Lectures on Dynamical Systems

Chapter 1. Introduction
Chapter 2. Spectral theory
Chapter 3. Calculus in Banach space
Chapter 4. Attracting fixed points

Oscar E. Lanford III
Mathematics Department
ETH-Zürich
8092 Switzerland

September 2, 1997

This text is a slightly revised version of part of a set of lecture notes distributed to the students in my course on dynamical systems at the ETH Zürich in the Winter Semester of 1991–92. I have had a number of requests for copies of these notes, and this has encouraged me to think that it might be useful to make them more generally available. The reader should be warned, however, that, since I am working on a major rewriting of this set of notes, I did not want to take the time to polish this older version carefully (although there is a good deal which needs to be done.) I will be grateful for comments about these notes and particularly grateful to have any errors called to my attention.

Oscar E. Lanford III (lanford@math.ethz.ch)

Chapter 1

Introduction

The fundamental motivation for the study of dynamical systems is the desire to understand, qualitatively, the long-time behavior of solutions of differential equations, even if these differential equations cannot be solved explicitly. This subject can be studied on a number of different levels of generality and abstraction. We can situate, roughly, the subject of these lectures by saying that, in contrast to ergodic theory, we will emphasize the *differentiable* character of the objects being studied but that, on the other hand, we will not investigate the properties of Hamiltonian motion which derive specifically from its conservative character. It is customary to refer to the general subject to be addressed in these lectures as the theory of *dissipative differentiable dynamical systems*, where the word "dissipative" is to be taken to mean merely "non-Hamiltonian"; it is rare that one can ever exploit the genuinely dissipative features of systems with friction.

Let us consider, then, a differential equation

$$\frac{dx}{dt} = X(x). \tag{1.1}$$

Note that we are supposing that our differential equation is *autonomous*, i.e., that t does not appear explicitly on the right-hand side. We will refer to the set of possible x's as the *state space*, and generally denote it by M; it might be Euclidean space, or an open set in Euclidean space, or an open set in a Banach space (infinite system of coupled equations, or partial differential equation) or it might be a manifold, or a Banach manifold. The natural framework for most of our investigations will be manifolds, possibly infinite dimensional ones. Nevertheless, we will write all formulas for the simple case of differential equations on Euclidean or Banach spaces. The reason for doing this is that, on the one hand, non-trivial global properties of manifolds play almost no role in the investigations to be described here—most of these analyses, once carried out in Euclidean spaces, are easy to extend to manifolds—and, on the other hand, the systematic use of the formalism of differentiable manifold theory, with co-ordinate patches, tangent bundles, etc., is a serious distraction from the real issues. This expositional strategy doesn't quite work, since we will want sometimes to assume that M is compact, but I don't intend to worry about this; the audience will be assumed to be translating everything into the language of differentiable manifold theory.[1]

This may be a good place to comment on the philosophy to be adopted in the lectures about whether or not to assume that the state space is finite dimensional. The first point to be made is that many of the most penetrating analyses of finite dimensional dynamical systems rest on the use of non-linear functional analysis on infinite spaces, so, even if we are really only interested in finite dimensional systems, we need to develop a fair amount of infinite-dimensional theory. Once that theory is at hand, it turns out that a great deal of dynamical systems theory can be done, at little or no extra cost, without assuming finite dimensionality. Indeed, as is often the case, the discipline of avoiding unnecessary assumptions of finite dimensionality frequently leads to better proofs. I have therefore adopted the point of view that, whenever finite-dimensionality assumptions can be avoided, without making an argument *essentially* more difficult, I avoid them.

One should not, however, overstate what is accomplished by avoiding unnecessary finite-dimensionality assumptions. It does **not**, in itself, make the theory applicable to interesting infinite dimensional systems like the Navier-Stokes equation. The problem with the Navier-Stokes equation, and most other partial-differential evolution equations, is that it is definitely not a smooth dynamical system. If one writes it in the schematic form 1.1, the $X(x)$ appearing on the right-hand side is not a smooth function of x but involves, e.g., unbounded operators. To a certain extent, the lack of smoothness of $X(x)$ can be circumvented by exploiting the fact that the equation of motion, like the heat equation, produces smoothing when run forward in time, so time evolution operators for positive times may be smooth even if the infinitesimal time evolution operators are not. Analyzing the dynamical properties of such evolution equations as the Navier-Stokes equation will certainly need to make use of the methods of infinite dimensional dynamical systems theory, but will require in addition a detailed study of the particular evolution equation. That kind of

[1] Mike Shub has suggested that most theorems about differentiable dynamical systems are proved assuming the state space to be a compact Euclidean space.

analysis of particular partial differential equations will not be given here.

If $X(x)$ in 1.1 is Lipschitz continuous, then solutions exist, at least locally, for arbitrary initial data, and they are unique. As simple and apparently well-behaved an equation as

$$\frac{dx}{dt} = x^2$$

shows that solutions need not exist for all times. In any event, we introduce the associated *solution mapping* f^t, defined as follows: $f^{t_0}(x_0)$ is defined if and only if the differential equation admits a solution $x(t)$ defined for all t between 0 and t_0, and its value in this case is $x(t_0)$. By the uniqueness of solutions of the differential equation,

$$f^{t_1}\left(f^{t_2}(x)\right) = f^{t_1+t_2}(x),$$

whenever the left-hand side is defined. This condition is expressed by saying that f^t is a *semi-flow*. If $f^t(x)$ is defined for all (x,t), then each f^t is a diffeomorphism (i.e., a smooth mapping of M onto itself with a smooth inverse) and $f^{t_1}f^{t_2} = f^{t_1+t_2}$. Furthermore, $f^t(x)$ is jointly smooth in (t,x). We thus say that f^t is a *smooth flow*. Conversely, any smooth flow f^t is the solution flow for the differential equation with

$$X(x) = \left.\frac{df^t}{dt}(x)\right|_{t=0}.$$

A flow is a one-parameter group of diffeomorphisms with parameter $t \in \mathbb{R}$. Often, one considers instead *discrete* one-parameter groups, i.e, the set of all iterates f^n of a single diffeomorphism. Sometimes, moreover, one drops the condition that f be invertible, i.e., one considers the positive iterates of an arbitrary smooth map.

There are a number of reasons for considering maps rather than flows. One is the fact that most theorems about maps have analogues for flows, and vice versa, but the versions for maps is often technically simpler and exhibits the main ideas in a less encumbered form. Secondly, if f^t is a flow, then the large-t behavior of f^t is essentially determined by the behavior of iterates of any single f^{t_0} with $t_0 \neq 0$, e.g., $t_0 = 1$. Thus any theorem about maps gives information about flows. A third reason is the "surface of section" idea of Poincaré which, in some cases, reduces the study of a flow in an m-dimensional space to that of a mapping in an $m-1$-dimensional space. The idea is as follows: Let f^t be a flow in an m-dimensional state space M and let Σ be a hypersurface (codimension-one submanifold) of M which is, at each of its points, transverse to the generator $X(x)$ of f^t. If the forward orbit of $x \in \Sigma$ crosses Σ again, define

$\Phi(x)$ to be the first such crossing point (and leave $\Phi(x)$ undefined otherwise.) The mapping Φ so defined is called the *first return map* or *Poincaré map* for Σ. It follows easily from the Implicit Function Theorem that the domain of Φ is an open subset of Σ and that Φ is smooth on its domain. It is often the case that, if Σ is judiciously chosen, quite a lot of the dynamics of f^t is captured by the dynamics of the iterates of Φ. For example, a fixed or periodic point x_0 for Φ corresponds to a closed orbit (periodic solution) for f^t, and the long-time behavior of orbits with initial point near x_0 can be determined from the behavior of iterates of Φ near x_0. Fourthly, iteration of maps provides a natural way to analyze the behavior of solutions of differential equations which are not autonomous but periodic. Consider a differential equation

$$\frac{dx}{dt} = X(t, x),$$

where X is periodic in t: $X(t + \tau, x) = X(t, x)$. If solutions exist for all time, then we can construct the corresponding solution mapping f^t. This will *not* be a flow—the group property will fail—but it will remain true that $f^{n\tau} = (f^\tau)^n$. Thus, the iterates of the time τ mapping essentially determine the long-term behavior of solutions of the differential equation. Finally, interesting phenomena occur in state spaces of smaller dimension for mappings than for flows and may therefore be more comprehensible. Roughly speaking, flows on spaces of dimension less than three cannot exhibit any interesting asymptotic behavior, while invertible maps can already do interesting things in two dimensions and non-invertible maps already in one dimension.

Chapter 2

Rudiments of spectral theory for general bounded operators on Banach spaces

2.1 Orientation

Any normal operator A on a finite dimensional (complex) Hilbert space $\mathcal{X}$ has a very simple structure: If the eigenvalues of A are $\lambda_1, \ldots, \lambda_m$, and if $\mathcal{X}_i$ is the set of vectors ξ such that $A\xi = \lambda_i \xi$ (i.e., the eigenvectors with eigenvalue λ_i together with the zero vector), then $\mathcal{X}$ is the direct sum of the $\mathcal{X}_i$.

This simple structure theory can be extended satisfactorily in two directions. The first extension is to *general* linear operators on *finite-dimensional* vector spaces. The principal complication here is the necessity to take into account the possible existence of *generalized eigenvectors*, i.e., vectors annihilated by some power of $A - \lambda\mathbf{1}$, λ a complex number, but not necessarily by $A - \lambda\mathbf{1}$ itself. If, again, we let the eigenvalues of A be $\lambda_1, \ldots, \lambda_m$, and if we now let $\mathcal{X}_i$ denote the set of all vectors ξ such that $(A - \lambda_i\mathbf{1})^j\xi = 0$ for some integer $j \geq 1$ then, again, each $\mathcal{X}_i$ is a linear subspace mapped into itself by A, and $\mathcal{X}$ is the direct sum of the $\mathcal{X}_i$. To complete the structure theory it is only necessary to work out what A restricted to any one of the $\mathcal{X}_i$ can look like, and this is a relatively straightforward task.

The second extension is to *normal* operators on infinite dimensional Hilbert spaces. Leaving aside problems concerning unbounded operators, the main complication here is the fact that such an operator need not have any eigenvectors at all. The direct-sum decomposition into eigenspaces in the finite dimensional case is replaced in this more general situation by an integral representation $A = \int \lambda P(d\lambda)$ where $P(d\lambda)$ is a projection valued measure on the complex plane. The intuitive interpretation of this projection valued measure is that the range of projection $P(X)$ corresponding to a Borel set X is the proper infinite-dimensional analogue of the linear span of the eigenvectors with eigenvalues in X in the finite-dimensional case.

For a general bounded operator A on a Banach space $\mathcal{X}$, our knowledge is much less complete. There is, however, a good definition of the spectrum in this general situation, and one can construct useful spectral projections $P(X)$ associated with subsets X of the spectrum which are at a non-zero distance from the rest of the spectrum.

2.2 The spectrum

From now on in this chapter, A will denote a bounded linear operator on a Banach space $\mathcal{X}$. By the *resolvent set* of A we mean the set of complex numbers λ such that $\lambda\mathbf{1} - A$ has a bounded (everywhere defined) inverse. By the closed graph theorem, this is the same as saying that $\lambda\mathbf{1} - A$ is injective—i.e., that λ is *not* an eigenvalue of A—and that the range of $\lambda\mathbf{1} - A$ is all of $\mathcal{X}$. The inverse $(\lambda\mathbf{1} - A)^{-1}$ is called the *resolvent* (of A, or of λ, or of (A, λ), depending on the context.) The *spectrum* of A means the complement of the resolvent set; we will denote it by $\sigma(A)$.

One of the principal tools of elementary spectral theory is the systematic exploitation of the *Neumann series*. This means the following: Let U be an operator with a bounded inverse, and V some other operator. Then, formally, the infinite series

$$U^{-1} - U^{-1}VU^{-1} + \ldots + (-1)^n \left(U^{-1}V\right)^n U^{-1} + \ldots \tag{2.1}$$

should give the inverse of $U + V$. It is easy to check that this is true in the sense that, whenever the series (2.1) converges in norm, $U + V$ is invertible and the sum of the series is $(U + V)^{-1}$. Furthermore, the series is easily seen to converge if $\|V\| < \left(\|U^{-1}\|\right)^{-1}$ or, more generally, if $\left\|\left(U^{-1}V\right)^m\right\| < 1$ for some positive integer m. This remark shows, for example that the resolvent set for any A is open, and that the resolvent $(\lambda\mathbf{1} - A)^{-1}$ varies analytically with λ on the resolvent set, i.e., that it is given locally in the neighborhood of each point by a norm-convergent power series. Since the spectrum is defined to be the complement of the resolvent set, it is closed.

It is quite a non-trivial fact that the spectrum of an operator A cannot be empty. To prove it, assume the contrary, i.e., that the resolvent set of some operator A is the whole complex plane. Then $(\lambda\mathbf{1} - A)^{-1}$ is an entire operator valued function, which furthermore is easily seen—using the Neumann series—to go to zero at infinity. Thus, for any ξ in $\mathcal{X}$ and any ϕ in the dual of $\mathcal{X}$, $\phi\left((\lambda\mathbf{1} - A)^{-1}\xi\right)$ is an entire complex valued analytic function vanishing at infinity which therefore, by Liouville's Theorem, vanishes identically. Since this is true for all ξ, ϕ, it follows that $(\lambda\mathbf{1} - A)^{-1}$ vanishes for all λ, and this is manifestly impossible.

The argument in the preceding paragraph illustrates a general method of extending results from complex analysis to vector- and operator-valued functions. Many important results are easy to extend in this way, among them the Cauchy Integral Theorem and Cauchy Integral Formula; the integrals involved can be interpreted simply as Riemann integrals.

By elementary properties of the norm, $\|A^{n+m}\| \leq \|A^n\|\,\|A^m\|$. From this it follows easily that $\|A^n\|^{1/n}$ converges as n goes to infinity, and that the limit is actually the infimum of the $\|A^n\|^{1/n}$. The limit is called the *spectral radius* of A; we will denote it by $\rho(A)$. The terminology is justified by the fact that the spectral radius turns out to be equal to the supremum of $|\lambda|$ over the spectrum of A. To see this, we first invoke the Neumann series to write

$$(\lambda\mathbf{1} - A)^{-1} = \sum_{n=0}^{\infty} A^n / (\lambda)^{n+1} \tag{2.2}$$

wherever the right-hand side converges. Since the right-hand side certainly converges for $|\lambda| > \rho(A)$, this shows that the supremum of $|\lambda|$ over the spectrum of A is at most equal to $\rho(A)$. On the other hand, integrating (2.2) term-by-term gives

$$A^n = \frac{1}{2\pi i}\oint_{\gamma_R} \lambda^n\,(\lambda\mathbf{1} - A)^{-1}\,d\lambda,$$

where γ_R is the circle $\{|\lambda| = R\}$, traversed in the counterclockwise sense, and R is any positive number larger than $\rho(A)$. By the Cauchy Integral Theorem we can reduce the radius R to any number strictly larger than the supremum of $|\lambda|$ over the spectrum. Estimating the integral in a straightforward way gives a bound $\|A^n\| \leq B_R R^n$ for any such R and any positive integer n, and from this bound it follows that $R \geq \rho(A)$; Hence, the supremum of $|\lambda|$ over the spectrum cannot be strictly smaller than $\rho(A)$, so equality is established.

Since $\rho(A) = \inf_n \|A^n\|^{1/n}$, we always have $\rho(A) \leq \|A\|$. The 2×2 matrix

$$\begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}$$

shows that this inequality cannot always be made an equality, even with the freedom to replace the norm by an equivalent one. However, we can do almost as well:

Proposition 2.2.1 *If $R > \rho(A)$, there exists a norm $|||\ \ |||$, equivalent to the original norm $\|\ \ \|$, such that $|||A||| \leq R$*

Proof. Since $R > \rho(A)$, $\|(A/R)^n\|$ goes to zero as n goes to infinity. Hence, if we put

$$|||\xi||| = \sup_{n \geq 0} \|(A/R)^n \xi\|\,,$$

we get

$$\|\xi\| \leq |||\xi||| \leq \left(\sup_{n \geq 0} \|(A/R)^n\|\right) \|\xi\|,$$

i.e., the norm $|||\ \ |||$ is equivalent to the norm $\|\ \ \|$. On the other hand,

$$|||A\xi||| = \sup_{n \geq 0} \|(A/R)^n A\xi\| = R \sup_{n \geq 0} \left\|(A/R)^{n+1}\xi\right\| \leq R|||\xi|||,$$

so $|||A||| \leq R$, as desired. □

2.3 Spectral projections

A *projection* will mean a bounded operator P such that $P^2 = P$, i.e., such that $P\xi = \xi$ for all ξ in the range of P. Note that P is a projection if and only if $\mathbf{1} - P$ is. To specify a projection is essentially the same thing as to specify a direct-sum decomposition of $\mathcal{X}$ into two closed complementary subspaces. In one direction: Given a projection P, we let $\mathcal{X}_1$ denote the range of P (or, equivalently, the null space of $\mathbf{1} - P$) and $\mathcal{X}_2$ the null space of P. Only the zero vector is in both $\mathcal{X}_1$ and $\mathcal{X}_2$. If ξ is any vector, then $P\xi \in \mathcal{X}_1$ and $\xi - P\xi \in \mathcal{X}_2$; hence, $\mathcal{X}_1$ and $\mathcal{X}_2$ span $\mathcal{X}$. Thus $\mathcal{X} = \mathcal{X}_1 \oplus \mathcal{X}_2$. In the other direction: Given a direct sum decomposition, the mapping $\xi_1 \oplus \xi_2 \longmapsto \xi_1$ is a projection and the Closed Graph Theorem implies that it is continuous if $\mathcal{X}_1$ and $\mathcal{X}_2$ are closed. In this case, we speak of the projection *onto* $\mathcal{X}_1$ *along* $\mathcal{X}_2$.

Note that it it necessary to specify both the range and the null space in order to specify a projection. This is contrary to experience with spectral theory for normal operators, where the projections encountered are all self-adjoint or *orthogonal* projections, i.e., ones for which the null space is the orthogonal complement of the range; evidently, in this case, it suffices to specify one of them. It is easy to see that a projection P commutes with a bounded operator A if and only if both the range and null space of P are mapped into themselves by A (i.e., are "invariant under A"). Again, the invariance of the null space follows from that of the range and vice versa if A is a normal operator and P an orthogonal projection, but the two conditions are independent in the general case.

The main piece of spectral theory which persists in the general case we are considering is the following: To any splitting $\sigma(A) = S_1 \cup S_2$ into two *disjoint closed* subsets (which, by the compactness of the spectrum, must be at a non-zero distance from each other), there corresponds a direct sum splitting $\mathcal{X} = \mathcal{X}_1 \oplus \mathcal{X}_2$ of the space into two closed invariant subspaces such that the spectrum of the restriction of A to $\mathcal{X}_1$ is S_1 and that of the restriction of A to $\mathcal{X}_2$ is S_2. Furthermore, the splitting is invariant not only under A itself, but under any operator which commutes with A, and this condition, together with what was said above, determines it uniquely.

The construction of this splitting is trivial in the case of a diagonalizable operator on a finite-dimensional space—$\mathcal{X}_1$ is the linear span of the eigenvectors with eigenvalues belonging to S_1 and $\mathcal{X}_2$ the linear span of the eigenvectors with eigenvalues belonging to S_2. Described in this way, the construction does not generalize; the key to generalizing it is the observation that the projection onto $\mathcal{X}_1$ along $\mathcal{X}_2$ can be written as

$$\frac{1}{2\pi i}\oint_\gamma (\lambda\mathbf{1} - A)^{-1}\, d\lambda,$$

where γ is a contour enclosing all the points of S_1 and none of the points of S_2. To see that the above integral formula does give the indicated projection, it is only necessary to verify that the integral is the identity on eigenvectors with eigenvalues in S_1 and zero on eigenvectors with eigenvalues in S_2, and this follows easily, for diagonalizable operators, from the Cauchy Integral Formula. We should note here that the description of the conditions on the "contour" γ were a little imprecise. What we need, precisely, is that γ is a *cycle*, i.e., a formal sum of closed paths, in the complement of the spectrum, which winds once around each point of S_1 and zero times around each point of S_2, i.e., $(1/2\pi i)\oint_\gamma (\lambda - \lambda')^{-1}\, d\lambda$ is one for λ' in S_1 and zero for λ' in S_2.

