As noted previously, a connected isolated set of zeros is an isolated invariant set. Its index is determined by the sign of g on either side of C, and it is a simple matter to check that Σ^0, Σ^1, and $\bar{0}$ are the only possibilities. To compute the index of $S_\alpha(C)$, choose a constant path $\bar{\lambda}$. Then the flow on $X\times[-1,2]$ is a product flow, and it is well-known that the index $h(S_{\bar{\lambda}}(C))$ is given by $h(S)\wedge h(C)$. But the index is continuation-invariant, and $S_{\bar{\lambda}}(C)$ continues to $S_\alpha(C)$ over $\mathcal{P}(\Lambda)\times\mathcal{C}_g$, so the index is computed by $h(S)\wedge h(C)$ over $\mathcal{P}(\Lambda)\times\mathcal{C}_g$. □

Since the index of $S_\alpha(C)$ is smash product of the indices of S in X and C in $[-1,2]$, the homology index is just the tensor product:

$$CH_*(S_\alpha(C))\cong CH_*(S)\otimes CH_*(C).$$

Of course, if $h(C)=\bar{0}$ then $CH_*(S_\alpha(C))=CH_*(C)=0$. If $h(C)=\Sigma^n$, then $CH_{k+n}(S_\alpha(C))\cong CH_k(S)$. It will be important for our purposes to describe this suspension isomorphism precisely. Choose a point $c\in C$ and a path α. If we form the constant path $\overline{\alpha(c)}$, then $S_{\overline{\alpha(c)}}(C)=S_{\alpha(c)}\times C$ and there is a suspension isomorphism

$$CH_k(S_{\alpha(c)}) \xrightarrow{\otimes\sigma_n} CH_k(S_{\alpha(c)}) \otimes CH_n(C) \xrightarrow{\times} CH_{n+k}(S_{\alpha(c)} \times C) = CH_{n+k}(S_{\overline{\alpha(c)}}(C)),$$

where σ_n is the generator of $CH_n(C)$. Further, if we take any path in $\mathcal{P}(\Lambda)$ from the constant path $\overline{\lambda(c)}$ to α, continuation along the corresponding path from $(\overline{\lambda(c)}, g)$ to (α, g) induces an isomorphism

$$F_{\alpha,\overline{\alpha(c)}}: CH_k(S_{\overline{\alpha(c)}}(C)) \to CH_k(S_\alpha(C)).$$

We will denote the composition $\Sigma(S): CH_k(S_c) \to CH_{k+n}(S_\alpha(C))$ as the *index suspension isomorphism*. (The various embellishments with parentheses and subscripts will be omitted when clear from the context, or when their presence distracts more than it clarifies.) This isomorphism has the following important continuation property.

COROLLARY 4.4. *Suppose that S is an isolated invariant set which continues over Λ, and that C is a closed subset of $[-1,2]$ with a finite number of components. Further, suppose that $\mathcal{C}$ is a component of $\mathcal{G}(C)$ such that $h(C) = \Sigma^n$ on $\mathcal{C}$.*

Then, for every $c \in C$ and every $(\alpha, g) \in \mathcal{P}(\Lambda) \times \mathcal{C}$, the index suspension isomorphism

$$\Sigma(S): CH_k(S_{\alpha(c)}) \to CH_{k+n}(S_\alpha(C))$$

commutes with the continuation isomorphisms. That is, suppose that $\hat{\alpha}$ is a path in $\mathcal{P}(\Lambda) \times \mathcal{C}$ such that $\hat{\alpha}(s) = (\alpha_s, g_s)$. Let $\beta \in \mathcal{P}(\Lambda)$ be the path defined by $\beta(s) = \alpha_s(c)$. Then there is a commutative diagram

$$\begin{array}{ccc} CH_k(S_{\beta(0)}) & \xrightarrow{\Sigma(S)} & CH_{k+n}(S_\beta(0)) \\ \downarrow F_\beta(S) & & \downarrow F_{\hat{\alpha}}(S) \\ CH_k(S_{\beta(1)}) & \xrightarrow{\Sigma(S)} & CH_{k+n}(S_\beta(1)). \end{array}$$

Proof. First, suppose that each α_s is a constant path. Then $F_{\hat{\alpha}}(S) = F_\alpha(S) \times \mathrm{id}$ and the commutativity is trivial. For a general $\hat{\alpha}$, note that $\hat{\alpha}$ is homotopic to the path $(\overline{\alpha_s(c)}, g)$, with the homotopy running through the paths from each α_s to the constant paths $\alpha_s(c)$. But this homotopy is exactly what is used to define $\Sigma(S)$, so the commutativity of the diagram for the family of paths $\overline{\alpha_s(c)}$ implies the commutativity for $\hat{\alpha}$. □

