CHAPTER 5

System Restoration through Managed Regime Shifts

The defining tenet of ecological resilience [55] is that collapse is inevitable, but that buried within that collapse are the resources needed for eventual restoration of system function. The restoration of a collapsed system is therefore a multi-stage process in which the resources of the collapsed system must first be *reorganized* to a point from which full *recovery* can readily and naturally proceed. In particular, restoration requires the system to traverse a sequence of *alternative states* [108] before full system function is restored. If the traverse of that sequence is left to chance, then the reorganization phase may take a prohibitively long time to complete. If, on the other hand, an environmental resource manager can artificially trigger the regime shift between alternative states, then recovery may occur at a faster and more predictable pace. This line of thinking suggests that one way to enhance a system's ecological resilience is through the active management of the system's regime shifts.

This reorganization phase can be illustrated through the tritrophic food web from Fig. 9(a). Recall that this food web consists of a producer (algae), primary consumer (Daphnia), and secondary consumer (fish - crappie). The system's nominal regime is taken as a state where most of the system's biomass is held in the two consumer compartments, as illustrated by the bar graph on the left side of Fig. 18. A collapse of this nominal regime may occur if some environmental event (i.e. storm, epidemic, overfishing) depletes the population of these consumer guilds. The collapse of the consumer guilds will reduce predation pressure on the producers, thereby resulting in explosive growth in that guild as shown in the *collapsed system's* bar graph. System reorganization involves the eventual rebuilding of the primary consumer compartment. This can occur naturally, or its growth can be hastened through intentional restocking of that guild. Once this is done, then the system has reached a stage from which the secondary consumer guild can readily re-establish itself without any further external assistance.

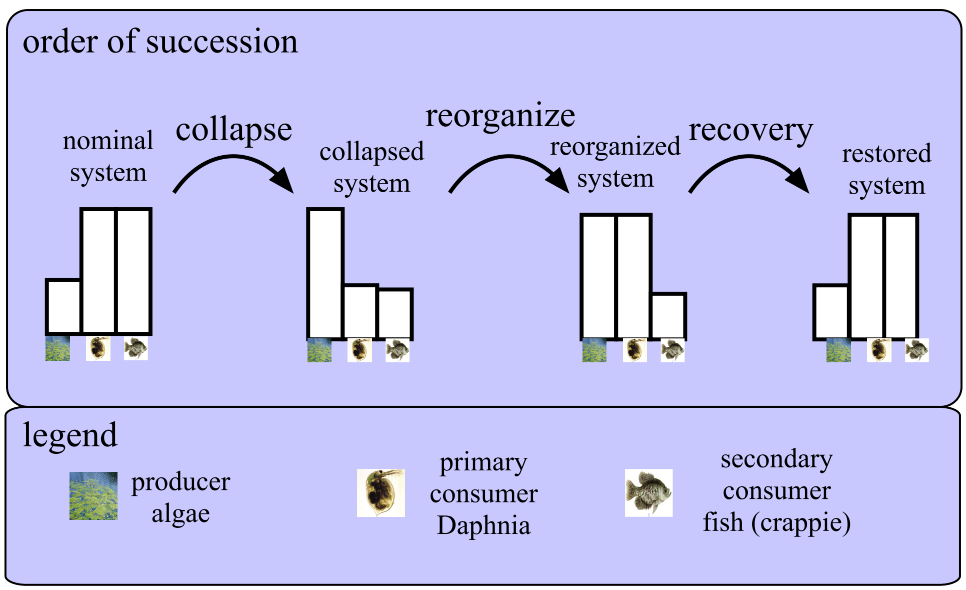

FIGURE 18. Order of Succession and Restoration Plans

The necessity of following the order of "collapse" to "reorganize" to "recovery" may be seen by trying another "common sense" approach that short circuits this order. From the standpoint of ecosystem services, it is the secondary consumer guild (fish) that is of greatest interest to the environmental resource manager. If the trophic relationships between these guilds are ignored, the resource manager may simply decide to restock the secondary consumer compartment after the collapse. That strategy, however, will fail since restocking the secondary consumers increases predation pressure on the primary consumers, thereby depressing primary consumer production. Without a sufficiently large primary consumer guild, the secondary consumer guild cannot maintain itself as a minimum level of resources is required for continued growth. In other words, successful recovery follows a well defined *order of succession* which rebuilds (reorganizes) guilds at lower trophic levels, before attempting to recover guilds at higher trophic levels. Plans that ignore this order of succession have a high probability of failure.

Traversing the order of succession also requires that the system meet threshold conditions to trigger the sequence of regime shifts leading to system restoration. The example in Fig. 18 has only one threshold to meet; namely the requirement that the primary consumer guild was rebuilt to a critical level from which no further intervention was required. If one had stopped rebuilding the primary consumers before that critical level was achieved, the system would simply slip back into its collapsed state. This sequence of regime shifts represent a path or plan for full system restoration. Ecologically resilient systems will eventually follow the path, but that traversal may take a very long time. Path traversal can be sped up through interventions that more quickly meet the thresholds marking the regime shift. The key pieces of information needed to affect this speed up are 1) knowledge of the pathways leading to recovery, 2) knowledge of the thresholds that have to be met to traverse the pathways, and 3) understanding which actions can be taken to speed up meeting those thresholds. The objective of this chapter is to discuss ways for finding these pathways from the mechanistic model we have for the system.

In practice, biologists use their knowledge of trophic interactions and experimental field work to propose restoration plans. This work is done with the help of computer simulation models in which Monte Carlo runs are used to evaluate the effectiveness of a chosen restoration plan. Prior successes with this approach, unfortunately, have been limited. A prior study [63] found that only a third of completed real-world restoration project were successful with at least another third being classified as abject failures. Given that the computer model was based on a mechanistic system model, one might ask whether appealing to formal methods, rather than Monte Carlo methods, would provide a more efficient way for developing restoration plans with more predictable real-world outcomes. The following sections show how this might be done. In particular, we show how algorithmic methods from dynamical system analysis [78, 65, 36, 64] can be used to construct a *regime transition system* (RTS) that characterizes all possible shock-induced regime shifts. The regime transition system provides a complete picture of the pathways that can be followed for system restoration.

