NEW GEOMETRIC AND TOPOLOGICAL METHODS OF ANALYSIS OF THE GLOBAL NONLINEAR CONTROL PROBLEM

by

Efthimios Kappos

Memorandum No. UCB/ERL M96/42

1 July 1996

NEW GEOMETRIC AND TOPOLOGICAL METHODS OF ANALYSIS OF THE GLOBAL NONLINEAR CONTROL PROBLEM

by

Efthimios Kappos

Memorandum No. UCB/ERL M96/42

1 July 1996

ELECTRONICS RESEARCH LABORATORY

College of Engineering
University of California, Berkeley
94720

New Geometric and Topological Methods of Analysis of the Global Nonlinear Control Problem

Lecture Notes for a Series of
Talks on Nonlinear Control
given at the EECS Department,
U.C. Berkeley, May–June 1996.

Efthimios Kappos*

* EECS Department,
University of California, Berkeley CA 94720,
on leave from
the School of Mathematics and Statistics,
University of Sheffield, England.
This work was supported by the UK EPSRC Grant GR/K83205.

July 1, 1996

Contents

1 Introduction

We begin by giving the form of the nonlinear control problem that we intend to study. It is expressed in terms of a pair, consisting of a vector field and a distribution in a manifold. It is then possible to state the aim of this work, namely to give an analysis of the global dynamics that arise as controlled dynamics of nonlinear control systems.

A short section deals with the comparison between differential-geometric approaches (using the Lie bracket as the main tool) and our more geometric/topological one. It is noted that, away from equilibria, the controllability question is far from solved in the differential-geometric setting.

We close this introductory part with a brief description of the novel tools we shall employ and an outline of the main results.

1.1 Preliminaries: Control Pairs and Controlled Dynamics

In control theory, we study *control-affine systems* of the form

$$\dot{x} = f(x) + g(x)u = f(x) + \sum_{i=1}^{m} u_i g_i(x) \tag{1}$$

on an n-dimensional *state space*; this is usually $\mathbf{R}^n$, but is often a manifold, say M^n. The control actions u_i are assumed unconstrained, for the moment. Let us also assume that the span of the control vector fields $g_1, \ldots, g_m$ is m throughout.

We can move to a more compact global, invariant form for the above control system by making the following definition:

Definition 1 *A* **control pair** (X, D) *consists of a smooth vector field* $X \in \mathcal{X}(M^n)$, *called the* **state dynamics**, *together with a constant rank subbundle* $D \subset TM^n$ *called the* **control distribution.**

(By *'smooth'*, we shall usually mean C^2 –twice continuously differentiable.)

We recall next the definitions of some elementary objects from global analysis; this will serve mainly the purpose of fixing notation.

1.2 Manifolds, Vector Fields and Flows

A **manifold** M^n is a Hausdorff topological space that is locally homeomorphic to $\mathbf{R}^n$ (via charts (U_i, ϕ_i)) and is such that the composition of charts is smooth. The **tangent space** TM^n of M^n is defined in a number of ways, all equivalent; in the simplest case, one imagines the manifold to be embedded in some $\mathbf{R}^N$, in which case the tangent space T_xM^n at every point is easy to define. One then forms TM^n set-theoretically as the disjoint union of the point-wise tangent spaces and gives it a natural topology making it a manifold of dimension $2n$. In fact TM^n is a **vector bundle** over M^n (a sort of collection of vector spaces smoothly parametrised by the points of the manifold.) We write $\pi : TM^n \to M^n$ for the natural projection: $(x, v) \mapsto x$.

A **vector field** X on M^n is a smooth *section* of TM^n, in other words a smooth assignment of an element of T_xM^n for every $x \in M^n$ (the condition for a section is that X maps $M^n \to TM^n$ such that $\pi \circ X(x) = x$ for all x.) One writes $\mathcal{X}(M^n)$ for the set of vector fields on M^n. The notation $\Gamma(TM^n)$ is also common and is more general.

Vector fields correspond to systems of first-order ordinary differential equations. In fact, one proves the existence of integral curves of vector fields, i.e of curves $c : [t_0, t_1] \to M^n$ such that

$$\dot{c}(t) = X(c(t))$$

for all $t \in [t_0, t_1]$ by using local charts and appealing to the fundamental existence theorem for odes.

If we assume that solution curves exist for all times (if X is *complete*), we can then define the **flow** $\phi(t, x)$ as the solution curve at time t, starting from the point x at time 0. It is convenient sometimes to write $\phi_t(x)$ for this or even $\phi^X(t, x)$ if more than one flow is being considered and we want to make clear the vector field that gives rise to the flow.

We have the fundamental group property

$$\phi_{t_1}(\phi_{t_1}(x)) = \phi_{t_2}(\phi_{t_1}(x)) = \phi_{t_1+t_2}(x)$$

which implies, since $\phi_0 = \mathrm{id}$, that the maps $\phi_t : M^n \to M^n$ are diffeomorphisms.

Definition 2 *A* **sub-bundle** *(or* **distribution***) $D \subset TM^n$ of constant rank $m \leq n$ is a collection of m-dimensional subspaces of the tangent spaces $D_x \subset T_xM^n$ that depends smoothly on x. (It is possible to choose, locally, a smooth basis for D consisting of m pointwise linearly independent vector fields.)*

The distribution D is in fact another *vector bundle* over M^n. The set of all (smooth) sections of D is denoted by $\Gamma(D)$.

1.3 Formulation of the Main Aim of Global Control Design

In the context of a control pair, $\Gamma(D)$ is the set of all smooth, **global feedback controls.** We write

$$U : M^n \to D$$

for such a feedback control, so U maps $x \mapsto (x, U(x))$, with $U(x) \in D_x$.

Definition 3 *For every choice of $U \in \Gamma(D)$, the vector field $X + U$ will be called the* **controlled dynamics** *corresponding to U.*

One should compare at this point the traditional viewpoint of choosing smooth functions $u_i(x)$ and forming the system

$$\dot{x} = f(x) + \sum_i u_i(x) g_i(x).$$

(We implicitly assume that U is complete as well, so that $X + U$ is complete and it defines a global flow.)

Now as we move in the space of sections $\Gamma(D)$, we obtain dynamics whose topological type (in the sense of definition 23 below) varies, in general. If for some U, the controlled dynamics is structurally stable (see definition 24), then nearby controls will not change the topological type. However, 'large' changes in the control will, typically, lead to global changes in the dynamics.

The most general question one can ask about control systems is then to classify the equivalence classes of dynamics achievable through control and to obtain partitions of the set of all controls $\Gamma(D)$ according to the topological type of the flow of $X + U$. This is an even more difficult problem, in principle, than the *fundamental problem of dynamical system theory*, which seeks to classify all classes of flows on any given manifold.

The problem in the generality just given is fortunately of little immediate relevance to control practitioners. In applications, one asks a simpler question, namely: is it possible to achieve controlled dynamics that have certain dynamical features? In the simplest case, we ask whether it is possible to 'simplify' the dynamics so that there is a simgle asymptotically stable equilibrium, *at a fixed location, or neighborhood, in the state space.*

We can summarise this discussion by stating what we take to be the main object of the qualitative theory of control systems:

Main Aim of Geometric Control Theory: *Study the controlled dynamics $X + U$ as we vary the control U and, in particular, examine whether it is* *possible to select a U so as to achieve dynamics that have certain desirable features (such as global stability in the simplest case.)*

In fact the specification of less trivial dynamics is an interesting problem that we shall analyse in more detail later.

Warning: We are primarily interested in the case where the state dynamics X is *not identically zero!* If indeed the control pair degenerates to just a specification of a control distribution, then our subject has a very different flavour.

In the smooth, or continuous case, this is essentially a problem in differential geometry and the control theoretic dimension is not very noticeable. Relevant aspects are:

1. **Involutivity/Integrability: (Frobenius' Theorem and Foliations)** Given vector fields (controls) $X, Y \in \Gamma(D)$, their Lie algebra bracket $[X, Y]$ is of relevance to control because even if $[X, Y] \notin \Gamma(D)$, control laws can be found that push the state along the direction of the bracket. Thus, as is well known, one is interested in forming the smallest Lie algebra generated by $\Gamma(D)$, call it $\mathcal{C}$. This is then an involutive distribution and, assuming its rank is constant, gives rise to a **foliation** of state space. (The rank neeed not be constant, of course, in which case we appeal to Sussmann's generalization of Chow's theorem.)

   The consequence for control is that starting from a point on some leaf of the foliation, the state is contrained to lie on that leaf for all times.

2. **Relaxing the notion of control:** In much of classical control theory, one finds that it is essential that we specify the space of our control functions. In the case of the control distribution D, we find that we must allow, say, piecewise smooth, or analytic controls.

Let us give an example at this point ilustrating the usefulness of the Lie bracket in the context of zero state dynamics.

Example 1: (due to Brockett, see [3], p.182)

$$\begin{aligned} \dot{x}_1 &= u_1 \\ \dot{x}_2 &= u_2 \\ \dot{x}_3 &= x_2 u_1 - x_1 u_2 \end{aligned}$$

Claim 1 *There is no continuous stabilizing feedback law near* 0. *However, the local controllability condition is satisfied and every state can be driven to the origin using a smooth control function.*

The nonexistence of a continuous stabilizing feedback control follows from the Krasnosel'skii-Brockett's necessary condition; this is because, by continuity of the distribution $D = \text{span}\{g_1, g_2\}$ and since at 0 there is no x_3-component, the Gauss map cannot be onto in a small neighborhood of the origin (see Section 7 for an explanation of the terms used.)

To see that the local controllability condition is satisfied, rewrite the system as

$$\begin{pmatrix} \dot{x}_1 \\ \dot{x}_2 \\ \dot{x}_3 \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \\ x_2 \end{pmatrix} u_1 + \begin{pmatrix} 0 \\ 1 \\ -x_1 \end{pmatrix} u_2 = g_1 u_1 + g_2 u_2 \qquad (2)$$

and note that $[g_1, g_2] = -2\frac{\partial}{\partial x_3}$.

It is not difficult to visualize how it is possible to choose smooth feedback control in the open sets $\{x_3 > 0\}$ – (x_3 axis) and $\{x_3 < 0\}$ – (x_3 axis) to drive a state towards the origin monotonically in x_3. However, one sees that then on the $x_3 = 0$ plane (minus the origin, of course), we must have chattering, since either g_1 or g_2 has a nonzero x_3 component. If we give up the requirement that the control be of feedback type, we can of course steer a state on the x_3–axis to 0 simply by first pushing it off this plane and then proceeding as described above. Also note that, on the x_3-axis, to drive a state to zero we must first drive it off the axis, then let the previous control act. (Try to sketch a picture of this.) Finally, it is worth pointing out that we made no real use of the extra direction provided by the bracket *locally*; globally though, we used the fact that we can reduce the x_3 component at the same time as reducing some 'norm' in the (x_1, x_2)–plane (distance from the x_3-axis.)

1.4 Limitations of the Lie bracket in the case of nontrivial state dynamics

In the presence of nontrivial state dynamics (X not identically zero), the flavour of our subject changes considerably.

Definition 4 *The affine subspace $X(x) + D_x \subset T_x M^n$ will be called the* **control indicatrix** *at x.*

This is sometimes called the set of *associated vector fields* of the control pair –see [22], p.74. (Note that Assumption 3.1(a) in [22], p.73 is not satisfied if $X(x_0) \neq 0$ and $X(x_0) \notin D_x$. It is satisfied, though, at a zero of the state dynamics, which is the case that is treated in most presentations of the controllability problem using a *differential geometric* approach.)

The object we obtain by pasting together these spaces over all the points of M^n is *not a vector bundle*. The *fibre* is either a vector subspace or an affine one; we may choose to call such an object an *affine bundle*; however, one does not have available a well-developed theory of affine bundles and we can only usefully categorise it as a **fibration**. (A *fibration* is a map $p: E \to M$ of topological spaces that is onto and has the *Homotopy Lifting Property* for all topological spaces Y that are mapped to the base through a map f: $Y \xrightarrow{f} M$ –see [26], p.29.) In Part II, we shall present another fibration associated to that of the control indicatrix that is of crucial importance to the solution of the general problem of control.

Now it is useful to have available explicit **'negative' results** concerning nonlinear controllability/accessibility. This is especially important since we are trying to motivate an approach to nonlinear control theory that is not dependent on the formation of Lie brackets. The theorems below will assert that, if we are at a general *nonzero* point for the state dynamics, then achieving Lie bracket directions is constrained by the fact that we cannot go backwards along the state dynamics.

The first result points out a crucial difference between systems where the state dynamics is zero and those with nontrivial state dynamics. It says that a *very large control action is necessary if we are to have a chance to go backwards along the state vector field* **locally** (assuming this is possible, of course.) In other words —and this is the form the theorem takes, if the control action is bounded (as is the case in practice), then we cannot locally reach an open subset of some neighborhood of our starting point.

In the case of an *integrable control distribution* this subset can be given explicitly: its boundary is a union of leaves of the foliation of D. This is our second result. If D is of codimension one, then the boundary is exactly the leaf of D through the starting point. If the codimension is greater than one, then there is more freedom in choosing the boundary. (We assume M^n is Riemannian, i.e. it has been provided with some Riemannian metric; thus it makes sense to talk about norms and orthogonality.)

Theorem 1 *Suppose that at a point $x \in M^n$ the state dynamics is nonzero, $X(x) \neq 0$ and is such that $X(x) \notin D_x$ (thus also $D_x \neq T_xM^n$.) Suppose given a compact subset $\mathcal{U} \subset \mathbf{R}^m$.*

Then, if we constrain the controls u to lie in $\mathcal{U}$ for all x (equivalently, if we take a subset $\mathcal{G}$ of $\Gamma(D)$ such that U has sup norm bounded by some K for *all $U \in \mathcal{G}$), there is a neighborhood N of x and an open subset $N' \subset N$ such that $x \in \overline{N'}$ and any trajectory of controlled dynamics $X + U$ with $u(x) \in \mathcal{U}$ that is contained in the neighborhood N misses the open subset N'.*

The second theorem does not impose a contraint on the control, but requires D to be integrable.

Theorem 2 *Suppose that at a point $x \in M^n$ the state dynamics is nonzero, $X(x) \neq 0$ and is such that $X(x) \notin D_x$. Also suppose D is integrable.*

Then there is some neighborhood N of x and a local hypersurface $\Sigma \subset N$ through x that is invariant under D and divides N into two open sets N_+ and N_- with common boundary Σ (so that $N = \overline{N_+} \cup \overline{N_-}$) and are such that N_+ is positively invariant for all controlled trajectories contained in N.

Remark: The conclusions of these theorems hold *irrespective of whether any Lie algebraic rank condition holds, such as the condition that the span of the $\{\mathrm{ad}_X^k D\}_{k\geq 0}$ is of rank n at x.*

Proof of Theorem 1:

The first step is to set up a nonzero vector field and obtain a flow box for it. Decompose $TM^n = D \oplus D^{\perp}$; hence write $X = X_1 + X_2$, with $X_1 \in D$ and $X_2 \in D^{\perp}$. By assumption, $X(x) \notin D_x$, so $X_2(x) \neq 0$ and hence is nonzero locally in some neighborhood N_0 of x.