We will be particularly interested in studying the behavior on $\mathfrak{D}([0,1])$ and $\mathfrak{D}^+$. As noted before, $\mathfrak{D}^+$ is a connected dense open subset of $\mathfrak{D}$; $\mathfrak{D}([0,1])$ is also connected, and has empty interior. Of course, on $\mathfrak{D}([0,1])$, the flow on $X \times [0,1]$ over (α, g) is simply the restriction of the flow on $X \times \Lambda$ to the one-parameter family of flows picked out by α. That is, $\mathfrak{D}([0,1])$ is our device for picking out one-parameter families of flows. The isolated invariant sets of interest are $S_\alpha([0,1]) = \bigcup_{0 \le s \le 1} S_{\alpha(s)}$. However, $h([0,1]) = \bar{0}$ across $\mathcal{G}([0,1])$, so $h(S_\alpha([0,1])) = \bar{0}$ across $\mathfrak{D}([0,1])$. Thus, applying the index theory directly to the one-parameter families will tell us nothing.

However, since $\mathfrak{D}([0,1])$ lies in the closure of $\mathfrak{D}^+$, we can apply the singular theory of Section 3. If $\mathcal{U}$ is the maximal open set in $\mathfrak{D}$ over which $S_\alpha([0,1])$ continues, let $\mathcal{U}^+ = \mathcal{U} \cap \mathfrak{D}^+$. Then $\mathcal{U}^+$ is open and dense in $\mathcal{U}$.

LEMMA 4.5. *$\mathcal{U}^+$ is path-connected.*

Proof. By construction, $\mathcal{U}$ is connected. It thus suffices to show that, if ω is a path in $\mathcal{U}$ with endpoints in $\mathcal{U}^+$, then there is a path ω^+ in $\mathcal{U}^+$ that is endpoint-homotopic to ω.

Let $0 < t < 1$ be the least t such that $\omega(t) = (\alpha_t, g_t) \notin \mathcal{U}^+$. Then there is a neighborhood $N \subset X \times [-1, 2]$ that is isolating for all (α, g) near (α_t, g_t). In particular, there is some $g' \in \mathcal{G}^+$ such that $(\alpha_{t'}, (1-\tau)g_{t'} + \tau g') \in \mathcal{U}$ for all t' near t and for all $0 \le \tau \le 1$. Clearly, this allows us to construct an $\omega' \simeq \omega$ so that $\omega'(t) \in \mathcal{U}^+$ for all t' near t. By compactness, this procedure can be applied across the interval $0 \le t \le 1$ to construct the required path. □

Over $(\alpha, g) \in \mathcal{U}^+$, the maximal invariant set no longer possesses the form $\bigcup_{0 \le s \le 1} S_{\alpha(s)}$ (as the g-flow on $[0,1]$ is no longer trivial), and in general depends on both α and g. We will denote the isolated invariant set over (α, g) by $S_{(\alpha,g)}([0,1])$, or more simply by $S_{(\alpha,g)}$. Since the flow on $[0,1]$ has 0 as an attractor and 1 as a repeller, $S_{(\alpha,g)}([0,1])$ has $S_{\alpha(0)}$ as an attractor and $S_{\alpha(1)}$ as a repeller. These features will be present for all $g \in \mathcal{G}^+$, but the structure of the connecting orbit set will vary with g. Our goal is to glean information about the one-parameter family $S_\alpha([0,1])$ from the behavior of the connecting orbit set as $\|g\| \to 0$.

5. Singular Transition Matrices

Transition matrices were introduced in [14] as a device for studying one-parameter families of flows. The setting for this study is the existence of an isolated invariant set that continues over the parameter space (without loss of generality, the unit interval). It is (at least tacitly) assumed that the flows at the two ends are well understood—in the sense of the existence of a Morse decomposition and knowledge of the connection matrices at these parameter values—and that we seek information about the parameter values in between. To obtain this information, a drift flow is put on the parameter space, so that 0 is an attractor and 1 is a hyperbolic repeller (i.e. $g \in \mathcal{G}^+$). Then, for small drift flows, the Morse decompositions $\mathcal{M}_0$ and $\mathcal{M}_1$ together form a Morse decomposition for the flow on $X \times [0,1]$. Connection matrices for this Morse decomposition are then computed. Such matrices have the form

$$\begin{bmatrix} \Delta_0 & T \\ 0 & \Delta_1 \end{bmatrix},$$

with Δ_0 and Δ_1 connection matrices for $\mathcal{M}_0$ and $\mathcal{M}_1$. The "off-diagonal" block T is a *transition matrix.* Nonzero entries in T detect connections from

$\mathfrak{M}_1$ to $\mathfrak{M}_0$. As the drift flow is slowed to zero, these connections limit to a sequence of connecting orbits for parameter values between 0 and 1.