With this simple example as motivation, we formulate the following more general result:

Proposition 2.3.1 *Let A be a bounded operator on the Banach space $\mathcal{X}$, $\sigma(A) = S_1 \cup S_2$ a decomposition of the spectrum of A into two closed disjoint subsets, and γ a cycle in $\sigma(A)$ winding once around each point of S_1 and zero times around each point of S_2. Then*

$$P = \frac{1}{2\pi i}\oint_\gamma (\lambda \mathbf{1} - A)^{-1}\, d\lambda$$

is a projection which commutes with every bounded operator commuting with A (so in particular commutes with A itself). The integral depends only on S_1, S_2, not on γ. Furthermore, the spectrum of the restriction of A to the range of P is S_1 and the spectrum of the restriction of A to the null space of P is S_2. The properties of P listed determine it uniquely.

We will refer to P as the *spectral projection* of A corresponding to the subset S_1 of its spectrum.

Proof. We first remark that it follows easily from a sufficiently general form of the Cauchy Integral Theorem (see, e.g., L. Ahlfors, *Complex Analysis*, 2nd. ed., Theorem 18 of Chapter 4) that the integral does not depend on γ. Also, in the case where S_2 is empty and γ is a sufficiently large circle centered at the origin, we have already noted that integrating the Neumann series term-by-term gives $P = \mathbf{1}$. From this and the fact that P does not depend on γ, it follows that interchanging the roles of S_1 and S_2 replaces P by $\mathbf{1} - P$.

We now want to show that P, as defined by the contour integral, is a projection. For this purpose, we will need the following topological remark: Let γ be as in the statement of the proposition. Then there exists another cycle γ_1, which also winds once around each point of S_1 and zero times around each point of S_2, and which is "inside" γ in the sense that γ winds once around each point of γ_1 but γ_1 winds zero times around each point of γ. One way to guarantee this is to take γ_1 to lie in a small neighborhood of S_1. Then

$$\begin{aligned}
P^2 &= \frac{1}{2\pi i}\oint_\gamma (\lambda \mathbf{1} - A)^{-1}\, d\lambda \frac{1}{2\pi i}\oint_{\gamma_1} (\lambda_1 \mathbf{1} - A)^{-1}\, d\lambda_1 \\
&= \left(\frac{1}{2\pi i}\right)^2 \oint_\gamma \oint_{\gamma_1} (\lambda \mathbf{1} - A)^{-1} (\lambda_1 \mathbf{1} - A)^{-1}\, d\lambda d\lambda_1 \\
&= \left(\frac{1}{2\pi i}\right)^2 \oint_{\gamma_1} (\lambda_1 \mathbf{1} - A)^{-1} \oint_\gamma \frac{1}{\lambda - \lambda_1} d\lambda d\lambda_1 \\
&\quad - \left(\frac{1}{2\pi i}\right)^2 \oint_\gamma (\lambda \mathbf{1} - A)^{-1} \oint_{\gamma_1} \frac{1}{\lambda_1 - \lambda} d\lambda_1 d\lambda,
\end{aligned}$$

where we have used the identity

$$U^{-1} - V^{-1} = U^{-1}(V - U)V^{-1}$$

with $U = \lambda\mathbf{1} - A$ and $V = \lambda_1\mathbf{1} - A$ to write

$$(\lambda\mathbf{1} - A)^{-1}(\lambda_1\mathbf{1} - A)^{-1} = (\lambda - \lambda_1)^{-1}\left((\lambda_1\mathbf{1} - A)^{-1} - (\lambda\mathbf{1} - A)^{-1}\right).$$

Finally, since γ_1 is inside γ, we have

$$\oint_\gamma \frac{1}{\lambda - \lambda_1} d\lambda = 2\pi i; \quad \oint_{\gamma_1} \frac{1}{\lambda_1 - \lambda} d\lambda_1 = 0,$$

and hence $P^2 = P$, as desired.

It is clear from the definition that P commutes with any bounded operator which commutes with A. In particular, it commutes with A, so the range and null space of P are invariant subspaces for A. We want next to show that the spectrum of A restricted to the range of P is contained in S_1, i.e., that, if $\lambda \in S_2$, then $(\lambda\mathbf{1} - A)$ is invertible on the range of P. We prove this by constructing its inverse explicitly as a contour integral. Let γ be as in Proposition 1; let $\lambda \in S_2$; and let

$$B = \frac{1}{2\pi i}\oint_\gamma \frac{1}{\lambda - \lambda_1}(\lambda_1\mathbf{1} - A)^{-1} d\lambda_1.$$

Then B commutes with A; hence also with P, and

$$\begin{aligned}
(\lambda\mathbf{1} - A)B &= \frac{1}{2\pi i}\oint_\gamma \frac{\lambda - \lambda_1 + (\lambda_1 - A)}{\lambda - \lambda_1}(\lambda_1\mathbf{1} - A)^{-1} d\lambda_1 \\
&= \frac{1}{2\pi i}\oint_\gamma \left[(\lambda_1\mathbf{1} - A)^{-1} + \frac{1}{\lambda - \lambda_1}\right] d\lambda_1 \\
&= P + \frac{1}{2\pi i}\oint_\gamma \frac{1}{\lambda - \lambda_1} d\lambda_1 \\
&= P,
\end{aligned}$$

since $\lambda \in S_2$ and γ winds zero times around S_2. Thus, if $P\xi = \xi$, $(\lambda\mathbf{1} - A)B\xi = \xi$, and hence, since A and B commute, the restriction of B to the range of P is a two-sided inverse to the restriction of $(\lambda\mathbf{1} - A)$ to the same subspace.

Interchanging the roles of S_1 and S_2 shows that the spectrum of the restriction of A to the null space of P is contained in S_2. But since the union of the spectral of the restrictions of A to S_1 and S_2 must fill up the whole spectrum of A, it follows that the spectrum of the restriction of A to the range of P (respectively the null space of P) must be all of S_1 (respectively, all of S_2.

It remains to prove the uniqueness of P. Thus, let $\hat{P}$ be another projection which commutes with every operator commuting with A and such that the spectrum of A restricted to the range (null space) of $\hat{P}$ is S_1 (S_2). Then P and $\hat{P}$ commute, so $P(\mathbf{1} - \hat{P})$ is again a projection, and it commutes with A since both P and $\hat{P}$ do. Since its range is contained in both the range of P and the null space of $\mathbf{1} - \hat{P}$, the spectrum of A restricted to its range is contained both in S_1 and in S_2. Since S_1 and S_2 are disjoint, this range must be zero-dimensional—Recall that a bounded operator on a Banach space of non-zero dimension has non-empty spectrum—so $P - P\hat{P} = 0$, i.e., $P = P\hat{P}$. Interchanging the roles of P and $\hat{P}$ gives $\hat{P} = \hat{P}P$, so $P = \hat{P}$, as desired.

General spectral theory gives a bit more information if the operator A is *compact*, i.e., if the closure of the image of the unit ball under A is compact in the norm topology. In this situation it can be shown that the spectrum of A is a finite or countable set with no accumulation point except, perhaps, zero. Thus, there is a spectral projection for each non-zero point of the spectrum. Furthermore, each of these spectral projections is finite-dimensional, i.e., if $\lambda \neq 0$ is in the spectrum of the compact operator A, and if P is the associated spectral projection, then the range of P is a finite-dimensional space invariant for A. The restriction of A to this invariant space is thus a linear operator on a finite-dimensional space whose spectrum consists of the single point λ; it thus has a Jordan Canonical Form with all diagonal elements equal to λ.

It is illuminating to see what the structure of one-dimensional spectral projections is. Let P be a projection with one-dimensional range, and let ξ_0 be a vector spanning its range. The null space of P is a codimension-one subspace of $\mathcal{X}$, i.e., a hyperplane. Any hyperplane is the null space of some linear functional which is uniquely determined up to multiplication by a non-zero constant. Let ϕ_0 be a linear functional whose null space is the null space of P. Since ξ_0 is not in the null space of P, $\phi_0(\xi_0) \neq 0$; by multiplying ϕ_0 by an appropriate constant, we can arrange to have $\phi_0(\xi_0) = 1$. We now claim that, for a general vector ξ, $P\xi = \phi_0(\xi)\xi_0$. The proof of this is almost immediate; $\xi \mapsto \phi_0(\xi)\xi_0$ defines a linear operator which is easily checked to be a projection whose range is the set of multiples of ξ_0, i.e., the range of P, and whose null space is the null space of ϕ_0, i.e., the null space of P.

Now suppose that P is a spectral projection for some operator A. Then, since the range of P is invariant for A, ξ_0 must be mapped by A to a multiple of itself, i.e., must be an eigenvector of A with some eigenvalue λ. We define the *adjoint* A^* of A to be the mapping of the dual $\mathcal{X}^*$ of $\mathcal{X}$ to itself defined by $(A^*\phi)(\xi) = \phi(A\xi)$ for all $\xi \in \mathcal{X}$ and all $\phi \in \mathcal{X}^*$. Since the null space of ϕ_0 is the same as the null space of P and hence is mapped into itself by A, it follows that $A^*\phi_0$ vanishes on the null space of ϕ_0 and hence that $A^*\phi_0$ is a multiple of ϕ_0, i.e., that ϕ_0 is an eigenvector of A^*. Furthermore,

$$\lambda = \lambda\phi_0(\xi_0) = \phi_0(A\xi_0) = A^*\phi_0(\xi_0),$$

which implies that the corresponding eigenvalue is again λ. Summing up the results of the above arguments: If λ is an isolated point of the spectrum of A, and if the spectral projection corresponding to $\{\lambda\}$ is one dimensional, then that spectral projection may be written as

$$\xi \mapsto \phi_0(\xi)\xi_0,$$

where ξ_0 is an eigenvector of A with eigenvalue λ, ϕ_0 is an eigenvector of A^* with the same eigenvalue, and these eigenvectors are normalized to that $\phi_0(\xi_0) = 1$. One of the lessons to be drawn from this is that the eigenvectors of A^* are as important as those of A in determining the "spectral representation" of A.

2.4 Real vector spaces

There is one more small twist needed for most applications to dynamical systems theory. In the above discussion, we assumed tacitly that $\mathcal{X}$ is a Banach space over the *complex* numbers; even in the finite-dimensional case, an operator on a real vector space need not have any spectrum on the real axis. In dynamical systems theory, on the other hand, we usually want to work in *real* vector spaces, but would still like to be able to use simple spectral theory.

One satisfactory way to develop spectral theory for operators on a real vector space $\mathcal{X}$ is to work systematically with an associated complex vector space, called the *complexification* $\mathcal{X}_{\mathbb{C}}$ of $\mathcal{X}$. As an additive group, $\mathcal{X}_{\mathbb{C}}$ is simply $\mathcal{X} \times \mathcal{X}$, and multiplication by complex numbers is defined by

$$(\alpha + i\beta)(\xi_1, \xi_2) = (\alpha\xi_1 - \beta\xi_2, \alpha\xi_1 + \beta\xi_2).$$

It is easy to check that this does make $\mathcal{X} \times \mathcal{X}$ into a complex vector space. A norm on $\mathcal{X}$ can be extended to a norm on $\mathcal{X}_{\mathbb{C}}$ by

$$\|(\xi_1, \xi_2)\| = \sup_{\theta} \|\cos\theta\xi_1 + \sin\theta\xi_2\|;$$

again it is easy to check that this does define a norm on the complex vector space $\mathcal{X}_{\mathbb{C}}$ and that the complexified space is complete if the initial real space is.

If A is any real-linear operator on $\mathcal{X}$, we construct a corresponding complex linear operator $A_{\mathbb{C}}$, called the complexification of A, on $\mathcal{X}_{\mathbb{C}}$, by $A_{\mathbb{C}}(\xi_1, \xi_2) =$

$(A\xi_1, A\xi_2)$; again, it is necessary—but trivial—to check that this does define a complex-linear operator and that its norm is the same as the norm of A. With the construction available, we can then interpret any statement about spectral properties of A as referring instead to $A_{\mathbb{C}}$.

Although this device permits us to apply spectral theory to operators on real-linear vector spaces, it is not entirely satisfactory because it entails working in the auxiliary space $\mathcal{X}_{\mathbb{C}}$, and this is often inconvenient. There are some statements in spectral theory, however, which, although most conveniently proved by passing through the complexification, can be formulated directly in the original space $\mathcal{X}$. The following proposition gives one such statement:

Proposition 2.4.1 *Let A be a bounded linear operator on the real Banach space $\mathcal{X}$. Then the spectrum of A (i.e., the spectrum of $A_{\mathbb{C}}$) is invariant under complex conjugation. Let $\sigma(A) = S_1 \cup S_2$ be a decomposition of the spectrum of A into two disjoint closed subsets invariant under complex conjugation. Then there is a unique projection P on $\mathcal{X}$, commuting with every operator which commutes with A, such that the spectrum of the restriction of A to the range of P is S_1 and the spectrum of the restriction of A to the null space of P is S_2.*

Proof. Not every complex-linear operator on $\mathcal{X}_{\mathbb{C}}$ is the complexification of a real-linear operator on $\mathcal{X}$. It is not difficult to see, in fact, that an operator is a complexification if and only if it commutes with the complex conjugation operator $K : (\xi_1, \xi_2) \mapsto (\xi_1, -\xi_2)$. Furthermore, $A_{\mathbb{C}} - \lambda\mathbf{1}$ is invertible if and only if $K^{-1}(A_{\mathbb{C}} - \lambda\mathbf{1})K = A_{\mathbb{C}} - \bar{\lambda}\mathbf{1}$ is, so the spectrum of $A_{\mathbb{C}}$ is invariant under complex conjugation.

Now let S_1, S_2 be as in the statement of the proposition, and let $P_{\mathbb{C}}$ denote the spectral projection belonging to the subset S_1 of the spectrum of $A_{\mathbb{C}}$. The uniqueness statements in Proposition 3.1 imply that $K^{-1}P_{\mathbb{C}}K = P_{\mathbb{C}}$, i.e., that $P_{\mathbb{C}}$ is, as its name suggests, the complexification of a real-linear operator P on $\mathcal{X}$. It is nearly immediate that this operator has the properties asserted in the proposition. To prove uniqueness, we note that the complexification of any operator with the asserted properties has the properties which characterize the spectral projection $P_{\mathbb{C}}$. □

A linear operator is A said to be *hyperbolic* if its spectrum does not intersect the unit circle.

Corollary 2.4.2 *Let A be a hyperbolic linear operator on the real Banach space $\mathcal{X}$. Then there is a splitting $\mathcal{X} = \mathcal{X}_s \oplus \mathcal{X}_u$ into subspaces $\mathcal{X}_s$, $\mathcal{X}_u$ each of which is invariant under all operators commuting with A, so that the spectrum of the restriction of A to $\mathcal{X}_s$ is inside the unit circle and the spectrum of the restriction of A to $\mathcal{X}_u$ is outside the unit circle.*

We will refer to $\mathcal{X}_s$, $\mathcal{X}_u$ respectively as the *stable* and *unstable* subspaces of A.

Exercises.

1. Evidently, a complex-linear vector space $\mathcal{X}$ is also real-linear and so can re-complexified. This doubles its dimension if it is finite dimensional. Similarly for complex-linear operators. What is the relation between the spectrum and spectral projections of a complex linear operator A and of its re-complexification $A_{\mathbb{C}}$?

2. Again let $\mathcal{X}$ be a complex-linear space, A a complex-linear operator on it. Suppose $\mathcal{X}$ is equipped with a complex conjugation, i.e., a real-linear operator K such that $K^2 = \mathbf{1}$ and $K\lambda\xi = \bar{\lambda}K\xi$. What is the relation between the spectra of the complex-linear operator A and the real-linear operator KA?

3. In the notation of Corollary 2.4.2, show that $\mathcal{X}_s$ can be characterized as the set of vectors ξ such that $A^n\xi$ is uniformly bounded in n. Give an analogous characterization of $\mathcal{X}_u$ which does not assume that A is invertible.

Chapter 3

Differential calculus on Banach spaces

Much of elementary multidimensional differential calculus extends in a straightforward way to Banach spaces. In this section, we sketch this extension, concentrating mostly on pointing out the few places where the standard arguments do *not* work.

3.1 Differentiable mappings

Let $\mathcal{X}$, $\mathcal{Y}$ be Banach spaces, f a mapping defined on some subset $\mathcal{D}$ of $\mathcal{X}$ and taking values in $\mathcal{Y}$. We say that f is *differentiable* at a point x of $\mathcal{D}$ provided that x is an interior point of $\mathcal{D}$ and that there is a bounded linear operator $L : \mathcal{X} \longmapsto \mathcal{Y}$ such that

$$f(x+\xi) - f(x) - L\xi = o(\|\xi\|) \quad \text{as } \xi \to 0$$

i.e., such that for any $\epsilon > 0$ there is a $\delta > 0$ such that

$$\|f(x+\xi) - f(x) - L\xi\| \leq \epsilon\|\xi\| \quad \text{when } \|\xi\| \leq \delta.$$

If this is the case, the linear operator L is unique, and is called the *derivative* of f at x; we will denote it by $Df(x)$.

The kind of differentiability we have defined above is sometimes called *Fréchet* differentiability, to distinguish it from the weaker notion of *Gateaux* differentiability, which means that $t \mapsto f(x+t\xi)$ is differentiable at 0 (or perhaps just right-differentiable at 0) for each ξ in $\mathcal{X}$, but without any uniformity in ξ. There is also a notion of *weak Gateaux* differentiability, which means that $t \mapsto \phi\left(f(x+t\xi)\right)$ is differentiable at 0 for each ϕ in the dual of $\mathcal{Y}$ and each ξ in $\mathcal{X}$. The difference between the various notions is not very great. For example, if, f is Gateaux differentiable at each point of some non-empty open set; if, for each x in that set,

$$\xi \mapsto L_x(\xi) \equiv \left.\frac{df(x+t\xi)}{dt}\right|_{t=0}$$

is linear and bounded and if $x \mapsto L_x \in \mathcal{L}(\mathcal{X},\mathcal{Y})$ is continuous (in the norm topology), then f is also Fréchet differentiable and $L_x = Df(x)$. We will not take the time here to investigate the extent to which the results we discuss can be generalized to Gateaux differentiable functions; Fréchet differentiability generally suffices for the purposes of dynamical systems theory.

We have explained what is meant by the derivative of a function at a point. The *function* Df is defined at exactly those points of $\mathcal{X}$ where f is differentiable, and its value at such a point x is $Df(x)$, which is in $\mathcal{L}(\mathcal{X},\mathcal{Y})$, the space of bounded linear mappings from $\mathcal{X}$ to $\mathcal{Y}$. Since $\mathcal{L}(\mathcal{X},\mathcal{Y})$ is itself a Banach space, it makes sense to speak of the derivative of the function Df, and this derivative is what is called the *second derivative* D^2f of f. As defined, this second derivative takes values in $\mathcal{L}\left(\mathcal{X},\mathcal{L}(\mathcal{X},\mathcal{Y})\right)$, but it is often useful to use the observation that this space is naturally identified with the space of all continuous bilinear mappings from $\mathcal{X}\times\mathcal{X}$ to $\mathcal{Y}$ and to think of the second derivative as taking values in this space of bilinear mappings. Higher derivatives are defined by an obvious extension of this definition, and we can think of D^nf as taking values in the space of bounded n-linear mappings from $\mathcal{X}\times\ldots\times\mathcal{X}$ (n times) to $\mathcal{Y}$. A function defined on an open set $\mathcal{U}$ whose n-th derivative exists at each point of $\mathcal{U}$, and such that, furthermore, $x \mapsto D^nf(x)$ is continuous, is said to be of class $\mathcal{C}^n$, and a function of class $\mathcal{C}^n$ for all n is said to be of class $\mathcal{C}^\infty$.