Pick a flow box $\psi : N_1 \to W \subset \mathbf{R}^n$ ($N_1 \subset N_0$) so that $\psi(x) = 0$ and $Y = \psi_*(X_2) = (0, \ldots, 0, 1)^T$. Choose an orthogonal basis for $\mathbf{R}^n$ so that $b_n = Y(0)$ and $\psi_*(D_x) = \text{span}\,\{z_1, \ldots, z_m\}$ (remember $X_2 \in D^{\perp}$.) Write $z_1, \ldots, z_n$ for the coordinates in thisbasis. Thus $dz_n(Y)(0) = 1$. Then the one-form dz_n pulls back to an exact form$\alpha = d\beta$ in N_1 and $\alpha(X_2)(x) = 1$ ($\beta = \psi^*(z_n)$.) The level sets z_n constant pull back to leaves of the local foliation of α.

We compute $\mathcal{L}_{X+U}\beta$:

$$\mathcal{L}_{X+U}\beta = \alpha(X + U) = \alpha(X) + \alpha(U). \tag{3}$$

Now $\alpha(X)(x) = \alpha(X_2)(x) = 1$ and $\alpha(U)(x) = 0$ since $D_x \subset \ker\, d\beta$ by construction. Thus $\mathcal{L}_{X+U}\beta(x) = 1 > 0$.

Claim 2 *There is some neighborhood N of x such that $\alpha(U)(x) > 0$ for all $x \in N$.*

This is because $\alpha(D_x) = 0$ and the compact K–ball in the fibres of D gives a continuous non-negative function $\gamma : N_1 \to \mathbf{R}$: $y \mapsto \sup\{|\alpha(v)|\ ;\ v \in B_K(0) \subset D_y\}$. One then takes $N = \gamma^{-1}[0, 1) \cap N_1$.

For $y \in N$ we have $\alpha(U)(y) < 1$ and so $\mathcal{L}_{X+U}\beta > 0$. This completes the proof of Theorem 1, since one easily checks that the set N' in the statement of the theorem is the set where $\beta < 0$. $\square$

Remark: As we take larged and larger controls, the neighborhood N will shrink and it is possible that controls will exist driving us downwards with respect to the local Lyapunov function β. Let us look at Example 1 again, this time introducing some nontrivial state dynamics.

Example 1(b): Suppose we have the system

$$\begin{pmatrix} \dot{x}_1 \\ \dot{x}_2 \\ \dot{x}_3 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} + \begin{pmatrix} 1 \\ 0 \\ x_2 \end{pmatrix} u_1 + \begin{pmatrix} 0 \\ 1 \\ -x_1 \end{pmatrix} u_2 = f + g_1 u_1 + g_2 u_2 \tag{4}$$

so we have constant state dynamics pointing 'upwards.'

Following the procedure given in the above proof, we take the function $\beta = x_3$. Then $\alpha = \begin{pmatrix} 0 & 0 & 1 \end{pmatrix}$ and we compute $\mathcal{L}_{X+U}\beta = \dot{x}_3 = 1 + x_2 u_1 - x_1 u_2$. Taking our starting point to be $x = \begin{pmatrix} 0 & 0 & 1 \end{pmatrix}^T$, we have $\dot{\beta}(x) = 1 > 0$, which stays positive for nearby points provided the control is bounded.

Note that the 'unbounded' feedback control $u_1 = 0$, $u_2 = \frac{1}{x_1}$ gives $\dot{\beta} = 0$ and so we can achieve $\dot{\beta} < 0$ by an even larger control. The reader should compare this situation with that of Example 1 itself, where a 'small' control action was enough to steer downwards with respect to the x_3 direction.

Proof of Theorem 2:

The basic idea is the same. We decompose $X = X_1 + X_2$, with $X_1 \in D$ and $X_2 \in D^\perp$ and use a local flow box $\psi : N_0 \to W \subset \mathbf{R}^n$ for X_2 near x. The m-dimensional foliation of the integrable distribution D pushes forward to a foliation of W. Now the orthogonal complement $D^\perp$ is a distribution —not integrable in general, but which can give us local transverse sections for the foliation of D.

Thus, consider Σ_0 to be a local $(n-m)$-dimensional manifold transverse to the distribution D and orthogonal to it at x, so that $X_2(x) \in \Sigma_0$. It intersects the $z_n = 0$ plane along an $(n-m-1)$-dimensional manifold Σ_1 and the leaves through Σ_1 give the desired D–invariant Σ.

Theorem 2 follows, since $dz_n(Y) > 0$ in W. $\square$

Remarks: Theorems 1 and 2 say that, locally near a point where $X \neq 0$, we cannot use the extra directions arising from the Lie brackets of X and the basis vector fields of D. The question is then: under what conditions are these extra directions available?

One instance where this is possible is in the vicinity of an equlibrium point (again we emphasize that then Theorems 1 and 2 do not apply.) The reader should perhaps go back at this point to the instances where Lie bracket conditions lead to local accessibility and notice that they either apply in a neighborhood of an equilibrium point or to the case when $X \in D$ (e.g Assumption 3.1(a) in [22], p.72.)

Given our geometric viewpoint, however, we can give an appreciation of the *global problem* we shall be concerned with: even though *locally* we are constrained by the geometry of the control pair, *globally* these objects have nontrivial 'twisting'. In the case of an equilibrium, for example, the state dynamics will be used to control in the complement of the control distribution. This can be best seen, at this point, in the following linear example.

Example 2:

Consider the two-dimensional linear system

$$\begin{aligned} \dot{x}_1 &= -x_1 + u \\ \dot{x}_2 &= x_1 + x_2 \end{aligned}$$

It is controllable, since

$$Ab = \begin{pmatrix} -1 \\ 1 \end{pmatrix}.$$

However, let us pick a point, such as $x = (0,1)^T$, where $Ax \neq 0$ and $Ax \notin D = \text{span}\{b\}$. It should be clear that the 'extra direction' given by the Lie bracket is not available at x. In fact, as we expected from theorem 2, in a neighborhood of x, it is only possible to move 'upwards' –i.e. no trajectory can enter the lower half of a local plane.

We know it should be possible to steer x to zero, but the question is how? In this simple two-dimensional case, it is easy to see that what we should do is steer to the left until the state dynamics starts pointing downwards and then drive towards the origin. Thus any stabilized trajectory through our chosen x must take a large 'detour' before heading towards zero. An explanation of this, together with a re-interpretation of the linear stabilizability/controllability conditions will be given in a later section.

We close this section by giving another example where the formation of the accessibility Lie algebra $\mathcal{C}$ is not meaningful when the state dynamics is nonzero.

Example 3:

Consider the planar system

$$\begin{aligned}\dot{x}_1 &= x_2^2 \\ \dot{x}_2 &= u\end{aligned}$$

The vector fields f and g are $f = x_2^2 \frac{\partial}{\partial x_1}$ and $g = \frac{\partial}{\partial x_2}$ and so forming brackets we get

$$[f,g] = -2x_2 \frac{\partial}{\partial x_1}$$

and

$$[[f,g],g] = 2\frac{\partial}{\partial x_1}$$

and so the system satisfies the local accessibility condition for any x_0. However, we can see easily that, since $\dot{x}_1 \geq 0$ everywhere, if we start at the point, say $(1,1)$, the reachable set is a half-plane and thus the system is not stabilizable to the origin.

Finally, here is a simple exercise:

Exercise 1 *Given two vector fields X and Y =constant, suppose that the direction given by the vector field $\frac{\partial}{\partial x_1}$ is not in the span of X and Y at a point x_0.*

Give conditions on the components of the vector fields X and Y so that their bracket $[X,Y]$ has a nonzero $\frac{\partial}{\partial x_1}$ component at x_0. What can you say about the germ of the vector field X at x_0? (e.g. can X be zero in any neighborhood of x?)

2 The Relevance of Global Nonlinear Dynamical Control

Despite the fact that so far there are few instances where nonlinear control theorists are asked to design dynamics more complicated that a single asymptotic attractor, global feedback control is potentially useful in several areas, some of which are presented below.

1. *Using control to build systems with multiple equilibria or limit cycles.*

   Given a nonlinear system in state space form such that certain components admit 'control action' (say a nonlinear circuit or a power system with control), it may be desirable to design nontrivial dynamics such as, for example, multiple equilibria or limit cycles. One may think of the multiple equilibria as 'memory locations' and of stable limit cycles as oscillators.

The theory we shall present is ideally suited to such problems since it is based on a description of dynamics through functions and hence can be used as a design method for achieving dynamics of a particular qualitative type.

2. *Control of complicated dynamics and the use of control to 'simplify' them, eg to turn a chaotic attractor into a stable limit cycle.*

   The converse problem, in a sense, is that of altering given dynamics because their behaviour is undesirable. In the case of a chaotic attractor, for example, we may wish to use some 'control action' to get rid of the chaotic behaviour altogether.

   The well-known pioneering work of Grebogi, Ott and Yorke has led to an explosion of interest in this area recently. However, most approaches to date are not truly nonlinear and they certainly do not take into account the full global geometry of the dynamics involved. The approach we develop could hopefully lead to more elaborate nonlinear, global control-of-chaos techniques.

3. *Large deviation theory.*

   The context here is, I think, of great importance to control theory. Not only does it offer the nonlinear control theorist a large area for novel research, but it gives a solid reason for enlarging the scope of nonlinear control beyond the familiar stabilization and tracking problems. The connections and similarities with physical theories (as in the work of B. Simon, Jona-Lasinio and others on quantum-mechanical tunnelling) should not be overlooked.

   Briefly, one considers the stochastic diffusion process with small noise

$$dx_t^\epsilon = b(x_t^\epsilon)dt + \epsilon\sigma(x_t^\epsilon)dw_t \tag{5}$$

   in the limit as $\epsilon \to 0$. For concreteness, let us consider the case where the drift dynamics b is gradient-like (see next section for definition.) Since the noise is small, starting from some state in the region of attraction of a stable equilibrium, we flow down following closely a trajectory of the drift vector field. The small noise can, however, lead to large excursions away from attractors. Under uniform ellipticity of the diffusion matrix σ, the theory of stochastic stability will in fact tell us that we shall leave the region of attraction for any nonzero noise intensity (any nonzero ϵ.)

What is interesting, however, is the paths we are likely to follow; these, for small noise, turn out to be those of the minimal-energy steering of the state of the *associated control problem* below, obtained by replacing the 'white noise' $\xi_t = \frac{dw_t}{dt}$ by a control action

$$\dot{x}_t = b(x_t) + \sigma(x_t)u, \tag{6}$$

using the cost functional $\int_0^T |u|^2 dt$ and requiring that we start at an attractor and control the state to some point in its region of attraction. Note that ϵ does not appear in the above equation.

Under certain conditions, we end up with an optimal feedback control that in some sense *inverts the gradient dynamics* of b: the attractor becomes a repeller and the saddles switch index from, say, k to $n-k$. This, by the way, was the context in which I first came up with the idea that it is enough to be able to control through certain 'Lyapunov functions' instead of requiring the 'reachable set' to be 'large' –thus switching from an essentially local, strong form of controllability (involving derivatives) to a weaker, global one which is more easily verified and is enough for the purposes of control design (involving functions.)

3 Outline of the Main Methods and Results

In this subsection, we try to sketch answers to the questions:

1. What is the novelty of the tools we are about to develop?
2. What are the advantages in using them? In particular, what results do we obtain that are not derivable by other, more classical/differential methods?
3. *'Enrichment':* Are these tools useful in areas sharing features with control theory?

3.1 New Methods, Tools and Concepts

- **From Dynamical System Theory:** Our understanding of global dynamics has improved vastly in the past three decades. Little of this understanding has trickled through to nonlinear systems theory.

  We claim that, by being specific about the classes of dynamics we are interested in, we can make our control task easier and benefit from the special features of the chosen class (eg gradient-like systems.)

  Concepts that will find application in our approach include: structural stability, notions of equivalence (topological orbital equivalence) of dynamics, qualitative classification of dynamics, bifurcation theory, decomposition of flows into gradient-like and chain-recurrent parts (Conley's theorem), existence of Lyapunov global functions and consequences of their existence etc.

- **From Geometry and Topology:** Here, in addition to using modern toplogical concepts to 'modernize' our approach to nonlinear control, we shall specifically make use of basic notions such as: homotopy invariance of certain maps and objects (leading to degree/index theory and more generally to algebraic topology), genericity/transversality, enabling a powerful transition from *local* conditions (such as rank conditions of Jacobians and counting number of independent equations defining a geometric object) to *global* objects (transverse manifolds). Also, we shall find that many familiar control-theoretic aspects are best understood using tools such as the *Gauss map* and the existence of topological invariants. Finally, we must note that a proper geometric setting makes it easier to formulate questions in a coordinate-free way, as we have seen in the first section.

- **From Other Areas (e.g. Singularity Theory, Algebra):** There will be instances where an awareness of the progress made in areas such as the ones mentioned can lead to improved understanding of some subtle aspects of nonlinear systems (this is so especially in the case where vector fields etc are modules over a polynomial ring.)

  Finally, we hope that useful applications can be found for Computational Algebraic Geometry (which provides nice answers to problems about parametrization of solutions, implicitization and elimination of variables concerning systems of polynomial equations), even though this set of tools is not yet fully applicable to systems over the real field.

3.2 Advantages of our Methods and Main Results

- We have already argued that simply narrowing attention to a specific class of dynamics can go a long way towards making the task of the control theorist easier. This is especially clear in the case of systems with global Lyapunov functions which are almost everywhere strict. In the next section, we shall present the main results from the genral theory of global dynamics and try to make explicit the information contained in a global Lyapunov function.

- **Necessary Conditions:** Our use of Lyapunov functions and the Gauss map permits a streamlined presentation of necessary conditions for achieving dynamics of a particular type (e.g stabilization); we see that the Krasnosel'skii-Brockett-Coron conditions are simple consequences of the geometry of the flow near an asymtpotic attractor. We give generalizations to other dynamics (saddles and global gradient-like dynamics for example) and to more complicated isolated invariant sets.

  The geometry of the Lyapunov level sets gives stronger conditions that avoid the transition to algebraic-topological invariants. In the particular case of a stable limit cycle, we give a condition that says that the image of the Gauss map cannot be 'too small' in the sence that it cannot be contained in a half-space.

- **Objects of Interest to Control:**

  We start by giving a new interpretation of the linear stabilizability and controllability conditions. Instead of using Lie brackets, we show that these conditions simply mean that there exists a subspace (manifold in the nonlinear case) complementary to D and on which the state dynamics is 'stable' (see Part II for details.)

  This motivates a discussion of geometric objects of interest to control: singular sets (for example, the set where the state dynamics is in D or the set where D is in the kernel of dh, where h is a Morse-Lyapunov function.) Here transversality is heavily used and generic dimensions obtained.

- **Constant Control Distributions:** Here we can give a more or less complete theory. We can give explicit conditions for stabilization and other types of dynamics; we can in fact give a description of *all* Lyapunov functions, at least in the convex case.

  Several examples will be given to illustrate the procedure. In low dimension, this is surprisingly easy to visualize. In higher dimension, one still has a nontrivial problem to solve, yet the relation to the linear systems case is very clear.

3.3 Applications to Other Areas

Let us make a historical comment: differential geometric tools such as the ones that are used in control theory are a 19th century creation. Differential geometry has seen much progress, of course, but much of nonlinear control theory can be considered as a variation on rather classical material.