We now revisit this development in our more general framework, deriving transition matrices from singular connection matrices. To emphasize this development, we will refer to these objects as singular transition matrices. As in the previous section, most of these results are the obvious generalization of the corresponding results in [14]. The exception is Theorem 5.5, which establishes the connection between transition matrices and continuation, and which is the essential step in relating singular transition matrices and topological transition matrices.

Given a Morse decomposition $\mathfrak{M}$ that continues over $\alpha([0,1])$, the first thing we must do is determine the structure of the corresponding Morse decomposition in $X\times[0,1]$ over some $(\alpha,g)\in\mathcal{U}^+$.

PROPOSITION 5.1. *For every $(\alpha,g)\in\mathcal{U}^+$, $(S_{\alpha(0)},S_{\alpha(1)})$ is an attractor–repeller pair for $S_{(\alpha,g)}([0,1])$. Moreover, these attractor–repeller decompositions are related by continuation.*

Clearly, if $\mathfrak{M}$ is a Morse decomposition that continues over $\alpha(I)$, then $\mathfrak{M}_0$ and $\mathfrak{M}_1$ are Morse decompositions of $S_{\alpha(0)}$ and $S_{\alpha(1)}$. If $g\in\mathcal{G}^+$, then the index of 0 in $[-1,2]$ is Σ^0 while the index of 1 is Σ^1. Thus the indices of the elements of $\mathfrak{M}_1$ in $X\times[-1,2]$ are suspensions of their indices in X. We denote the sum of these suspended indices by $C_*\Delta^{\Sigma}(P)$, or by $C_*\Delta^{\Sigma}_{\alpha(1)}(P)$ when it is necessary to display the path.

PROPOSITION 5.2. *If $\mathfrak{M}$ is a Morse decomposition of S that continues over Λ, then for every $(\alpha,g)\in\mathfrak{D}([0,1])$, the collection $\{M_{(\alpha,g)}(p)\}$ is a Morse decomposition of $S_{(\alpha,g)}$.*

Combining these propositions, we have the following corollary.

COROLLARY 5.3. *If $\mathfrak{M}$ is a Morse decomposition of S that continues over Λ, then there is an open subset $\mathcal{V}^+\subset\mathcal{U}^+$, whose closure contains $\mathfrak{D}([0,1])$, such that $\mathfrak{M}_0\cup\mathfrak{M}_1$ continues as a Morse decomposition across $\mathcal{V}^+$.*

The next task is to understand the relationship between the partial order $<_{\alpha([0,1])}$ on P (i.e. the partial order, without the drift flow, that is valid across $\alpha([0,1])$) and the partial orders $<_{(\alpha,g)}$ on $P\times\{0,1\}$ (i.e. the partial orders that result when the drift flow is imposed). Actually, we want to compare $<_{\alpha([0,1])}$ to the drift partial order $<_{d,\alpha}$ on P, which is derived from the flow-defined orders $<_{(\alpha,g)}$ on $P\times\{0,1\}$ in $\mathcal{U}^+$. The drift partial order is obtained by first taking the inverse limit of the flow-defined orders, then projecting this onto a partial order on P. Namely, $p<_{d,\alpha}q$ if there exists an $i,j\in\{0,1\}$, a $g\in\mathcal{G}([0,1])$, and a sequence (α_n,g_n) in $\mathcal{U}^+$ converging to (α,g) such that $(p,i)<_{(\alpha_n,g_n)}(q,j)$ for every n. In other words, $p<q$ in the drift partial order if there is a connection (or sequence of connections) from a q-Morse set to a p-Morse set in every open family of slow drift flows.

LEMMA 5.4. *For any α, the flow-defined order on $\alpha([0,1])$, $<_\alpha$, refines the drift partial order $<_{d,\alpha}$.*

Singular transition matrices are defined in terms of $<_{d,\alpha}$-connection matrices on $\mathcal{M}_0 \cup \mathcal{M}_1$. If $\mathcal{M}$ is a Morse decomposition that continues over Λ, then for every path $\alpha \in \mathcal{P}(\Lambda)$ we have the set of singular connection matrices $\mathcal{C}M(\mathcal{M}_0 \cup \mathcal{M}_1, \alpha \times \mathcal{G}([0,1]), \mathcal{V}_\alpha^+)$. These are connection matrices with the drift partial order. For each connection matrix

$$\Delta(P \times \{0,1\}):$$
$$\bigoplus_{p \in P} C_*\Delta_{\alpha(0)}(P) \oplus \bigoplus_{p \in P} C_*\Delta^{\Sigma}_{\alpha(1)}(P) \to \bigoplus_{p \in P} C_*\Delta_{\alpha(0)}(P) \oplus \bigoplus_{p \in P} C_*\Delta^{\Sigma}_{\alpha(1)}(P),$$

there is an off-diagonal block

$$T = \Delta(P \times \{0\}, P \times \{1\}) : \bigoplus_{p \in P} C_*\Delta^{\Sigma}_{\alpha(1)}(P) \to \bigoplus_{p \in P} C_*\Delta_{\alpha(0)}(P).$$