It will be useful to define also non-integral degrees of differentiability. A function f defined on an open subset $\mathcal{U}$ of a Banach space $\mathcal{X}$, with values in a Banach space $\mathcal{Y}$, is said to be *Hölder continuous of order* α if there is a constant C such that $\|f(x_1)-f(x_2)\| \le C\,\|x_1-x_2\|^\alpha$ for all x_1, x_2 in $\mathcal{U}$. Here, in principle, α can be any strictly positive real number, but the notion is not very useful for $\alpha > 1$: A Hölder continuous function of order greater than one is everywhere differentiable, with derivative zero, and hence is locally constant. The term *Lipschitz continuous* is used as a synonym for Hölder continuous of order one.

Now let r be a strictly positive real number, and write $r = n+\alpha$, with n an integer and $0 \le \alpha < 1$. We then say that a function f defined on $\mathcal{U}$ is of class $\mathcal{C}^r$ if it is of class $\mathcal{C}^n$ and if D^nf is Hölder continuous of order α on a neighborhood of each point of $\mathcal{U}$.

As already noted, a great deal of standard multidimensional differential calculus extends to the infinite-dimensional situation in a completely straightforward way. Among the most important such results are the chain rule, Taylor's Theorem with remainder in integral form, and the fact that existence and continuity of partial derivatives implies differentiability. Another important result of this kind, to be discussed in the following subsections, is the implicit function theorem. There are, nevertheless, differences between the finite and infinite dimensional situations, and care is needed to avoid pitfalls. Perhaps the most obvious pitfall concerns results which depend on the local compactness of the domain space. For example, a continuous function on an infinite dimensional Banach space need not be bounded on bounded sets; in fact, even a function which is $\mathcal{C}^\infty$ on the whole space need not be bounded on bounded sets (See Exercise 1). Similarly, continuity apparently does not imply local uniform continuity, i.e., a continuous function apparently does not need to be uniformly continuous on a neighborhood of each point in its domain of definition.

A more subtle difference concerns the existence of smooth partitions of unity. *On a given Banach space, there need not exist non-zero smooth functions vanishing outside bounded sets.* Such functions of course exist on Hilbert spaces and other spaces whose norms are smooth away from zero. But many spaces arising naturally do not have smooth norms. The most complete results of which I am aware concern the sequence spaces ℓ^p. It is easy to see that $\|.\|_p^p$ is of class $\mathcal{C}^n$ on ℓ^p for any integer $n < p$. If p is an even integer, then $\|.\|_p^p$ is everywhere p times differentiable, and that the p-th derivative is constant, so $\|.\|_p^p$ is in fact of class $\mathcal{C}^\infty$. (If, on the other hand, p is an odd integer, the p-th derivative fails to exist on the dense set of points (x_n) for which $x_n = 0$ for some n.) If we exclude the trivial case of even p, it is not hard to see that $\|.\|_p^p$ is nowhere n-times differentiable for any $n > p$. The fact that the standard norm has only a limited number of derivatives does not in any obvious way imply that there are no smoother equivalent norms, to say nothing of nonconstant smoother function of bounded support. Nevertheless, it has been shown by R. Bonic and J. Frampton ("Differentiable functions on certain Banach spaces", *Bull. Amer. Math. Soc.*, **71**, (1965) 393-395.) that, on ℓ^p, p not an even integer, there exist no non-constant functions of class $\mathcal{C}^n$, $n > p$, which vanish outside the unit ball. They prove this as a corollary of the following remarkable "Maximum Modulus Principle":

Theorem 3.1.1 (Bonic and Frampton) *Let $p \geq 1$ be a real number which is not an even integer; let D be a connected bounded open set in ℓ^p, and let f be a real-valued function defined and continuous on the closure of D and n times differentiable on D itself for some integer $n \geq p$. Then the image under f of the boundary of D is dense in the image under f of D itself.*

Exercises

1. Construct a real-valued function f, defined and of class $\mathcal{C}^\infty$ everywhere on a Banach space $\mathcal{X}$, which vanishes outside the unit ball but which is nevertheless unbounded.

2. (Easy) If $\mathcal{X}$, $\mathcal{Y}$ are Banach spaces show that $A \mapsto A^{-1}$ is differentiable on the (open) set of invertible linear mappings $\mathcal{X} \to \mathcal{Y}$ with bounded inverses, and compute its derivative. The main point of this exercise is to get the "functorial" aspects right, i.e. the derivative is to be exhibited as a linear mapping from $\mathcal{L}(\mathcal{X}, \mathcal{Y})$ into $\mathcal{L}(\mathcal{Y}, \mathcal{X})$

3. We first have to establish some notation, and we do it only in a simple case: For n a non-negative integer, let $\mathcal{C}^n(-1,1)$ denote the space of all real valued functions of class $\mathcal{C}^n$ on $(-1,1)$ whose n-th derivatives are bounded on that interval. Equipped with the norm

$$\|f\|_{\mathcal{C}^n} \equiv \max\left(\sup_{-1<x<1} \{|f(x)|\}, \sup_{-1<x<1} \left\{|f^{(n)}(x)|\right\}\right)$$

$\mathcal{C}^n(-1,1)$ becomes a Banach space. (This construction can evidently be generalized to define Banach spaces of $\mathcal{C}^r$ mappings, r not necessarily an integer, from an open set in one Banach space into another.)

Now consider the mapping $f \mapsto f \circ f$ which sends the open unit ball of $\mathcal{C}^n(-1,1)$ into itself. Show that this mapping is **nowhere** differentiable from $\mathcal{C}^n(-1,1)$ to itself, but that it is differentiable—in fact, of class $\mathcal{C}^1$—from the open unit ball of $\mathcal{C}^n(-1,1)$ into $\mathcal{C}^{n-1}(-1,1)$.

The reason for the non-differentiability will probably strike you as an inessential pathology, but the non-differentiability of composition operators turns out to have serious consequences.

3.2 Fixed point theorems and Newton's method

The starting point for this subsection is the Contraction Mapping Principle:

Theorem 3.2.1 (Contraction mapping principle) *Let (X, ρ) be a complete metric space; f a mapping of X to itself such that, for some $\kappa < 1$,*

$$\rho(f(x_1), f(x_2)) \leq \kappa\rho(x_1, x_2) \quad \textit{for all } x_1,\ x_2 \textit{ in } X$$

Then f has exactly one fixed point $\bar{x} \in X$, and, for any $x_0 \in X$, $f^n(x_0)@ >> n \to \infty > \bar{x}$.

The essential observation is that

$$\rho(f^{n+1}(x_0), f^n(x_0)) \leq \kappa^n \rho(f(x_0), x_0).$$

From this it also follows that

$$\rho(x_0, \bar{x}) \leq (1-\kappa)^{-1} \rho(x_0, f(x_0)).$$

This estimate implies immediately the following corollary, which shows that the fixed point varies continuously with the map:

Corollary 3.2.2 *Let (X, ρ) be as in Theorem 2.1; f_1 and f_2 two mappings of X to itself, each contractive with contraction constant κ; $\overline{x_1}$ and $\overline{x_2}$ their respective fixed points. Then*

$$\rho(\overline{x_1}, \overline{x_2}) \leq (1-\kappa)^{-1} \rho(\overline{x_1}, f_2(\overline{x_1})) \leq (1-\kappa)^{-1} \sup_x \rho(f_1(x), f_2(x)).$$

For a differentiable mapping (on a convex set), the obvious way to check contractivity is to estimate the norm of the derivative:

Remark. *Let x_1, x_2 be two points of a Banach space $\mathcal{X}$, and write $[x_1, x_2]$ for the straight line segment joining them. Let f be defined and differentiable on $[x_1, x_2]$. Then*

$$\|f(x_1) - f(x_2)\| \leq \sup_{x \in [x_1, x_2]} \{\|Df(x)\|\} \|x_1 - x_2\|.$$

To apply the contraction mapping theorem to a differentiable mapping on a ball in a Banach space, it is not enough to know that the derivative of the mapping has norm less than one on that ball; one also needs to know that the ball is mapped into itself. The following extremely simple proposition provides a way to check this which is often useful in practice:

Proposition 3.2.3 *Let f be differentiable on the (closed) ball of radius ρ about x_0 in the Banach space $\mathcal{X}$, and suppose there is a constant $\kappa < 1$ such that $\|Df\| \leq \kappa$ on that ball. Then if*

$$\|f(x_0) - x_0\| \leq (1-\kappa)\rho$$

f has exactly one fixed point in that ball.

Proof. We have only to check that f maps the ball into itself. But $\|x-x_0\| \leq \rho$ implies

$$\|f(x)-x_0\| \leq \|f(x)-f(x_0)\| + \|f(x_0)-x_0\| \leq \kappa\rho + (1-\kappa)\rho = \rho,$$

as desired.

The Contraction Mapping Theorem is a simple and powerful tool, but it does have the disadvantage that it only works for contractions. By using the circle of ideas around Newton's method, its range of applicability can be extended considerably. Let us review the idea of Newton's method. Let f be a mapping defined on an open subset of a Banach space $\mathcal{X}$, with values in a Banach space $\mathcal{Y}$. To fix ideas, suppose that f is of class $\mathcal{C}^2$ and that $Df(x)$ is invertible for each x. We want to find a zero of f, i.e., a $\bar{x} \in \mathcal{X}$ such that $f(\bar{x}) = 0$. The idea is as follows: If we have an approximate zero x_0, then, near x_0, $f(x)$ may be approximated by the affine mapping $f(x_0) + Df(x_0)(x-x_0)$, which has a zero at $x_0 - Df(x_0)^{-1}f(x_0)$. We can then hope that this zero of the affine approximation to f is a better approximation to a zero of f itself than is the starting point x_0 and repeat the process starting there. In other words: We note that a zero of f is the same thing as a fixed point of

$$\Phi(x) = x - Df(x)^{-1}f(x),$$

and we seek a fixed point for Φ by iteration. If we suppose that $\bar{x}$ is a zero of f, we can see that this iteration must converge if we start near enough to $\bar{x}$ by computing the derivative of Φ:

$$\begin{aligned} D\Phi(x) &= \mathbf{1} - D\left(Df(x)^{-1}\right)f(x) + Df(x)^{-1}Df(x) \\ &= -D\left(Df(x)^{-1}\right)f(x), \end{aligned}$$

which vanishes at $\bar{x}$. Thus, simply by continuity, Φ is a contraction on a sufficiently small ball about the fixed point $\bar{x}$, so iterating Φ starting at any point in such a ball gives a sequence which converges to $\bar{x}$. Furthermore, the convergence is asymptotically extremely fast, provided that f is of class $\mathcal{C}^3$. In fact, if f is of class $\mathcal{C}^3$, then Φ is of class $\mathcal{C}^2$, and if $\|D^2\Phi\| \leq K$ on a ball of radius ρ about $\bar{x}$, then, if $\|x-\bar{x}\| \leq \rho$,

$$\|\Phi(x)-\bar{x}\| = \|\Phi(x)-\Phi(\bar{x})\| \leq (K/2)\|x-\bar{x}\|^2,$$

i.e., the error after $n+1$ iterations is of the order of the *square* of the error after n iterations, i.e., *the number of decimal places of accuracy doubles with each iteration.*

In the above discussion we have, for simplicity, started from the assumption of the existence of a zero of f. It is not difficult, however, to use the preceding proposition to give a concrete sufficient condition for the existence of a zero which also shows that it can be found by Newton's method. The formulation, admittedly a little awkward, is as follows:

Proposition 3.2.4 *Let f be of class $\mathcal{C}^2$, with bounded second derivative, on the ball of radius ρ about x_0. Assume that $Df(x_0)$ is invertible with bounded inverse. We will use the following notation: $M = \|Df(x_0)^{-1}\|$, $\epsilon = \|f(x_0)\|$, $K_1 = \sup_{\|x-x_0\|\leq\rho} \|Df(x)\|$, $K_2 = \sup_{\|x-x_0\|\leq\rho} \|D^2f(x)\|$, We now assume*

- $K_1\rho M < 1$
- $\kappa \equiv K_2 \left(\frac{M}{1-K_1\rho M}\right)^2 (\epsilon + K_1\rho) < 1$
- $M\epsilon \leq \rho(1-\kappa)$.

Then f has a unique zero in the ball of radius ρ about x_0.

Proof. The proof consists simply in checking that the various hypotheses suffice to guarantee that

$$\Phi(x) \equiv x - (Df(x))^{-1} f(x)$$

is defined on the ball of radius ρ about x_0 and satisfies the hypotheses of Proposition ??. We note, first of all, that for x in that ball, $\|Df(x) - Df(x_0)\| \leq K_1\rho$, and hence the Neumann series implies that $Df(x)$ is invertible and that

$$\|\,(Df(x))^{-1}\,\| \leq \frac{M}{1-K_1\rho M}.$$

Second, as we have already noted, $D\Phi(x) = -D\left(Df(x)^{-1}\right) f(x)$. This actually calls for a little interpretation. By definition, $D\Phi(x)$ is a linear operator from $\mathcal{X}$ to $\mathcal{Y}$, whereas $D\left(Df(x)^{-1}\right)$ is a linear mapping from $\mathcal{X}$ to $\mathcal{L}(\mathcal{Y},\mathcal{X})$; what the above equation means is that, for any $\xi \in \mathcal{X}$,

$$D\Phi(x)\xi = -\underbrace{\left(D\left(Df(x)^{-1}\right)(\xi)\right)}_{\in\mathcal{L}(\mathcal{Y},\mathcal{X})} f(x).$$

Again using the Neumann series, one checks easily that

$$D\left(Df(x)^{-1}\right)(\xi)(\eta) = -Df(x)^{-1}D^2f(x)\left(Df(x)^{-1}\xi, Df(x)^{-1}\eta\right).$$

Also, $\|f(x)\| \leq \epsilon + K_1\rho$. Combining all these remarks we see that $\|D\Phi(x_)\| \leq \kappa$ for all x with $\|x - x_0\| \leq \rho$. Furthermore, $\|\Phi(x_0) - x_0\| = \|Df(x_0)^{-1}f(x_0)\| \leq M\epsilon$, so our final assumption is just what is needed to guarantee that the final assumption of the Proposition ?? holds. □

The preceding proposition is inconvenient to apply because it requires estimating the second derivative. For this reason, it is sometimes advantageous to replace Newton's method by another approximation scheme which is easier to control. The idea is to redefine Φ by replacing $Df(x)$ by some fixed constant invertible operator Γ; if Γ is a good enough approximation to $Df(x)$ for all x in the ball of radius ρ about x_0, the redefined Φ is still a contraction on this ball. The precise conditions which have to be satisfied are quite simple:

Proposition 3.2.5 *Let $B_\rho(x_0)$ denote the closed ball with radius ρ centered at x_0; let f be a mapping defined on a neighborhood of $B_\rho(x_0)$ taking values in $\mathcal{Y}$ and differentiable at each point of $B_\rho(x_0)$; and let Γ be an invertible linear mapping $\mathcal{X} \to \mathcal{Y}$. If there exists $\kappa < 1$ such that*

1. $\left\|\Gamma^{-1}\left(Df(x) - \Gamma\right)\right\| \leq \kappa$ *for all* $x \in B_\rho(x_0)$.
2. $\left\|\Gamma^{-1}f(x_0)\right\| \leq (1-\kappa)\rho$

then f has a unique zero in $B_\rho(x_0)$.

In fact, the exact conditions of the preceding proposition are not often very important, but it is simpler to state them explicitly than to formulate qualitatively what it is about them which is actually used.

Proof. The proof is almost immediate. Put $\Phi(x) = x - \Gamma^{-1}f(x)$. Evidently, a fixed point of Φ is the same thing as a zero of f. By a very simple computation,

$$D\Phi(x) = \Gamma^{-1}\left(\Gamma - Df(x)\right),$$

so the two condition of Proposition ?? are simply reformulations of those of Proposition ?? □

It should be noted that, although this modified Newton's method is easier to control, and requires less regularity, than Newton's method itself, it does not give the extremely rapid convergence that Newton's method does (unless $\Gamma = Df(\bar{x})$, which is hard to arrange in practice).

3.3 The implicit function theorem

The ideas of the preceding section can be used to prove the implicit function theorem:

Theorem 3.3.1 (Implicit function theorem) *Let $\mathcal{X}$, $\mathcal{Y}$, $\mathcal{Z}$ be Banach spaces, f a function defined on a neighborhood $\mathcal{U}$ of 0 in $\mathcal{X} \times \mathcal{Z}$, with values in $\mathcal{Y}$. Assume:*

i) *For each $(x_0, z_0) \in \mathcal{U}$, $f(x, z_0)$ is differentiable in x on a neighborhood of x_0, and the partial derivative $D_1 f(x, z)$ varies continuously with $(x, z) \in \mathcal{U}$.*

ii) $f(0,0) = 0$.

iii) *$D_1 f(0,0)$ is invertible.*

Then

1) *If $\delta_x > 0$ is sufficiently small, there exists $\delta_z > 0$ such that, for each z with $\|z\| < \delta_z$, there is exactly one x with $\|x\| < \delta_x$ such that $f(x, z) = 0$. (Intuitively: For each small z, the equation $f(x, z) = 0$ has a unique small solution.)*

We will from now on suppose that δ_z, δ_x have been fixed once and for all so that the condition of i) holds, and are also small enough so that $D_1 f(x, z)$ is invertible for all (x, z) in

$$\Delta \equiv \{(x, z) : \|z\| < \delta_z \text{ and } \|x\| < \delta_x\} .$$

We will denote the x corresponding to z by $X(z)$

2) *If, in addition, f is differentiable on Δ, X is differentiable on*

$$\Delta_z \equiv \{z : \|z\| < \delta_z\} ,$$

and

$$DX(z) = -\left(D_2 f(X(z), z)\right)^{-1} D_1 f(X(z), z).$$

3) *If f is of class $\mathcal{C}^r$, $r \geq 1$ on Δ, then X is of class $\mathcal{C}^r$ on δ_z.*

Proof. To prove 1), we simply apply Proposition ?? to the mapping $f^{(z)} : x \mapsto f(x, z)$, with 0 as an approximate zero, $\Gamma = D_1 f(0,0)$, and $z \in \mathcal{Z}$ as a "parameter." We have to choose ρ and κ. At $z = 0$, we have $\Gamma^{-1}\left(Df^{(z)}(0) - \Gamma\right) = 0$ and $\Gamma^{-1} f^{(z)}(0) = 0$. By continuity, there exist $\delta_x, \delta_z > 0$ such that

$$\left\|\Gamma^{-1}\left(Df^{(z)}(x) - \Gamma\right)\right\| \leq 1/2 \text{for } \|x\| \leq \delta_x \text{ and } \|z\| \leq \delta_z.$$

By making δ_z smaller if necessary, we can also ensure that $\|\Gamma^{-1} f^{(z)}(0)\| \leq \delta_x/2$ if $\|z\| \leq \delta_z$. Proposition 2.5 now applies, with $\kappa = 1/2$ and $\rho = \delta_x$, for each z with $\|z\| \leq \delta_z$ and proves 1).

We now turn to the proof of 2). Note that the formula for $DX(z)$ would follow immediately by differentiating the identity $f(X(z), z) = 0$, if we knew that X was differentiable. Proving that X is differentiable, assuming that f is, is an instructive exercise in the definition of the derivative. To get started, we need to know that X is continuous. This follows from the proof of Proposition 2.5 together with Corollary 2.2. To prove differentiability, fix z, and, for sufficiently small δz, define $\delta X = X(z + \delta z) - X(z)$. From $f(X, z) = 0$, $f(X + \delta X, z + \delta z) = 0$, and the definition of derivative, we get

$$D_2 f(X, z)\delta z + D_1 f(X, z)\delta X = o(\|\delta z\| + \|\delta X\|).$$

Since we have set things up so that we know that $D_1 f(X, z)$ is invertible, we get

$$\delta X = -\left(D_1 f(X, z)\right)^{-1} D_2 f(X, z)\delta z + o(\|\delta z\| + \|\delta X\|).$$

A first application of this estimate together with the continuity of X shows that $\|\delta X\| = O(\|\delta z\|)$, and hence $o(\|\delta z\| + \|\delta X\|)$ can be replaced simply by $o(\|\delta z\|)$. A second application of the estimate shows, directly from the definition of derivative, that X is differentiable at z and that its derivative is as indicated in 2).