On the other hand, most of the work in the fields mentioned above has been recent. Dynamical system theory, in particular, was essentially started in the sixties. I hope that this belated but justified attempt to bring into nonlinear control theory more modern methods and concepts will inspire work in related fields, such as nonlinear circuit and device theory, power system analysis and control etc.

4 Dynamical Systems Fundamentals

In this section we shall present some fundamental, but not so well known results on global stable aspects of dynamical systems. The main aim will be to prove the Conley Flow Decomposition Theorem and to examine its consequences for control design. The crucial role played by global Lyapunov functions will be emphasized.

4.1 Basic definitions and Conley's theorem

Throughout this part, we shall take the state space to be a *compact* manifold; this will include the case of *dissipative dynamics*, in other words complete vector fields on a possibly non-compact manifold that admit a global, compact attractor (a set that is not a manifold, in general.) In this case, one considers the flow on the one-point compactification of the manifold obtained by taking the point 'at infinity' to be a repeller.

Definition 5 *Given points x_0 and x_N in M^n and a number $\epsilon > 0$, an* ϵ-**chain** *from x_0 to x_N is a finite sequence of points $x_1, \ldots, x_{N-1}$ and times $t_k > 0$ ($k \in \mathrm{N} - 1$) such that*

$$|x_{k+1} - \phi(t_k, x_k)| < \epsilon \, , \ \forall k \in \mathrm{N}. \tag{7}$$

(ϕ is the flow of X.)

Definition 6 *A point x is* **chain-recurrent** *for the flow ϕ if for any $\epsilon > 0$ and any $T > 0$, there is an ϵ-chain from x to itself such that $t_k \geq T$ for all k.*

We denote by $\mathcal{R}(\phi)$ the **chain-recurrent set**, *i.e. the set of all chain-recurrent points of the flow ϕ.*

One should compare this concept with the more familiar concepts of α– and ω–limit sets and of the *nonwandering set* of a flow. In general, $\Omega(X) \subset \mathcal{R}(X)$, but the chain recurrent set may be larger, as in the following example.
Example 1: Consider the vector field on the circle $S^1 = \{e^{i\theta} \; ; \; \theta \in [0, 2\pi]\}$, given by $X(\theta) = (\cos(\theta/2))^2 \frac{\partial}{\partial \theta}$. It has a single equilibrium point at $\theta = \pi$ and gives a flow that moves counter-clockwise around the circle, except at the point x given by $\theta = \pi$. It is easy to see that the point x is the unique α and ω–limit point of every point on the circle. Thus $\Omega(X) = \{x\}$.

On the other hand, we claim that $\mathcal{R}(X) = S^1$, the whole of the circle. This is because, for chain recurrence, we are allowing ourselves an 'error' of

ϵ at each stage. Thus, for any $\epsilon > 0$ and any $T > 0$, starting at any point y in the circle, we move towards x, where the flow 'slows down' because of the zero of the vector field; once we are near enough (this happens in a finite number of steps), we can 'jump over' x using the ϵ leeway. Finally, going backwards in time, we can compute a similar sequence of points from y to a point ϵ–close to the other side of x. This gives a finite ϵ–chain from y to itself.

It turns out that the chain-recurrent set contains even complicated sets such as strange attractors. Thus, from the dynamical viewpoint, it is a very useful concept (and was used to great effect by, for example, R. Bowen, in the study of so-called Axiom-A differomorphisms; the chain-recurrent set is then assumed to have a *hyperbolicity* property generalizing the hyperbolicity of equilibrium points.)

The chain-recurrent set is also the *right* concept for separating two parts in every flow: the chain recurrent part, where there exists this approximate 'periodicity' or recurrence and the 'dissipative' or gradient-like part, where a certain quantity strictly decreases along the flow. This quantity is of course easily compared with the traditional view of a Lyapunov function; Conley's theorem makes all this precise.

First, let us give a definition that will fix the relation between gradient-like systems and Lyapunov functions

Definition 7 *A flow ϕ on the compact manifold M is* **gradient-like** *if there exists a continuous function V on M, called a* **global Lyapunov function** *for ϕ, that is strictly decreasing along the trajectories of ϕ everywhere except at the set of equilibrium points of ϕ.*

The reader should spend a minute or two thinking about how a global Lyapunov function differs from the Lyapunov functions one sees in traditional control theory. In particular, Lyapunov functions exist near saddles or even repelling equilibrium points. The conventional way geometers think of Lyapunov functions is by visualizing them as 'height functions' in some landscape, or on a manifold embedded in Euclidean space. The height function on a torus lying at a small angle off the vertical on a table provides a conventional example (see Example 2 below.)

In the case of a gradient-like system, it is easy to see that $\mathcal{R}(X)$ is precisely the set of equilibrium points, in other words the complement of the set where the Lyapunov function decreases. This is true in general; we have the fundamental result

Theorem 3 (Conley's Flow Decomposition Theorem): *For any flow ϕ on a compact manifold M^n, there exists a continuous function V that is strictly decreasing along the trajectories of ϕ outside the chain recurrent set $\mathcal{R}(\phi)$ and is such that the image $V(\mathcal{R}(\phi))$ of $\mathcal{R}(\phi)$ under V is a nowhere dense subset of the compact interval $V(M^n) \subset \mathbf{R}$.*

The original proof of Theorem 3 appeared in the report [5], republished in [6]. We follow the proof given in [23].

4.2 Outline of the proof of Conley's theorem

The main idea is to first prove the existence of Lyapunov functions for a simpler object: an attractor-repeller pair. It then remains to connect the chain-recurrent set with the set of all possible attractor-repeller pairs.

Definition 8 *A subset $A \subset M^n$ is an* **attracting set** *for the flow ϕ if there exists a positively invariant set $U \supset A$ and a time $T > 0$ such that $\phi_T(\overline{U}) \subset U^o$ and $A = \cap_{t \geq 0} \phi_t(U)$.*

The set U is called a **trapping region** *for the attracting set A.*

A subset A^* is called a **complementary repelling set** *of A for the flow ϕ if it is attracting for the negative of the flow with respect to the trapping region $M^n \setminus U$.*

The pair (A, A^*) is called an **attractor-repeller pair** *for ϕ (abbreviated to 'A-R pair' from now on.)*

One sees easily that if $x \notin A \cup A^*$, then $\omega(x) \subset A$ and $\alpha(x) \subset A^*$. The following properties of attracting and repelling sets are straightfoward to prove.

Proposition 1 *Attracting and repelling sets are compact invariant subsets of the compact manifold M^n.*

Let

$$\mathcal{A} = \{(A, A^*) \ ; \ A \text{ and } A^* \text{ are an A-R pair for } \phi\} \tag{8}$$

be the set of all A-R pairs. The next proposition says that not only does every A-R pair contain the chain-recurrent set of ϕ, but $\mathcal{R}(\phi)$ is exactly the common intersection of all the A-R pairs.

Proposition 2 *Let ϕ be a flow on a compact Riemannian manifold. Then*

$$\mathcal{P} \doteq \bigcap_{(A,A^*) \in \mathcal{A}} (A \cup A^*) = \mathcal{R}(\phi). \tag{9}$$

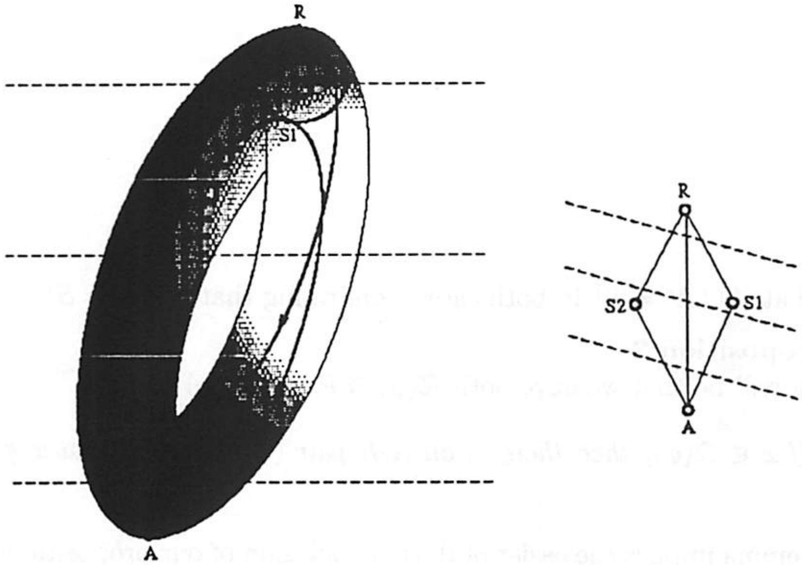

Figure 1: Flow on a two-torus

Before giving the proof, let us give some examples to get a feeling for why this is the case.

Example 2: Consider the gradient flow on the torus with respect to the height function when the torus is positioned as in Figure 1.

There are four equilibrium points: an attractor, p_1, two saddles, p_2 and p_3, and a repeller, p_4.

One choice of an attractor-repeller pair is to take

$$A = \{p_1\} \text{ and } A^* = \{p_2, p_3, p_4\} \cup W^s(p_2) \cup W^s(p_3) \tag{10}$$

(W^s stands for 'stable manifold.')

A different choice is

$$A = \{p_1, p_2\} \cup W^s(p_2) \text{ and } A^* = \{p_3, p_4\} \cup W^s(p_3) \tag{11}$$

Can you think how many A-R pairs exist in this case? (*Hint:* Choose a height h; if the h-level set of the height function V does not pass through an equilibrium point, then the set $U = \{V(x) \leq h\}$ is a trapping region and hence defines an A-R pair –the attractor A, for example, is the maximal invariant set in U; if there are no equilibria between two levels, the A-R pair does not change –all this is elementary Morse theory. The technique of slicing using Lyapunov level sets is a useful one; we shall meet it again when discussing Morse sets and Morse decompositions.)

The chain-recurrent set $\mathcal{R} = \{p_1, p_2, p_3, p_4\}$, as we knew from Conley's theorem and the fact that the flow is gradient-like –in fact gradient.

Example 1 (Contd.): In the case of the flow on the circle, there are no *nontrivial* attracting or repelling sets, according to our definition. The circle itself and the empty set are the trivial ones. Thus

$$\mathcal{A} = \{(S^1, \emptyset), (\emptyset, S^1)\} \tag{12}$$

and we see that $A \cup A^* = S^1$ in both cases, confirming that $\mathcal{R}(X) = S^1$.

Proof of Proposition 2:

The proof will be that we have both $\mathcal{R}(\phi) \supset \mathcal{P}$ and $\mathcal{R}(\phi) \subset \mathcal{P}$.

Lemma 1 *If $x \notin \mathcal{R}(\phi)$, then there is an A-R pair (A, A^*) such that $x \notin A \cup A^*$.*

Clearly, the lemma implies the easier of the two inclusion of our proposition: $\mathcal{R}(\phi) \supset \mathcal{P}$. The proof relies on the fact that, since $x \notin \mathcal{R}(\phi)$, there is an $\epsilon > 0$ such that there is no ϵ–chain from x to itself; this then is used to construct a trapping region U. We omit the details, which can be found in [23].

Lemma 2 **Existence of Lyapunov functions for A-R pairs:** *Let $(A, A^*) \in \mathcal{A}$. Then there exists a continuous function $V : M^n \to \mathbf{R}$ such that*

1. *$V(A) = 0$ and $V(A^*) = 1$,*
2. *For $x \notin A \cup A^*$, $0 < V(x) < 1$ and*
3. *V is strictly decreasing outside the set $A \cup A^*$.*

Proof: The proof follows lines similar to those used in proving the existence of a local Lyapunov function near an attractor. One tries an obvious guess as to what V should be and then modify the guess so that it satisfies properties (1) to (3).

Step 1: Let

$$V_0(x) = \frac{d(x, A)}{d(x, A) + d(x, A^*)} \tag{13}$$

where $d(x, y)$ is the Riemannian distance function in the compact M^n and we make use of the fact that $A \cap A^* = \emptyset$.

It is readily checked that V_0 is continuous, $V_0(A) = 0$, $V_0(A^*) = 1$ and $V(M^n \setminus (A \cup A^*)) = (0, 1)$. However, V_0 is not necesssarily decreasing along ϕ. **Step 2:** Let

$$V_1(x) = \sup_{t \geq 0}\{V_0(\phi_t(x))\}. \tag{14}$$

Since A and A^* are invariant, we still have $V_1(A) = 0$ and $V(A^*) = 1$. Now if $x \notin A \cup A^*$, we want to make sure that $V_1(x)$ is well defined. Since $d(\phi_t(x), A) \to 0$ as $t \to \infty$ and $V_0(x) \in (0, 1)$, we deduce that the supremum is attained, in other words there is a time $t(x)$ such that $V_1(x) = V_0(\phi_{t(x)}(x))$.

In fact the function $t(x)$ is continuous and nonnegative and so V_1 is continuous.

The function V_1 is also decreasing, i.e. $V_1(\phi_t(x)) \leq V_1(x)$, since the supremum on the left is over a smaller set. In order to make it strictly decreasing, we use a weighted average:

Step 3: Define

$$V(x) = \int_0^\infty e^{-s} V_1(\phi_s(x))ds. \tag{15}$$

It is obviously continuous.

Now for any $t \geq 0$ we have by the group property of the flow that

$$V(\phi_t(x)) = \int_0^\infty e^{-s} V_1(\phi_{s+t}(x))ds \leq \int_0^\infty e^{-s} V_1(\phi_s(x))ds \tag{16}$$

using the monotonicity of V_1. Thus V is decreasing.

To show it is *strictly* decreasing, we use the fact that $V(x)$ is the integral of the product of the weight function e^{-s} with a monotonically decreasing nonnegative function $V_1(\phi_s(x))$. Thus, if for some $x \notin A \cup A^*$ and some $t > 0$, $V(x) = V(\phi_t(x))$, then we must have that $V_1(\phi_{s+t}(x)) = V_1(\phi_s(x))$ *for all* s. But this is impossible, since $\omega(x) \subset A$ and so picking, say $s = nt$ for $n = 0, 1, 2, \ldots$ we get

$$V_1(x) = V_1(\phi_t(x)) = V_1(\phi_{2t}(x)) = \cdots = V_1(\phi_{nt}(x)) \tag{17}$$

and $V_1(\phi_{nt}(x)) \to 0$ as $n \to \infty$, a contradiction. Thus the lemma is proved. □

The next result proves the converse inclusion $\mathcal{R}(\phi) \subset \mathcal{P}$ and thus completes the proof of Proposition 2.

Lemma 3 *If* $(A, A^*) \in \mathcal{A}$, *then* $\mathcal{R}(\phi) \subset A \cup A^*$.

Hence every attractor-repeller pair contains the chain-recurrent set of ϕ. This will be useful later on, when we are trying to extract the information contained in a Lyapunov function, assuming we know nothing else about the flow.

Proof: (Outline) Take $x \notin A \cup A^*$; we must show $x \notin \mathcal{R}(\phi)$.

From the above lemma, we know there exists a Lyapunov function V for (A, A^*). Use the fact that V is strictly decreasing at x to show that it is not possible to find an ϵ–chain from x to itself, for some ϵ and some (large) $T > 0$, since the function V must decrease at each step and thus it prevents the return to x. □

Finally, we prove Conley's Theorem using Proposition 2. We need a technical lemma about the cardinality of the set $\mathcal{A}$.