These are the *singular transition matrices.* The interpretation of these matrices is that, if $T(p,q) = T((p,0),(q,1)) \neq 0$, then $p \leq_{d,\alpha} q$ and so $p \leq_\alpha q$. That is, somewhere along the path α, there is a connection from $M(q)$ to $M(p)$. In this interpretation, it is clear that what singular transition matrices are detecting is codimension-1 connections. (This should be contrasted with the interpretation of connection matrices, which detect connections that persist over an open set of parameter values.) This is the best that can reasonably be expected. If a connection occurs only for a set of parameter values with codimension 2 or more, then an arbitrarily small perturbation will perturb any path in Λ off of that set.

We know a certain amount about the structure of T. The crucial statement about this structure is the following.

THEOREM 5.5. *The connection homomorphism for the attractor-repeller decomposition $(S_{\alpha(0)}, S_{\alpha(1)})$ of $S_{(\alpha,g)}([0,1])$ is an isomorphism, which is computed by continuation of S across Λ. That is, there is a commutative diagram*

$$\begin{array}{ccc} CH_p(S_{\alpha(1)}) & \xrightarrow{\quad\Sigma\quad} & CH^{\Sigma}_{p+1}(S_{\alpha(1)}) \\ & {\scriptstyle F_\alpha^{-1}} \searrow \quad \swarrow {\scriptstyle \partial} & \\ & CH_p(S_{\alpha(0)}). & \end{array}$$

Proof. As before, we will first prove the result for a product flow. If α is a constant path, then $\{\alpha\} \times \mathcal{G}^+ \subset \mathcal{U}^+$ and $S_{(\alpha,g)}([0,1]) = S_{\alpha(0)} \times [0,1]$. If (N,L) is an index pair for $S_{\alpha(0)}$, then

$$(N \times [-\tfrac{1}{2}, \tfrac{3}{2}], (N \times ([-\tfrac{1}{2}, \tfrac{1}{2}] \cup \{\tfrac{3}{2}\})) \cup (L \times [-\tfrac{1}{2}, \tfrac{3}{2}]), (N \times \{\tfrac{3}{2}\}) \cup (L \times [-\tfrac{1}{2}, \tfrac{3}{2}]))$$

is an index triple for (S_0, S_1). This triple can be simplified to

$$(N \times [0,1], (N \times \{0,1\}) \cup (L \times [0,1]), (N \times \{1\}) \cup (L \times [0,1])).$$

The boundary map is then a map of the form

$$H_p(N,L)\otimes H_1([0,1],\{0,1\})\to H_p(N,L)$$

with $z\otimes o_1 \mapsto z$, where $o_1\in H_1([0,1],\{0,1\})$ is the generator. Further, in this setting Σ is $\times o_1$, so $\partial\circ\Sigma=\text{id}$. But in this setting, continuation along α is trivial, so $F_\alpha=\text{id}$.

The general result then follows by continuation over $\mathfrak{U}^+$. If $(\alpha,g)\in\mathfrak{U}^+$, choose a path $\omega=(\alpha_\tau,g_\tau)$ in (α,g) to $(\overline{\alpha(1)},g)$. Without loss, we may assume $\alpha_\tau(1)=\alpha(1)$ for all τ. Now consider continuation along ω. From Corollary 2.6, $\partial_\alpha=F_{\omega^{-1}}\circ\partial_{\overline{\alpha(1)}}\circ F_\omega$. By construction, $\Sigma_\alpha=F_{\omega^{-1}}\circ\Sigma_{\overline{\alpha(1)}}\circ F_\omega$. Clearly, the induced isomorphisms on $CH_*(S_{\alpha(0)})$ and $CH_*(S_{\alpha(1)})$ are F_α and id, respectively. Thus

$$\begin{aligned}\partial_\alpha\circ\Sigma_\alpha&=F_{\omega^{-1}}\circ\partial_{\overline{\alpha(1)}}\circ F_\omega\circ F_{\omega^{-1}}\circ\Sigma_{\overline{\alpha(1)}}\circ F_\omega\\&=F_{\alpha^{-1}}\circ\partial_{\overline{\alpha(1)}}\circ\Sigma_{\overline{\alpha(1)}}\\&=F_{\alpha^{-1}}.\end{aligned}$$