This chapter focuses on the construction and use of the regime transition system (RTS). Building this transition system first requires that we identify as many of the system's regimes as possible. This will be done by first finding *all* of the system's fixed point equilibria, using simulations to generate forward orbits from points in the equilibria's neighborhood, and finally using methods from algorithmic dynamics [78, 65, 36, 64] to construct cubical complexes that isolate the system's basic sets from each other. The regime transition system is then defined as a finite state machine whose states are these cubical complexes and whose state transitions are determined from transitions between complexes that were observed in the simulated orbits. The chapter closes by showing how shock-induced and bifurcation-induced regime shifts can be used to enhance the ecological resilience of the system. The methods are demonstrated on a four guild food web with intra-guild predation.

1. Finding All Equilibria of Consumer-Resource Systems

Construction of the regime transition system starts with finding *all* equilibria of the consumer-resource system. Since these equilibria are basic sets of the system, we use them as a starting point for generating simulated orbits that reach other basic sets (regimes). This provides a systematic way of searching for all of the basic sets that can be reached from the system's fixed points.

Traditional methods for finding system equilibria use successive approximation to find a *single* root of the algebraic equation,

$$0 = f(x; k_0)$$

When f is a vector of polynomials in $\mathbb{R}[x]$, then methods for algebraic geometry [17] can be used to find *all* roots of the system of equations. This section reviews these methods and then applies them to finding all equilibria of a four guild consumer-resource system that is used throughout this chapter as a running example.

Algebraic Geometry: This subsection reviews those concepts from algebraic geometry used in developing algorithms that find all system equilibria. Consider a subset $I \in \mathbb{R}[x]$ where $x = \{x_1, \dots, x_n\}$ is the indeterminate variable. This set is called an *ideal* if

- $0 \in I$
- for all $f, g \in I$, then $f + g \in I$
- for all $f \in I$ and $g \in \mathbb{R}[x]$, then $f \cdot g \in I$

Let $f_1, \dots, f_m$ be polynomials in $\mathbb{R}[x]$, then the set

$$\langle f_1, \dots, f_m \rangle := \left\{ \sum_{i=1}^{m} h_i(x) f_i(x), \quad \text{with } h_i \in \mathbb{R}[x] \right\}$$

can be shown to be an ideal. An ideal I is said to be finitely generated if there exist polynomials $f_1, \ldots, f_m \in \mathbb{R}[x]$ such that $I = \langle f_1, \ldots, f_m \rangle$ and we say that these polynomials form a *basis* for I. It can be shown [17] that every ideal of $\mathbb{R}[x]$ is finitely generated. The bases for an ideal I, in general, are not unique, but certain bases will be more useful than others.

When $f(x)$ is a polynomial in $\mathbb{R}[x]$, it can be written as

$$f(x) = \sum_{\alpha} k_\alpha x^{[\alpha]}$$

where $k_\alpha \in \mathbb{R}$ are real coefficients for monomial, $x^{[\alpha]}$, with multi-index $\alpha = (\alpha_1, \ldots, \alpha_n)$. There are many ways of explicitly ordering the terms in the above summation, depending on how we choose to order the multi-indices. This ordering is important in developing computationally efficient algorithms. The most commonly used order is the *standard lexicographic order*. In particular let $\alpha = (\alpha_1, \ldots, \alpha_n)$ and $\beta = (\beta_1, \ldots, \beta_n)$ be two multi-indices. We say $\alpha >_{\text{lex}} \beta$ if the the leftmost nonzero index in $\alpha - \beta$ is positive (not zero). We say $x^{[\alpha]} >_{\text{lex}} x^{[\beta]}$ if $\alpha >_{\text{lex}} \beta$.

Let $I = \langle f_1, \ldots, f_m \rangle$ denote an ideal of $\mathbb{R}[x]$ that is finitely generated by polynomials $f_1, \ldots, f_m$ in $\mathbb{R}[x]$. The set

$$\mathbb{V}(I) = \left\{ z \in \mathbb{C}^n \;\middle|\; f_i(z) = 0 \right\}$$

is called an *algebraic variety* of I generated by $\{f_i\}_{i=1}^m$. The variety is a subset of the complex field consisting of all zeros of the polynomials generating the ideal. The variety is said to be *zero dimensional* if it consists of a finite set of points in $\mathbb{C}^n$.

Let the ideal $I \subset \mathbb{R}[x]$ be zero dimensional then a basis

$$\mathcal{T} = \{f_1, f_2, \ldots, f_m\}$$

is said to be *triangular* if

- $f_j \in \mathbb{R}[x_{n-j+1}, \ldots, x_n]$ for any $j = 1, \ldots, m$ and

- the leading monomial in f_j with respect to the standard lexicographic order is of the form $x_{n-j+1}^{m_j}$ for some $m_j \geq 1$.

A list of triangular bases $\mathcal{T} = \{\mathcal{T}_1, \ldots, \mathcal{T}_\ell\}$ is called a *triangular decomposition* [22] of I if

$$\mathbb{V}(I) = \mathbb{V}(\mathcal{T}_1) \cup \cdots \cup \mathbb{V}(\mathcal{T}_\ell)$$

Since $\mathcal{T}_j$ is triangular, it has at least one univariate polynomial whose zeros/varieties can be readily computed using a traditional root finding algorithm. The computed zero is then back substituted into another polynomial of the basis to get another univariate polynomial. The back substitution process is then repeated for all polynomials of the basis, thereby obtaining a specific solution for the system of equations. The process is then repeated for all bases in the triangular decomposition to obtain a set of complex numbers satisfying the algebraic equations generating the ideal, I. Because I's affine variety, $\mathbb{V}(I)$, equals the union of the varieties of all bases in the triangular decomposition, the union of all solutions obtained for each basis in the decomposition must equal the set of *all* possible solutions to the system of polynomial equations.

Working Example: This subsection demonstrates how computer algebra platforms like SINGULAR [42] can be used to find all equilibria of a four guild system with intra-guild predation. This particular system will be used as a running example throughout the chapter.