Lemma 4 *The set $\mathcal{A}$ is at most countable.*

Proof: We use the fact that the compact manifold M^n is a metric space and hence has a countable basis. We would like to assert that a trapping region defines a unique A-R pair. Note that, in general, there is more than one attracting set inside a trapping region U; however, by also inclusing the complementary repeller, whose trapping region is $M^n \setminus U$ we get a unique object.

Thus it is appropriate to consider the product $M^n \times M^n$ and to map each A-R pair to the set $A \times A^*$ since then the correspondence with the set U is well-defined, namely the A-R pair corresponds to the open set $U^o \times (M^n \setminus U)^o \subset M^n \times M^n$. Now since the product has a countable basis, we conclude that $\mathcal{A}$ too is at most countable. □

(Note that of course many open sets may give the same A-R pair – remember the discussion in Example 2.)

To finish the proof of Conley's theorem, let us index the (at most) countably many Lyapunov functions we obtain for the A-R pairs in $\mathcal{A}$, say, writing V_k for each such function ($k \in \mathbf{Z}$.) Now let

$$V(x) = \sum_{k=1}^{\infty} \frac{2}{3^k} V_k(x). \tag{18}$$

The sum is absolutely convergent, so V is well defined and continuous; also $V(x) \in [0,1]$. It is decreasing since the functions V_k are decreasing.

To show it is strictly decreasing outside $\mathcal{R}$, we use the fact that if $x \notin \mathcal{R}(\phi)$, then, by Lemma 3, there is some A-R pair (A_k, A_K^*) (remember the indexing) such that $x \notin A_k \cup A_k^*$ and so at least one term $V_j(\phi_t(x)) \leq V_j(x)$ gives strict inequality.

The final part of the Theorem involves a subtlety: since the chain-recurrent set is contained in every attractor-repeller set, for $x \in \mathcal{R}(\phi)$, $V_j(x)$ is either 0 or 1. Thus the expression for such x is

$$V(x) = \sum_{k=1}^{\infty} \frac{2\delta_k}{3^k} \tag{19}$$

with δ_k zero or one. This is the ternary expansion of some number between zero and one and *since only the digits* 0 *and* 2 *are used*, it is only possible to

get in this way a subset of a nowhere dense Cantor set in the unit interval.
□

5 What do Lyapunov functions tell us about global dynamics?

In this section we do not assume any knowledge of the chain recurrent set of the flow ϕ; however, we shall assume given a Lyapunov function on a subset of the state space. The question we will try to answer in this section is: *what information about the global dynamics of ϕ is contained in the Lyapunov function and how do we extract it?* The obvious simple case of having a function in a neighborhood of a point such that the point is the unique minimum of the function is precisely Lyapunov's stability theorem, for asymptotic equilibria.

We suppose, as usual, that the flow ϕ (and the corresponding complete vector field X) is given on a compact manifold M^n.

KEY IDEA: **Use the gradient vector field of the Lyapunov function and notions of equivalence of dynamics –including the Conley index.**

Definition 9 *A continuous function V defined on an open subset $\mathcal{U}$ of M^n is a* **Lyapunov function** *for the dynamics X on the set $\mathcal{U}$ if V is strictly decreasing along the flow ϕ restricted to the set $\mathcal{U}$. We write $(V,\mathcal{U})$ for the pair consisting of the function and its domain.*

If V is C^1, we require $dV(X)(x) < 0$ for all $x \in U$.

Note: We reserve the name *global Lyapunov function* for a function as in the Conley Theorem, in other words a Lyapunov function in $M^n \setminus \mathcal{R}(X)$.

The following proposition is proved along the lines of Lemma 3 of Section 4.

Proposition 3 *Suppose we are given a Lyapunov function $(V,\mathcal{U})$ for ϕ. Then $\mathcal{R}(X|_{\mathcal{U}}) \subset M^n \setminus \mathcal{U}$.*

Remark: In general, of course, $\mathcal{R}(X|_{\mathcal{U}})$ is a proper subset of $M^n \setminus \mathcal{U}$, as we saw before in the discussion of attractor-repeller pairs. The content of the above proposition should therefore be taken to be that open sets can be 'thrown away' from state space in a search for the asymptotic or chain recurrent part of a flow. Caution should also be exercised since we are only asserting that the flow *restricted* to the open set has these properties. It can easily happen that, for dynamics more complex than gradient-like ones, a Lyapunov function is found that is defined on an open set that contains a part of a periodic orbit, for example. Any *global* Lyapunov function should of course be constant on any periodic orbit.

Now we may think that a lot of dynamical information can be extracted from a Lyapunov function if the set $\mathcal{U}$ is all of M^n except for a finite number of distinct points, say $E = \{e_1, \ldots, e_P\}$. However, even in this case we **cannot** conclude that $\mathcal{R}(X) = E$.

Example 1: Suppose the e_i are isolated points where X is in ker dV, but both X and dV are everywhere nonzero. Then $\mathcal{R} = \emptyset$, but $E \neq \emptyset$.

As a concrete example, take $V = x_1$ and $X = (x_1^2 + x_2^2, 1)^T$ in a subset of $\mathbf{R}^2$; then $E = \{(0,0)\}$ and there is a unique trajectory of X through the origin that is tangent to the level of V; however, X is trivial and V has no critical points.

What went wrong in the above example was that the gradient vector field of the Lyapunov function also had trivial dynamics: the points e_i were not equilibria of $-\nabla V$. Let us see how things improve when, together with the function V, we consider its gradient dynamics.

5.1 A generalized Lyapunov theorem

Definition 10 *Provide M^n with a Riemannian metric, $< \, , \, >$. Then for any C^1 function V defined on an open set $\mathcal{U} \subset M^n$, the* **gradient vector field** *of V with respect to the given Riemannian metric, $-\nabla V$ is the vector field satisfying*

$$dV(Y)(x) = < Y(x), \nabla V(x) > \tag{20}$$

The *equilibrium points* of any gradient vector field of V so defined coincide with the *critical points* of V. We write $cp(V) = E(-\nabla V)$. In fact, even though the gradient vector field seems to depend on the chosen metric for its definition, the dynamics so obtained are all equivalent (in a sense to be made precise later.) Moreover, the purpose of this section is to show the equivalence with *any other vector field for which V is a Lyapunov function.*

We collect some obvious facts on the relation between a function and its gradient vector field.

A critical point e of V is **nondegenerate** iff e is a nondegenerate equilibrium point of the gradient vectof field $-\nabla V$. In local coordinates, this amounts to having a nondegenerate Hessian at e; also note that nondegeneracy implies the equilibria (critical points) are *isolated*. Finally, since the Hessian is symmetric, in the case of a gradient vector field, nondegeneracy is equivalent to **hyperbolicity** (no eigenvalues on the imaginary axis since they are all real and nonzero.)

A function on a compact manifold whose critical points are all nondegenerate is called a **Morse function** (it therefore has only a finite number of them.) (One can formulate much of elementary algebraic topology in terms of Morse functions: for example, a homology theory can be defined in terms of the unstable manifolds of the gradient flow with a natural boundary map which can be shown to be equivalent to singular homology; it is called Witten–Floer homology. Also, Poincaré duality is easy to see by reversing the direction of the flow, i.e. letting the gradient push us up instead of down.) Nice references for Morse theory are: [19], [7], [24] and the wonderful article by Bott [2].

Let $E = \{e_1, \ldots, e_P\}$ be a finite set of points of M^n. The extent to which the dynamics of a vector field possessing a Lyapunov function resembles the dynamics of the gradient vector field of the function is made precise in the following theorem:

Theorem 4 (Generalized Lyapunov Theorem): *Suppose V is a C^1 function on M^n such that $(V, M^n \setminus |E|)$ is a Lyapunov function for the dynamics X on M^n. Suppose $E = cp(V)$. Then:*

1. $\mathcal{R}(X) = E$.
2. *for all $i = 1, \ldots P$, if the index of the equilibrium point e_i of $-\nabla V$ is* $\operatorname{ind} e_i = k_i$, $0 \leq k_i \leq n$, *then*

$$CI_X(e_i) = \Sigma^{k_i} = CI_{-\nabla V}(e_i) \tag{21}$$

where CI stands for the Conley index of the isolated invariant set e_i and Σ^{k_i} is a pointed k_i–sphere.

The next section explains these terms and their use in the study of global dynamics. We then procced with the proof of the theorem.

5.2 The Conley index

Note: The material in this section is too brief to do justice to this important topic. I recommend that you read first the introductory chapter of Conley's monograph and then refer to the survey articles [12] and [13].

Definition 11 *A (closed) invariant set S of the flow ϕ is* **isolated** *if it is the* maximal *invariant set in some neighborhood $\mathcal{U} \supset S$.*

This is equivalent to requiring $\mathcal{S} = \cap_{t \in \mathbf{R}} \phi_t(\mathcal{U})$. Nondegenerate equilibria are isolated as are limit cycles with nondegenerate Poincaré maps. We write $\mathcal{S} = I(\mathcal{U})$ and we say that $\mathcal{S}$ is an **isolated invariant set** (IIS for short.)

It will also be convenient to consider closed neighborhoods N of $\mathcal{S}$ (i.e. closed sets such that $\mathcal{S} \subset N^o$); these we shall call **isolating neighborhoods** or **isolating blocks**. An isolating neighborhood is thus defined by the requirements

1. N is a compact subset of the compact manifold M^n
2. the maximal invariant set in N is contained in the interior of N.

Thus, roughly, we are not allowing trajectories in N to have *'internal tangencies'* with the boundary of N.

MAIN IDEA OF THE CONLEY INDEX: **For a given isolating neighborhood N, examine its boundary ∂N and try and deduce something about the isolated invariant set $\mathcal{S}(N)$, the maximal IIS contained inside N.**

(The notations $N = I(\mathcal{S})$ and $\mathcal{S} = \mathcal{S}(N)$ will be used often and should be clear to the reader from the above definitions.) Thus, the Conley index is a **Black-box** approach, in the sense that the behaviour at the boundary is telling us something about the internal dynamics, even if these are not known explicitly. Of crucial importance in this black-box study is the concept of an *exit set.*

Definition 12 *The* **exit set** $N_+ \subset N$ *is the subset of N defined by the conditions:*

1. N_+ *is positively invariant for the flow* $\phi|_N$*; in other words, if* $x \in N_+$*, then* $\phi(t,x) \in N_+$ *for all* $t > 0$ *such that* $\phi([0,t],x) \subset N$.
2. *if for some* $x \in N$*, there is a time* $t_1 > 0$ *such that* $\phi(t_1,x) \notin N$*, then there is a time* $t_0 < t_1$ *such that* $\phi([0,t_0],x) \subset N$ *and* $\phi(t_0,x) \in N_+$.

Definition 13 *The pair* (N, N_+)*, where N is an isolating block and* N_+ *is an exit set for N is called an* **index pair** *for the IIS* $\mathcal{S}(N)$*, the maximal IIS in* N^o.

In general, given an isolated invariant set and given an arbitrary isolating block for it, it is possible to modify the isolating block and make a clever choice of exit set so as to obtain an index pair. Let us give some examples.

Examples:

1) **An attracting equilibrium** e: We can always choose a local Lyapunov function V so that $N = \{x \; ; \; V(x) \leq c\}$ is an isolating block for e and the exit set $N_+ = \emptyset$.

If, however, we are not careful and we pick a block whose boundary is not a Lyapunov level set, then we get a nonempty exit set; try to see how, by throwing out any piece of trajectory $\phi([-t_0, 0], x) \subset N$ with x a point of internal tangency, we can modify N so that is has empty exit set.

2) **A saddle equilibrium,** of index k: Here we use disk-diffeomorphs (sets diffeomorphic to disks) D^k and D^{n-k} for the unstable and stable manifolds using local Lyapunov functions and then we get a cylinder $D^k \times D^{n-k}$ as the isolating block, with exit set along the unstable directions, i.e. along a set $S^{k-1} \times D^{n-k}$.

3) **The empty set:** locally, by the flow box theorem, we can pick a disk neighborhood $N = D^n$ near a nonzero point for the vector field so that $N_+ = S_+^{n-1}$, a hemisphere.

The above examples seem to suggest that, in agreement with the elementary notion of index of an equilibrium, an index pair will have an exit set that has dimension equal to the intersection of the unstable manifold with a surrounding disk.

This is good, but has not given us much that is new; to bring out the full power of the Conley index approach, we must first realize that an index pair can be defined for an *arbitrary* IIS –an object for which the usual index is not defined. Furthermore, and this is the key, it is possible to construct a topological object out of an index pair that does not depend on the specific index pair, i.e. is the same for any choice of index pair for a fixed IIS. This object is the Conley index.

We give its definition, which depends on a number of elementary concepts from topology; a brief outline of these concepts then follows.

Definition 14 *Suppose an index pair* (N, N_+) *is given and let* $\mathcal{S} = \mathcal{S}(N)$ *be the maximal IIS in* N. *Then the* **Conley index** *of the IIS* $\mathcal{S}$ *is the topological equivalence class of the pointed space*

$$(N/N_+, [N_+]) \tag{22}$$

where N/N_+ *is the quotient space formed by identifying all the points in* N_+. *This class does not depend on the index pair chosen, in other words, if* (N', N'_+) *is another index pair for* $\mathcal{S}$, *then*

$$(N/N_+, [N_+]) \simeq (N'/N'_+, [N'_+]) \tag{23}$$

5.3 Some Concepts from Topology

To understand the above definition, we need the following concepts:

1. Quotient spaces
2. Pointed spaces
3. Topological equivalence of spaces

We explain each one of these terms briefly.

5.3.1 Quotient spaces

If a surjective map $p : \Omega \to Q$ from a topological space Ω to a set Q is given, the **quotient topology** on the set Q is defined by requiring that a subset $A \subset Q$ is open iff $p^{-1}(A) \subset \Omega$ is open.

The main example we have in mind is the case where an equivalence relation $\sim$ is given in the space Ω; this leads to a partiton of Ω into equivalence classes. Call Q the set of equivalence classes. Then the canonical map p is defined by assigning each point of Ω to its equivalence class under $\sim$. By the above, Q is a topological space with the quotient topology and the map p is continuous. One writes $Q = \Omega / \sim$.

Collapsing a subset to a point: In the definition of the Conley index, we had to form the quotient N/N_+. This is the quotient space under the equivalence relation: $x \sim y \Leftrightarrow x, y \in N_+$. Thus, the elements of N/N_+ are the points of $N \setminus N_+$ and the 'point' $[N_+]$ which is the collapsed exit set.

Example: Take the pair (D^k, S^{k-1}), a k-disk and its boundary sphere. Then $D^k/S^{k-1} \cong S^k$ (check this); the intuitive reason is that collapsing the boundary sphere to a point means drawing a string wrapped around the disk tight, thus forming a sphere of the same dimension as the disk.

5.3.2 Space pairs and pointed spaces

A *space pair* (Ω, A) consists of a space and a subset A of it. A *map of pairs* is a map

$$f : (\Omega, A) \to (\Omega', B) \tag{24}$$

such that $f(A) \subset B$. In homotopy theory, one is interested in particular in *based spaces* or **pointed spaces** (Ω, x_0). Maps between based spaces are thus required to send basepoints to basepoints.