□

In particular, since Σ and F_α are isomorphisms, so is ∂. If we insert this isomorphism into the commutative diagram (2.1), we obtain

$$\begin{array}{ccccccc}0\longrightarrow & H_k\Delta(I\times\{1\}) & \xrightarrow{[T]} & H_{k-1}\Delta(I\times\{0\}) & \longrightarrow 0\\ & \downarrow{\scriptstyle\phi(I\times\{0\})} & & \downarrow{\scriptstyle\phi(I\times\{0\})} & \\ 0\longrightarrow & CH_k^\Sigma(M_1(I)) & \xrightarrow{\partial(I)} & H_{k-1}(M_0(I)) & \longrightarrow 0.\end{array}$$

When we take $I=\{p\}$, this implies that $T(p,p)=\partial(p)=F_{\alpha^{-1}}\circ\Sigma^{-1}$ is an isomorphism. Combining this with the upper triangularity of T, we can codify our knowledge of singular transition matrices as follows.

Corollary 5.6. *A singular transition matrix T is an upper-triangular isomorphism with respect to the drift partial order $<_{d,\alpha}$. The diagonal entries of T are given by $T(p,p)=F_{\alpha^{-1}}(p)\circ\Sigma^{-1}(p)$.*

That is, T has the form

$$\begin{bmatrix}F_{\alpha^{-1}}\circ\Sigma^{-1} & & *\\ & \ddots & \\ 0 & & F_{\alpha^{-1}}\circ\Sigma^{-1}\end{bmatrix}.$$

There are some important gaps in our understanding of singular transition matrices. First, there is the fundamental computational issue: How does one ensure that all the singular transition matrices have been identified? Beyond this, there are questions about the relationship between the path α and the set of matrices $\mathfrak{J}$ over α. For example, what is the relationship between $\mathfrak{J}(\alpha)$ and $\mathfrak{J}(\alpha^{-1})$? If α_1 and α_2 are paths with $\alpha_1(1)=\alpha_2(0)$, what is the relationship between $\mathfrak{J}(\alpha_1)$, $\mathfrak{J}(\alpha_2)$, and $\mathfrak{J}(\alpha_1*\alpha_2)$? Finally, the interpretation of the results is slightly unsatisfactory in that the matrices give information about $<_{d,\alpha}$, rather than about $<_\alpha$ directly.

6. Topological Transition Matrices

In [10], an object was constructed with many of the same features as singular transition matrices.

Suppose $\mathcal{M}$ is a Morse decomposition that continues over Λ. Let $\Lambda' \subset \Lambda$ be the set of parameter values for which there are no connecting orbits:

$$\Lambda' = \left\{ \lambda \in \Lambda \mid S_\lambda = \bigcup_{p \in P} M_\lambda(p) \right\}.$$

For $\lambda \in \Lambda'$, there is a canonical isomorphism $\Phi_\lambda : C_*\Delta_\lambda(P) \to CH_*(S_\lambda)$. If $\lambda_0, \lambda_1 \in \Lambda'$, choose a path ω from λ_0 to λ_1. Then we can carry out continuation along ω in two ways: we can continue S along ω; and we can continue $\bigcup_{p \in P} M(p)$ along ω. Since the connecting orbits are precisely the complement $S \setminus \bigcup_{p \in P} M(p)$, the difference between these two continuations should reveal something about the connecting orbit set. This is made precise as follows. We can construct a diagram of isomorphisms

$$\begin{array}{ccc} C_*\Delta_{\lambda_1}(P) & \xrightarrow{F\Delta_{\lambda_0,\lambda_1}(P)} & C_*\Delta_{\lambda_0}(P) \\ \downarrow \Phi_{\lambda_1} & & \downarrow \Phi_{\lambda_0} \\ CH_*(S_{\lambda_1}) & \xrightarrow{F_{\lambda,\mu}} & CH_*(S_{\lambda_0}). \end{array}$$

This diagram is not commutative in general, and it is precisely this failure of commutativity that gives information about connecting orbits. To reduce this to a more manageable form, fix a field of coefficients for the homology groups, choose a basis on $C_*\Delta_{\lambda_1}(P)$, and use the continuation isomorphism $F\Delta_{\lambda_0,\lambda_1}(P)$ to carry this to a basis on $C_*\Delta_\lambda(P)$. Then the composition $\Phi_{\lambda_0}(P)^{-1} \circ F_{\lambda_0,\lambda_1} \circ \Phi_{\lambda_1}(P)$ can be represented as a matrix with respect to these bases. We denote this matrix $T^{\text{top}}_{\lambda_0,\lambda_1}$ and refer to it as the *topological transition matrix.*