The working example is a four guild system with intra-guild predation. The system was randomly generated using William's niche model [122] with allometrically scaled parameters [10]. The system equations for this food web are

$$
\begin{aligned}
\dot{x}_1 &= r_1 x_1(1-x_1) - \tfrac{x_1 x_2}{1+x_1} - \tfrac{43}{64}\tfrac{x_1 x_3}{1+x_1+x_2} - \tfrac{13}{64}x_1 \\
\dot{x}_2 &= -\tfrac{43}{64}\tfrac{x_2 x_3}{1+x_1+x_2} + \tfrac{x_1 x_2}{1+x_1} - \tfrac{97}{256}\tfrac{x_2 x_4}{1+x_2+x_3} - \tfrac{13}{64}x_2 \\
\dot{x}_3 &= -\tfrac{97}{256}\tfrac{x_3 x_4}{1+x_2+x_3} + \tfrac{43}{64}\tfrac{(x_1+x_2)x_3}{1+x_1+x_2} - \tfrac{35}{256}x_3 \\
\dot{x}_4 &= \tfrac{97}{256}\tfrac{(x_2+x_3)x_4}{1+x_2+x_3} - \tfrac{5}{64}x_4
\end{aligned} \tag{41}
$$

This food web consists of a single producer (x_1), two primary consumer guilds (x_2 and x_3), and a secondary consumer guild (x_4). The trophic diagram for this system is shown in Fig. 19(a) where we can see that both primary consumer guilds compete for the same resource (the producer) with x_3 predating on x_2. This system exhibited a range of steady state behaviors that are not restricted to fixed point convergence. Fig. 19(b) shows a simulated forward orbit in which the system's qualitative behavior is that of a limit cycle.

	equilibrium vector	stability index
A	$[\,0.64\ \ 0.26\ \ 0.00\ \ 0.62\,]^T$	$(0,0,4)$
B	$[\,0.65\ \ 0.22\ \ 0.04\ \ 0.48\,]^T$	$(1,0,3)$
C	$[\,0.69\ \ 0.00\ \ 0.26\ \ 0.46\,]^T$	$(2,0,2)$
D	$[\,0.26\ \ 0.00\ \ 1.01\ \ 0.00\,]^T$	$(1,0,3)$
E	$[\,0.25\ \ 0.68\ \ 0.00\ \ 0.00\,]^T$	$(2,0,2)$
F	$[\,0.80\ \ 0.00\ \ 0.00\ \ 0.00\,]^T$	$(2,0,2)$
G	$[\,0.00\ \ 0.00\ \ 0.00\ \ 0.00\,]^T$	$(1,0,3)$

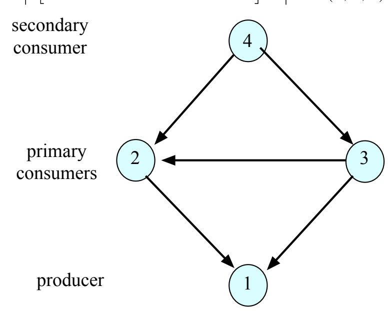

(a) Trophic relationships for 4 guild food web and system equilibria

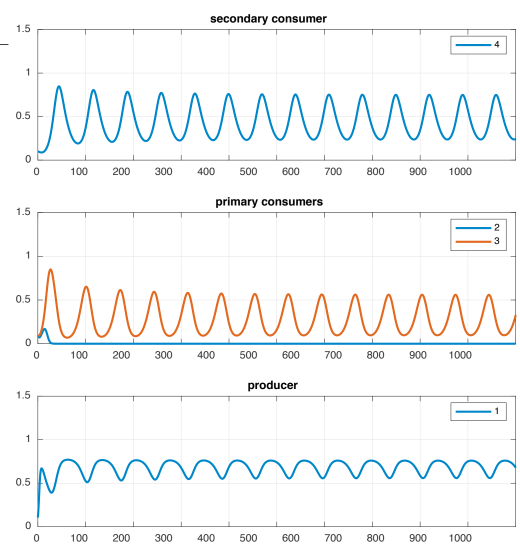

(b) State trajectories for a specified initial condition

FIGURE 19. Example Four Guild System

To use triangular decompositions to find all system equilibria, we first need to obtain a set of polynomial equations whose roots give the equilibria of the system in equation (41). This was done by introducing a new variable for each resource pool. The system in equation (41) has three distinct resource pools,

$$\mathcal{R}_1 = \{1\}, \quad \mathcal{R}_2 = \{1,2\}, \quad \mathcal{R}_3 = \{2,3\}$$

So there are three new variables $w = \{w_1, w_2, w_3\}$

$$w_1 = \frac{1}{1+x_1}, \quad w_2 = \frac{1}{1+x_1+x_2}, \quad w_3 = \frac{1}{1+x_2+x_3}$$

These variables are then used to transform equation (41) into a polynomial system in $\mathbb{R}[x,w]$. To complete the ideal, we then augment the preceding equations for the w variables to the original system equations. The ideal whose affine variety we wish to determine is therefore generated by the following polynomial equations

$$\begin{aligned}
0 &= F_1(x,w) = -\tfrac{13}{64}x_1 - x_1(x_1-1) - w_1x_1x_2 - \tfrac{43}{64}2_2x_1x_3 \\
0 &= F_2(x,w) = -\tfrac{13}{64}x_2 - \tfrac{43}{64}w_2x_2x_3 + w_1x_1x_2 - \tfrac{97}{256}w_3x_2x_4 \\
0 &= F_3(x,w) = \tfrac{43}{64}w_2x_3(x_1+x_2) - \tfrac{97}{256}w_3x_3x_4 - \tfrac{35}{256}x_3 \\
0 &= F_4(x,w) = \tfrac{97}{256}w_3x_4(x_2+x_3) - \tfrac{5}{64}x_4 \\
0 &= F_5(x,w) = w_1(1+x_1) - 1 \\
0 &= F_6(x,w) = w_2(1+x_1+x_2) - 1 \\
0 &= F_7(x,w) = w_3(1+x_2+x_3) - 1
\end{aligned} \tag{42}$$

Algorithms computing a triangular decomposition of a polynomial system's ideal [**79**] have been implemented in the computer algebra tool SINGULAR [**42**]. These codes were used to find all equilibria for the system (41) by finding the zeros of the algebraic equations (42). In particular, a MATLAB script was used to automate the generation of the system equations (41) and the associated polynomial equations (42). These equations were passed to SINGULAR, whose output was then evaluated in MATLAB to compute the stability index for the linearization at each equilibrium. This computation found seven non-negative real valued equilibria, labeled from A to G with their location in the state space and their stability indices given in the table in Fig. 19(a). These equilibria will be used in the next section as initial conditions for a search that builds the regime transition system for this particular system.