In the case of the Conley index, the pair $(N/N_+, [N_+ +])$ is a pointed (or based) space, since in the quotient, the exit set is considered as a single point.
Note: By convention, if the exit set is empty, one takes the base point to be a point disjoint from N.

5.3.3 Topological equivalence

Definition 15 *Two continuous maps $f_0, f_1 : \Omega \to \Omega'$ are* **homotopic** *if there is a continuous map*

$$F : [0,1] \times \Omega \to \Omega' \tag{25}$$

such that $F(0,x) = f_0(x)$ and $F(1,x) = f_1(x)$ for all $x \in \Omega$.

This says that we can continuously deform the one map to the other. We write $f_0 \simeq f_1$. In the case of maps $\gamma : S^1 \to T^2$, the two torus, one checks that any map wrapping once around the length of the torus is not homotopic to a map going round it once. The classification of homotopy classes of maps is one of the main aims of algebraic topology.

Definition 16 *Two (topological) spaces Ω, Ω' are* **homotopically equivalent** *if there exist maps*

$$\begin{aligned} f &: \Omega \to \Omega' \\ g &: \Omega' \to \Omega \end{aligned}$$

such that the compositions $g \circ f \simeq \mathrm{id}_\Omega$ and $f \circ g \simeq \mathrm{id}_{\Omega'}$.

We say that f and g are *homotopy inverses* of each other. It is often the case that Ω' is a subset of Ω, the map g is the inclusion map and f is a retraction, i.e a continuous map of the total space onto its subset (think of the punctured Euclidean space $\mathbf{R}^n - \{0\}$ being retracted to the unit sphere S^{n-1}.) Another example is the solid torus being homotopy equivalent to its core, a circle.

There exist *relative* versions of the above definitions, in other words applying to space pairs or to pointed spaces.

5.4 Examples of the Conley index

Recall that the Conley index of an IIS S is the topological type (i.e. the homotopy equivalence class) of the pointed space

$$CI(S) = [(N/N_+, [N_+])] \tag{26}$$

where (N, N_+) is any index pair and we work with the convention of appending a disjoint distinguished point $*$ if the exit set is empty.

In applications, it is not always possible to compute the Conley index, i.e. to establish the precise topological type of the quotient space N/N_+. One then uses a weaker form of the index, one that only takes into account an algebraic-topological object associated to a homotopy equivalence class of topological spaces. Information is lost in this process, but there is the advantage of having available powerful methods of computation (excision, Mayer-Vietoris sequene etc.)

Definition 17 *The* (co)homological Conley index $h_*(S)$ *(resp.* $h^*(S)$*) is the (co)homology* $H_*(CI(S)$ *(resp.* $H^*(S)$*) of the homotopy equivalence class of the pointed space* $(N/N_+, [N_+])$.

(By H_* (resp H^*) we mean the graded groups $\oplus_p H_p$ (resp. $\oplus_p H^p$), where H_p is the pth–*singular homology group* and H^p is the pth–cohomology ring of the class $CI(S)$.)

The use of algebraic topology can be avoided sometimes, especially if only gradient-like systems are being considered, but is mandatory in many other cases, as we shall see below (as soon as we allow limit cycle behavior, for example.) The reader unfamiliar with these concepts can skip these points –they are marked by the sign $**$. (I cannot resist referring you to, for example, Fomenko [8], for a beautiful elementary treatment of algebraic topology.)

5.4.1 Attractors

We saw that, for an attractor, it is always possible to choose an index pair with empty exit set.

Thus in the case of an attracting *equilibrium point* e, the isolating block is a disk $D^n \simeq \text{pt}$ so that the Conley index is

$$CI(e) \simeq \Sigma^0.$$

For more general attractors, the situation is very similar: an isolating neighborhood with empty exit set exists, so that we can retract to the attractor itself. An example is a *stable limit cycle* γ. If $\gamma \subset \mathbf{R}^n$ (thus γ is an embedded closed orbit in $\mathbf{R}^n$), then an isolating block has the type of a solid torus $S^1 \times D^{n-1}$, which is contractible to its core $S^1 \times \{0\}$ and so the Conley index is

$$CI(\gamma) = [(\gamma, *)].$$

More generally, if K is a compact attractor that admits a 'tubular' isolating neighborhood, then

$$CI(K) \simeq (K, *).$$

5.4.2 Saddles

A saddle *equilibrium point* s of index k has Conley index

$$CI(s) \simeq \Sigma^k.$$

This is because, as we saw, we can isolate using the pair $D^k \times D^{n-k}$ with exit set $S^{k-1} \times D^{n-k}$. Collapsing the exit set gives the same space as $D^k/S^{k-1} \simeq S^k$, since the 'stable' directions do not play any role (we can make the disk very thin by squeezing down along the $W^s(s)$ before we form the quotient.) ** A hyperbolic limit cycle γ of index k has locally the structure of a vector bundle $E^n(\gamma) \to \gamma$ with fibre $\mathbf{R}^n$ that decomposes into the product $W^u_{\mathrm{loc}}(\gamma) \times W^s_{\mathrm{loc}}(\gamma)$. By 'squeezing in' the stable directions, we end up with a k-bundle $E^k(\gamma) \to \gamma$ (i.e. only the unstable manifold matters); this may be twisted, i.e. not a product. The *exit set* for this 'disk' bundle is the **sphere bundle** $S^{k-1}(\gamma)$.

The quotient $E^k(\gamma)/S^{k-1}(\gamma)$, by a remarkable coincidence, is precisely what is called the **Thom space** $\mathcal{T}(E^k(\gamma))$ of $E^k(\gamma)$ in algebraic topology (a much studied object, of relevance to areas such as cobordism theory and the Lefshetz fixed point formula –see [21], [11], [9].)

The Conley index is then the topological type of the Thom space

$$CI(\gamma) \simeq \mathcal{T}^k(\gamma)$$

This is *not* easy to compute –or even to visualize. Some nice algebraic-topological machinery comes to our aid, in the form of the so-called **Thom isomorphism.** Suppose $E^k \to B$ is a k–vector bundle over the base space B and suppose we have a Riemannian metric given on E. Then the *disk* *bundle* $D^k(B)$ and *sphere bundle* $S^{k-1}(B)$ over B can be defined (with fibre the unit closed disk and its bounding sphere respectively.) The *Thom space* $\mathcal{T}(E^k)$ is now defined, as above, to be the quotient $D^k(B)/S^{k-1}(B)$. It is to be considered as a *pointed space*, with distinguished point $t_0 = [S^{k-1}(B)]$.

Theorem 5 (The Thom Isomorphism): *There exists a group homomorphism between the integral homology groups*

$$\tau : H_{p+k}((\mathcal{T}(E^k), t_0); \mathbf{Z}) \to H_p(B, \mathbf{Z})$$

for all p, called the Thom morphism; *this is in fact an isomorphism of groups.*

A similar statement holds for cohomology. So we can compute the (co)homology of the Thom space from the (co)homology of the base.

In the case of a limit cycle γ, $H_p(\gamma) = \mathbf{Z}$ for $p = 0, 1$ and zero otherwise. Thus $H_k((\mathcal{T}(\gamma), t_0)) \cong \mathbf{Z}$ and $H_{k+1}((\mathcal{T}(\gamma), t_0)) = \mathbf{Z}$ and is zero for all $p > k+1$. In $\mathbf{R}^3$ for example, with $k = 1$, we get H_1 and H_2 equal to $\mathbf{Z}$ and $H_p = 0$ for $p > 2$. In fact in this special case we can see that the Thom space is actually the wedge sum $S^1 \vee S^2$, confirming the computation of homology (can you visualize this?)

5.4.3 Normally hyperbolic invariant sets

For IISs more general than equilibrium points or limit cycles, a common assumption is that in directions *normal* to the IIS, we have *hyperbolicity* (and so roughly local stable and unstable manifolds.) We then proceed along lines similar to the case of a limit cycle:

We 'ignore' the stable directions and consider the Thom space obtained by collapsing the bounding sphere bundle of the unstable manifold to a point. The homological Conley index can then be computed using the Thom isomorphism. As we do not have in mind applications where it is desirable to design dynamics with IIS more complicated than equilibria and limit cycles, we do not pursue this general case any further. **

5.5 Sums and products of Conley indices

Let (Ω, x_0), (Ω', x_1) be two pointed spaces. Sometimes we shall write Ω for a pointed space (the distinguished point being understood.)

Definition 18 *The* **wedge sum** $\Omega \vee \Omega'$ *of the two pointed spaces is the space*

$$\Omega \vee \Omega' = (\Omega \amalg \Omega'/_{x_0 \sim x_1}, [x_0]) \tag{27}$$

obtained by identifying the distinguished points in the disjoint union of the spaces.

Thus, for example, $\Sigma^1 \vee \Sigma^1$ is a *'figure-eight'* with the node point being distinguished.

Let $A = \Omega \times \{x_1\} \cup \{x_0\} \times \Omega'$.

Definition 19 *The* **wedge product** *of the pointed spaces* Ω *and* Ω' *is the pointed space*

$$\Omega \wedge \Omega' = (\Omega \times \Omega'/A, [A]). \tag{28}$$

Products of pointed spheres are particularly well behaved; we have

$$\Sigma^1 \wedge \Sigma^1 = \Sigma^2$$

and more generally

$$\Sigma^k \wedge \Sigma^l = \Sigma^{k+l}.$$

5.5.1 Applications to the computation of the Conley index

Sums: Let $\mathcal{S}_1$ and $\mathcal{S}_2$ be two disjoint isolated invariant sets. Suppose they are isolated using blocks N_1 and N_2 respectively, with $N_1 \cap N_2 = \emptyset$. Then the Conley index of the maximal IIS $\mathcal{S}(N_1 \cup N_2) = \mathcal{S}_1 \cup \mathcal{S}_2$ is

$$CI(\mathcal{S}_1 \cup \mathcal{S}_2) = CI(\mathcal{S}_1) \vee CI(\mathcal{S}_2)$$

Caution: Great care should be taken with choosing isolating blocks: remember that the Conley index takes into account the *maximal IIS* in N. Suppose the IISs above are two equilibrium points, say e_1 and e_2, the first of which is an attractor and the second a saddle of index 1. Then $CI(e_1 \cup e_2) = \Sigma^0 \vee \Sigma^1$. However, suppose the unstable manifold of e_2 intersects the stable manifold of e_2; then, if the isolating block encloses the intersection orbit, the resulting index is easily seen to be trivial (a point). This is because the maximal invariant set is $e_1 \cup e_2 \cup C(e_1, e_2)$, where $C(e_1, e_2)$ is the set of connecting orbits for the two equilibria, and this is not just the union of e_1 and e_2.

Products: The main application here is to *product flows*, i.e. those that can be written as

$$\begin{aligned} \dot{x}^1 &= f^1(x^1) \\ \dot{x}^2 &= f^2(x^2). \end{aligned}$$

For example, given a linear flow that can be written, in some coordinates as

$$\begin{aligned} \dot{x}_1 &= x_1 \\ \dot{x}_2 &= x_2 \\ &\cdots \\ \dot{x}_k &= x_k \\ \dot{x}_{k+1} &= -x_{k+1} \\ &\cdots \\ \dot{x}_n &= -x_n \end{aligned}$$

we conclude that, since along each of the coordinates, the index of the origin is either Σ^0 or Σ^1, the index of the origin in the product flow is

$$CI(0) = k\Sigma^1 \vee (n-k)\Sigma^0$$

which, by the properties of the product, is just Σ^k, as we expected.

Algebraic structures Let $\mathcal{P}$ denote the set of pointed spaces. The wedge sum and product are defined on $\mathcal{P}$; we can ask: in what sense are they like the usual algebraic sums and products?

It is obvious that the sum does not make $\mathcal{P}$ into a group: if either of the two pointed spaces Ω and Ω' is nontrivial (not a single point, denoted $*$), then $\Omega \vee \Omega' \neq *$. Thus, with the one-point space playing the role of zero, there are no inverses in $(\mathcal{P}, \vee)$.

On the other hand, the product $\wedge$ is more like the usual multiplication. The reader should check that we have the properties

$$* \wedge \Omega = * \tag{29}$$

$$\Sigma^0 \wedge \Omega = \Omega \tag{30}$$

so that $*$ is like multiplication by zero and the unit space is the zero-sphere Σ^0. These properties are important later on in the discussion of connections between saddles. Also note that it confirms the computation of the index in a product space with a hyperbolic saddle (the Σ^0's had no effect in the final index); this in some sense establishes the contention that it is only the unstable directions that are involved in the Conley index.

Finally, a remark about modifying the sum so that it is better behaved: considering for the moment only hyperbolic equilibria, can you think of a way to define a 'sum' of such equilibria, so that inverses exist? (*Hint:* Consider including connecting orbits.) We return to this point later.

5.6 Continuation

One of the most important properties of the Conley index is that it is a *'stable object'*, in the sense that it is invariant under small perturbations of the flow. This constituted, in fact, the main motivation of Conley: he wanted to develop a theory of the 'stable features' of dynamical systems (should we also perhaps use the term 'robust'?)

Consider an equilibrium point at the precise moment when it is undergoing a saddle-node bifurcation; it is obviously not (locally) stable, since its topological type changes by small changes in the bifurcation parameter. However, the index of the equilibrium in the sense of Conley **does not change** locally. This is because, in order to compute the index, we had to isolate the IIS using a block and this block *continues to isolate for nearby flows!* (Check that in $\mathbf{R}^2$ this index is trivial for the saddle-node bifurcating equilibrium.)

The continuation property can now be stated: suppose we consider a dynamical system that depends on a number of parameters

$$\dot{x} = f(x, \lambda) \tag{31}$$

where $\lambda \in \Lambda$, some set. One can define a natural 'product flow' in $M^n \times \Lambda$ (the dynamics in Λ being trivial) and thus we can talk about isolated invariant sets, isolating blocks etc. for the product flow.

Let N be an isolating block for the product flow; this means that it isolates $S(N)$, the maximal invariant set in N. If N is of the form $N_0 \times \Lambda_0$, where $N_0 \subset M^n$ and $\Lambda_0 \subset \Lambda$, then we call N a *product neighborhood*. Since the dynamics in parameter space are trivial, we see that $N_0 \times \lambda$ must be an IIS for the flow f_λ for all $\lambda \in \Lambda_0$. We say that N gives a **continuation** for the IIS $S(\lambda)$ of the flow, in the subset Λ_0 of parameter space.

In the case of our bifurcating equilibrium (say at 0), we can take $\Lambda_0 = [-\epsilon, \epsilon] \subset \mathbf{R}$ and N_0 any (closed) neighborhood of the origin that misses the other equilibria of f_λ, for all $\lambda \in [-\epsilon, \epsilon]$. We then obtain a continuation between the empty IIS and the IIS consisting of two equilibria together with their connecting orbit (which, as far as the outside is concerned, 'looks like' the empty set, in other words as a Conley index it is the same space as that for the empty set IIS.)

The notion of continuation, as we saw, is not sensitive to *local bifurcations*, but it can be used to produce a classification of IIS according to whether one IIS can bifurcate to the other. For example, to start with the most elementary case, the *empty set* can continue to a saddle-node type equilibrium or to the IIS described above, consisting of two equilibria and a connecting orbit. It cannot continue (i.e. there can be no sequence of bifurcations not involving *other* IIS) to, for example, a hyperbolic equilibrium of arbitrary index or to an IIS consisting of a limit cycle of any stability type etc. The reader should notice how this concept compares with the elementary theory of two-dimensional dynamics and in particular the Poincaré-Bendixson theory. (What is really special about planar systems is that compact IISs can only be equilibria, limit cycles and unions of equilibria with connecting orbits.)