Once the formula for $DX(z)$ is established, it is straightforward to prove—by differentiating that formula repeatedly—that X is n times differentiable if f is. This process also leads to an explicit (if unwieldy) formula for the n-th derivative of X in terms of the partial derivatives (of total order not greater than n) of f. From this formula it follows at once that if the partial derivatives of f of order n are Hölder continuous of order α, then so is the n-th derivative of X. □

Corollary 3.3.2 (Inverse function theorem) *Let f be a continuously differentiable mapping defined on a neighborhood of $x_0 \in \mathcal{X}$, taking values in $\mathcal{Y}$. Assume that $Df(x_0)$ is invertible. Then f maps any sufficiently small ball about x_0 homeomorphically onto a neighborhood of $f(x_0)$. The inverse homeomorphism is continuously differentiable, and of class $\mathcal{C}^r$ if f is. We have*

$$Df^{-1}(f(x)) = \left(Df(x)\right)^{-1}.$$

Proof. To prove the local invertibility of f, we need to show that $f(x) = y$, regarded as an equation to be solved for x with y fixed, has a unique solution near x_0 for each y sufficiently near to $f(x_0)$. This follows from the implicit function theorem applied to $(y, x) \mapsto f(x) - y$. Differentiability of the inverse function also follows from the implicit function theorem; the formula for the derivative can be derived from the formula given above for the derivative of an implicitly defined function, but it is easier simply to differentiate the identity $f^{-1}(f(x)) = x$ and apply the chain rule. □

Theorem 3.3.1 is frequently referred to as the “soft” implicit function theorem, to distinguish it from the much deeper “hard” implicit function theorems of Nash and Moser.

Chapter 4

Attracting fixed points

4.1 Orientation

Let f be a map. A *periodic point* of f means a point x_0 such that $f^p(x_0) = x_0$ for some strictly positive integer p. The smallest (strictly positive) such p is called the *period* of f. A *periodic cycle* of f means a sequence $x_0, f(x_0), \ldots, f^{p-1}(x_0)$, where x_0 is a periodic point of period p. Study of the behavior of a map near a periodic point can be reduced to study of the behavior near a fixed point by replacing f by f^p. Somewhat more precisely: The motion induced by f near a periodic cycle of period p is, in a natural sense, the semi-direct product of a cyclic permutation of order p with the motion induced by f^p near any point of the cycle. The cyclic permutation if "trivial"; we therefore focus our attention on the case of a fixed point.

Thus, let f be a differentiable mapping defined on an open set in a Banach space $\mathcal{E}$, with a fixed point x_0. Near this fixed point, f is well approximated by the *linear* mapping $x \mapsto f(x_0) + Df(x_0)\,(x - x_0)$. It is thus natural to ask: To what extent is the dynamics of f near the fixed point like the dynamics of $Df(x_0)$? The point of view here is that the dynamics of *linear* mappings is well-understood and simple—at least in the finite-dimensional case—and may provide valuable guidance for the understanding of the non-linear case.

One nice situation, which happens reasonably frequently (at least in the finite-dimensional case) is that, near the fixed point, f *differs from* $Df(x_0)$ *only by a change of coordinates.* More precisely, what we mean is that there is an invertible mapping ϕ, defined on a neighborhood of 0 and mapping 0 to x_0, such that $f = \phi Df(x_0)\phi^{-1}$ on a neighborhood of 0. We speak of such a ϕ as a *linearization* of f (near x_0), or we say that ϕ *conjugates* $Df(x_0)$ to f

(near x_0). It matters a lot how much regularity ϕ is required to have. In order to be of any interest at all, a linearization must be at least a homeomorphism, in which case we speak of a *topological* linearization. As we shall see later on, topological linearizations exist quite often—not always—but the information they give about the dynamics near the fixed point is too coarse to be very useful.

Linearizations with some degree of differentiability are more useful than ones which are merely topological, but conditions for their existence are more complicated. For example, among maps of class $\mathcal{C}^\infty$, "most" (in a sense which we will not define here) admit $\mathcal{C}^\infty$ linearizations, but a dense set of them do not. Because the existence or non-existence of a smooth linearization depends on delicate properties of the map, these linearizations are not very effective tools for analyzing general smooth maps. Nevertheless, it can be a helpful *heuristic* device to imagine that linearizations exist. Judiciously chosen "obvious" facts about linear maps often extend to give useful and less trivial properties of general maps. For example: Suppose f admits a $\mathcal{C}^1$ linearization, with $\mathcal{C}^1$ inverse, at the fixed point x_0. Without loss of generality, we can assume that $D\phi(0) = \mathbf{1}$.[1] Now if $\mathcal{E}_1$ is a linear subspace of $\mathcal{E}$ invariant under $Df(x_0)$, $\phi\mathcal{E}_1$ is an invariant manifold for f. (Actually, since we are only supposing that ϕ is locally defined, this manifold is in general only *locally invariant* in an appropriate sense. This manifold passes through x_0, and its tangent space there is exactly the subspace $\mathcal{E}_1$. Thus, if a smooth linearization exists, invariant subspaces for $Df(x_0)$—about which spectral theory provides detailed information—give rise to locally invariant manifolds for f. It is then natural to ask to what extent this correspondence can be established *without* assuming the existence of a smooth linearization. It turns out, in fact, that a strict correspondence between invariant subspaces and invariant manifolds breaks down in general, for roughly the same reasons as does the existence of smooth linearizations, but that there are some "dynamically defined" invariant subspaces for $Df(x_0)$ to which there do correspond invariant manifolds. The study of these questions will be the main topic of this chapter, but, before beginning on it, we will look at a simpler relation between the dynamics of f near a fixed point x_0 and the dynamics of $Df(x_0)$, manifested in what is called "linear stability analysis".

[1] In general, we get from the chain rule that $Df(x_0) = D\phi(0)Df(x_0)D\phi(0)^{-1}$, i.e., $D\phi(0)$ commutes with $Df(x_0)$. Hence, replacing ϕ by $\phi \circ [D\phi(0)^{-1}]$ gives a new linearization whose derivative at 0 is the identity.

4.2 Linear stability analysis

A fixed point x_0 for a mapping f will be said to be *attracting* if there exists a neighborhood $\mathcal{U}$ of x_0 whose images under f shrink down uniformly to $\{x_0\}$ in the sense that, for any neighborhood $\mathcal{V}$ of x_0, $f^n\mathcal{U} \subset \mathcal{V}$ for all sufficiently large n. (Although we have not explicitly assumed that $\mathcal{U}$ is mapped into itself by f, we can add this assumption without making the definition more restrictive; simply replace $\mathcal{U}$ by $\bigcup_{n=0}^{\infty} f^n\mathcal{U}$.)

This definition calls for two comments on terminology. The first comment is that what we call an attracting fixed point is commonly called a *stable* or (in the older literature) *asymptotically stable* fixed point. The reason for this terminology is that, if we imagine perturbing an attracting fixed point slightly, we get another initial point whose orbit converges back to that of the unperturbed fixed point. The term "stable" is, however, overused in dynamical systems theory, and I have preferred to reserve it primarily to refer to *properties of a system which are preserved under small changes in that dynamical system* (as opposed to properties of an orbit which are preserved under small changes in initial condition). The second comment is that the above definition of attracting fixed point is never satisfied by a Hamiltonian dynamical system (at least on a finite-dimensional state space); such systems, because they preserve volumes in state space, cannot map a set $\mathcal{U}$ of nonzero volume into an arbitrarily small subset of itself. There is another definition of stability more appropriate for Hamiltonian systems: A fixed point x_0 is said to be *Lyapunov stable* if, for any neighborhood $\mathcal{U}$ of x_0, there exists a (smaller) neighborhood $\mathcal{V}$ such that $f^n\mathcal{V} \subset \mathcal{U}$ for all $n \geq 0$, i.e., any orbit starting sufficiently near to x_0 stays near forever.

Whether or not a fixed point is attracting can often be determined from the spectrum of the derivative at that fixed point.

Proposition 4.2.1 *Let f be a mapping of class $\mathcal{C}^1$; x_0 a fixed point of f.*

1. *If the spectral radius ρ of $Df(x_0)$ is strictly less than one, x_0 is attracting.*
2. *If $Df(x_0)$ is a compact operator, and if the spectral radius of $Df(x_0)$ is strictly greater than one, x_0 is not attracting.*

Thus, at least in the case where $Df(x_0)$ is compact—and in particular in the finite-dimensional situation which is our principal interest—the spectral radius of $Df(x_0)$ determines whether or not x_0 is attracting unless it is exactly one. I do not know whether the second assertion is true without some assumption like compactness to permit the splitting off of the part of the spectrum with largest modulus.

Proof. We will prove only the first assertion. We choose We choose a norm such that $\|Df(x_0)\| < 1$; then, by continuity of the derivative, there exist $\kappa < 1$ and $\rho > 0$ such that $\|Df(x)\| < \kappa$ for all x with $\|x - x_0\| < \rho$. We take $\mathcal{U}$ to be the ball of radius ρ about x_0; then for $n = 1, 2, \ldots$, $f^n\mathcal{U}$ is contained in the ball of radius $\kappa^n\rho$ about x_0. Any neighborhood $\mathcal{V}$ of x_0 contains the ball of radius $\kappa^n\rho$ for all sufficiently large n. □

The analogue of the above considerations for flows is as follows: A *stationary solution* of a differential equation

$$\frac{dx}{dt} = X(x)$$

means an x_0 such that $X(x_0) = 0$. Then x_0 is a fixed point for the solution flow f^t for all t. (Recall: It is our one of our standing hypotheses for flows that $X(x)$ possesses enough regularity to guarantee uniqueness of solutions.) A stationary solution x_0 is *attracting* if there is a neighborhood $\mathcal{U}$ of x_0 such that, for every neighborhood $\mathcal{V}$ of x_0, $f^t(\mathcal{U}) \subset \mathcal{V}$ for all sufficiently large t. Under very mild technical assumptions, this is equivalent to attractivity of the stationary solution for the time-one map f^1, so the study of attracting stationary solutions for differential equations essentially reduces to that of attracting fixed points for maps.

For example, if there exists a k such that

$$\|X(x)\| \leq k\|x\| \quad \text{for small } x,$$

then

$$\|f^t x\| \leq e^{kt}\|x\| \quad \text{for } x \text{ small and } 0 < t < 1.$$

Then if

$$f^n(\mathcal{U}) \subset e^{-k}\mathcal{V} \quad \text{for } n \geq n_0,$$

we get

$$f^{n+\alpha}(\mathcal{U}) \subset \mathcal{V} \quad \text{for } n \geq n_0 \text{ and } 0 \leq \alpha < 1,$$

i.e.,

$$f^t(\mathcal{U}) \subset \mathcal{V} \quad \text{for } t \geq n_0.$$

It $X(x)$ is continuously differentiable, the spectrum of the derivative of the time-one solution mapping at a stationary solution x_0 is also easy to compute—given the spectrum of the linearization $DX(x_0)$ of the differential equation—using the *first variational equation*. The idea is as follows: Let $x(t)$, $0 \leq t \leq t_0$ be a solution of the differential equation. Write $V(t)$ for $Df^t(x(0))$. Then $V(t)$ satisfies the linear differential equation

$$\frac{dV}{dt} = V(t)DX(x(t)) \quad \text{with initial condition } V(0) = \mathbf{1}.$$

This equation, furthermore, determines $V(t)$. In the simple case where $x(t) \equiv x_0$, we get

$$V(t) = e^{tDX(x_0)},$$

and hence we have

Corollary 4.2.2 *A stationary solution x_0 of the differential equation*

$$\frac{dx}{dt} = X(x(t)),$$

(with $X(x)$ continuously differentiable) is attracting if the spectrum of $DX(x_0)$ is contained in the open left half-plane; in finite dimensions, such a stationary solution is not attracting if the spectrum of $DX(x_0)$ is not contained in the closed left half-plane.

Lectures on Dynamical Systems

Chapter 5. Stable manifolds: introduction
Chapter 6. Stable manifolds: smoothness
Chapter 7. Stable manifolds: dynamics
Chapter 8. Unstable manifolds
Chapter 9. Invariant manifolds: miscellany

Oscar E. Lanford III
Mathematics Department
ETH-Zürich
8092 Switzerland

September 2, 1997

This text is a slightly revised version of part of a set of lecture notes distributed to the students in my course on dynamical systems at the ETH Zürich in the Winter Semester of 1991–92. I have had a number of requests for copies of these notes, and this has encouraged me to think that it might be useful to make them more generally available. The reader should be warned, however, that, since I am working on a major rewriting of this set of notes, I did not want to take the time to polish this older version carefully (although there is a good deal which needs to be done.) I will be grateful for comments about these notes and particularly grateful to have any errors called to my attention.

Oscar E. Lanford III (lanford@math.ethz.ch)

Chapter 5

Stable manifolds: introduction

5.1 Introduction

We will consider here a mapping f, at least once continuously differentiable, defined on an open set in a Banach space $\mathcal{E}$. A fixed point z_0 for f will be said to be *hyperbolic* if the linear operator $Df(z_0)$ is hyperbolic in the sense defined in §2.4, i.e., if its spectrum is disjoint from the unit circle. We will fix our attention here on a particular hyperbolic fixed point z_0. By Corollary 2.4.2, there is a splitting of $\mathcal{E}$ as a direct sum of two closed linear subspaces $\mathcal{E}_s$ and $\mathcal{E}_u$, each invariant under $Df(z_0)$, such that the spectrum of the restriction of $Df(z_0)$ to $\mathcal{E}_s$ is entirely inside the unit circle and that of the restriction to $\mathcal{E}_u$ entirely outside. Furthermore, $\mathcal{E}_s$ can be characterized by

$$\mathcal{E}_s = \{z \in \mathcal{E} : \Lambda^n z \longrightarrow 0 \quad \text{as } n \to \infty.\}$$

In this and the following chapters, we are going to establish the existence of a "non-linear version" of $\mathcal{E}_s$. Loosely formulated, the idea is as follows: Let W^s denote the set of points z such that

$$f^n(z) \text{ remains near } z_0 \text{ for all } n \geq 0 \text{ and } \to z_0 \text{ as } n \to \infty.$$

So defined, W^s is just a set, manifestly mapped into itself by f and manifestly containing z_0, but which a priori might contain nothing else. We are going to show, in fact, that

W^s is a smooth submanifold of $\mathcal{E}$ and the tangent space to W^s at z_0 is $\mathcal{E}^s$[1].

This invariant manifold is known as the *stable manifold* for f at x_0. We will concentrate, for the first part of the analysis, on the behavior in the immediate vicinity of the fixed point, and will construct only a little piece of invariant manifold. In nice cases, it is possible to extend the small invariant manifold we construct here in a natural way to a *global* invariant manifold. Thus, what we are concerned with here is more properly called a *local stable manifold.* There is an analogous non-linear version of $\mathcal{E}_u$; it is called the *unstable manifold* and will be denoted by W^u. If f is invertible, the unstable manifold for f can be constructed as the stable manifold for f^{-1}. As we will see later on, however, unstable manifolds can also be constructed without assuming invertibility of f.

Stable and unstable manifolds are objects of fundamental importance in the study of hyperbolic fixed points and in dynamical systems theory in general. Fortunately, they are very well-behaved objects, and in particular there is no loss of differentiability in passing from f to the corresponding stable and unstable manifolds: If f is of class $\mathcal{C}^r$, $1 \leq r \leq \infty$, then W^s and W^u are submanifolds of class $\mathcal{C}^r$[2]. The theory of these objects, while not exceptionally difficult, is of substantial depth, and a circumspect indirect approach seems to be needed to make the proofs reasonably simple, especially if the objective is to obtain results with weakest possible hypotheses. In broad outline, we are going to proceed as follows:

- At the outset, we forget about the fact that W^s is supposed to be the set of points whose orbits converge to z_0 and concentrate on finding an invariant manifold tangent to $\mathcal{E}^s$ at z_0.

- It is in fact convenient to elarge the framework a little and to look for an invariant set with a bit less regularity that necessary for a submanifold.

- We convert the search for an invariant set to the search for a fixed point of an auxiliary operator on an appropriate function space, and we show that this operator is contractive and hence has a unique fixed point.

- With the fixed point—i.e., the invariant set—already in hand, we proceed to prove first that the invariant set is a $\mathcal{C}^1$ manifold; then that it has at least as much additional smoothness as f itself.

---

[1]This means in particular that the dimension of W^s is equal to that of $\mathcal{E}^s$

[2]The converse doesn't hold: A mapping of finite differentiability can easily have—by accident, so to speak—$\mathcal{C}^\infty$ stable and unstable manifolds

– Finally, we show that our invariant manifold does indeed consist exactly of those points whose orbits remain near the fixed point and converge to it.

5.2 Notation and preliminary reductions

Our basic notation and standing hypotheses are as follows: We are considering a mapping f defined on an open set in a Banach space $\mathcal{E}$. When nothing is said to the contrary, it is always to be understood that f is at least continuously differentiable. We investigate f in the neighborhood of a fixed point, which we will take—for notational simplicity—to be the origin in $\mathcal{E}$. We will denote $Df(0)$ by Λ, and we assume that Λ is a hyperbolic operator. It follows that $\mathcal{E}$ has a representation as

$$\mathcal{E} = \mathcal{E}_s \oplus \mathcal{E}_u,$$

and that, in this representation, Λ has the form

$$\begin{pmatrix} \Lambda_s & 0 \\ 0 & \Lambda_u \end{pmatrix},$$

where

1. the spectral radius of Λ_s is < 1, and
2. Λ_u is invertible and the spectral radius of $(\Lambda_u)^{-1}$ is < 1.

When referring to points of $\mathcal{E} = \mathcal{E}_s \oplus \mathcal{E}_u$, we will generally use

– x (with subscripts, accents, etc.) to refer to the $\mathcal{E}_s$ co-ordinate,

– y to refer to the $\mathcal{E}_u$ co-ordinate, and

– z to refer to the pair (x, y).

We choose norms on $\mathcal{E}_s$ and $\mathcal{E}_u$ so that

$$\|\Lambda_s\| < 1 \quad \text{and} \quad \|(\Lambda_u)^{-1}\| < 1, \tag{5.1}$$

and we equip $\mathcal{E}$ with the norm

$$\|(x, y)\| = \max\{\|x\|, \|y\|\}.$$

We now separate off the linear part of f near the origin by writing:

$$f(x, y) = (\Lambda_s x + f_s(x, y), \Lambda_u y + f_u(x, y)), \tag{5.2}$$

where f_s and f_u are continuously differentiable and vanish, together with their first derivatives, at the origin.

We will refer to the f_s and f_u as the *nonlinear parts* of f. A first very important observation is that, because we want to prove a result which is local at the fixed point, we can without loss of generality assume that f_s and f_u are "small". This goes as follows: By assumption, the f_s and f_u are continuously differentiable and vanish, together with their first derivatives, at $(0,0)$. Hence, the norms of their first derivatives can be made as small as we like on the ball of radius ϵ about $(0,0)$ by making ϵ small. We can then replace f by $z \mapsto \frac{1}{\epsilon} f(\epsilon z)$ to get a new mapping with the same linear parts as f but whose non-linear parts have first derivatives which are uniformly as small as we like on the unit ball. That is: Working on the unit ball with the nonlinear parts small in the $\mathcal{C}^1$ sense is equivalent to working on a small ball for a general $\mathcal{C}^1$ mapping f. We will refer to this reduction as "magnification".