2. Regime Transition System

The regime transition system (RTS) may be informally described as a finite state machine whose logical states represent global invariant structures in the state space and whose edges represent the orbits connecting these structures. In this regard, the RTS may be seen as a discrete abstraction that maps out all of the system's possible regime shifts. The fundamental problem we face is the inference of this discrete abstraction for *global* behavior from a set of differential equations that are *local* representations of the system.

We know a great deal about a system's global behavior when its orbits are compact. The compactness assumption is not unreasonable for ecological systems since we know all compartmental states must be bounded. So let (X, ϕ) be a smooth dynamical system and assume there exists a compact invariant set, $S \subset X$. As discussed in section 4 of chapter 2, this system's chain recurrent set, $R(S)$, can be partitioned into a collection of mutually disjoint *basic sets*, $\{[R(S)]_i\}_{i=1}^m$. Or more formally

$$R(S) = \bigcup_{i=1}^{m} [R(S)]_i \text{ where } [R(S)]_i \cap [R(S)]_j = \emptyset \text{ for } i \neq j$$

Recall that we used this fact in Chapter 2 to define regime shifts as a transition between two different basic sets.

Perhaps the most important thing we know about smooth systems with compact orbits is that their basic sets can be partially ordered with respect to the connecting orbits between these basic sets. This last fact is a consequence of Conley's decomposition theorem [**14**]. The following statement of this theorem follows from [**77**]. Formal proofs for this theorem will be found in Conley's original monograph [**14**] and more recent textbooks on smooth dynamical systems [**30, 60**]

THEOREM 5. **(Conley's Decomposition Theorem)** *Let S be a compact invariant set for the smooth dynamical system (X, ϕ) and let $\{[R(S)]_i\}_{i=1}^n$ denote the basic sets of this system's chain recurrent set. Then there exists a function $V : S \to [0, 1]$ such that*

- *there exist constants $\sigma_i \in [0, 1]$ ($i = 1, 2, \ldots, n$) for which $V(x) = \sigma_i$ for all $x \in [R(S)]_i$.*
- *and $V(\phi(t; x)) < V(x)$ for all $x \in S - R(S)$ and for any $t > 0$.*

Note that the function, V, in theorem 5 is a Lyapunov-like function. Since this function is decreasing along forward orbits starting outside of the chain recurrent set, $R(S)$, we can classify the global behavior of these orbits in one of two ways. Consider an orbit $x(\cdot; p)$ with starting state p.

- If p is in one of the basic sets, $[R(S)]_i$, then the orbit $x(t; p) \in [R(S)]_i$ for all $t \in \mathbb{R}$. In other words, the basic sets are "traps" for orbits starting in the basic set.
- If $p \in S - R(S)$ (starts outside of the chain recurrent set) then there are basic sets $[R(S)]_i$ and $[R(S)]_j$ where $i \neq j$ such that $x(t; p) \to [R(S)]_i$ as $t \to -\infty$ and $x(t; p) \to [R(S)]_j$ as $t \to \infty$. In other words, the orbit for any p outside of a basic set connects two different basic sets. The set of all *connecting orbits* from basic set $[R(S)]_i$ to basic set $[R(S)]_j$ will be denoted as $C(i, j)$.

We therefore see that smooth compact systems have two distinct qualitative behaviors; they are either trapped in a basic set (regime) or they are moving between two regimes along a connecting orbit.

This decomposition of the orbits also establishes a partial order on the basic sets. In particular we say that $[R(S)]_i >_c [R(S)]_j$ if and only if $C(i,j)$ is non empty (i.e. there is a connecting orbit from $[R(S)]_i$ to $[R(S)]_j$). One can readily use the Lyapunov-like character of V in theorem 5 to prove that the binary relation $>_c$ partially orders the basic sets. Moreover, since there is a finite number of basic sets, there is also a "minimal" element of the set.

We now use these results to define a *regime transition system* to characterize the system's logical transitions between different regimes (basic sets). In particular, a regime transition system is the ordered triple $\Sigma = (\mathcal{B}, Q, \mathcal{B}_m)$ whose "logical " state space is $\mathcal{B} = \{B_1, B_2, \ldots, B_n\}$ and where the ith logical state (regime) is $B_i = [R(S)]_i$, the ith basic set. The map $Q : \mathcal{B} \to 2^{\mathcal{B}}$ is a set valued map such that $Q(B_i) \subset \mathcal{B}$ and $B_j \in Q(B_i)$ if and only if $C(i,j)$ is nonempty (i.e. there is a connecting orbit from $[R(S)]_i$ to $[R(S)]_j$). The set $\mathcal{B}_m \subset \mathcal{B}$ is a set of logical states (regimes) that are minimal with respect to the partial order $>_c$ induced by the original system's connecting orbits. A *regime sequence* is any finite sequence, $\{B_{i_j}\}_{j=0}^{n}$, of logical states in $\mathcal{B}$ such that $B_{i_{j+1}} \in Q(B_{i_j})$ for $j = 0, 1, \ldots, n-1$. The regime sequence will be said to be *accepted* by the regime transition system Σ if and only if the last regime (logical state) in the sequence, B_{i_n}, is a minimal logical state in $\mathcal{B}_m$.

The regime transition system Σ may be seen as accepting all regime sequences that can be generated through shock-induced regime shifts. In particular, if we know that the system's actual state, $x(t_0)$, at time t_0 is in basic set $[R(S)]_i$, we know it will stay in that basic set forever unless an impulsive disturbance "shocks" the system state out of this basic set. In particular, that shock will cause the system state to jump onto a connecting orbit between $[R(S)]_i$ and some other basic set $[R(S)]_j$. If this shock occurs, then the system state will converge to $[R(S)]_j$. Reaching the basic set $[R(S)]_j$ is the "threshold" that needs to be met before applying another shock that drives the system state onto another connecting orbit. In other words, we can use shocks to trigger regime shifts that force the system's "logical" state to traverse a path accepted by the regime transition system Σ. In this regard, the regime transition system provides a roadmap that an environmental resource manager can follow to steer a system towards alternative regimes.