The following *nontriviality theorem* is indicative:

Theorem 6 *Suppose the Conley index of some index pair (N, N_+) is not trivial. Then there is a complete trajectory of the flow in N, i.e. $\mathcal{S}(N) \neq \emptyset$.*

5.7 Morse decompositions, orbit diagrams and Lyapunov functions

Recall that the set of all attractor-repeller pairs $\mathcal{A}$ is at most countable infinite. A finite object can be defined that captures interesting detail about a flow; this object, called a Morse decomposition, is a generalization of the more classical (ca 1960) one of the *orbit diagram* of a Morse-Smale flow.

In this section, we shall present the rudiments of Morse decompositions and their uses and also make the crucial connections with global Lyapunov functions.

We start with a compact isolated invariant set $\mathcal{S}$, which we may think of as either the whole state space, if compact, or as the compact global attractor of some dissipative dynamics.

Definition 20 *A* **Morse decomposition** *$\mathcal{M}$ of the IIS $\mathcal{S}$ is a finite collection of IIS, $\{M_1, \ldots, M_N\}$, $M_i \subset \mathcal{S}$, such that for any $x \notin \cup_i M_i$, the α- and ω-limit sets of x, $\alpha(x)$ and $\omega(x)$, are subsets of distinct members of the decomposition, say $\alpha(x) \subset M_i$ and $\omega(x) \subset M_j$, $i \neq j$.*

The IIS M_i, $i = 1, \ldots, N$ of the decomposition are called **Morse sets.** Roughly, each Morse set consists of components of the chain recurrent set, together with sets of *connecting orbits*. It is a more general object than an A-R pair, which consists of only two Morse sets. Also note that we have that $\mathcal{R}(\phi) \subset \cup_i M_i$ and hence there exist Lyapunov functions that are strict away from the set $\cup_i M_i$. In fact, it is not difficult to see that there are Lyapunov functions that are *constant* on each M_i and strict elsewhere (in

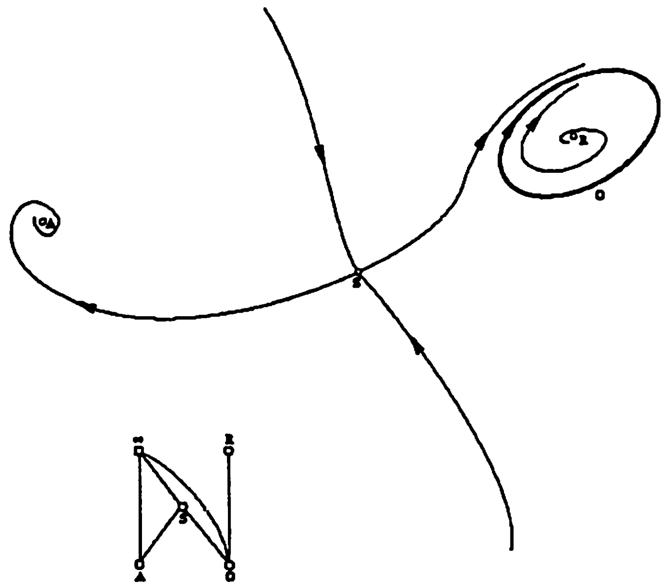

Figure 2: Orbit diagram of planar flow

the construction of global Lyapunov functions take only those attractor-repeller pair that do not break up any of the Morse sets in the given Morse decomposition.)

Definition 21 *Given an ordered pair of distinct Morse sets (M_i, M_j), the* **set of connecting orbits** *$C(M_i, M_j)$ is defined as the set of all orbits whose α–limit set is in M_i and whose ω–limit set is in M_j.*

Morse decompositions are related to a simpler object, defined for **Morse-Smale flows**; these are flows with a finite number of hyperbolic equilibria and limit cycles which exactly coincide with the chain recurrent set of the flow and whose stable and unstable manifolds intersect transversely.

Definition 22 *The* **orbit** or **Smale diagram** *of a Morse-Smale flow is the directed graph obtained as follows:*

1. *The vertices are the equilibria and limit cycles of the flow.*
2. *A directed arrow from a vertex I_i to a vertex I_j exists iff there is an orbit whose α–limit set is exactly I_i and its ω–limit set is exactly I_j.*

Note that it is possible, and convenient for drawing an orbit diagram, to distinguish **levels** from 0 (bottom) to n (top), according to the *index* of the hyperbolic limit set. One easily sees, from the transversality condition, that all arrows then point 'downwards.' An example of an orbit diagram is given in Figure 2.

Let us isolate three crucial features of orbit diagrams; these are:

1. **Indexing:** It is possible to classify limit sets according to their classical (Morse) index –the dimension of the unstable manifold.

2. **Transversality:** This, as we saw, was equivalent to having the arrows point down, from limit sets of high index to limit sets of lower index.

3. **Lyapunov functions:** It is possible to choose a Lyapunov function V such that, for any limit set I_i, there are values of the Lyapunov function V such that I_i is the maximal invariant set in the set lying between the two levels of V; in other words

$$\mathcal{S}(\{x\,;\ c_1 \leq V(x) \leq c_2\}) = I_i. \tag{32}$$

   This is a slightly more general statement than that in elementary Morse theory referring to gradient flows (and it needs proof, which the reader should look up in the standard references on Morse theory.) Also note that the *exit set* for this neighborhood is exactly the lower of the two levels, $\{x\,;\ V(x) = c_1\}$.

   Finally, any Lyapunov function gives a total order of the IISs of the Morse-Smale flow that is consistent with the partial order of Definition 22.

We ask the question: how do these features generalize in the case of an arbitrary flow?

We saw that it is always possible to find Morse decompositions –since A-R pairs are (trivial) Morse decompositions. Morse decomposition are *not unique*, though, in general.

Let us look at the above three attributes of orbit diagrams in turn and see how we can go about building an 'orbit diagram' from a Morse decomposition. The obvious starting point is to take the Morse sets as vertices and to connect them with arrows whenever there is a trajectory whose α–limit set is in one Morse set and its ω–limit set lies in the other.

Indexing: This is obviously generalized by the *Conley index*; note that we have lost the simple indexing by levels, since the Conley index is a topological space and there is in general no way of ordering these spaces by the integers (can we use some other totally ordered set? If so, how?)

A related, important question we should be asking at this point is: **what stability information is contained in the Conley index?**

Two things are clear: the correspondence between IISs and their Conley index is not one-to-one, except by restriction to simple IISs, such as equilibrium points. Despite this non-uniqueness, however, we may choose to classify IISs by their index and call an IIS an 'attractor,' for example, if it admits an isolating block with empty exit set. This is a crude classification, but is useful in some cases (as in the discussion of bifurcations above.)

(In the case of isolated equilibrium points, it should be clear that each Conley index Σ^k corresponds to a unique topological type of equilibrium point.)

Transversality: The condition that we imposed on points outside the support of the Morse sets going to distinct Morse sets in negative and positive time, together with the fact that Lyapunov functions exist there, implies the **no-cycle condition,** i.e. that we cannot have a closed sequence of arrows all pointing in the same direction along the loop.

However, this condition is *not equivalent*, in this case, to transversality of stable and unstable manifolds (assuming they exist): consider the case of a two planar saddles connected by a trajectory, i.e. such that $W^u(s_1) \cap W^s(s_2) \neq \emptyset$, but $W^s(s_1) \cap W^u(s_2) = \emptyset$. In this case, one can make the saddles s_1 and s_2 be Morse sets of a Morse decomposition and a global Lyapunov function can be found; yet transversaility fails.

We conclude that Morse decompositions are clearly a generalization of orbit diagrams, even if we consider systems with only simple chain recurrent set (equilibria and limit cycles); in particular, a Morse decomposition exists for certain systems that fail the transversality condition of Morse-Smale flows. Be careful, though: if the two saddles are connected by two nontransverse orbits, then this saddle connection is part of the chain-recurrent set and hence cannot be broken down into simpler IISs.

Lyapunov functions: This is an important aspect of the theory we are about to develop, so we start with some preliminaries.

The Morse sets $M_1, \ldots, M_N$ of a Morse decomposition are **partially ordered** by the relation

$$M_i \succ M_j \Longleftrightarrow \exists x : \alpha(x) \subset M_i \,, \omega(x) \subset M_j. \tag{33}$$

This relation is well defined by what we have said and it means that we can assign 'arrows' between vertices in the 'orbit diagram' of a Morse decomposition. (We continue using the term 'orbit diagram' since it is clear what we mean and it aids visualization.)

If a Lyapunov function V is given for the flow ϕ, we have a **total ordering** of the Morse sets by the relation

$$M_i \succeq M_j \Longleftrightarrow \exists I_i, I_j \in \mathcal{R}(\phi), I_i \subset M_i, I_j \subset M_j : V(I_i) \geq V(I_j). \tag{34}$$

(Remember that V must be constant on each connected component of the chain recurrent set.) The relation is well defined (check) and moreover **must be consistent** with the partial order of the Morse decomposition described above!

5.8 Proof of the generalized Lyapunov theorem

Part (a) is easy: We know, since the Lyapunov function is strict away from E, that $\mathcal{R}(X) \subset E$. Each point in E must be an equilibrium of X; this may sound obvious (it was not part of the theorem's assumptions), but is only true because the points of E are nondegenerate critical points of the Lyapunov function.

Indeed, suppose $X(e_i) \neq 0$, $e_i \in E$. By a flow-box argument, there is a neighborhood N of e_i such that the image of the Gauss map restricted to N is a contractible neighborhood of $G_X(e_i)$. However, again locally, the Lyapunov function V can be written, by the Morse lemma (see Milnor [19]), as a 'sum of squares,' i.e. as

$$V(y) = y_1^2 + \ldots + y_k^2 - y_{k+1}^2 - \ldots - y_n^2.$$

Thus the Gauss map on a level set is onto, contradicting the fact that the image of the Gauss map of X is contractible. (This is more or less the argument one uses to prove the Krasnosel'skii–Brockett necessary condition, of course.)

Part (b) now follows from the assertion that any equilibrium point of a gradient-like flow can be isolated using Lyapunov function intervals and level sets and, as a consequence, the index of the equilibrium is determined by the Conley index of the index pair $(V^{-1}([c_1, c_2]), V^{-1}(c_1))$. In case the value of V is the same for two equilibria, we must first perturb V to make it generic, i.e. to ensure that $V(e_i) \neq V(e_j)$, for any distinct e_i, e_j in E.

Comment: Note that the theorem cannot assert that $-\nabla V$ and X have the same orbit diagram; this is false, in general, as Example 1 showed.

Exercise 2 *1. Show that if e_i, e_j are equilibrium points of a Morse Smale flow such that $C(e_i, e_j) \neq \emptyset$ and $|\mathrm{inde}_i - \mathrm{inde}_j| = 1$, then*

$$CI(\overline{C(e_i, e_j)}) = * \neq CI(e_i) \vee CI(e_j)$$

2. Show that in the absence of the transversality condition, in other words if two equilibrium points have a one-dimensional intersection of their *stable and unstable manifolds, then*

$$CI(\overline{C(e_i, e_j)}) = CI(e_i) \vee CI(e_j)$$

It is thus possible for the Conley index to be 'blind' to the presence of connecting orbits. A subtler tool, called a **connected simple system,** *can distinguish between the qualitatively different perturbations of this nongeneric situation (see [12] for an example.)*

6 Global Control Design Through Morse-Lyapunov Functions

The aim of these notes is to present a way of using global feedback controls to design nontrivial global dynamics. The specification of the desirable dynamics is for now only a *qualitative* one. Thus in the case of stabilization, we specify a rather simple type of dynamical behaviour (one with a single asymptotically stable equilibrium point) and we do not address issues such as transient performance, robustness etc. In this section we define a way of specifying global dynamics to be achieved through control. This method uses what is in some sense the complement of the concept of a Morse decomposition, which we shall call a *Morse specification*. The analysis of dynamics presented in Section 4 will naturally play a crucial role.

Let us first make precise the notion of equivalence as used in the theory of dynamical systems.

Definition 23 *Suppose two flows, ϕ^1 and ϕ^2 are given on the compact manifold M^n. We say that ϕ^1 and ϕ^2 are* **topologically orbitally equivalent (TOE)** *if there exists a continuous function $\tau : M^n \times \mathbf{R} \to \mathbf{R}$ and a homeomorphism $h : M^n \to M^n$ such that*

1. *$\tau(x, 0) = 0$ for all $x \in M^n$ and τ is strictly monotone increasing as a function of t for all x.*
2. *For all $x \in M^n$ and all $t \in \mathbf{R}$,*

$$h(\phi^1(t,x)) = \phi^2(\tau(x,t), h(x)), \tag{35}$$

in other words, the diagram

$$\begin{array}{ccc} M^n & \xrightarrow{\phi^1(t,\cdot)} & M^n \\ h\downarrow & & \downarrow h \\ M^n & \xrightarrow{\phi^2(\tau(\cdot,t),\cdot)} & M^n \end{array} \tag{36}$$

commutes for all $t \in \mathbf{R}$.

This definition is weak enough to get rid of so-called *moduli*, for example the eigenvalues of the linearization at an equilibrium point would be invariants if we insisted on h being a *diffeomorphism*, yet strong enough to guarantee that oriented trajectories of one system are mapped one-to-one to trajectories of the second system (and so in particular TOE dynamics have chain recurrent sets that are mapped componentwise to each other and also have the same orbit diagram.) It is not true, however, that fixing the orbit diagram of a flow fixes the TOE class even for Morse-Smale flows.

For the purposes of control, we have a choice as to what definition of equivalence to use. If we specify that we want to achieve dynamics in a given equivalence class, we come up against problems in order to guarantee the desirable global behaviour. This is because we shall use functions and as we saw the dynamics for which the given function is a Lyapunov function does not have a fixed orbit diagram, in general. It is worth remembering that a global Lyapunov function, according to our theorem, does not fix the orbit diagram, but only the chain recurrent set components and their Conley index.

Let us make these issues more precise and give a description of the topological way of specifying global dynamics through functions.

6.1 Morse specifications of desirable dynamics

A qualitative description of dynamics will be of little use if the dynamics we specify are not structurally stable. The concept is a basic one in the theory of dynamical systems.

A topology can be defined on the space of all vector fields of a fixed differentiability, say the C^k vector fields on M^n (see Hirsch for details [10].). We write the metric space thus obtained $\mathcal{X}^k(M^n)$ (for M^n compact.)

Definition 24 *A flow ϕ is called (C^k-)* **structurally stable** *if there is some open neighborhood of it in the space $\mathcal{X}^k(M^n)$ consisting of topologically orbitally equivalent flows.*

A local definition is possible, leading to the conclusion that hyperbolic equilibria and limit cycle are structurally stable. Globally, transverse intersections of stable and unstable manifolds are structurally stable. More general is the concept of a *normally hyperbolic chain recurrent set.*

The early hope that structurally stable systems are 'dense' in the set of all vector fields proved false with the discovery by Smale of open sets consisting of structurally unstable systems (in four dimensions; later, three dimensional examples were found by Newhouse.)