The relevant properties of topological transition matrices are:

(1) If continuation is path-independent (in particular, if $\pi_1(\Lambda) = 0$), then T^{top} is path-independent.
(2) If $\lambda_0, \lambda_1, \lambda_2 \in \Lambda'$, then $T^{\text{top}}_{\lambda_2,\lambda_0} = T^{\text{top}}_{\lambda_2,\lambda_1} \circ T^{\text{top}}_{\lambda_1,\lambda_0}$.
(3) The diagonal entries $T^{\text{top}}_{\lambda_1,\lambda_0}(p,p) : CH_*(M_{\lambda_0}(p)) \to CH_*(M_{\lambda_1}(p))$ are represented by the identity map. If λ_0 and λ_1 are in the same path component of Λ', then $T^{\text{top}}_{\lambda_1,\lambda_0} = \text{id}$ (i.e., all off-diagonal entries are zero).

Thus, if $T^{\text{top}}_{\lambda_1,\lambda_0} \neq \text{id}$, then every path from λ_0 to λ_1 in Λ must pass through a parameter value κ such that $S_\kappa \neq \bigcup_{p \in P} M_\kappa(p)$, a parameter value such that S_κ has connecting orbits. So essentially T^{top} detects codimension-1 families of connecting orbits. This can be refined further to:

(4) If $T^{\text{top}}_{\lambda_1,\lambda_0}(p,q) \neq 0$, then $p <_\Lambda q$. Further, if ω is a path from λ_0 to λ_1, there is a sequence $s_1 < \cdots < s_n$ in $[0,1]$ and a sequence $q = p_0, p_1, \ldots, p_{n-1}, p_n = p$ in P such that the connecting orbit set $C(M_{\omega(s_i)}(p_{i-1}), M_{\omega(s_i)}(p_i))$ is nonempty.

Thus topological transition matrices, like singular transition matrices, are upper-triangular matrices with "id" down the diagonal and with off-diagonal entries indicating the presence of connecting orbits. In comparison with singular transition matrices, topological transition matrices have the obvious disadvantage that they can only be defined between parameter values with no connecting orbits. However, when they are defined, they have several useful properties. First, there is only a unique topological transition matrix between any two parameter values in Λ' (in fact, between any two path components of Λ'). Second, since the product of topological transition matrices is a transition matrix, it is easy to compute all possible matrices once a few are known. In particular, $T_{\lambda_0,\lambda_1}^{-1} = T_{\lambda_1,\lambda_0}$.

In simple examples, it is observed that topological transition matrices and singular transition matrices have the same matrix representation. Given this, together with the similarity in the properties of the two matrices, it is natural to ask if the two definitions coincide. We now turn to this question.

7. Equivalence of Singular and Topological Transition Matrices

Before discussing the relationship between topological transition matrices and singular transition matrices, we must point out that the question "Is T^{top} a singular transition matrix?" does not really make sense as stated. Singular transition matrices are derived from drift flows, in which the index of the repeller has been suspended. However, we are in a position to reconcile this difference. We have established the existence of a canonical suspension isomorphism Σ. Using this, we can canonically relate indices in $X_{\alpha(1)}$ to indices in $X \times [-1, 2]$. With that relation in place, we can now state our main result.

THEOREM 7.1. *Suppose $\mathcal{M} = \{M_p\}_{p \in P}$ is a Morse decomposition that continues over Λ. If $\lambda_0, \lambda_1 \in \Lambda'$, then for every $(\alpha, g) \in \mathcal{U}^+$ with $\alpha(i) = \lambda_i$, the unique singular transition matrix is $T_{\text{top}} \circ (\bigoplus_{p \in P} \Sigma(p))^{-1}$.*

Proof. It suffices to show that, if $T(\alpha, g)$ is a singular transition matrix, then $T(\alpha, g) \circ \bigoplus_p \Sigma_p = T_{\text{top}}$. This equality follows from the commutativity of the following diagram:

$$
\begin{array}{ccccc}
C_k\Delta_1(p) & \xrightarrow{\bigotimes_{p\in P}\Sigma_p} & C_{k+1}\Delta_1^{\Sigma}(P) & \xrightarrow{T(\alpha,g)} & C_{k+1}\Delta_0(P) \\
\downarrow{\scriptstyle \Phi_1(P)} & & \downarrow{\scriptstyle \Phi_1^{\Sigma}(P)} & & \downarrow{\scriptstyle \Phi_0(P)} \\
CH_k(S_1) & \xrightarrow{\Sigma(P)} & CH_{k+1}(S_1) & \xrightarrow{\partial} & CH_k(S_0) \\
\downarrow{\scriptstyle \text{id}} & & & & \downarrow{\scriptstyle \text{id}} \\
CH_k(S_1) & & \xrightarrow{\qquad F_{01}(S) \qquad} & & CH_k(S_0).
\end{array}
$$