Viewing the regime transition system in this way also provides a formal definition for ecological resilience. Informally, we said a system is ecologically resilient if it eventually reorganizes and recovers it nominal regime after a collapse. So let $B^* \in \mathcal{B}$ denote the logical state associated with this nominal regime. The system is ecologically resilient if B^* is the only minimal logical state in $\mathcal{B}_m$ and for every logical state $B_0 \in \mathcal{B}$ there is a regime sequence starting at B_0 that is accepted by the regime transition system Σ.

3. Algorithmic Construction of Regime Transition System

This section examines an *algorithmic* approach for constructing the regime transition system, $\Sigma = (\mathcal{B}, Q, \mathcal{B}_m)$, from the differential state equations, $\dot{x} = f(x; k)$, for a compact smooth dynamical system. Algorithmic methods use numerical integration of the system's state equations to detect global invariant structures [**36**] in the system's state space and then characterize those structures as computer data structures that can be more easily worked with. In our case, these algorithmic methods are used to detect basic sets that can be reached by shock-induced regime shifts from the system equilibria. The detected sets are then represented as cubical complexes that are *isolating neighborhood* for the basic set. An isolating neighborhood of an invariant set, U, is any set N such that U is the largest invariant set in N. The connecting orbits between these isolating neighborhoods are generated by numerically integrating the system equations to obtain sampled orbits originating in the isolating neighborhood of one basic set and terminating in the isolating neighborhood of another basic set. These isolating neighborhoods and their connecting orbits can be efficiently represented and manipulated in the computer as directed graphs, thereby allowing one to systematically construct a regime transition system whose regime sequences are all rooted in the system's equilibria.

For the regime transition system, Σ, constructed in this manner, the logical states $\mathcal{B}$ are derived from isolating neighborhoods, $\{N_1, \ldots, N_m\}$, of the basic sets $\{[R(S)]_1, \ldots, [R(S)]_m\}$. The opening description we gave for Σ identified the logical states with the basic sets. There are several reasons why constructing the regime transition using isolating neighborhoods is a better thing to do than trying to directly find the basic sets.

In the first place, the topology of a basic set is inherently sensitive to computer precision error. Isolating neighborhoods, on the other hand, are relatively insensitive to computational error. Moreover, the topology of the isolating neighborhood provides a lower bound (with respect to Betti numbers of the cubical complex) on the basic set's topology. Therefore the transition systems built from isolating neighborhoods will have topologies will be numerically stable objects.

In the second place, the transition system characterizing transitions between isolating neighborhoods is bisimilar to the transition system for the basic sets (provided the gridding is chosen correctly). This means that one can identify the regime transition system, $(\mathcal{B}, Q, \mathcal{B}_m)$, from the transition system discovered using the isolating neighborhoods. The algorithm that is used to discover the neighborhood's transition system is simply a graph search; that starts at the isolating neighborhoods for the equilibria and then systematically explores to find the isolating neighborhoods of other basic sets. Because we have methods from section 1 that can find *all* system equilibria, we can use this method to discover a large part of the regime transition system.

Finally, while the algorithmic approach is based on a gridding of the system's state space, the actual data that is stored are only those grid elements discovered by the computed orbits. In particular, the only information that is actually retained is the cubical complex representing the isolating neighborhood. This means the space complexity of the scheme is rooted in the underlying complexity of that basic set's topology; something that can be controlled through a judicious selection of the grid size.

There are many ways of constructing isolating neighborhoods for the basic sets. The easiest way is to introduce a cubical gridding of the state space and use sampled orbits of the system (obtained by numerical integration of the system equations) to generate a *symbolic trajectory* in this gridded space. This symbolic trajectory is represented as a directed graph whose vertices are drawn from the grid elements of the space and whose edges are the observed transitions between distinct grid elements. The isolating blocks are the cubical complexes associated with the largest strongly connected components of that graph.

This algorithmic approach was used with the four guild system in Fig. 19 to find isolating blocks and the connecting orbits between them. In particular, each system state, $x \in S$, was associated with a *cubical grid element*

$$\mathbf{g}_x = (g_1, g_2, \ldots, g_n)$$

whose components are integers

$$g_i = \left\lfloor \frac{x_i + dx_i/2}{dx_i} \right\rfloor$$

for $i = 1, 2, \ldots, n$ and where dx_i is the length of the ith side of the grid element. The system equations are numerically integrated from a specified initial conditions to create a *sampled* forward orbit, $\{x(kh)\}_{k=0}^{\infty}$, of the system. Since each state of that orbit has an associated grid element we can construct a sequence of grid elements $\left\{\mathbf{g}_{x(kh)}\right\}_{k=0}^{\infty}$. There is a lot of redundant information in this sequence, so we identify a subsequence $\{k_j\}_{j=0}^{\infty}$ of non-negative integers such that $k_0 = 0$ and $k_j = k$ if $\mathbf{g}_{x(kh)} \neq \mathbf{g}_{x((k-1)h)}$. This subsequence marks time instants when the sampled orbits shifts between two different cubical grid elements. We refer to the sequence $\left\{\mathbf{g}_{k_j}\right\}_{j=0}^{\infty}$ as a *symbolic orbit* of the system.

The symbolic orbit $\left\{\mathbf{g}_{k_j}\right\}_{j=0}^{\infty}$ is conveniently represented as a weighted directed graph $G = (V, E, w)$ whose vertex set, V, consists of the distinct cubical grid elements of the orbit. The edge set, $E \subset V \times V$, contains distinct single step transitions along that orbit. The weighting function $w : E \rightarrow [0, 1]$ maps each edge onto the fraction of time that the particular edge was traversed by the orbit. Fig. 20(a) shows the symbolic orbit generated by a sampled orbit of our four guild system. This sampled orbit $\{x(kh)\}$ is shown in blue and the grid elements comprising the symbolic orbit are the black cubes shown in the picture. The figure only plots the orbit in a phase space spanned by the two most active consumer guilds. But what this shows is that the sampled orbit is contained within the cubical complex (i.e. the set formed from the union of the cubical sets shown by the black squares). In other words, the cubical complex isolates the orbit from the rest of the flow if the grid size is sufficiently small.