6.1.1 Gradient-like dynamics

Fix a number of distinct points in M^n, $e_1, \dots, e_P$ and integers k_i, $0 \le k_i \le n$, $i = 1, \dots, P$. The e_i are the desirable equilibria and the k_i are their indices. We define the set

$$E = \{(e_i, k_i) \; ; \; i = 1, \dots, P\}$$

and use the symbol $|E|$ (for the *'support of E'*) to denote the set $\{e_1, \dots, e_P\}$.

Naturally, there are topological limitations on the number and configuration of these 'equilibria' (by the Morse inequalities, for example, which require the number of equilibria at each index level to exceed the corresponding *Betti number*.) A simpler way to guarantee that the specification makes sense dynamically is to also assume given a Morse function h_0, called the **validating function**, having the property that its gradient vector field has the set E as its chain recurrent set (with the stability type given by the indices k_i.)

Definition 25 *A* **Morse specification** $\mathcal{M}$ *of gradient-like dynamics consists of a set E as above, together with a validating Morse function h_0.*

The validating function h_0 has an orbit diagram that implies a partial order on the 'equilibrium points' in $|E|$.

Now consider Morse functions whose gradient dynamics are top. orb. equivalent to those of h_0 and whose critical set is precicely $|E|$. We call these the **Morse-Lyapunov functions** for the Morse specification $\mathcal{M}$ and we denote the set by $\mathcal{F}(\mathcal{M})$. The name is chosen to indicate that these 'global' functions are the analogue of the more familiar 'local' Lyapunov function candidates.

In this setting, we say the dynamics with Morse specification $\mathcal{M}$ has been achieved if there is a section U in $\Gamma(D)$ and an element $h \in \mathcal{F}(\mathcal{M})$ such that

$$dh(X + U)(x) < 0 \tag{37}$$

for all $x \in M^n \setminus |E|$.

We see that this is an **existence problem** that has two components:

1. Finding an 'appropriate' Morse-Lyapunov function h. 'Appropriate' in this context means exactly one for which a suitable feedback control exists.

2. Finding an explicit feedback control, $U \in \Gamma(D)$.

It is reasonably well-understood by now that the first problem is the crucial one; this is clear, for example, in the **control Lyapunov function** approach of Artstein–Sontag (see Bacciotti [1].).

It is worthwhile to point out the main differences between this approach and ours. To start with, the CLF approach is primarily local, in that it deals only with stabilization. This is not too important. The main difference is that the main purpose of the CLF approach is to give a controller in a nice form (a *universal controller*) that can be applied once the Lie derivatives of the Lyapunov function in the state and control directions are known. Thus, one makes it a starting point that there exists a function V such that

$$\inf_{u} \mathcal{L}_{f(x,u)} < 0$$

for all x in a neighborhood of the equilibrium point to be stabilized. (Here the control system is $\dot{x} = f(x,u)$.)

The core of the theory we are developing is to solve the much harder first part of the existence problem. This is indeed a difficult problem and we shall bring a lot of machinery to bear on its solution. We also hope to convince the reader that the topological methods we shall use are necessary and well worth mastering.

Let us hasten to add that we do not have a complete solution at hand (otherwise nonlinear control theory would have been 'solved,' to a large extent!) We do have a reasonably complete theory in the case of constant control distribution. We also have an *obstruction theory* that exploits the various algebraic topological methods available and promises to be even more profitable in the future.

6.1.2 More general invariant sets

Leaving the safe and well-understood world of equilibrium points behind, we must proceed with caution.

Limit cycles present new problems, but these are essentially easy to handle. They of course correspond to requiring the controlled dynamics to have *'oscillatory' behaviour*; for practical applications, **stable** limit cycles are the only types to consider, but doing the general case involves no additional effort.

We must first decide the extent to which we want the limit cycles to be localized in space and also the way the closed orbits are embedded in state space: are we interested in knotted or unknotted orbits; is there a reason to consider orbits that are period-doubled bifurcations of simpler orbits etc.

The case of stabilizing a saddle periodic orbit of a chaotic attractor illustrates these issues (the *'control of chaos'* case.)

Now the definition of a **Morse specification** of Morse–Smale dynamics proceeds in a way similar to that of the previous subsection. Morse-Lyapunov functions are still useful and orbit diagrams are obtained just as in the case of gradient-like systems. Note that Morse-Lyapunov functions are constant on limit cycles.

Finally, to deal with more general invariant sets or with the use of control to alter given chain recurrent sets, we shall prefer a more indirect path. Since we have chosen an approach through Morse functions and the Conley index –both objects that are meant for the study of the 'gradient-like' part of a flow and not for the chain-recurrent part, we 'localize' by isolating a region of state space with the right Conley index for the IIS under consideration; this means that the invariant set isolated has the same Conley index as the invariant set we aim to achieve and the two are related by continuation.

The approach we give is thus similar to the study of bifurcations (local and global), except that control enters in the given affine way instead of through the standard bifurcation normal form.

A simple example, that of 'doubling the period' of an oscillator through small control action, has been considered in [14]. We leave this topic for further study in a later section.

6.2 Global aspects and limitations

We defined the class of Morse–Lyapunov functions for a given Morse decomposition $\mathcal{M}$ to be the class of Morse functions giving TOE gradient dynamics.

On the other hand, the condition for achieving the control objective of equation 37 only requires that the Morse–Lyapunov function h is a global Lyapunov function for the gradient-like dynamics $X + U$. We have seen through examples that the Lyapunov function does not fix the orbit diagram of the flow.

In this subsection, we ask the question:

WHAT ARE THE POSSIBLE ORBIT DIAGRAMS OF FLOWS WITH LYAPUNOV FUNCTIONS IN THE CLASS $\mathcal{F}(\mathcal{M})$?

[This section willbe completed later.]

7 Necessary Conditions for Global Control Design

Our approach to necessary conditions is again based on the dynamical systems and topological ideas presented so far; more specifically, we shall use the existence of (strict) global Lyapunov functions away from the chain recurrent set. In this way, not only do we manage to recover the rather elementary results known for stabilization (Krasnosel'skii–Brockett, Coron), but we are able to give general necessary conditions of a global nature and for arbitrary invariant sets (for example, new results on limit cycles are obtained.) We also strive to make clear the limitations of the set of necessary conditions we derive and the difficulties in applying them to concrete control systems. The treatment is still preliminary and these notes should be read as a draft rather than as the final exposition.

We take the opportunity to give a capsule review of some methods of algebraic topology. There is a good reason for doing this: in order to understand the various necessary conditions and the differences between them, we need to make clear the relations between *homotopy equivalence, index theoretic results such as the Hopf theorems and the transition to the algebraic groups of homology, homotopy etc.*

7.1 Some background and methods of algebraic topology

A thumbnail sketch of a number of important concepts and methods of computation of algebraic topology will be given. There is no effort to make anything rigorous; we simply aim to give the basic keys to the user, so that algebraic topological objects can be computed in applications.

Singular Homology The easiest algebraic object to handle is the graded abelian group $H_*(X)$, the **singular homology group** of a topological space X (with integer coefficients.) Powerful methods for its computation exist; we outline the Mayer–Vietoris sequence and explain the concept of a long exact sequence of a pair and its relation to excision.

The pth-homology group $H_p(X;\mathbf{Z})$ measures in some sense the 'holes' of X that are like p-spheres. The 0th-group $H_0(X)$ is always equal to $\mathbf{Z}$, if X is connected. There is a way of defining **reduced homology groups** $\tilde{H}_k$ so that $\tilde{H}_0(X) = 0$ for any connected space and all higher groups coincide. Let us give some examples.

The group $H_1(S^1)$ is isomorphic to $\mathbf{Z}$ (which is also the same as $\pi_1(S^1)$, the fundamental group of the circle.) This group, therefore, contains the degree information that says that maps from the circle to itself are classified, up to homotopy, by the net number of times they wind around the circumference.

For the sphere of arbitrary dimension, S^m, $m \geq 1$, the only nontrivial homology group is the mth one, $H_m(S^m) \simeq \mathbf{Z}$. So far, we have restated the results on the homotopy groups π_k since, by the *Hurewicz isomorphism theorem*, homotopy and homology groups are isomorphic at the first level (greater than zero) where one, and hence both, are nontrivial (the abelianization of the fundamental group is to be considered for the case $m = 1$.)

A difference between homotopy and homology groups arises first in the case of the two-torus T^2. We have that $H_2(T^2) = \mathbf{Z}$, even though $\pi_2(T^2) = 0$! (The reader may spend a minute or two pondering the difference.)

Mayer–Vietoris Sequence Suppose a space X is the union of two open subsets, $X = A \cup B$, A, B open and $A \cap B \neq \emptyset$. Then there is *long exact sequence*

$$\ldots \rightarrow H_k(A \cap B) \rightarrow H_k(A) \oplus H_k(B) \rightarrow H_k(X) \rightarrow H_{k-1}(A \cap B) \rightarrow \ldots \quad (38)$$

This a surprisingly powerful tool of computation. A first application is to spheres, where one decomposes S^m into two hemispheres overlapping a little, so that A and B are disks and $A \cap B$ is homotopically equivalent to a sphere S^{m-1}. The basis of an inductive argument then gives, using Mayer–Vietoris:

$$\ldots \rightarrow H_k(S^{m-1}) \rightarrow H_k(A) \oplus H_k(B) \rightarrow H_k(S^m) \rightarrow H_{k-1}(S^{m-1}) \rightarrow \ldots \quad (39)$$

yielding, for $k = m$,

$$\ldots \rightarrow 0 \rightarrow 0 \oplus 0 \rightarrow H_m(S^m) \rightarrow H_{m-1}(S^{m-1}) \simeq \mathbf{Z} \rightarrow 0 \ldots \quad (40)$$

and hence $H_m(S^m) \simeq \mathbf{Z}$, by induction, since exactness means the map in the middle is one-to-one and onto, i.e. is an isomorphism.

Long exact sequence of a pair and excision Another very useful method of computation of homology is the one involving the pair (X, A), where $A \subset X$. One gets a long exact sequence

$$\ldots \rightarrow H_k(A) \rightarrow H_k(X) \rightarrow H_k(X, A) \rightarrow H_{k-1}(A) \rightarrow \ldots \quad (41)$$

where the groups $H_k(X, A)$ are the *relative homology groups* of X relative to A. Without getting into the precise definition, we mention that in many important cases these relative groups are isomorphic to the homology groups of the quotient space X/A and thus the long exact sequence of a pair is extremely useful for the computation of the Conley index.

As an example, let us show that D^n/S^{n-1} has the homology of the sphere S^n. For the pair (D^n, S^{n-1}) (the sphere S^{n-1} is the boundary sphere of the disk D^n), we have

$$\ldots \to H_k(D^n) \to H_k(D^n, S^{n-1}) \to H_{k-1}(S^{n-1}) \to H_{k-1}(D^n) \to \ldots \quad (42)$$

and so, setting $k = n$, we get

$$\ldots \to 0 \to H_n(D^n, S^{n-1}) \to \mathbf{Z} \to 0 \to \ldots \quad (43)$$

and hence $H_n(D^n, S^{n-1}) \simeq H_n(D^n/S^{n-1}) \simeq \mathbf{Z}$. Similarly, one finds that, for $k \neq n$ (and nonzero), $H_k(D^n/S^{n-1}) = 0$.

Maps and homomorphisms Given a continuous map $f : X \to Y$, one gets an induced map in homology, which we denote either by f_* or by $H_*(f)$

$$H_*(f) : H_*(X) \to H_*(Y)$$

This is a homomorphism of groups. One understands this to mean that there are maps at each level of homology, i.e. there are homomorphisms $H_k(f) : H_k(X) \to H_k(Y)$ for all k.

The usefulness of homology groups derives from the following basic results

Proposition 4 *If the maps f and g are homotopic, then $f_* = g_*$ (actually equal as homomorphisms.)*

Proposition 5 *If the two spaces X and Y are homotopy equivalent and f has a homotopy inverse, then f_* is an isomorphism and so the homology groups $H_*(X)$ and $H_*(Y)$ are isomorphic.*

Corollary 1 *If f is a homeomorphism of spaces, then f_* is an isomorphism.*

** It is worth pointing out the abstract *homological algebra* viewpoint on these matters, since it, arguably, clarifies the topological approach. The notation $H_*(f)$ was chosen because H_* is a **functor** between two **categories**, the category **hTop** of homotopy equivalence classes of topological spaces and the category **Ab** of abelian groups. (The **objects** of **hTop** are equivalence classes of topological spaces and the **morphisms** are the homotopy equivalence classes of maps. The **objects** of the category **Ab** are abelian groups and the **morphisms** are group homomorphisms.) The above propositions are used to prove the functorial properties in this case (if you know what they are, check them.) **

When the space X is a finite-dimensional manifold, the homology groups $H_k(X)$ are finitely generated; there is in this case the following basic structure theorem for finitely generated abelian groups

Theorem 7 (Structure Theorem) *Any finitely generated abelian group G decomposes uniquely as*

$$G = F \oplus \tau,$$

where F is **free** *and τ is* **torsion.**

In fact, one knows τ in more detail (see for example, Lang [18].) The dimension of the free part of the homology group H_k is called the kth–**Betti number**, $b_k = \dim H_k(X)$. The **Euler characteristic** $\chi(X)$ is defined as the alternating sum of the Betti numbers

$$\chi(X) = \sum_k (-1)^k b_k.$$

7.2 Collections of necessary conditions

We shall need the definition of a Gauss map. We assume that $M^n = \mathbf{R}^n$ or is some open subset of it.

Definition 26

1. *Suppose X is a vector field that is nowhere zero in M^n. The* **Gauss map** *$G_X : M^n \to S^{n-1}$ is defined by $x \mapsto X(x)/|X(x)|$.*

2. *Suppose $N^{n-1} \subset M^n$ is a submanifold such that the vector field X is nowhere zero on N. Then the* **Gauss map** *$G_X|_N : N^{n-1} \to S^{n-1}$ is given by restricting G_X to N^{n-1}. The important thing about this Gauss map is that it is a map between spaces of the same dimension, one of which is a sphere.*

3. *If, as above, $N^{n-1} \subset M^n$ is an oriented submanifold of M, we also define the* **Gauss map** *$G_N : N^{n-1} \to S^{n-1}$ by: $x \mapsto n(x)$, where n is the vector field of 'outward' unit normal vectors to N. This is also a map between manifolds of the same dimension.*

7.2.1 Index–theoretic results

We give an overview of results using the classical topological index of equilibrium points; the presentation is not for the most part novel, but an effort is made to use a modern viewpoint.

If $e \in M^n$ is an isolated equilibrium point of the vector field X, take a 'ball neighborhood' U of e (an open set homeomorphic to a ball) such that e is the only equilibrium of X in U and its boundary $N = \partial U$ is a submanifold homeomorphic to a sphere. Then the Gauss map $G_X|_N$ extends to a map from S^{n-1} to itself

$$S^{n-1} \xrightarrow{h} N \xrightarrow{G_X|_N} S^{n-1},$$

where h is the homeomorphism for the ball neighborhood U.

At the level of homology, we thus get a homomorphism $\psi = G_X|_N \circ h$ from $H_{n-1}(S^{n-1})$ to itself; since this group is isomorphic to $\mathbf{Z}$, we get a homomorphism from $\mathbf{Z} \to \mathbf{Z}$. Such maps are specified by the image of the generator, say α of $H_{n-1}(S^{n-1})$. If $\psi(\alpha) = k\alpha$, then k is the **(topological) index** of the equilibrium e. It does not depend on the precise U chosen.