We next want to give a concrete analytical formulation of the geometrical notion of manifold tangent to $\mathcal{E}_s$ at 0. It both corresponds to our geometrical intuition and can easily be proved from the standard definitions in manifold theory that a submanifold containing 0 is tangent to $\mathcal{E}_s$ there if and only if a sufficiently small piece of this manifold around 0 can be represented as the graph $\Gamma(w)$ of a mapping w, defined on a neighborhood of 0 in $\mathcal{E}_s$ and taking values in $\mathcal{E}_u$, such that

$$w(0) = 0 \quad \text{and} \quad Dw(0) = 0. \tag{5.3}$$

(The first condition simply says that the origin lies on the submanifold; the second expresses tangency.) Furthermore, the smoothness of the submanifold—or rather, of the small piece of it so represented—is exactly the smoothness of w. In our original—unmagnified—problem, what we are looking for is a function, defined on a small neighborhood of 0 and satisfying (5.3), whose graph is mapped into itself by f. After magnification, we can assume that w is defined on the open unit ball. Thus, what we want to prove is something along the following lines:

Preliminary formulation of desired result. *Let Λ_s and Λ_u be as above. Then for any pair of functions f_s, f_u, defined and continuously differentiable on the open unit ball of $\mathcal{E}_s \oplus \mathcal{E}_u$, vanishing together with their first derivatives at the origin and having sufficiently small $\mathcal{C}^1$ norm on the unit ball, the mapping*

$$f : (x, y) \longmapsto (\Lambda_s x + f_s(x, y), \Lambda_u y + f_u(x, y))$$

leaves invariant a set of the form $\Gamma(w)$, where w is a $\mathcal{C}^1$ function from the unit ball of $\mathcal{E}_s$ into $\mathcal{E}_u$, vanishing together with its first derivative at the origin.

With this preliminary formulation in mind, we now ask when the graph of a function w is mapped into itself by f of the form (5.2). To answer this question, we write out, in our detailed notation for f, the image under f of a general point of the graph of w:

$$f(x, w(x)) = (\Lambda_s x + f_s(x, w(x)), \Lambda_u w(x) + f_u(x, w(x)).$$

This will again be in the graph of w if and only if

$$\Lambda_u w(x) + f_u(x, w(x)) = w\left(\Lambda_s x + f_s(x, w(x))\right), \tag{5.4}$$

i.e., if and only if

$$w(x) = \Lambda_u^{-1}\left(w\left(\Lambda_s x + f_s(x, w(x))\right) - f_u(x, w(x))\right). \tag{5.5}$$

Now the right-hand side of (5.5) can be taken as defining a (non-linear) operator on functions w, i.e., (5.5) can be interpreted as the fixed-point equation $w = \mathcal{F}w$ where

$$(\mathcal{F}w)(x) = \Lambda_u^{-1}\left(w\left(\Lambda_s x + f_s(x, w(x))\right) - f_u(x, w(x))\right). \tag{5.6}$$

The final result, then, is:

The graph of w is mapped into itself by f if and only if w is a fixed point for $\mathcal{F}$ as defined in (5.6).

(Included in the fixed-point condition is the requirement that the right-hand side of (5.6) is defined on all of the domain of w, i.e., that $\Lambda_s x + f_s(x, w(x))$ is in the domain of w for all x in the domain of w.)

5.3 Lipschitz functions

The idea now is to choose a function space on which the operator $\mathcal{F}$ is a contraction. Although many choices of function space can be made to work under sufficiently strong hypotheses on f, some care is needed to get the sharpest possible results. A space which turns out to work well is *a space of Lipschitz continuous functions, equipped with a judicious modification of the supremum norm.* The technical advantage here in working with with Lipschitz functions is that

- for our estimates, a Lipschitz condition is almost as good as a bound on the first derivative, while
- Lipschitz continuity, unlike differentiability, is well-behaved under uniform—or even pointwise—limits.

We need to introduce a little formalism. If (X_1, ρ_1), (X_2, ρ_2) are metric spaces, we say that a mapping $f : X_1 \to X_2$ is *Lipschitz continuous* if there exists a constant k such that

$$\rho_2(f(x), f(x')) \leq k\rho_2(x, x') \quad \text{for all } x, x' \in X_1.$$

The set of k's for which this inequality holds is closed; its infimum will be denoted by $\mathrm{Lip}(f)$. If the range space is a normed vector space, then the set of Lipschitz continuous functions is a vector space and $\mathrm{Lip}(f)$ is almost a norm on this vector space. Although it fails to be a norm because it vanishes on constant functions, we will nevertheless refer to $\mathrm{Lip}(f)$ as the *Lipschitz norm* of f.

The Lipschitz norm is particularly well-behaved with respect to composition of functions:

Proposition 5.3.1 *Let (X_1, ρ_1), (X_2, ρ_2), (X_3, ρ_3) be metric spaces, $f : X_1 \to X_2$ and $g : X_2 \to X_3$ Lipschitz continuous mappings. Then $g \circ f$ is Lipschitz continuous and*

$$Lip(g \circ f) \leq Lip(g)\, Lip(f).$$

Proof. For x, x' in X_1,

$$\rho_3(g(f(x)), g(f(x'))) \leq \mathrm{Lip}(g)\rho_2(f(x), f(x')) \leq \mathrm{Lip}(g)\mathrm{Lip}(f)\rho_1(x, x').$$

□

Once we have decided to look for the invariant set as the graph of a Lipschitz continuous function, it is no longer necessary to assume that f is continuously differentiable. Instead of requiring that f_s and f_u be continuously differentiable on the unit ball with uniformly small first derivatives, it suffices to assume instead that they are Lipschitz continuous with small Lipschitz norm. The condition that their derivatives vanish at the origin no longer makes sense in this more general framework, but we retain the assumption that $f_s(0) = 0$ and $f_u(0) = 0$.

We can now formulate our first existence and uniqueness result:

Proposition 5.3.2 *Let Λ_s, Λ_u^{-1} be operators of norm < 1 on Banach spaces $\mathcal{E}_s$, $\mathcal{E}_u$ respectively, and let f_s, f_u be Lipschitz-continuous mappings from the unit ball of $\mathcal{E}_s \oplus \mathcal{E}_u$ into $\mathcal{E}_s$, $\mathcal{E}_u$ respectively which vanish at the origin. If the Lipschitz norms of f_s, f_u are small enough, then there is a unique mapping w_s from the unit ball of $\mathcal{E}_s$ into $\mathcal{E}_u$, Lipschitz continuous with Lipschitz norm ≤ 1, whose graph is mapped into itself by*

$$f : \quad (x, y) \longmapsto (\Lambda_s x + f_s(x, y), \Lambda_u y + f_u(x, y))$$

Proof. Let

$$\mathcal{X}_{\mathcal{F}} := \{w : \mathcal{B}_s \to \mathcal{E}_u, w(0) = 0 \quad \text{and} \quad \mathrm{Lip}(w) \le 1\}$$

(where we have written $\mathcal{B}_s$ for the open unit ball in $\mathcal{E}_s$.) The idea is going to be to show that $\mathcal{F}$ as defined in (5.6) maps $\mathcal{X}_{\mathcal{F}}$ contractively into itself. For the remainder of this proof, w will always denote an element of $\mathcal{X}_{\mathcal{F}}$. We will frequently use without comment the simple estimates

$$\|w(x)\| \le \mathrm{Lip}(w)\|x\| \le \|x\| < 1$$

for all $x \in \mathcal{B}_s$.

To simplify the formulas arising in the course of the proof, it is convenient to introduce a streamlined notation. We rewrite the definition (5.6) of the operator $\mathcal{F}$ as

$$(\mathcal{F}w)(x) = F(x, w(x), w(v(x, w(x)))) \tag{5.7}$$

where

$$F(x, y, y') = \Lambda_u^{-1}\,[y' - f_u(x, y)] \quad \text{and} \quad v(x, y) = \Lambda_s x + f_s(x, y). \tag{5.8}$$

Note that—from the vanishing of f_s, f_u, and w at the origin—

$$v(0, 0) = 0 \quad \text{and} \quad F(0, 0, 0) = 0.$$

When we consider functions of several variables, like F and v, we will write Lip_i for the Lipschitz norm with respect to the i-th variable with the other variables held fixed (i.e., in the differentiable case, the supremum norm of the i-th partial derivative.) From the definitions, we have

$$\begin{array}{lcll}
\mathrm{Lip}_1(F) & \le & \|\Lambda_u^{-1}\|\,\mathrm{Lip}_1(f_s) & \text{(arbitrarily small)},\\
\mathrm{Lip}_2(F) & \le & \|\Lambda_u^{-1}\,\|\mathrm{Lip}_1(f_u) & \text{(arbitrarily small)},\\
\mathrm{Lip}_3(F) & = & \|\Lambda_u^{-1}\|, & \\
\mathrm{Lip}(v) & \le & \|\Lambda_s\| + \mathrm{Lip}(f_s) & (< \|\Lambda_s\| + \epsilon),\\
\mathrm{Lip}_2(v) & \le & \mathrm{Lip}_2(f_s) & \text{(arbitrarily small)}.
\end{array} \tag{5.9}$$

The last column in the above table—in parentheses—indicates how small the quantity in question can be made by making the Lipschitz norms of f_s and f_u small.

We note further that, because of our choice of norm on $\mathcal{E}_s \oplus \mathcal{E}_u$, the mapping $x \mapsto (x, w(x))$ has Lipschitz norm one, so

$$\mathrm{Lip}(x \mapsto v(x, w(x))) \le \mathrm{Lip}(v).$$

We will present the proof of the proposition as a seqence of steps, each a simple computation, showing successively that

1. $\mathcal{F}w$ is defined everywhere on $\mathcal{B}_s$
2. $\mathcal{F}w(0) = 0$
3. $\mathcal{F}$ preserves the condition $\mathrm{Lip}(w) \leq 1$
4. $\mathcal{F}$ is contractive

In the course of these computations we are led to impose a number of bounds on F and v; it follows easily from the estimates in the last column of (5.9) that each of these bounds can be made to hold by taking the Lipschitz norms of f_s and f_u small enough.

Step 1. *If*

$$Lip(v) \leq 1 \tag{5.10}$$

then $\mathcal{F}w$ *is defined on all of* $\mathcal{B}_s$ *for all* $w \in \mathcal{X}_{\mathcal{F}}$.

In terms of the notation introduced above, $\mathcal{F}w$ will be defined on all of $\mathcal{B}_s$ provided

$$\|v(x, w(x))\| < 1 \quad \text{when } \|x\| < 1.$$

Since $v(0, w(0)) = v(0, 0) = 0$, this will follow if we can arrange that the mapping $x \mapsto v(x, w(x))$ has Lipschitz norm ≤ 1. We have already remarked that this Lipschitz norm is $\leq \mathrm{Lip}(v)$. □

From now on, we assume that $\mathcal{F}w$ is defined everywhere in the open unit ball.

Step 2. $\mathcal{F}w(0) = 0$ *for all* $w \in \mathcal{X}_{\mathcal{F}}$.

This follows at once from $w(0) = 0$, $v(0, 0) = 0$, and $F(0, 0, 0) = 0$. □

Step 3. *If*

$$Lip_1(F) + Lip_2(F) + Lip_3(F)Lip(v) \leq 1, \tag{5.11}$$

then

$$Lip(\mathcal{F}w) \leq 1 \quad \textit{for all } w \in \mathcal{X}_{\mathcal{F}}.$$

From the definition (5.8) of F,

$$\begin{aligned} \mathrm{Lip}(\mathcal{F}w) &\leq \mathrm{Lip}_1(F) + \mathrm{Lip}_2(F)\mathrm{Lip}(w) + \mathrm{Lip}_3(F)\mathrm{Lip}(w)\mathrm{Lip}(v) \\ &\leq \mathrm{Lip}_1(F) + \mathrm{Lip}_2(F) + \mathrm{Lip}_3(F)\mathrm{Lip}(v). \end{aligned}$$

□

Next comes the questions of contractivity of $\mathcal{F}$. *We do not prove contractivity with respect to the Lipschitz norm, but with respect to a much weaker one.* The norm we choose is

$$\|w\| := \sup_{x \neq 0} \frac{\|w(x)\|}{\|x\|}. \tag{5.12}$$

It is easy to see that

- $\|w\| \leq \mathrm{Lip}(w)$ if $w(0) = 0$.
- $\| \ \|$ is a norm on the space of continuous mappings $w : \mathcal{B}_s \to \mathcal{E}_u$ for which $\|w\| < \infty$, and this space is complete with respect to $\| \ \|$.
- If w_n is a sequence in $\mathcal{X}_{\mathcal{F}}$ which converges with respect to $\| \ \|$, then the limit is again in $\mathcal{X}_{\mathcal{F}}$. (In fact: Pointwise convergence is enough.) Thus, $\mathcal{X}_{\mathcal{F}}$ is complete with respect to $\| \ \|$.

Step 4. *If*

$$Lip_2(F) + Lip_3(F)\, Lip(v) + Lip_3(F)\, Lip_2(v) < 1, \tag{5.13}$$

then $\mathcal{F}$ is contractive on $\mathcal{X}_{\mathcal{F}}$ with respect to $\| \ \|$.

Let w_1, w_2 be in $\mathcal{X}_{\mathcal{F}}$ and $x \in \mathcal{B}_s$, $x \neq 0$. Then

$$\begin{aligned} \mathcal{F}w_1(x) - \mathcal{F}w_2(x) &= F(x, w_1, w_1(v_1)) - F(x, w_2, w_1(v_1)) \\ &\quad + F(x, w_2, w_1(v_1)) - F(x, w_2, w_2(v_1)) \\ &\quad + F(x, w_2, w_2(v_1)) - F(x, w_2, w_2(v_2)) \end{aligned}$$

where we have suppressed many arguments, using the convention that w_i without argument means $w_i(x)$ and that v_i means $v(x, w_i(x))$. We then use the estimates:

$$\|v_1 - v_2\| \leq \mathrm{Lip}_2(v) \|w_1 - w_2\| \, \|x\|$$

and

$$\|w_2(v_1) - w_2(v_2)\| \leq \|w_1 - w_2\| \, \mathrm{Lip}(v) \, \|x\|$$

to get

$$\begin{aligned} &\|\mathcal{F}w_1(x) - \mathcal{F}w_2(x)\| \\ &\qquad \leq [\mathrm{Lip}_2(F) + \mathrm{Lip}_3(F)\, \mathrm{Lip}(v) + \mathrm{Lip}_3(F)\, \mathrm{Lip}_2(v)] \, \|w_1 - w_2\| \, \|x\|. \end{aligned}$$

Dividing by $\|x\|$ and taking the supremum over x, we see that $\mathcal{F}$ is a contraction on $\mathcal{X}_{\mathcal{F}}$ provided

$$\mathrm{Lip}_2(F) + \mathrm{Lip}_3(F)\,\mathrm{Lip}(v) + \mathrm{Lip}_3(F)\,\mathrm{Lip}_2(v) < 1,$$

as asserted. □

Putting all the pieces together: By the contraction mapping principle, $\mathcal{F}$ has one and only one fixed point in $\mathcal{X}_{\mathcal{F}}$ provided that (5.10), (5.11), (5.13) all hold, and, since

$$\|\Lambda_s\| < 1 \quad \text{and} \quad \|(\Lambda_u)^{-1}\| < 1,$$

the estimates in (5.9) show that all three of these conditions hold if the Lipschitz norms of f_s and f_u are small enough. Thus, the proof of Proposition 5.3.2 is complete. □

5.4 Generalizations

For the proof of Proposition 5.3.2 we need only estimates (5.10), (5.11) and (5.13), and these can be made to hold provided

$$\|\Lambda_s\| < 1 \quad \text{and} \quad \|\Lambda_s\|\,\|\Lambda_u^{-1}\| < 1.$$

The condition $\|\Lambda_u^{-1}\| < 1$ was never needed. This simple observation leads to the following useful generalization: Let $\rho \leq 1$ be such that

$$\sigma(Df(z_0)) \cap \{|\lambda| = \rho\} = \emptyset,$$

but such that

$$S_- := \sigma(Df(z_0)) \cap \{|\lambda| < \rho\} \quad \text{and} \quad S_+ := \sigma(Df(z_0)) \cap \{|\lambda| > \rho\}$$

are both non-empty. There is then an invariant direct sum decomposition of $\mathcal{E}$ as $\mathcal{E}_- \oplus \mathcal{E}_+$, where $\mathcal{E}_-$ (respectively $\mathcal{E}_+$) is the spectral subspace for $Df(z_0)$ corresponding to S_- (respectively S_+.) In this representation for $\mathcal{E}$, $Df(z_0)$ has block-diagonal form

$$Df(z_0) = \begin{pmatrix} \Lambda_- & 0 \\ 0 & \Lambda_+ \end{pmatrix}$$

and by proper choice of norm we can arrange that

$$\|\Lambda_-\| < \rho \quad \text{and} \quad \|\Lambda_+^{-1}\| < \rho^{-1}$$

The analysis of the preceding section applies in this more general situation and proves the existence of a Lipschitz function, defined on an open ϵ-ball about z_0 in $\mathcal{E}_-$ and taking values in $\mathcal{E}_+$, whose graph is invariant under f. The proof of the smoothness of the function—to be given in the next chapter—and of the vanishing of its derivative at z_0 also extend to this more general situation. Thus we get an invariant manifold tangent to $\mathcal{E}_-$. Such an invariant manifold is called a *strong stable manifold*; it can be characterized dynamically as the set of points whose orbits stay near to z_0 and converge to it faster than ρ^n (asymptotically.)

It is also useful—but a bit more difficult—to make a further extension to allow for $\rho > 1$. The following situation is particularly important in practice: Suppose that the spectrum of $Df(z_0)$ can be split into two parts:

$$\sigma(Df(z_0)) = S_{cs} \cup S_u$$

where

$$\sup\{|\lambda| : \lambda \in S_{cs}\} = 1 \quad \text{and} \quad \inf\{|\lambda| : \lambda \in S_u\} > 1.$$

There is again a corresponding splitting $\mathcal{E} = \mathcal{E}_{cs} \oplus \mathcal{E}_u$, etc., and we would like to show that there is an invariant manifold tangent to $\mathcal{E}_{cs}$. Such an invariant manifold is called a *center-stable manifold.*

The trouble in extending the above proof is in getting started, i.e., in proving that $\mathcal{F}w$ is defined on the whole unit ball. The estimates leading to preservation of $\mathrm{Lip}(w) \leq 1$ and to contractivity rely only on

$$\|\Lambda_{cs}\| \, \|\Lambda_u^{-1}\| < 1$$

(with a self-explanatory notation), and this bound can be made to hold without difficulty. To get $\mathcal{F}w$ defined on the whole unit ball, we needed $\mathrm{Lip}(v) \leq 1$, and for this the strict inequality

$$\|\Lambda_s\| < 1$$

is essential.

We can get around this difficulty—to some extent—by using the following idea: We cut off the non-linear terms in f away from the fixed point to get a mapping which is defined everywhere in $\mathcal{E}$ and *globally* close, in a Lipschitz sense, to a linear mapping. We then apply the analysis of the preceding section to prove the existence of a *global* invariant Lipschitz manifold W for the cut-off mapping. Because everything is defined everywhere, the problem of getting a large enough domain for $\mathcal{F}w$ disappears. The cutting-off is done in such a way as to make the cut-off mapping agree exactly with f on a neighborhood of the fixed point; then the manifold W which is globally invariant for the cut-off mapping is *locally invariant* for f in the sense that, for $\mathcal{U}$ a small enough neighborhood of the fixed point, if $z \in \mathcal{U} \cap W$, and $f(z) \in \mathcal{U}$, then $f(z) \in W$.