The commutativity of the bottom square was established in Theorem 5.5. The commutativity of the other two squares is established by the following lemmas. □

LEMMA 7.2. *For every* (α, g) *such that* $\alpha(1) \in \Lambda'$, *the suspension isomorphisms commute with the* Φ*s. That is, there is a commutative diagram*

$$\begin{array}{ccc} C_k\Delta_1(P) & \xrightarrow{\bigotimes_{p\in P}\Sigma_p} & C_{k+1}\Delta_1^{\Sigma}(P) \\ \downarrow \Phi_1(P) & & \downarrow \Phi_1^{\Sigma}(P) \\ CH_k(S_1) & \xrightarrow{\Sigma(P)} & CH_{k+1}(S_1). \end{array}$$

Proof. Once again, start with the case of the constant path $\overline{\alpha(1)}$. If $(N(p), L(p))$ is a disjoint collection of index pairs for the sets $M_1(p)$ in $X_{\alpha(1)}$, then $(\bigcup_{p\in P} N(p), \bigcup_{p\in P} L(p))$ is an index pair for S_1. Further,

$$(N(p), L(p)) \times ([1-\epsilon, 1+\epsilon], \{1-\epsilon, 1+\epsilon\})$$

is an index pair for $M_1(p)$ in $X_{\alpha(1)}$ and

$$\left(\bigcup_{p\in P} N(p), \bigcup_{p\in P} L(p)\right) \times ([1-\epsilon, 1+\epsilon], \{1-\epsilon, 1+\epsilon\})$$

is an index pair for $S_\alpha(1)$ in $X\times[-1,2]$. If $o_1 \in H_1([1-\epsilon, 1+\epsilon], \{1-\epsilon, 1+\epsilon\})$ is a generator, then all of the suspension isomorphisms are given by $\times o_1$. It is clear then that $\Phi_1^{\Sigma}(P)\circ(\bigoplus_{p\in P}\Sigma(p) = \Sigma\circ\Phi_1(P))$.

In general, if $\alpha(1)\in\Lambda'$, then there is an ϵ such that $\alpha(s)\in\Lambda'$ for every $1-\epsilon \le s \le 1+\epsilon$. If $A\colon \alpha \simeq \bar{\alpha}(1)$ is the straight-line homotopy then, for continuation along A, the maps Φ_1 and Φ_1^{Σ} commute with the continuation isomorphisms. The suspension isomorphisms have already been established to commute with the continuation isomorphisms. Thus the equality $\Phi_1^{\Sigma}(P)\circ(\bigoplus_{p\in P}\Sigma(p)) = \Sigma\circ\Phi_1(P)$ continues from $\bar{\alpha}(1)$ to α. □

LEMMA 7.3. *For every* (α, g), *the boundary homomorphisms commute with the* Φs. *That is, there is a commutative diagram*

$$\begin{array}{ccc} C_{k+1}\Delta_1^{\Sigma}(P) & \xrightarrow{T(\alpha,g)} & C_k\Delta_0(P) \\ \downarrow \Phi_1^{\Sigma}(P) & & \downarrow \Phi_0(P) \\ CH_{k+1}^{\Sigma}(S_1) & \xrightarrow{\partial} & CH_k(S_0). \end{array}$$

Proof. Since $S_i = \bigcup_{p\in P} M_i(p)$, for any $g\in\mathcal{U}^+$ we have

$$C(S_1, S_0) = \bigcup_{p,q\in P} C(M_1(q), M_0(p)).$$

Since the sets $C(M_1(q), M_0(p))$ form a separation of $C(S_1, S_0)$, the additivity of the connection homomorphism [9] implies that $\partial = \sum_{p,q\in P}\partial_{p,q}$ as maps from $CH_*(S_1)$ to $CH_*(S_0)$. Since $T(\alpha, g)$ is the matrix with pq entry $\partial_{p,q}$, $T(\alpha, g)$ represents $\sum_{p,q\in P}\partial_{p,q}$ with respect to $C_*\Delta(I)$. Thus $\Phi_0(I)\circ T(\alpha,g)\circ\Phi_1^{\Sigma}(P)^{-1}$ represents the sum with respect to $CH_*(S)$, and the commutativity of the diagram follows. □

As noted in the introduction, Theorem 7.1 has several consequences for our interpretation and understanding of singular transition matrices. Foremost among these is the uniqueness of the matrix, and the independence of the construction that this implies.