As mentioned above the symbolic orbit, $\{\mathbf{g}_{k_j}\}_{j=0}^{\infty}$, constructed from the sampled orbit generates a directed graph on the grid elements shown in Fig 20(a). The strongly connected components of this graph are readily computed using Tarjan's algorithm [**114**] and the cubical complex associated with the largest connected component is an isolating neighborhood for

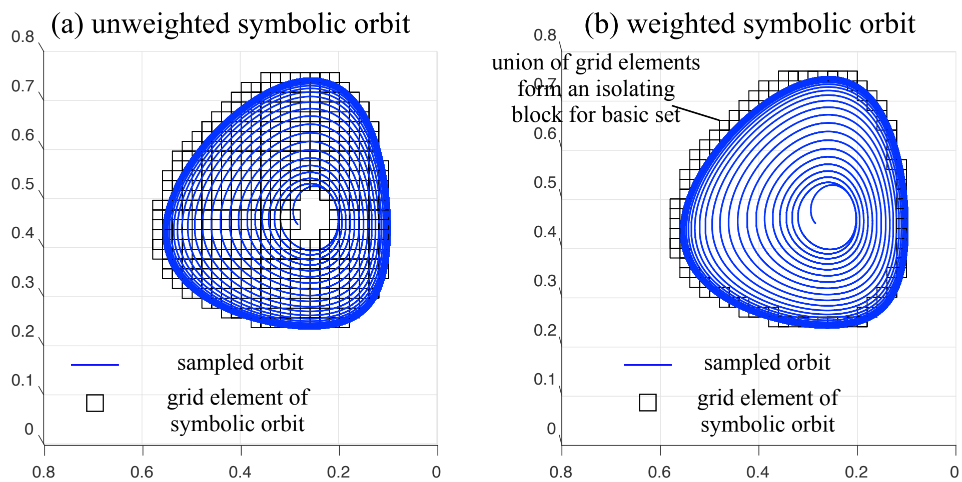

FIGURE 20. Symbolic orbits of four guild system and cubical complex isolating one the system's positive limit sets

a basic set of the system. This cubical complex is shown in Fig 20(b). This figure clearly shows that the cubical complex isolates the "limit cycle" that this orbit asymptotically approaches. The results portrayed in this figure therefore suggest that we can use these algorithmic methods to efficiently construct a finite length representation of the basic set whose size is governed in large part by the complexity of the basic set's topology.

Remark: Tarjan's algorithm finds all strongly connected components of the directed graph. We took the largest component as the one isolating the basic set. In practice, this may not be the best choice. Another possible approach would be to select that component whose topology (as measured by its Betti numbers) is persistent over a range of grid sizes.

The preceding discussion demonstrated how one can algorithmically identify cubical complexes that isolate the basic sets of a given system. In general constructing isolating blocks for all of the system's basic sets is difficult. But if the affine variety of the system is zero dimensional, then we can find *all* equilibria. Each of these equilibria is a basic set. Since Conley's theorem asserts the basic sets are partially ordered with respect to an order defined by their connecting orbits, this suggests we can use sampled orbits to find other basic sets by simply generating a sampled orbit, constructing its symbolic orbit, and finding the largest connected component for the symbolic orbit's directed graph. This procedure would identify all of those basic sets that can be reached through shock-induced regime shifts from the equilibrium points.

When this procedure was followed for the four guild system it generated the results shown in Fig. 21. In particular, Fig. 21(a), shows the cubical complexes isolating all basic sets of this system in a phase space that neglects the producer. This figure shows that there are eight isolating blocks. Seven of these blocks isolate the system equilibria and the eighth block isolates a limit cycle for the system. The directed graph associated with the transition system discovered by this algorithm is shown in Fig. 21(b). This figure readily shows that the transition system has only two minimal regimes in $\mathcal{B}_m$. These minimal logical states are the states that the system will eventually fall into. These are regimes A, which is a stable fixed point in which one of the primary consumers, x_3, is zero. The other minimal regime is the limit cycle regime H in which the other primary consumer, x_2, is zero.

Is this system "ecologically resilient"? If we take the minimal regimes A and H as being the "nominal" regimes that we wish to preserve, then the answer is yes. Because no matter how the system is perturbed, it will eventually return to one of these two regimes. Let us assume, however, that only one of these regimes, say A, is considered to be nominal. In that case, the system is not ecologically resilient because after collapse the system will return to A or H. In particular, for this system, there is a greater likelihood that regime H, rather than regime A, will be restored.

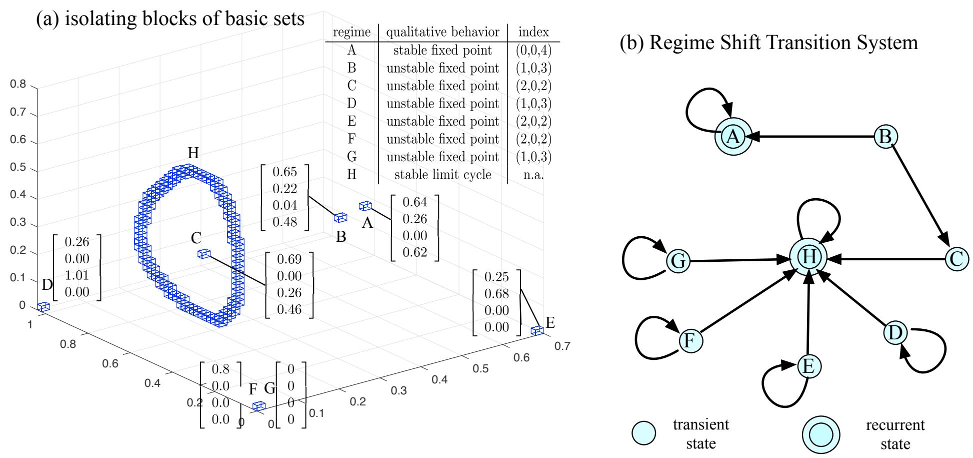

FIGURE 21. Isolating blocks for 4 guild system and its regime shift transition system

4. Managing Regime Shifts for System Restoration

A system's regime transition system, $(\mathcal{B}, Q, \mathcal{B}_m)$ is a finite abstraction of the original system. It is bisimilar to the original system in the sense that every symbolic orbit accepted by the system maps onto orbits in the original system, and every orbit of the original system has a symbolic orbit. Since the discrete states of the regime transition system represent distinct basic sets of the original system, this means that supervisory control of the transition system essentially manages regime shifts of the original process. This section considers the development of supervisory policies that ensure the original process has ecological resilience. In particular, we briefly describe two such policies; a shock-induced regime shift strategy and a bifurcation-induced regime shift strategy.