In more detail, the idea is as follows: We consider, in an obvious extension of earlier notation, mappings f of the form

$$f(x,y) = (\Lambda_{cs}x + f_{cs}(x,y), \Lambda_u y + f_u(x,y),$$

where f_{cs} and f_u are defined on the unit ball in $\mathcal{E}$, vanish at the origin, and have small Lipschitz norms.

We choose a Lipschitz-continuous real-valued function ψ on $\mathcal{E}_{cs}$ which is identically equal to one on a neighborhood of 0 and identically zero for $\|x\| \geq 1-\epsilon$ for some strictly positive ϵ. Such a function can easily be constructed by composing the norm with a smooth function on the positive real axis which is one on a neighborhood of 0 and 0 to the right of $1-\epsilon$. (Note, however, that by the remarks in §3.1, it may not be possible to take ψ to be differentiable if $\mathcal{E}_{cs}$ is infinite dimensional.) Define

$$\begin{aligned}
\widetilde{f_{cs}}(x,y) &= \psi(x) f_{cs}(x,y) \\
\widetilde{f_u}(x,y) &= \psi(x) f_u(x,y) \\
\widetilde{f}(x,y) &= \Big(\Lambda_{cs}x + \widetilde{f_{cs}}(x,y), \Lambda_u y + \widetilde{f_u}(x,y)\Big)
\end{aligned}$$

Although f might not have been defined on all of $\mathcal{E}$, we can evidently extend $\widetilde{f_{cs}}$, $\widetilde{f_u}$ to vanish when $\|x\| \geq 1$ and extend $\widetilde{f}$ correspondingly. By making $\mathrm{Lip}(f_{cs})$ and $\mathrm{Lip}(f_u)$ small on the unit ball, we can guarantee that the Lipschitz norms of $\widetilde{f_{cs}}$ and $\widetilde{f_u}$ *taken over the set of all* (x,y) *with* $\|y\| \leq \|x\|$ *with no restriction at all on* x are as small as we like[3]. Dropping the tildes, we are led to consider a *global* version of the invariant manifold problem, i.e., to consider a mapping f which is defined and uniformly near to linear in the sense of Lipschitz norm on—essentially—all of $\mathcal{E}_{cs} \oplus \mathcal{E}_u$ and to look for a function w defined and satisfying the condition $\mathrm{Lip}(w) \leq 1$ on all of $\mathcal{E}_{cs}$ whose graph is mapped into itself by f. Reexamination of the argument in the first few paragraph of this section shows that it works perfectly in this situation, without requiring $\mathrm{Lip}(v) \leq 1$, provided that all the Lipschitz norms are understood to be taken over the appropriate domains, e.g., the $\mathrm{Lip}_i(F)$ is to be understood

[3]We could have taken a cutoff function ψ which cuts off in both x and y, in which case the condition $\|y\| \leq \|x\|$ would be unnecessary. Cutting off only in x has the advantage that it may be easier to find smooth cutoff functions on $\mathcal{E}_{cs}$ than on $\mathcal{E}$ (particularly if $\mathcal{E}_{cs}$ is finite-dimensional and $\mathcal{E}_u$ infinite-dimensional.)

as taken over the domain

$$\{(x,y,y') : \|y\| \leq \|x\| \quad \text{and} \quad \|y'\| \leq \|x\| \quad \text{with no restriction on } x.\}$$

5.5 A more general setup

Thus, to deal prove the existence of both stable and center-stable manifolds, we need two versions of the basic existence theorem, one which requires strict contractivity of the "less expansive" part of $Df(z_0)$ and works on the unit ball, and a second which does not require this strict contractivity and works on the whole space. Once an appropriate cutoff of the non-linear terms has been made, the proofs in the two cases are essentially identical. It therefore seems efficient to carry out, simultaneously, two versions of the argument, adapted to the two situations. The idea of cutting off the nonlinear terms can evidently be made to work in more general circumstances than that described above, and we may as well carry out the proofs with this extra generality. We are thus led to carry out the analysis which will follow—notably, the proof of smoothness of the invariant manifolds—in the following somewhat abstract general setting:

Notation and standing hypotheses: *We consider two Banach spaces $\mathcal{E}_-$ and $\mathcal{E}_+$, equipped with bounded operators Λ_- and Λ_+ respectively. We denote by $\mathcal{E}$ the direct sum $\mathcal{E}_- \oplus \mathcal{E}_+$, and we equip $\mathcal{E}$ with the norm*

$$\|(x,y)\| = \max\{\|x\|, \|y\|\}.$$

We define

$$\rho_- := \sup\{|\lambda| : \lambda \in \sigma(\Lambda_-)\} \quad \text{and} \quad \rho_+ := \inf\{|\lambda| : \lambda \in \sigma(\Lambda_+)\}. \tag{5.14}$$

Our fundamental hypothesis is that

$$\rho_- < \rho_+,$$

so that the open annulus $\{\lambda : \rho_- < |\lambda| < \rho_+\}$ is disjoint from the spectrum of

$$\Lambda := \begin{pmatrix} \Lambda_- & 0 \\ 0 & \Lambda_+ \end{pmatrix} \quad \text{on } \mathcal{E}.$$

In particular, Λ_+ is invertible. We will always assume that the norms on $\mathcal{E}_\pm$ are chosen so that

$$\|\Lambda_-\| \, \|\Lambda_+^{-1}\| < 1;$$

we may later need to impose further restrictions on how the norm is chosen.

We will consider in parallel the following two situations:

– *(contractive case): $\rho_- < 1$; $\mathcal{B}_-$ denotes the open unit ball in $\mathcal{E}_-$; and $\mathcal{B}$ denotes the open unit ball in $\mathcal{E}$. The norm on $\mathcal{E}_-$ is chosen so that*

$$\|\Lambda_-\| < 1.$$

– *(non-contractive case): No assumption on ρ_-; $\mathcal{B}_-$ denotes $\mathcal{E}_-$; and $\mathcal{B}$ denotes an open convex neighborhood of $\{(x,y) \in \mathcal{E} : \|y\| \leq \|x\|\}$.*

In both cases, we consider mappings of the form

$$f(x,y) = (\Lambda_- x + f_-(x,y), \Lambda_+ y + f_+(x,y)), \tag{5.15}$$

with

$$f_\pm : \mathcal{B} \longrightarrow \mathcal{E}_\pm,$$

and

$$f_\pm(0,0) = 0.$$

When we speak of Lipschitz norms of function of x (respectively (x,y)), we mean Lipschitz norms over $\mathcal{B}_-$ (respectively $\mathcal{B}$.)

Summarizing—and extending slightly—the results of this chapter, we see that we have proved the following fundamental existence and uniqueness result:

Proposition 5.5.1 *If the $Lip(f_\pm)$ are small enough, there is a unique*

$$w : \mathcal{B}_- \to \mathcal{E}_+ \quad \text{with } w(0) = 0 \text{ and } Lip(w) \leq 1$$

whose graph is mapped into itself by f. In terms of the notation (5.8), a set of sufficient conditions for the existence and uniqueness of w is:

$$Lip(v) \leq 1 \tag{A}$$

$$Lip_1(F) + Lip_2(F) + Lip_3(F) Lip(v) \leq 1 \tag{B}$$

$$Lip_2(F) + Lip_3(F)\, Lip(v) + Lip_3(F)\, Lip_2(v) < 1 \tag{C}$$

In the non-contractive case, (A) is not needed.

The reason for including the—not very illuminating—explicit sufficient conditions (ABC) in the statement of this proposition is so that they will be available later for comparision with the conditions needed to deduce smoothness of w from smoothness of f. Indeed, since the condition (ABC) (respectively (BC)) are needed for our proof of the existence of the invariant set, it will be convenient to add them to our standing hypotheses:

Standing hypotheses 2: *From now on, we assume the above estimates (B) and (C) hold. In the contractive case, we assume further that (A) holds.*

Although it is efficient to consider the contractive and non-contractive cases together, the reader should be aware that the final results are considerably less satisfactory in the non-contractive case:

First: In the infinite-dimensional situation, there is the problem of whether the cut-off function ψ can be taken to be smooth, or, more generally, whether we can find any way of modifying the initial mapping outside a small neighborhood of 0 to produce a smooth global mapping with uniformly small nonlinear part. Unless this can be done, there would seem to be no hope of proving smoothness of the local invariant manifold.

Second: We will see in the next chapter that, if the $f_{\pm}$ are of class $\mathcal{C}^r$ with first derivatives vanishing at the origin, *and if*

$$(\rho_-)^r < \rho_+$$

then w is also $\mathcal{C}^r$. This latter condition is a consequence of $\rho_- < \rho_+$ in the contractive case, but not in the strongly non-contractive case ($\rho_- > 1$.) Examples show that it is really necessary. Thus, in the strongly non-contractive case, there is an intrinsic finite limit on the smoothness of the invariant manifold even for very smooth f's; no such limitation exists in the contractive case. (For the borderline case of the center-stable manifold, i.e., $\rho_- = 1$, there are no problems for any finite r, but we will give later an example of an analytic mapping with no $\mathcal{C}^\infty$ center-stable manifold.)

Third: As we shall see, the invariant manifold is unique in the contractive case but not—in general—in the non-contractive case. This may seem paradoxical in view of the uniqueness statement in Proposition 5.5.1. The explanation for the apparent paradox is that, to get from the original mapping to the f to which Proposition 5.5.1 applies, we need, in the non-contractive case, to cut off the non-linear terms. It is not surprising that the invariant manifold should depend on the cutoff. What may be surprising is that changing the cutoff can change the invariant manifold *even very near to the fixed point, where the cutoff does nothing.* That is: In the non-contractive case, local invariance is not really a local property; the invariant manifold can contain orbits which, although starting out very near to the fixed point, eventually get well away from it.

Chapter 6

Stable manifolds: smoothness

6.1 Differentiability of w.

We will consider here the general setup of §5.5. Proposition 5.5.1 then assures us of the existence—and uniqueness in an appropriate class—of a function whose graph is mapped into itself by f. For purposes of this chapter, w, without any further label, will always denote this function. We now add to our standing hypotheses that

Standing hypotheses 3: *the $f_\pm$ are continuously differentiable and their derivatives vanish at the origin.*

Our objective here is to prove:

Proposition 6.1.1 *With our current standing hypotheses, w is continuously differentiable on $\mathcal{B}_-$ and $Dw(0) = 0$.*

We will later investigate to what extent further smoothness of f implies further smoothness of w. It is useful to observe that this result requires no new "smallness" conditions on the $f_\pm$ beyond those needed to make the original existence proof (Proposition 5.5.1) work.

Proof. We will use the notation of the preceding chapter: The function w satisfies the equation

$$w(x) = F(x, w(x), w(v(x, w(x))))$$

where

$$F(x, y, y') = \Lambda_+^{-1}\left[y' - f_+(x, y)\right] \quad \text{and} \quad v(x, y) = \Lambda_- x + f_-(x, y).$$

With our stronger assumptions, F and v are continuously differentiable, and there are simple expressions for their derivatives at the origin. We will generally supress arguments when they are clear from the context. Thus, for example, we would normally write the right-hand side of the above functional equation for w simply as $F(x, w, w(v))$.

The general plan of attack is as follows:

1. We differentiate formally the functional equation satisfied by w to get the functional equation

$$Dw(x) = D_1 F + D_2 F Dw + D_3 F Dw(v)[D_1 v + D_2 v Dw]$$

for its derivative. Although we do not yet know that w is differentiable, we do know that, if it is differentiable, its derivative satisfies this equation.

2. We show that the above equation "can be solved uniquely for Dw." That is: We define an operator $\mathcal{K}$ acting on a space of mappings from $\mathcal{E}_-$ to $\mathcal{L}(\mathcal{E}_-, \mathcal{E}_+)$ by

$$\mathcal{K}\sigma(x) = D_1 F + D_2 F \sigma + D_3 F \sigma(v)[D_1 v + D_2 v \sigma], \tag{6.1}$$

and show that this operator has a unique fixed point in an appropriate function space. Again, we denote this unique fixed point simply by σ. Thus, without knowing that w is differentiable, we know that σ is the only possible candidate for its derivative.

3. We then use the functional equations for w and σ to verify that σ does have the defining property of the derivative of w, i.e., that

$$\|w(x + \delta x) - w(x) - \sigma(x)\delta x\| = o(\|\delta x\|),$$

and hence that w is indeed differentiable.

We will work in the space $\mathcal{X}_\mathcal{K}$ of continuous mappings σ from $\mathcal{E}_-$ to $\mathcal{L}(\mathcal{E}_-, \mathcal{E}_+)$ with $\sigma(0) = 0$ and $\|\sigma(x)\| \leq 1$ for all x. (This last bound is chosen to correspond to the condition $\mathrm{Lip}(w) \leq 1$.) We equip $\mathcal{X}_\mathcal{K}$ simply with the supremum norm. From the definition (6.1) of $\mathcal{K}$, and from $v(0, 0) = 0$, $w(0) = 0$, and $D_1 F(0, 0, 0) = 0$, it follows that $\mathcal{K}\sigma(0) = 0$ if $\sigma(0) = 0$. Also, simply by taking the norm of the defining equations, and using estimates like $\|D_1 F\| \leq \mathrm{Lip}_1(F)$, we see that $\|\mathcal{K}\sigma(x)\| \leq 1$ for all x provided $\|\sigma(x)\| \leq 1$ for all x and provided

$$\mathrm{Lip}_1(F) + \mathrm{Lip}_2(F) + \mathrm{Lip}_3(F)\mathrm{Lip}(v) \leq 1.$$

This latter inequality is exactly condition (B) in our standing hypotheses, and this is not an accident; (B) was introduced in order to guarantee that $\mathcal{F}$ preserves the condition $\text{Lip}(w) \leq 1$.

In the above we used the estimate:

$$\|D_1 v + D_2 v\sigma\| \leq \text{Lip}(v).$$

To prove this, let (x_0, y_0) be a point of $\mathcal{B}$, and let σ_0 be an operator from $\mathcal{E}_-$ to $\mathcal{E}_+$ with norm ≤ 1. Then $x \mapsto v(x, y_0 + \sigma_0(x - x_0))$ has Lipschitz norm not greater than that of v. Thus, the derivative of this mapping at x_0 has norm not greater than $\text{Lip}(v)$, i.e.,

$$\|D_1 v(x_0, y_0) + D_2 v(x_0, y_0)\sigma_0\| \leq \text{Lip}(v).$$

To find conditions which make $\mathcal{K}$ a contraction, we use the definition to write

$$\begin{aligned} \mathcal{K}\sigma_1(x) - \mathcal{K}\sigma_2(x) &= D_2 F(\sigma_1 - \sigma_2) + D_3 F(\sigma_1(v) - \sigma_2(v))[D_1 v + D_2 \sigma_1] \\ &\quad + D_3 F \sigma_2 D_2 v(\sigma_1 - \sigma_2). \end{aligned}$$

Taking norms and estimating in a completely straightforward way shows that $\mathcal{K}$ is a contraction provided

$$\text{Lip}_2(F) + \text{Lip}_3(F)\text{Lip}(v) + \text{Lip}_3(F)\text{Lip}_2(v) < 1,$$

and this is exactly condition (C), i.e., the bound needed to make $\mathcal{F}$ contractive.

We thus now let σ denote the unique fixed point for $\mathcal{K}$ in $\mathcal{X}_{\mathcal{K}}$. We want to show that $\sigma(x)$ is the derivative of w at x for each point x. For this purpose, let

$$M := \sup_x \limsup_{\delta x \to 0} \frac{\|w(x + \delta x) - w(x) - \sigma(x)\delta x\|}{\|\delta x\|}.$$

If follows from $\text{Lip}(w) \leq 1$ and $\|\sigma(x)\| \leq 1$ that $M \leq 2$. The idea now will be to use the functional equations to show that

$$M \leq \kappa M \quad \text{for some } \kappa < 1.$$

Since M is finite, this will imply $M = 0$, which is exactly the statement that $\sigma(x)$ is the derivative of w at x for all x.

To shorten the formulas, we introduce the abbreviations:

$$\delta w = w(x + \delta x) - w(x)$$

$$\delta v = v(x + \delta x, w + \delta w) - v(x, w)$$

From the functional equation for w, and the assumed differentiability of F,

$$\delta w = D_1 F \delta x + D_2 F \delta w + D_3 F[w(v+\delta v) - w(v)] + o(\delta x).$$

We then multiply δx by the expression given by the functional equation for $\sigma(x)$ and subtract from the above equation; this gives

$$\begin{aligned}\delta w - \sigma\delta x \quad = \quad & D_2 F\,[\delta w - \sigma \delta x] + D_3 F[w(v+\delta v) - w(v) - \sigma(v)\delta v] \\ & + D_3 F \sigma(v)\,[\delta v - D_1 v \delta x - D_2 v \delta w] \\ & + D_3 F \sigma(v) D_2 v\,[\delta w - \sigma \delta x] + o(\delta x).\end{aligned} \tag{6.2}$$

By the differentiability of v,

$$\delta v - D_1 v \delta x - D_2 v \delta w = o(\delta x).$$

Also,

$$\|\delta v\| \le \mathrm{Lip}(v)\|\delta x\|,$$

and hence

$$\limsup_{\delta x \to 0} \frac{\|w(v+\delta v) - w(v) - \sigma(v)\delta v\|}{\|\delta x\|} \le \mathrm{Lip}(v) M.$$

Taking norms of both sides of (6.2), dividing by the norm of δx, letting δx tend to 0, and inserting the preceding two estimates gives

$$\limsup_{\delta x \to 0} \frac{\|\delta w - \sigma(x)\delta x\|}{\|\delta x\|} \le \left[\mathrm{Lip}_2(F) + \mathrm{Lip}_3(F)\mathrm{Lip}(v) + \mathrm{Lip}_3(F)\mathrm{Lip}_2 v\right] M,$$

so w is differentiable, with derivative σ, provided

$$\mathrm{Lip}_2(F) + \mathrm{Lip}_3(F)\mathrm{Lip}(v) + \mathrm{Lip}_3(F)\mathrm{Lip}_2 v < 1,$$

and this is our condition (C). Thus, the proof of Proposition 6.1.1 is complete.
□

6.2 Hölder continuity of first derivatives

Proposition 6.2.1 *Assume, in addition to our standing hypotheses, that*

- *the $f_\pm$ are of class $\mathcal{C}^{1+\alpha}$ with $0 < \alpha < 1$*
- $Lip_2(F) + Lip_3(F)(Lip(v))^{1+\alpha} + Lip_3(F)Lip_2(v) < 1$ (C[1+α]).

Then w is of class $\mathcal{C}^{1+\alpha}$. A similar assertion holds for $\alpha = 1$; we assume that the $Df_\pm$ are Lipschitz continuous and prove that Dw is Lipschitz continuous.

The inequality (C[1+α]) follows from our standing hypotheses (C) and (A) in the contractive case. In the non-contractive case, it can be made to hold—first by choosing the norms properly and then by making the Lip($f_\pm$ small enough—if

$$(\rho_-)^{1+\alpha} < \rho_+.$$

Proof. We recall that a mapping g between metric spaces:

$$g : (X_1, d_1) \longrightarrow (X_2, d_2)$$

is said to be *α-Hölder continuous* if there is a constant k such that

$$d_2(g(x), g(x')) \leq k d_1(x, x')^\alpha \quad \text{for all } x, x' \text{ in } X_1.$$

We will let $\mathrm{Lip}^{(\alpha)}(g)$ denote the smallest such constant k, and we will refer to $\mathrm{Lip}^{(\alpha)}(\,.\,)$ as the *α-Hölder norm.*

We will present the proof of this proposition as a general lemma. This both simplifies the notation and permits us to avoid having to repeat essentially the same argument when we investigate higher derivatives. We showed above that Dw is the unique fixed point of a contractive operator $\mathcal{K}$ an appropriate function space. We will write the operator as

$$(\mathcal{K}\sigma)(x) = G(x, \sigma(x), \sigma(\bar{v}(x))) \tag{6.3}$$

where

$$G(x, s, s') = (D_1 F) + (D_2 F)s + (D_3 F)\, s' \left[(D_1 v) + (D_2 v)s\right] \tag{6.4}$$

(with the usual rule for filling in supressed arguments) and

$$\bar{v}(x) = v(x, w(x)). \tag{6.5}$$

The variable x takes values in $\mathcal{B}_-$, the variables s, s' in the closed unit ball in $\mathcal{L}(\mathcal{E}_-, \mathcal{E}_+)$. With our assumptions, G is α-Hölder continuous in its first argument (uniformly in its second and third arguments.) It is in fact quadratic—hence, $\mathcal{C}^\infty$—in its second and third arguments, but all we will use here is Lipschitz continuity in these arguments. The space on which $\mathcal{K}$ is contractive can be taken to be the set of all countinuous mappings from $\mathcal{B}_-$ into the closed unit ball of $\mathcal{L}(\mathcal{E}_-, \mathcal{E}_+)$, equipped with the supremum norm. (In the above, we imposed also the condition $\sigma(0) = 0$, but this condition was not needed in the proof of contractivity.) $\bar{v}$ is a mapping of class $\mathcal{C}^1$ from $\mathcal{B}_-$ into itself.