The classical theorem of Hopf describes maps from the sphere to itself.

Theorem 8 (Hopf's Classification Theorem) *Homotopy equivalence classes of maps $S^{n-1} \to S^{n-1}$ are in a one-to-one correspondence with the integers. For each integer k, the class of maps corresponding to it is called the class of maps of* **degree** k.

For a hyperbolic equilibrium point of (stability) index k, the topological index is equal to $(-1)^{n-k}$. Degree k maps are relatively easy to visualize (consider the equilibrium at the origin of the system written in complex form as $\dot{z} = z^k$.)

It is, moreover, true that maps to the sphere are classified by degree *even when the domain is not a sphere*; this is the following generalization of the Hopf theorem (see Whitehead [26], p.244)

Theorem 9 (Hopf–Whitney) *The homotopy classes of maps of an $(n-1)$-dimensional manifold N^{n-1} to S^{n-1} are in a one-to-one correspondence with the elements of the cohomology group $H^{n-1}(N^{n-1}; \mathbf{Z})$.*

Corollary 2 *If N is orientable, then the maps from N^{n-1} to S^{n-1} are in one-to-one correspondence with the integers; they are thus classified by 'degree.'*

This is because for every orientable manifold, $H^{n-1}(N^{n-1}) \simeq \mathbf{Z}$. If N is **not orientable**, then this group is $\mathbf{Z}_2$ and two maps are homotopic if and only if they have the same mod–2 degree.

The global version of the index result is the *Poincaré–Hopf theorem*:

Theorem 10 (Poincaré–Hopf) *1) Suppose $W^n \subset \mathbf{R}^n$ is a compact subset with nonempty interior such that its boundary is an $(n-1)$-dimensional submanifold of $\mathbf{R}^n$. Suppose X is a vector field on $\mathbf{R}^n$ that is nowhere vanishing on ∂W. Then*

$$\deg G_X|_{\partial W} = \sum_{e_i \in E} \operatorname{ind} e_i. \tag{44}$$

2) Suppose M^n is a compact manifold and X a vector field on M^n with isolated equilibria. If the boundary of M^n is not empty, we require the vector field to point inwards at all points. Then, we have

$$\sum_{e_i \in E} \operatorname{ind} e_i = (-1)^n \chi(M^n) \tag{45}$$

where $\chi(M^n)$ is the Euler characteristic of the manifold M^n and E is the (finite) set of equilibrium points of X.

In particular, the sum of the indices is a 'topological invariant' and does not depend on the vector field chosen.

3) Suppose W^k is any submanifold of $\mathbf{R}^n$ $(0 \leq k \leq n-1.)$ Consider a 'tubular neighborhood' $N_\epsilon(W^k)$ so that $\partial N_\epsilon(W^k)$ is an $(n-1)$-dimensional submanifold of $\mathbf{R}^n$. If X is any vector field on $\mathbf{R}^n$ such that, on W^k, X has nondegenerate equilibria, then

$$\sum_{e_i \in E} \operatorname{ind} e_i = \deg G_X|_{\partial N_\epsilon(W)} \tag{46}$$

We have given different versions of this important Theorem to help the reader find the most convenient way of extracting information in any particular example. Milnor [20] proves versions (2) and (3) and contains a nice discussion.

Remarks: 1) The index already contains considerable topological information since, as we saw, maps to the sphere are classified by degree. For the case of an **attracting equilibrium**, we know that the index is $(-1)^n$, which means that the generator of $H_{n-1}(S^{n-1})$ is mapped to its negative. As a result, the Gauss map gives an **isomorphism** in homology and we

also conclude that the Gauss map must be onto. This is essentially all the information in the *Krasnosel'skii–Brockett–Coron necessary conditions.*
2) It must be emphasized that the index is *'blind'* to all other dynamical features; only the equilibria affect it. Looking at it from the opposite side, only the topological type of the enclosing surface affects the configuration of equilibria inside (to the extent of fixing their index sum.)

A few examples to illustrate these points:

Examples:

1. In $\mathbf{R}^n$, a ball with a vector field pointing inwards at the bounding sphere must contain equilibria whose index sum is $(-1)^n$. If these are all hyperbolic, then the options are
   - A single attracting equilibrium.
   - Two attractors and a one-saddle.
   - If n is even, a single repeller is not ruled out; we notice that the two cases can be distinguished using the Conley index, since the exit set is empty.
   - Other configurations with the above net index sum.
2. In $\mathbf{R}^3$, an embedded torus T^2 gives possible Gauss maps of arbitrary degree. If we know that there are no equilibria enclosed, as for example in the case where the torus isolates a limit cycle, then the degree is zero, by part (1) of the Theorem, *independently of the stability type of the limit cycle.* This means that the Gauss map is homotopic to the constant map in this case. In the particular case of an attracting limit cycle, we thus see that the index does not give any information. We shall give alternative conditions for limit cycles in the next subsection.
3. On the torus T^2, we have, by version (2), that any vector field must have total index sum equal to zero, since $\chi(T^2) = 0$. Thus, we allow vector fields with no equilibria, or one attractor and one saddle, or one attractor, one repeller and two saddles (are these all possible?)

Necessary conditions using the topological index: It should now be clear how to derive necessary conditions for achieving global dynamics using the topological index.

Suppose given a Morse specification of gradient-type $((E, h_0), \mathcal{F}(\mathcal{M}))$. In the state space, any choice of oriented hypersurface that avoids $|E|$ gives an index sum determined by the equilibria enclosed. This then determines the degree of the Gauss map on the boundary. A necessary condition on the control dynamics is then that there should exist sections that have a well defined Gauss map on the boundary and whose degree is as required by the Poincare–Hopf Theorem.

In particular, if the index sum is equal to plus or minus one, then the Gauss map must be *onto*. Other values of the index sum, even zero, may give nontrivial necessary conditions.

Let us remark that the conditions obtained can either be used *locally* to check, for example, stabilizability by requiring the Gauss map to be onto for an arbitrarily small ball surrounding the equilibrium, or *globally*, since the only relevant information is that on the equilibria and any Morse specification leads to index sum conditions that must be true for any compact subset that avoids $|E|$.

These necessary conditions are not always easy to verify, since one needs to check that sections with specified Gauss map degree exist. It is often easier to translate them into algebraic terms (even then we must make explicit methods of computation.) This is attempted in the following section.

7.2.2 Homotopy equivalence and homotopic results

The fundamental result that forms the key to our approach is given below; it is an elemantary fact about some given dynamics and the gradient dynamics of any of its Lyapunov function.

Theorem 11 *Let $\mathcal{V}^{n-1}$ be a compact level set of some Lyapunov function V for the dynamics X on $M^n \subset \mathbb{R}^n$. Then the Gauss maps $G_X|_\mathcal{V}$ and $G_{-\nabla V}|_\mathcal{V}$ are homotopy equivalent.*

Proof: Let X_n be the projection of X onto the span of ∂V; on $\mathcal{V}$, X_n is never zero, by assumption, since V is a Lyapunov function.

Consider the isotopy of vector fields

$$Y_t(x) = (1-t)(X_n(x)) + tX(x)\ ,\ 0 \leq t \leq 1 \tag{47}$$

We have that $Y_0 = X_n$ and $Y_1 = X$.

Now notice that this gives an isotopy in terms of Gauss maps as well. This is because $Y_t(x) \neq 0$ for all t and all $x \in \mathcal{V}$. To see this, write Y_t as

$$Y_t = X_n + t(X - X_n)$$

and notice that $X - X_n$ is orthogonal to X_n, which is nonzero everywhere.

Now define the Gauss maps

$$G_t : \mathcal{V}^{n-1} \to S^{n-1}\ ,\ x \mapsto \frac{Y_t(x)}{|Y_t(x)|}. \tag{48}$$

Since $Y_t(x)$ is nowhere zero, this is well defined, and gives an isotopy between $G_0 = \frac{X_n}{|X_n|} = \frac{-\partial V}{|-\partial V|} = G_{-\partial V}$ and $G_1 = \frac{X}{|X|} = G_X$. □

For reference purposes, let us denote the set of homotopy equivalence classes of maps between two spaces Ω and Ω' by

$$[\Omega, \Omega']$$

If f is a map $f : \Omega \to \Omega'$, we write $[f]$ for its equivalence class. We thus have that $[G_X] = [G_{-\partial V}]$ in $[\mathcal{V}^{n-1}, S^{n-1}]$.

Relations to the index Since the spaces involved are of the same dimension $(n-1)$ and the maps are into the sphere, we have, by the Hopf theorems, that the homotopy equivalence classes are classified by degree.

[This section will be completed at a later draft.]

Limit cycles In the case of a limit cycle γ, we saw that the Gauss map always has degree zero.

Additional necessary conditions are obtained by examining the Gauss map in more detail.

Theorem 12 *Suppose γ is a limit cycle for the dynamics X on $\mathbf{R}^n$. Then*

1. *For any $\epsilon > 0$, there is a neighborhood $N_\delta(\gamma)$ such that $G_X(N_\delta(\gamma)) \subset N_\epsilon(G_X(\gamma))$.*
2. *$G_X(\gamma)$ cannot be contained in a half space of the sphere; equivalently, for any hyperplane $\mathcal{P} \subset \mathbf{R}^n$, $G_X(\gamma) \cap \mathcal{P} \neq \emptyset$. Moreover, for generic $\mathcal{P}$, $|G_X(\gamma) \cap \mathcal{P}|$ is even.*

Proof: The first part is by continuity and the long flow box.

The second part is by contradiction. Suppose there is a hyperplane $\mathcal{P}_a = \{v \in \mathbf{R}^n\ ;\ a(v) = 0\}$, $a \in (\mathbf{R}^n)^*$ such that $G_X(\gamma) \cap \mathcal{P}_a = \emptyset$.

Since $\mathcal{P}_a$ separates S^{n-1} into two parts, we must have that $a(G_X(x))$ is of uniform sign, say negative, for all $x \in \gamma$.

Choose a basis $b_1, \ldots, b_n$ of $\mathbf{R}^n$ such that $a(b_1) = 1$ and $a(b_k) = 0$ for $k > 0$. Write $x_1, \ldots, x_n$ for the coordinates in this basis.

Claim 3 *The function $V(x) = \frac{1}{2}x_1^2$ is a Lyapunov function for X in some open neighborhood of γ.*

To see this, compute $\frac{dV}{dt}|_\gamma$. We have

$$\frac{dV}{dt} = (x_1\, 0 \cdots 0) \cdot \dot{\gamma}$$

and, since $G_X = \frac{X}{|X|}$, this is just $a(X) < 0$.

The claim gives the desired contradiction, since $G_X(\gamma)$ is a closed curve. The last part also follows from this fact. □

The Theorem thus says that, even though the image of the Gauss map of a limit cycle is 'thin,' it still must curve sufficiently so that intersects all possible hyperplanes.

As for the Lyapunov level sets near a limit cycle, we have

Theorem 13 *Suppose γ is a stable limit cycle for some controlled dynamics. Then on each level set of a Lyapunov function near γ each direction appears twice, in other words, for each $v \in S^{n-1}$,*

$$|G_X^{-1}(v)| \geq 2.$$

Thus, even though the Gauss map of the gradient vector field of Lyapunov functions is of degree zero, as it should be by Poincaré–Hopf, it covers the unit sphere twice. In this case we see, therefore, that members of the same homotopy equivalence class can have widely different Gauss images. The proof is omitted.

7.3 The road forward

We have presented ways of deriving necessary conditions for achieving global dynamics and we also pointed out some of their limitations. The basic aim of our analysis, of course, has to be *constructive*. This will be done in the Second Part of the Notes, where the existence problem of sections in the appropriate homotopy equivalence class is addressed. The dynamical tools of the First Part will still be useful. New tools will have to come into play, however, and we hope to give a complete analysis of the obstructions to the existence of sections.

The case of constant control distribution will occupy a large part of the Notes, since the analysis is more satisfactory in this case.

References

[1] Andrea Bacciotti: Local Stabilizability of Nonlinear Control Systems, World Scientific, 1992.

[2] Raoul Bott: Morse Theory Indomitable, Inst. des Hautes Etudes Sc. Publ. Math. **68**, 1988.

[3] Roger W. Brockett: Asymptotic Stability and Feedback Stabilization, in Brockett et. al. eds: Differential Geometric Control Theory, Birkhäuser, 1983, pp.181–191.

[4] Alberto Isidori: Nonlinear Control Systems, 3rd Edition, Springer, 1994.

[5] Charles Conley: IBM Research Center Report, reprinted in [6].

[6] Ergodic Theory and Dynamical Systems, Volume **8***, 'Conley Memorial Volume,' 1988.

[7] Anatoly Fomenko: Differential Geometry and Topology, Gordon Breach, 1985.

[8] Anatoly Fomenko: Visual Geometry and Topology, Springer, 1995.

[9] M. Greenberg and J. Harper: Algebraic Topology: a First Course, Benjamin–Cummings, 1981.

[10] Morris Hirsch: Differential Topology, Springer, 1977.

[11] Dale Husemoller: Fibre Bundles, 3rd ed., GTM **20**, Springer, 1994.

[12] E. Kappos: The Conley Index and Global Bifurcations, Part I: Concepts and Theory, Int.J. Bif. and Chaos, 5–4, 1995, pp.937–953.

[13] E. Kappos: The Conley Index and Global Bifurcations, Part II:Illustrative Applications, Int.J. Bif. and Chaos, to appear, December 1996.

[14] E. Kappos: Global Feedback Dynamics: Some Preliminary Results, IFAC NOLCOS, Capri, 1988.

[15] E. Kappos: A Global Geometrical Input-Output Linearization Theory, IMA J. Math. Con and Inf., 9–1, 1992, pp.1–21.

[16] E. Kappos: Necessary Conditions for Global Feedback Control, Proc. of NOLTA '95, Las Vegas, December 1995.

[17] E. Kappos: The Rolé of Morse-Lyapunov Functions in the Design of Nonlinear Global Feedback Dynamics, Chapter 12, pp.249–267, in A. Zinober ed.: Variable Structure and Lyapunov Control, Springer, 1995.

[18] Serge Lang: Algebra, Addison–Wesley, (several editions), eg 1970.

[19] John Milnor: Morse Theory, Ann. of Math. **51**, PUP, 1963.

[20] John Milnor: Topology from the Differentiable Viewpoint, The University Press of Virginia, Charlottesville, 1965.

[21] John Milnor and J. Stasheff: Characteristic Classes, Ann.Math. **76**, PUP, 1974.

[22] Henk Nijmeijer and Arian van der Schaft: Nonlinear Dynamical Control Systems, Springer, 1990.

[23] Clark Robinson: Dynamical Systems, CRC Press, 1995.

[24] Dietmar Salamon: Connected Simple Systems and the Conley Index of Isolated Invariant Sets, Trans. AMS, **291**–1, pp.1–41. 1985.

[25] Héctor J. Sussmann: Subanalytic Sets and Feedback Control, J.Diff.Eq., **31**, pp.31–52, 1979.

[26] George W. Whitehead, Elements of Homotopy Theory, Springer, GTM 61, 1978.