Shocked-Induced System Restoration: Consider the regime transition system, $\Sigma = (\mathcal{B}, Q, \mathcal{B}_m)$. Let $B^* \in \mathcal{B}$ denote the nominal logical state (regime) of the system. The original system (X, ϕ) is ecologically resilient if and only if the $\mathcal{B}_m = \{B^*\}$ and for any logical state $B_0 \in \mathcal{B}$, there is a regime sequence accepted by Σ. As noted above if $\mathcal{B}_m$ contains B^* as well as other logical states, then the system is not ecologically resilient. But if we can control the system's regime shifts so that only those regime sequences terminating in B^* are actually generated by the original system, then we can force the restoration of the nominal regime.

Let us see how this might be done for the four guild system in equation (41). The system's regime transition system, Σ, was constructed in the preceding section and is shown in Fig. 21(b). One can readily see that Σ's set of minimal logical states is $\mathcal{B}_m = \{A, H\}$ in which A is a stable fixed point and H is a limit cycle. Let us assume that the A is chosen as the nominal regime because it is a stable fixed point. This system will not be ecologically resilient because H is also in $\mathcal{B}_m$.

This system is not ecologically resilient, but we can force the restoration of the nominal regime by adopting a rather simple policy that waits until the system state reaches one of the logical states in $\mathcal{B}_m$. If that terminal state is the nominal regime, then we are done and the system has been restored. If that terminal state is not the nominal regime, then we intervene using a shock-induced regime shift that causes the system state to jump into the ROA of the basic set containing the nominal regime.

This policy is rather easy to enforce for our example system. Assuming the system starts in its nominal regime, A, we force a system collapse at $t = 200$ with an external disturbance that causes guild 2 to collapse. The collapse of guild 2 also causes a collapse in guild 4, so the system transitions from regime A to regime F, after which the system enters the limit cycle regime H as shown in Fig. 22.

Let us now consider two different restoration policies. The first policy simply restocks guild 2 at $t = 600$ after the system has entered the minimal

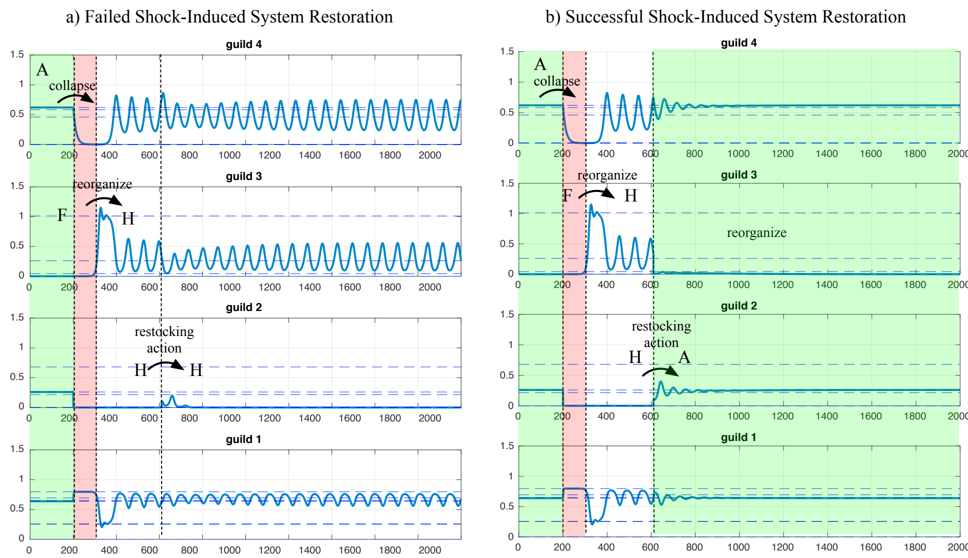

FIGURE 22. Example of Successful and Unsuccessful Shock-Induced Restoration of System

regime H. The outcome of this policy is shown in Fig. 22(a) where we see that the policy fails to restore the nominal regime A. The reason for this failure is that the restocking operation did not completely force the system into a logical state from which it could reach the nominal regime.

So we now consider a second policy that does restore the nominal regime. This policy restocks guilds 2 and 3 after the system has entered regime H. The outcome for this policy is shown in Fig. 22(b) which shows recovery of the nominal regime. The reason why the second restoration policy was successful was because its restocking action forced the system state into the isolating neighborhood for regime A.

This example shows that it is possible to use managed shock-induced regime shifts to restore a system's nominal regime, even if that system is not ecologically resilient. The necessary precondition for this approach is that the nominal regime must lie in terminating set $\mathcal{B}_m$ of the regime transition system, Σ. There are other policies that can be used as well. For instance, if the system had collapsed to regime, B, then the directed graph for Σ shown in Fig. 21 shows that we can reach the nominal regime by simply adopting policies that prevent transitions from B to C. In other words, we are adopting a simple supervisory control scheme for the transition system Σ to prevent transitions from which the nominal regime cannot ever be reached. In other words, the problem of regime shift management has been reduced to a well studied problem in the supervisory control of discrete-event systems [**91, 68**].

Bifurcation-Induced System Restoration: From the preceding discussion, it should be clear that the success of a shock-induced restoration policy relies on the desired nominal regime being in $\mathcal{B}_m$. If this is not the case then the system cannot be rendered ecologically resilient through shock-induced regime shifts. This subsection shows that it may still be possible to restore the nominal regime, but this would be done bifurcation-induced regime shifts.