For later use, we consider the following setup:

– *$\mathcal{S}$ denotes a closed ball in a Banach space $\mathcal{Z}$.*

– *G denotes a continuous mapping from $\mathcal{B}_- \times \mathcal{S} \times \mathcal{S}$ to $\mathcal{S}$ which is Lipschitz with respect to its second and third variables.*

– *$\bar{v}$ denotes a Lipschitz mapping from $\mathcal{B}_-$ to itself.*

We assume that

$$Lip_2(G) + Lip_3(G) < 1 \tag{6.6}$$

from which it follows readily that the operator $\mathcal{K}$ defined by

$$(\mathcal{K}\sigma)(x) = G(x, \sigma(x), \sigma(\bar{v}(x)))$$

maps the space of all continuous functions from $\mathcal{B}_-$ to $\mathcal{S}$—equipped with the supremum norm—contractively into itself.

Lemma 6.2.2 *Let $\mathcal{S}$, G, $\bar{v}$ be as above, and assume further that G is α-Hölder continuous with respect to its first variable, uniformly with respect to the other variables, and that*

$$Lip_2(G) + Lip_3(G)\,(Lip(\bar{v}))^\alpha < 1. \tag{6.7}$$

Then the unique fixed point of $\mathcal{K}$ is α-Hölder continuous.

Proof of Proposition 6.2.1 from Lemma 6.2.2. The assumption on the α-Hölder continuity of G with respect to the first variable means that

$$\mathrm{Lip}_1^{(\alpha)}(G) := \sup\left\{\frac{\|G(x_1, s, s') - G(x_2, s, s')\|}{\|x_1 - x_2\|^\alpha} : x_1 \neq x_2 \in \mathcal{B}_-, s, s' \in \mathcal{S}\right\}$$

is finite. If we take G, $\bar{v}$ as defined in (6.4) and (6.5), it is easy to see that $\mathrm{Lip}_1^{(\alpha)}(G)$ is finite and that

$$\mathrm{Lip}_2(G) \leq \mathrm{Lip}_2(F) + \mathrm{Lip}_3(F)\mathrm{Lip}_2(v)$$

$$\mathrm{Lip}_3(G) \leq \mathrm{Lip}_3(F)\mathrm{Lip}(v)$$

$$\mathrm{Lip}(\bar{v}) \leq \mathrm{Lip}(v).$$

Hence, (6.7) follows from

$$\mathrm{Lip}_2(F) + \mathrm{Lip}_3(F)(\mathrm{Lip}(v))^{1+\alpha} + \mathrm{Lip}_3(F)\mathrm{Lip}_2(v) < 1,$$

i.e., from our assumption (C[1+α]). Thus, α-Hölder continuity of Dw follows. That (C[1+α]) follows from (C) if $\mathrm{Lip}(v) \leq 1$ is obvious from inspection of the formulas. That (C[1+α]) can be made to hold if

$$(\rho_-)^{1+\alpha} < \rho_+$$

follows easily from the table of estimates (5.9) □

Proof of Lemma 6.2.2. The idea will be as follows: We will look for a bound on $\mathrm{Lip}^{(\alpha)}(\sigma)$ which is preserved by $\mathcal{K}$, i.e., for a number B such that

$$\sigma : \mathcal{B}_- \longrightarrow \mathcal{S} \quad \text{and} \quad \mathrm{Lip}^{(\alpha)}(\sigma) \leq B \quad \text{implies} \quad \mathrm{Lip}^{(\alpha)}(\mathcal{K}\sigma) \leq B.$$

Once we have found such a B, it is easy to finish the proof of the Hölder continuity of Dw: The set of mappings

$$\sigma : \mathcal{B}_- \longrightarrow \mathcal{S} \quad \text{with} \quad \mathrm{Lip}^{(\alpha)}(\sigma) \leq B$$

is closed and hence complete with respect to the supremum norm. It is mapped into itself by $\mathcal{K}$, and $\mathcal{K}$ is contractive in the supremum norm. Hence, $\mathcal{K}$ has a fixed point in this set. But $\mathcal{K}$ has only one fixed point, so the unique fixed point must satisfy

$$\mathrm{Lip}^{(\alpha)}(\sigma) \leq B < \infty.$$

To show that such a constant B exists, we let σ be any α-Hölder continuous mapping from $\mathcal{B}_-$ to $\mathcal{S}$, and we estimate the α-Hölder norm of $\mathcal{K}\sigma$. For any $x_1 \neq x_2$, we have

$$\begin{aligned}\mathcal{K}\sigma(x_1) - \mathcal{K}\sigma(x_2) &= G(x_1, \sigma_1, \sigma(\bar{v}_1)) - G(x_2, \sigma_2, \sigma(\bar{v}_2)) \\ &= G(x_1, \sigma_1, \sigma(\bar{v}_1)) - G(x_2, \sigma_1, \sigma(\bar{v}_1)) \\ &\quad + G(x_2, \sigma_1, \sigma(\bar{v}_1)) - G(x_2, \sigma_2, \sigma(\bar{v}_1)) \\ &\quad + G(x_2, \sigma_2, \sigma(\bar{v}_1)) - G(x_2, \sigma_2, \sigma(\bar{v}_2)),\end{aligned}$$

Here we have written subscripts 1 and 2 to indicate objects to be evaluated at $x = x_1$ and $x = x_2$ respectively. We make the following straightforward estimates:

$$\|G(x_1, \sigma_1, \sigma(\bar{v}_1)) - G(x_2, \sigma_1, \sigma(\bar{v}_1))\| \leq \mathrm{Lip}_1^{(\alpha)}(G)\|x_1 - x_2\|^\alpha.$$

$$\|G(x_2, \sigma_1, \sigma(\bar{v}_1)) - G(x_2, \sigma_2, \sigma(\bar{v}_1))\| \leq \mathrm{Lip}_2(G)\mathrm{Lip}^{(\alpha)}(\sigma)\|x_1 - x_2\|^\alpha$$

$$\begin{aligned}&\|G(x_2, \sigma_2, \sigma(\bar{v}_1)) - G(x_2, \sigma_2, \sigma(\bar{v}_2))\| \leq \mathrm{Lip}_3(G)\|\sigma(\bar{v}_1) - \sigma(\bar{v}_2)\| \\ &\qquad \leq \mathrm{Lip}_3(G)\mathrm{Lip}^{(\alpha)}(\sigma)\|\bar{v}_1 - \bar{v}_2\|^\alpha \\ &\qquad \leq \mathrm{Lip}_3(G)\mathrm{Lip}^{(\alpha)}(\sigma)\left(\mathrm{Lip}(\bar{v})\right)^\alpha \|x_1 - x_2\|^\alpha\end{aligned}$$

Inserting into the above formula for $\mathcal{K}\sigma(x_1) - \mathcal{K}\sigma(x_2)$ gives

$$\|\mathcal{K}\sigma(x_1) - \mathcal{K}\sigma(x_2)\| \le \|x_1 - x_2\|^\alpha \times \\ \times \Big(\mathrm{Lip}_1^{(\alpha)}(G) + [\mathrm{Lip}_2(G) + \mathrm{Lip}_3(G)(\mathrm{Lip}(\bar{v}))^\alpha]\,\mathrm{Lip}^{(\alpha)}(\sigma)\Big)$$

Dividing by $\|x_1 - x_2\|$ and taking the supremum over $x_1 \neq x_2$ gives an estimate of the form

$$\mathrm{Lip}^{(\alpha)}(\mathcal{K}\sigma) \le A + \kappa\,\mathrm{Lip}^{(\alpha)}(\sigma),$$

where

$$A := \mathrm{Lip}_1^{(\alpha)}(G) \quad \text{and} \quad \kappa := \mathrm{Lip}_2(G) + \mathrm{Lip}_3(G)(\mathrm{Lip}(\bar{v}))^\alpha.$$

By (6.6), $\kappa < 1$, so, if we take

$$B := A/(1-\kappa),$$

we get

$$\mathrm{Lip}^{(\alpha)}(\sigma) \le B \quad \text{implies} \quad \mathrm{Lip}^{(\alpha)}(\mathcal{K}\sigma) \le B,$$

as desired □

6.3 Higher derivatives

Proposition 6.3.1 *Assume, in addition to our standing hypotheses, that*

- *the $f_\pm$ are of class $\mathcal{C}^r$, $1 < r < \infty$,*
- *$Lip_2(F) + Lip_3(F)(Lip(v))^r + Lip_3(F)Lip_2(v) < 1$* (C[r]).

Then w is of class $\mathcal{C}^r$. In the contractive case, if the $f_\pm$ are of class $\mathcal{C}^\infty$, then w is of class $\mathcal{C}^\infty$.

Proof. The assertion for $1 < r < 2$, and a partial result for $r = 2$, has already been proved in Proposition 6.2.1. The complete proof will go by induction on the integer part of r, and part of the induction step has already been proved (Lemma 6.2.2). The careful and complete formulation of the induction argument is a little heavy, and it is therefore perhaps useful to say at the outset what the idea is, and to explain where the condition (C[r]) comes from. For this purpose, let us consider only integer r. Suppose that the $f_\pm$ are of class $\mathcal{C}^r$ and that, by the induction step, we already know that w is of class $\mathcal{C}^{r-1}$.

The idea is to imitate the proof of differentiability of w. Differentiating the functional equation for w r times (formally) gives an equation of the form

$$\begin{aligned} D^r w(x) &= D_2 F D^r w + D_3 F D^r w(v)(Dv)^r \\ &\quad + D_3 F D w(v) D_2 v D^r w + \text{lower order terms}, \end{aligned}$$

where the "lower order terms" are ones not involving $D^r w$ and hence are under control by the induction hypothesis. (Dv is used here as an abbreviation for the first derivative of $x \mapsto v(x, w(x))$.) We can regard this formula as an equation to be solved for $D^r w$ in terms of the already-known lower order derivatives. In contrast to the situation with the first derivative, right-hand side of this equation is *affine* (i.e., linear plus constant), and the equation will have a unique solution if the linear operator on the right has norm < 1. This will be the case if

$$\mathrm{Lip}_2(F) + \mathrm{Lip}_3(F)(\mathrm{Lip}(v))^r + \mathrm{Lip}_3(F)\mathrm{Lip}_2(v) < 1.$$

and this is exactly (C[r]). Thus, if (C[r]) holds, we can compute what $D^r w$ must be before we know that it exists. We can then show, by roughly the same method as for the first derivative, that the solution to this equation is indeed the r-th derivative of w.

Although the proof of this last point is similar in spirit for $r = 1$ and $r > 1$, the details in the two cases are sufficiently different so that it does not seem to be practical to invent hypotheses general enough to encompass them both. What we will in fact do is to repeat the argument of §6.1 with the appropriate changes. We will do this is the general framework introduced in the preceding section which is adapted to formalizing the induction argument. The technical result needed is as follows:

Lemma 6.3.2 *Let $\mathcal{S}$, G, $\bar{v}$ be as in Lemma 6.2.2, and assume in addition to the hypotheses of that lemma that*

- $G \in \mathcal{C}^1$, *and*
- $\mathit{Lip}_2(G) + \mathit{Lip}_3(G)\mathit{Lip}(\bar{v}) < 1$.

Then σ is of class $\mathcal{C}^1$ and $D\sigma$ satisfies

$$D\sigma(x) = D_1 G + D_2 G D\sigma + D_3 G D\sigma(\bar{v}) D\bar{v} \tag{6.8}$$

With the original choices for G, etc., this lemma shows that w is of class $\mathcal{C}^2$ if the $f_\pm$ are and if (C[2]) is satisfied. The key to the induction on r is the observation that the equation for $D\sigma$ has the same general form as the equation for σ and so the lemmas can be made to apply successively to D^2w, D^3w, etc. We will first give the proof of the lemma, then explain the induction argument.

Proof of Lemma 6.3.2. Differentiating the equation for σ, we see that the derivative of σ, if it exists, must satisfy (6.8). We define a linear operator $\hat{\mathcal{K}}$ on the space $\hat{\mathcal{X}}$ of all bounded continuous functions from $\mathcal{B}_-$ to $\mathcal{L}(\mathcal{E}_-, \mathcal{Z})$ (Reminder: $\mathcal{Z}$ denotes the Banach space in which the values of σ lie) by

$$(\hat{\mathcal{K}}\tau)(x) = D_2G\,\tau(x) + D_3G\,\tau(\bar{v}(x))\,D\bar{v}(x),$$

so the equation for $D\sigma$ reads

$$D\sigma = D_1G + \hat{\mathcal{K}}(D\sigma).$$

We equip the space $\hat{\mathcal{Z}}$ with the supremum norm; then

$$\|\hat{\mathcal{K}}\| \leq \mathrm{Lip}_2(G) + \mathrm{Lip}_3(G)\mathrm{Lip}(\bar{v})$$

Call the quantity on the right κ. By assumption $\kappa < 1$, and it follows easily that the mapping

$$\tau \longmapsto D_1G + \hat{\mathcal{K}}\tau$$

maps the ball $\hat{\mathcal{S}}$ of radius $B := \mathrm{Lip}_1(G)/(1-\kappa)$ in $\hat{\mathcal{Z}}$ contractively into itself. Thus, this mapping has a unique fixed point, and we will from now on let τ denote this fixed point.

To show that τ is the derivative of σ, we introduce, as before,

$$M := \sup_x \limsup \frac{\|\sigma(x+\delta x) - \sigma(x) - \tau\delta x\|}{\|\delta x\|}.$$

By Lemma 6.2.2, with $\alpha = 1$, we see that σ is Lipschitz continuous, and from this it follows that M is finite. We are going to prove

$$M \leq \kappa M,$$

with κ as above, and this will finish the proof.

To prove this estimate, we fix x and define

$$\delta\sigma := \sigma(x+\delta x) - \sigma(x),$$

and

$$\delta v := \bar{v}(x+\delta x) - \bar{v}(x)$$

Then

$$\begin{aligned} \delta\sigma &= G(x+\delta x, \sigma+\delta\sigma, \sigma(\bar{v}+\delta v)) - G(x,\sigma,\sigma(\bar{v})) \\ &= G_1\delta x + G_2\delta\sigma + G_3(\sigma(\bar{v}+\delta v) - \sigma(\bar{v})) + o(\|\delta x\|), \end{aligned}$$

where we have used Lipschitz continuity to see that

$$\delta v \text{ and } \sigma(\bar{v}+\delta v)-\sigma(\bar{v}) \text{ are both } O(\|\delta x\|).$$

Multiplying the functional equation for τ by δx and subtracting gives

$$\begin{aligned}\delta\sigma - \tau\delta x &= G_2[\delta\sigma - \tau\delta x] \\ &\quad + G_3[\sigma(\bar{v}+\delta v) - \sigma(\bar{v}) - \tau(\bar{v})\delta v + o(\|\delta x\|).\end{aligned}$$

Since

$$\|\delta v\| \le \mathrm{Lip}(\bar{v})\|\delta x\|,$$

we have

$$\limsup_{\delta x\to 0} \frac{\|\sigma(\bar{v}+\delta v)-\sigma(\bar{v})-\tau(\bar{v})\delta v\|}{\|\delta x\|} \le \mathrm{Lip}(\bar{v})M.$$

Combining the various estimates, we get

$$\limsup_{\delta x\to 0}\frac{\|\delta\sigma - \tau\delta x\|}{\|\delta x\|} \le \mathrm{Lip}_2(G)\, M + \mathrm{Lip}_3(G)\mathrm{Lip}(\bar{v})\, M = \kappa M,$$

as desired. □

Lemma 6.3.3 *Let $\mathcal{S}$, G, $\bar{v}$ be as in Lemma 6.2.2, let $1 < r < \infty$, and assume in addition that G and $\bar{v}$ are of class $\mathcal{C}^r$ and that*

$$Lip_2(G) + Lip_3(G)(Lip(\bar{v}))^r < 1. \tag{6.9}$$

Then the unique fixed point σ is of class $\mathcal{C}^r$.

Proof. In the proof of Lemma 6.2.2 we saw that σ is differentiable and that $\tau = D\sigma$ satisfies the equation

$$\tau(x) = \hat{G}(x, \tau(x), \tau(\bar{v}(x))),$$

where

$$\hat{G}(x,t,t') := D_1G + D_2G\, t + D_3G\, t'\, D\bar{v},$$

and where the variables t, t' take values in a sufficiently large ball $\hat{\mathcal{S}}$ in the Banach space $\hat{\mathcal{Z}} := \mathcal{L}(\mathcal{E}_-, \mathcal{Z})$. If G and $\bar{v}$ are of class $\mathcal{C}^r$, then $\hat{G}$ is of class $\mathcal{C}^{r-1}$. It also follows immediately from the formula that

$$\mathrm{Lip}_2(\hat{G}) = \mathrm{Lip}_2(G) \quad \text{and} \quad \mathrm{Lip}_3(\hat{G}) \le \mathrm{Lip}_3(G)\,\mathrm{Lip}(\bar{v}).$$

Putting these considerations together we see that, if G, $\bar{v}$, $\mathcal{S}$ fulfill the original assumptions—in particular that

$$\text{Lip}_2(G) + \text{Lip}_3(G) < 1$$

—and if in addition G, $\bar{v}$ are of class $\mathcal{C}^r$, $r > 1$, and satisfy (6.9):

$$\text{Lip}_2(G) + \text{Lip}_3(G)(\text{Lip}(\bar{v}))^r < 1,$$

then $\hat{G}$, $\bar{v}$, $\hat{\mathcal{S}}$ also fulfill the original assumptions, $\hat{G}$, $\bar{v}$ are of class $\mathcal{C}^{r-1}$, and satisfy

$$\text{Lip}_2(\hat{G}) + \text{Lip}_3(\hat{G})(\text{Lip}(\bar{v}))^{r-1} < 1$$

which is just (6.9) with r replaced by $r - 1$. From these observations and Lemmas 6.2.2 and 6.3.2, it is easy to see by induction that $D\sigma$ is of class $\mathcal{C}^{r-1}$, i.e., that σ is of class $\mathcal{C}^r$. □

Proof of Proposition 6.3.1. As already noted, Lemma 6.3.2 with the original choice of G, etc., completes the proof for $r = 2$, so we need only consider $r > 2$, and we know that $\sigma = Dw$ satisfies the standard equation with

$$G(x, s, s') = (D_1F) + (D_2F)s + (D_3F)\, s' \left[(D_1v) + (D_2v)s\right]$$

and

$$\bar{v}(x) = v(x, w(x)).$$

Although it follows immediately from the smoothness assumption of the proposition that $F(x, y, y')$ and $v(x, y)$ are of class $\mathcal{C}^r$, the results of the preceding paragraph do not quite apply because w appears in the expressions for G and $\bar{v}$—e.g., D_1F is an abbreviation for $(D_1F)(x, w(x), w(v(x, w(x))))$—so we cannot directly conclude that G, $\bar{v}$ are of class $\mathcal{C}^{r-1}$. To get around this we need to make another small induction argument: We write the r of the proposition as $n + \alpha$ with n an integer and $0 < \alpha \leq 1$ and argue by induction on n. As we have already remarked, the assertion is true for $n = 1$. We thus assume the proposition for n and prove it for $n + 1$. By the induction hypothesis, w is of class $\mathcal{C}^{r-1}$; from this it follows at once that G and $\hat{v}$ are of class $\mathcal{C}^{r-1}$, then from the preceding paragraph that Dw is of class $\mathcal{C}^{r-1}$, i.e., that w is of class $\mathcal{C}^r$.