COROLLARY 7.4. *If $\lambda_0, \lambda_1 \in \Lambda'$, then there is a unique transition matrix T_{λ_0,λ_1}, from λ_1 to λ_0, that is independent of both the path α connecting λ_0 and λ_1 and of the drift flow g.*

Next, since singular transition matrices can be defined in terms of topological transition matrices, they inherit the composition formulas of topological transition matrices.

COROLLARY 7.5. *If $\lambda_0, \lambda_1, \lambda_2 \in \Lambda'$, then:*

(1) *$T_{\lambda_0\lambda_1}$ is a diagonal matrix if λ_0, λ_1 are in the same path component of Λ';*
(2) *$T_{\lambda_0\lambda_2} = T_{\lambda_0\lambda_1} \circ (\bigoplus_{p\in P} \Sigma_{\lambda_1}(p)) \circ T_{\lambda_1\lambda_2}$; and*
(3) *$T_{\lambda_1\lambda_0} = T_{\lambda_0\lambda_1}^{-1}$.*

Finally, this relation helps to bridge the gap between the drift partial order $<_{d,\alpha}$ and the flow-defined partial order over $\alpha([0,1])$, $<_\alpha$. While it does not guarantee that the two partial orders coincide, it does mean that they are effectively equal. We already knew that if $T_{\lambda_0\lambda_1}(p,q) \neq 0$ then $p <_{d,\alpha} q$. Now we see that if $T_{\lambda_0\lambda_1}(p,q) \neq 0$ then $p <_\alpha q$ for every path α from λ_0 to λ_1. That is, the two partial orders may not coincide, but the portions of the orders that are detectable by index techniques do coincide. In particular, if there is a connection in the flow-defined order that is not the limit of drift-flow connections, then it will be "invisible" to all index calculations.

Of course, all of these results depend on the corresponding results for topological transition matrices, which in turn depended on the particular structure assumed for the parameter space Λ'. We do not expect these results to generalize verbatim to the general setting (i.e. to transition matrices between parameter values where there are connecting orbits), because the uniqueness statement is false in general. However, it seems likely that these results are true in some form for all singular transition matrices. This is an open question for future research.

References

[1] C. Conley, *Isolated invariant sets and the Morse index,* CBMS Regional Conf. Ser. in Math., 38, Amer. Math. Soc., Providence, RI, 1978.
[2] M. Eidenschink and K. Mischaikow, *A numerical algorithm for finding isolating neighborhoods* (in progress).
[3] R. Franzosa, *Index filtrations and the homology index braid for partially ordered Morse decompositions,* Trans. Amer. Math. Soc. 298 (1986), 193–213.
[4] ———, *The continuation theory for Morse decompositions and connection matrices,* Trans. Amer. Math. Soc. 310 (1988), 781–803.
[5] ———, *The connection matrix theory for Morse decompositions,* Trans. Amer. Math. Soc. 311 (1989), 561–592.
[6] R. Franzosa and K. Mischaikow, *The connection matrix theory for semiflows on (not necessarily locally compact) metric spaces,* J. Differential Equations 71 (1988), 270–287.

[7] ———, *Algebraic transition matrices,* preprint.

[8] H. Kokubu, K. Mischaikow, and H. Oka, *Existence of infinitely many connecting orbits in a singularly perturbed ordinary differential equation,* preprint.

[9] C. McCord, *The connection map for attractor-repeller pairs,* Trans. Amer. Math. Soc. 307 (1988), 195–203.

[10] C. McCord and K. Mischaikow, *Connected simple systems, transition matrices and heteroclinic bifurcations,* Trans. Amer. Math. Soc. 333 (1992), 397–422.

[11] K. Mischaikow and M. Mrozek, *Chaos in the Lorenz equations: a computer assisted proof,* Bull. Amer. Math. Soc. (to appear).

[12] K. Mischaikow, M. Mrozek, and J. Reineck, *Singular index pairs,* preprint.

[13] R. Moeckel, *Morse decompositions and connection matrices,* Ergodic Theory Dynamical Systems 8 (1988), 227–250.

[14] J. Reineck, *Connecting orbits in one-parameter families of flows,* Ergodic Theory Dynamical Systems 8 (1988), 359–374.

[15] K. P. Rybakowski, *The homotopy index and partial differential equations,* Springer, New York, 1987.

[16] D. Salamon, *Connected simple systems and the Conley index of isolated invariant sets,* Trans. Amer. Math. Soc. 291 (1985), 1–41.

C. K. McCord
Department of Mathematics
University of Cincinnati
Cincinnati, OH 45221

K. Mischaikow
Center for Dynamical Systems
and Nonlinear Studies
School of Mathematics
Georgia Institute of Technology
Atlanta, GA 30332