Recall that a bifurcation-induced regime shift occurs when a change in the system parameters triggers a local bifurcation in one of the system equilibria. Since we know the equilibria for the nominal system, we can use the methods in chapter 3 to determine how large of a perturbation in the system parameters would trigger a bifurcation which makes a desired "nominal" regime recurrent. We demonstrate how this might be done for our example system.

So consider the four guild system in equation (41) and let us assume that the desired nominal regime is regime C. The system's transition system in Fig. 21 shows that regime C is not in $\mathcal{B}_m$ and so the system's nominal regime cannot be restored through shock-induced regime shifts. A robust stability analysis of all equilibria shows that a small reduction of the producer's carrying capacity will destroy regime B's equilibrium and change the stability index of regime C's equilibrium to that of a stable fixed point. Fig. 23(a) shows the isolating blocks for the perturbed system from which we see that the change in the stability indices of the equilibria causes the disappearance of regime's B and regime H (limit cycle). This change in stability indices would mean that regime C is recurrent and must be in the terminating set of states, $\mathcal{B}_m$. This assertion is readily verified by computing the regime transition system for the perturbed system. The directed graph for this perturbed transition system is shown in Fig. 23(b) where we can readily see that regime C is the only logical state in $\mathcal{B}_m$. So not only does a bifurcation-induced regime shift restore the nominal regime, but it also renders the entire system ecologically resilient.

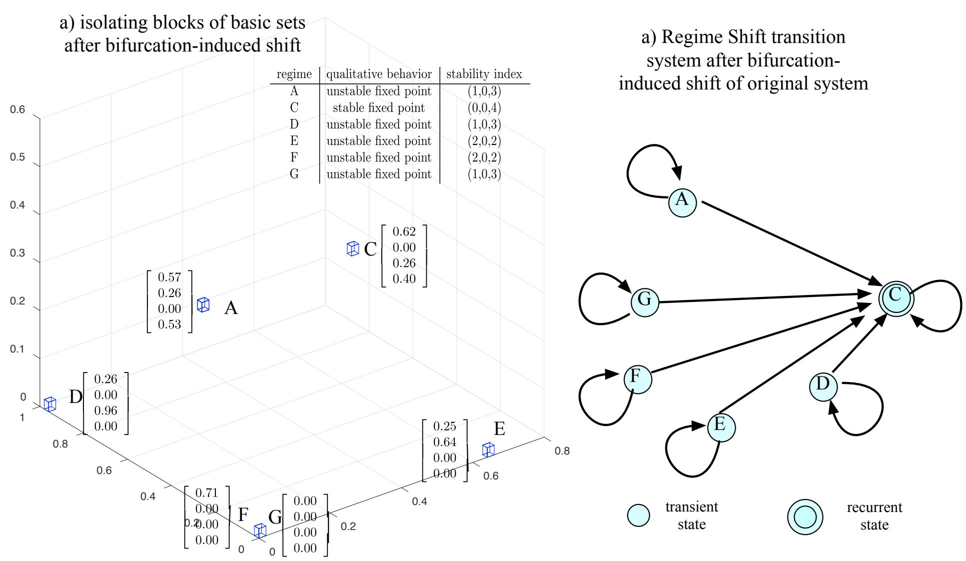

FIGURE 23. Example of Bifuracation-induced Restructuring of Regime Shift Transition System

5. Summary and Further Readings

This chapter provided a formal definition of ecological resilience in terms of the regime transition system $\Sigma = (\mathcal{B}, Q, \mathcal{B}_m)$ of a smooth dynamical system with compact orbits. The transition system is a discrete abstraction of the original system's global behavior with respect to the possible transitions between the system's basic sets. It essentially provides a roadmap of the pathways a system can follow through shock-induced regime shifts and in this way provides a basis for the supervisory control of ecological resilience. This chapter showed that the regime transition system can be constructed from the system's state equations through the use of algorithmic methods of used for computer analysis of dynamical systems. We saw that Conley's decomposition theorem guarantees the existence of such a transition system and that once determined the transition system could be used to restore a collapsed system through supervisory control of its shock-induced and bifurcation-induced regime shifts.

Restoration ecology is a branch of ecology that focuses on developing plans for the restoration of degraded ecosystems. Much of the work in restoration ecology focuses on specific restoration projects for degraded terrestrial systems. The poor track record of these restoration project was documented in [**63**]. Recent work has begun developing a theoretical framework for restoration ecology that restores degraded ecosystems by following a sequence of alternative states. [**53, 108, 52, 109**]. The results in this chapter should be relevant to that recent work by providing a systematic method for identifying the sequence of alternative states leading to restoration of the nominal system.

Algorithmic methods for the analysis of dynamical systems uses the computer to help prove properties about a dynamical system. The method used in this chapter to detect basic sets of the system is similar to that used in

[36]. The formal foundations establishing the correctness of this method using chain recurrence concepts was laid out in [78, 65]. The decision to use cubical complexes, rather than simplicial complexes, was based on experiments we did using both that supported the arguments made in [64] regarding cubical complexes.

The justification for the use of algorithmic methods draws heavily on results providing global characterizations of the orbits generated by smooth dynamical systems. The concept of chain recurrence is central to these stating and understanding these results. Further readings on chain recurrence may be found in [30] for discrete-time systems and [60]. Both of these books may be seen as expanding on Conley's original monograph [14] that extended Morse theory [76] to dynamical systems. The underlying algorithmic nature of chain recurrent concepts was detailed in [65]. Conley's decomposition theorem originates in [14], but there are a number of different statements of his main result in [30, 60, 77]. This chapter used the statement in [77] for its emphasis on the Lyapunov-like function V.

The idea of building discrete-abstractions characterizing the logical behavior of dynamical systems has its origins in the study of *hybrid systems* [44]; systems that combine both smooth dynamical systems and discrete-event systems. Hybrid systems may also be seen as a special class of *cyber-physical system* [3]. The idea that a smooth system could be abstracted into a finite state machine was introduced in [2, 106]. The use of such abstractions in the supervisory control [91] of smooth systems was reviewed in [67].