This completes the proof of the assertions for finite r. The assertion for $r = \infty$ applies only in the contractive case, and our standing assumptions then include the bound $\text{Lip}(v) \leq 1$. Thus, if (C) holds, (C[r]) holds for all $r > 1$, so w is of class $\mathcal{C}^r$ for all r, i.e. is of class $\mathcal{C}^\infty$. □

6.4 Some examples

6.4.1 Non-existence of smooth generalized stable manifolds

We will give here an example to show that the condition $\rho_-^r < \rho_+$ appearing in Proposition 6.3.1 for the existence of a $\mathcal{C}^r$ generalized stable manifold (in the case where $\rho_- > 1$) is, in at least one instance, sharp. Let $\lambda > 1$, and let

$$f(x,y) = (\lambda x, \lambda^2 y + x^2).$$

The origin is a fixed point, and

$$Df(0,0) = \begin{pmatrix} \lambda & 0 \\ 0 & \lambda^2 \end{pmatrix}.$$

We split $\mathbb{R}^2$ as the direct sum of the x and y axes, i.e., we take the x-axis for $\mathcal{E}_-$ and the y-axis for $\mathcal{E}_+$. The theory applies with $\rho_- = \lambda$ and $\rho_+ = \lambda^2$. Since

$$\rho_-^r = \lambda^r < \rho_+ = \lambda^2 \quad \text{if and only if } r < 2,$$

our general theory asserts that there is, for each $r < 2$, a locally invariant manifold of class $\mathcal{C}^r$ tangent to the x axis. We will show that there is no such manifold of class $\mathcal{C}^2$.

A manifold tangent to the x axis at $(0,0)$ is, locally near $(0,0)$, the graph of a function $y = w(x)$. The condition for invariance is easy to work out directly: The image under f of a point $(x, w(x))$ of the graph is $(\lambda x, \lambda^2 w(x) + x^2)$; this is again in the graph if and only if

$$w(\lambda x) = \lambda^2 w(x) + x^2; \tag{$*$}$$

hence, the graph is locally invariant if and only if this equation holds for all sufficiently small x. It is nearly trivial to see that no function which is twice-differentiable on a neighborhood of 0 can satisfy (*) on a—possibly smaller—neighborhood of 0: Assuming that w is twice differentiable and satisfies (*), we can differentiate twice to get

$$\lambda^2 w''(\lambda x) = \lambda^2 w''(x) + 2;$$

then put x equal to 0 to get the contradictory equation

$$\lambda^2 w''(0) = \lambda^2 w''(0) + 2.$$

6.4.2 Non-uniqueness of the center-stable manifold

Consider the (uncoupled) pair of differential equations

$$\frac{dx}{dt} = x^2, \quad \frac{dy}{dt} = (\log 2)y.$$

The general solution is easy to write down; it has the form

$$x(t) = \frac{x_0}{1 - tx_0}, \quad y(t) = 2^t y_0$$

and is defined for $(-\infty < t < 1/x_0)$ if $x_0 > 0$ and for $(1/x_0 < t < \infty)$ if $x_0 < 0$. The time-one solution mapping

$$f : (x, y) \longmapsto (\frac{x}{1-x}, 2\,y,)$$

is defined and analytic in a neighborhood of the origin, and has the origin as a fixed point. The derivative of f at the origin is

$$\begin{pmatrix} 1 & 0 \\ 0 & 2 \end{pmatrix},$$

so the center-stable subspace is the x-axis and the unstable subspace the y-axis. Consider now a solution curve for the differential equation with $x_0 > 0$. At $t \to -\infty$, $y(t)$ goes (exponentially) to 0 and $x(t)$ goes to zero like $-1/t$. Hence, the solution curve is asymptotic to the origin and has there a tangency of infinite order with the x-axis. The part of the solution curve near the origin is evidently locally invariant for the time-one map f. The union of such a solution curve with the negative x-axis (including the origin) is thus a locally-invariant $\mathcal{C}^\infty$ manifold tangent to the x-axis at the origin, i.e., is a center-stable manifold. Since there are continuously many pairwise disjoint solution curves, there are continuously many center-stable manifolds. It may nevertheless be remarked that there is only one of these—the x-axis—is an analytic manifold. This is true in general: In the $\mathcal{C}^\infty$ center-stable case, all the derivatives of w at the origin are uniquely determined by the functional equation for w in terms of the derivatives of f at the origin. Hence, there can be no more than one analytic center-stable manifold. Unfortunately, as the following example shows, there may be none at all.

6.4.3 An analytic mapping with no $\mathcal{C}^\infty$ center-stable manifold.

We start by generalizing the example given above of a mapping with no $\mathcal{C}^2$ invariant manifold tangent to $\mathcal{E}_-$. Consider a one-parameter family of mappings of the form

$$f_\lambda(x,y) = (\lambda x, 2y + h(x)),$$

where h is analytic in a neighborhood of 0, vanishes to second order at 0, and is *not* a polynomial. Applying an idea due to Ruelle and Takens, we regard the whole family of two-dimensional mappings as a single three-dimensional mapping by treating the parameter as a third coordinate with trivial evolution, and we look at center manifolds of this augmented system[1]. Specifically, we consider the mapping

$$F : (\lambda, x, y) \longmapsto (\lambda, \lambda x, 2y + h(x)).$$

Each point of the form $(\lambda, 0, 0)$ is a fixed point for this mapping, but we want to look in particular at the fixed point with $\lambda = 1$. The derivative of F at this fixed point is

$$\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 2 \end{pmatrix}$$

Thus, the center-stable subspace is the (λ, x) plane and the unstable subspace is the y-axis. A center-stable manifold is a locally invariant manifold of the form

$$\{y = w(\lambda, x)\}$$

Using the form of F we see—as in subsection 6.4.1—that the local invariance condition can be expressed as the requirement that the equation

$$w(\lambda, \lambda x) = 2w(\lambda, x) + h(x), \qquad (*)$$

hold in a neighborhood of $(1, 0)$. We can now generalize the argument of subsection 6.4.1 to show:

Let λ_n denote $2^{1/n}$. If $(*)$ holds in a neighborhood of $(\lambda_n, 0)$, and if $h^{(n)}(0) \neq 0$, then w cannot be $\mathcal{C}^n$ at $(\lambda_n, 0)$.

The proof is very simple: If w were $\mathcal{C}^n$, we could differentiate $(*)$ n times with respect to x to get

$$2\frac{\partial^n w}{\partial x^n}(\lambda_n, x) + h^{(n)}(x) = \frac{\partial^n w}{\partial x^n}(\lambda_n, \lambda_n x)(\lambda_n)^n.$$

[1]We will make more serious use of this idea when we discuss bifurcation theory. The technical point is that an invariant manifold for the augmented system gives a parametrized family of invariant manifolds for the mappings depending on a parameter, and the device of Ruelle and Takens permits us to deduce smoothness of the invariant manifold *with respect to the parameter* from results about individual mappings.

Using $(\lambda_n)^n = 2$ and setting $x = 0$ leads to the contradictory equation

$$2\frac{\partial^n w}{\partial x^n}(\lambda_n, 0) = 2\frac{\partial^n w}{\partial x^n}(\lambda_n, 0) + h^{(n)}(0).$$

By assumption, $h(x)$ is not a polynomial, i.e., there is an infinite sequence n_j going to infinity such that $h^{(n_j)}(0) \neq 0$ for all j. The corresponding λ_{n_j}'s converge to 1, so the conclusion is

There is no $\mathcal{C}^\infty$ *function* w *satisfying (*) in a neighborhood of* $(1, 0)$.

(This example is modelled on one due to S. J. van Strien, "Center manifolds are not C^∞", *Math. Z.* **166** 143-145 (1979).)

Chapter 7

Stable manifolds: dynamics

We initially introduced the stable manifold for a hyperbolic fixed point z_0 as the set of points z with forward orbit $(f^n(x) : n = 0, 1, \ldots)$ remaining near z_0 for all n and converging to z_0 as $n \to \infty$. We then constructed the stable manifold using primarily the requirement of invariance. We now want to show that the manifold we constructed really can be characterized as suggested above. This characterization has two aspects: We need to show both that orbits on the stable manifold do converge to the fixed point and that orbits not on the stable manifold do not. In both cases, there are generalizations to other kinds of generalized stable manifolds. We will work in the following general framework: The main estimates will be made in our usual magnified—and, in the non-contractive case, cut-off and globalized—context, i.e., working on the unit ball or the whole of $\mathcal{E}_-$ and imposing smallness conditions on the non-linear terms. Although this is not really necessary, we will assume here that the $f_\pm$ are continuously differentiable. Then, if for example we have established that some orbit converges to the fixed point, and we want to make sharp estimates on the *asymptotic rate* of convergence, we can assume—without changing the map, just by further magnification—that the non-linear terms are as small as we like.

For our purposes here we can forget most of the intricate notation of the preceding two chapters. We will consider mappings of the standard form:

$$f(x, y) = (\Lambda_- x + f_-(x, y), \Lambda_+ y + f_+(x, y));$$

we take the $f_\pm$ to be $\mathcal{C}^1$ and vanishing together with their first derivatives at the origin; and we will always assume that the norms are chosen so that

$$\|\Lambda_-\| \, \|\Lambda_+^{-1}\| < 1.$$

We further assume that there is a $\mathcal{C}^1$ function w, vanishing together with its first derivative at the origin, with $\mathrm{Lip}(w) \leq 1$, whose graph is mapped into itself by f; the main content of the two preceding chapters was to find conditions which guarantee the existence of such a w.

7.1 Convergence to the fixed point

We define

$$r_- = \|\Lambda_-\| + \mathrm{Lip}(f_-).$$

Proposition 7.1.1 *Assume that $r_- < 1$. Then for any z in the graph of w,*

$$\|f^n(z)\| \leq (r_-)^n \|z\|.$$

Proof. Because of the way we chose the norm on $\mathcal{E}_- \oplus \mathcal{E}_+$, and because

$$\|w(x)\| \leq \mathrm{Lip}(w)\|x\| \leq \|x\|,$$

we have

$$\|(x, w(x))\| = \|x\|.$$

The x component of $f(x, w(x))$ is

$$\Lambda_- x + f_-(x, w(x)),$$

so—by the invariance of the graph of w—for any $z = (x, w(x))$ on the graph of w

$$\begin{aligned} \|f(z)\| &= \|\Lambda_- x + f_-(x, w(x))\| \\ &\leq \|\Lambda_-\| + \mathrm{Lip}(f_-)\|x\| \\ &= r_- \|z\|. \end{aligned}$$

The assertion of the proposition follows at once by iterating this estimate. □

We can sharpen this result a little.

Proposition 7.1.2 *Under the same hypotheses, for any z in the graph of w,*

$$\limsup_{n \to \infty} \frac{1}{n} \log(\|f^n(z)\|) \leq \log \rho_-.$$

Proof. The preceding proposition says

$$\limsup_{n\to\infty} \frac{1}{n} \log(\|f^n(z)\|) \leq \log r_- .$$

Once we know that the orbit of z converges to 0, we know that, for any $\epsilon > 0$, the orbit is eventually in a ball where $\mathrm{Lip}(f_-) < \epsilon$. Thus,

$$\limsup_{n\to\infty} \frac{1}{n} \log(\|f^n(z)\|) \leq \log \|\Lambda_-\|.$$

The left-hand side of this inequality is unchanged if the norm is replaced by any equivalent norm, so we can minimize the right-hand side over all norms; this gives the desired estimate. □

7.2 Expanding wedges

The second part of the argument is to see what happens to orbits which are not on the invariant manifold. The situation here depends on the size of ρ_+:

- If $\rho_+ > 1$, then orbits not on the invariant manifold cannot stay near the fixed point for all time.
- If $\rho_+ \leq 1$, orbits not on the invariant manifold may nevertheless remain near the fixed point for all time and even converge to it, but they converge *less rapidly* than orbits on the invariant manifold.

To prove these assertions, we use a very simple version of a set of ideas idea which have extensive ramifications in the analysis of the very important phenomenon of exponential separation of orbits. As a first introduction to the idea, we consider for concreteness the standard hyperbolic case, and we assume that the norms are chosen so that

$$\|\Lambda_s\| < 1 \quad \text{and} \quad \|\Lambda_u^{-1}\| < 1.$$

The first inequality means that Λ_s is contractive; the second that Λ_u is expansive, as

$$\|\Lambda_u y\| \geq \left(\|\Lambda_u^{-1}\|\right)^{-1} \|y\| \quad \text{for all } y.$$

Consider a pair of points $z_1 = (x_1, y_1)$ and $z_2 = (x_2, y_2)$, and let $z_i' = (x_i', y_i')$ denote $f(z_i)$. If the non-linear terms were not present at all, and assuming that $y_1 \neq y_2$, we would have

$$\|y_2' - y_1'\| = \|\Lambda_u(y_2 - y_1)\| > \|(y_2 - y_1)\|.$$

Iterating, we find that the separation between the orbits of z_1 and z_2 grows exponentially.

The non-linear terms will generally spoil the simplicity of the above argument. For example, it can easily happen, even with very small non-linear terms, that $y_1 \neq y_2$ but $y'_1 = y'_2$. *This can however only happen if $x_1 - x_2$ is much larger than $y_1 - y_2$.* As we shall shortly see, if—for example—$\|y_1 - y_2\| > \|x_1 - x_2\|$, it then follows that $\|y'_1 - y'_2\| > \|y_1 - y_2\|$, provided that the non-linear terms are small enough. Not only that, but it also follows that $\|y'_1 - y'_2\| > \|x'_1 - \|x'_2\|$ so that the argument can be iterated. The result is that the orbits of z_1 and z_2 separate exponentially as long as they both remain near enough to the fixed point so that the non-linear terms remain small enough. Hence, they cannot both remain near the fixed point forever.

We now proceed to formalize this idea. It is useful to do this in the general context rather than just the original one of a hyperbolic fixed point. We let

$$r_- := \|\Lambda_-\| + \mathrm{Lip}(f_-) \quad \text{and} \quad r_+ := \|\Lambda_+^{-1}\|^{-1} - \mathrm{Lip}(f_+).$$

If the $\mathrm{Lip}(f_\pm)$ are small enough, then $r_+ > r_-$; in the standard hyperbolic situation we can similarly arrange $r_+ > 1 > r_-$, etc. We also say that the separation $z_2 - z_1$ between the two points z_1, z_2 is *predominantly expansive* if $\|y_2 - y_1\| > \|x_2 - x_1\|$.

Lemma 7.2.1 *If $r_+ > r_-$, and if the separation between z_1 and z_2 is predominantly expansive, then*

- $\|f(z_2) - f(z_1)\| \geq r_+\|z_2 - z_1\|$
- *the separation between $f(z_1)$ and $f(z_2)$ is predominantly expansive.*

If, conversely, the separation between $f(z_1)$ and $f(z_2)$ is not predominantly expansive, then

- *the separation between z_1 and z_2 is not predominantly expansive*
- $\|f(z_2) - f(z_1)\| \leq r_-\|z_2 - z_1\|$.

Proof. Write, as above, $z'_i = (x'_i, y'_i)$ for $f(z_i)$, and assume first that the separation between z_1 and z_2 is predominantly expansive. Then, because we defined

$$\|(x, y)\| = \max\{\|x\|, \|y\|\},$$

we have

$$\|z_2 - z_1\| = \|y_2 - y_1\|.$$

Then

$$\begin{aligned}\|y_2' - y_1'\| &= \|\Lambda_-(y_2 - y_1) + f_+(z_2) - f_+(z_1)\| \\ &\geq \|\Lambda_-(y_2 - y_1)\| - \mathrm{Lip}(f_+\|z_2 - z_1\| \\ &\geq \left((\|\Lambda_+^{-1}\|)^{-1} - \mathrm{Lip}(f_+)\right)\|z_2 - z_1\| \\ &= r_+\|z_2 - z_1\|,\end{aligned}$$

i.e.,

$$\|y_2' - y_1'\| \geq r_+\|z_2 - z_1\|. \tag{$*$}$$

Similarly—but without assuming the the separation between z_1 and z_2 is predominantly expanding—we have

$$\begin{aligned}\|x_2' - x_1'\| &\leq \|\Lambda_-\|\|x_2 - x_1\| + \mathrm{Lip}(f_-)\|z_2 - z_1\| \\ &\leq (\|\Lambda_-\| + \mathrm{Lip}(f_-))\,\|z_2 - z_1\| \\ &= r_-\|z_2 - z_1\|,\end{aligned}$$

i.e.

$$\|x_2' - x_1'\| \leq r_-\|z_2 - z_1\| \tag{$\dagger$}$$

From (*) it follows that

$$\|z_2' - z_1'\| \geq r_+\|z_2 - z_1\|.$$

From (*) and (†) together, and using also $r_+ > r_-$, we see that $\|y_2' - y_1'\| > \|x_2' - x_1'\|$, i.e., that the separation between z_1' and z_2' is also predominantly expansive.

It remains to deal with the assertions for the case where the separation between $f(z_1)$ and $f(z_2)$ is not predominantly expansive. From the preceding results, it follows at once that the separation between z_1 and z_2 cannot be predominantly expansive. Furthermore, from the definition of "not predominantly expansive",

$$\|z_2' - z_1'\| = \|x_2' - x_1'\|,$$

so the last assertion follows from (†). □

Proposition 7.2.2 *Assume that $r_+ > r_-$, and let z in the unit ball of $\mathcal{E}$ but not on the graph of w. Then*

1. *if $r_+ > 1$, there is an n such that $f^n(z)$ is not in the unit ball of $\mathcal{E}$.*
2. *if $r_+ \leq 1$, then either there is an n such that $f^n(z)$ is not in the unit ball of $\mathcal{E}$, or*

$$\liminf_{n\to\infty} \frac{1}{n} \log \|f^n(z)\| \geq \log r_+.$$

Proof. Write $z = (x, y)$, and let $\hat{z} = (x, w(x))$. Then, manifestly, the separation between z and $\hat{z}$ is predominantly expansive. By the preceding proposition, the separation between $f^n(z)$ and $f^n(\hat{z})$ remains predominantly expansive, and

$$\|f^n(z) - f^n(\hat{z})\| \geq (r_+)^n \|z - \hat{z}\|.$$

Now assume that $r_+ > 1$, so that the separation between $f^n(z)$ and $f^n(\hat{z})$ grows exponentially as long as both remain in $\mathcal{B}$. If $f^n(\hat{z})$ remains in the unit ball—as will be the case if $r_- < 1$—the exponential growth of the separation implies immediately that $f^n(z)$ cannot remain in the unit ball for all n.

We can avoid the assumption that $f^n(\hat{z})$ remains in the unit ball as follows: Assume that $f^n(z)$ remains in the unit ball for all n. Write (x_n, y_n) for $f^n(z)$ and $(\hat{x}_n, \hat{y}_n)$ for $f^n(\hat{z})$ Using the invariance of the graph of w and $\mathrm{Lip}(w) \leq 1$, we get

$$\|\hat{y}_n\| \leq \|\hat{x}_n\|. \qquad (\ddagger)$$

On the other hand, the proof of ?? shows that

$$\|x_n - \hat{x}_n\| \leq \alpha \|y_n - \hat{y}_n\| \quad \text{where } \alpha := r_-/r_+ < 1.$$

Inserting the assumption that $\|(x_n, y_n)\| \leq 1$ for all n and using (‡) give the bound

$$\|\hat{y}_n\| \leq \alpha \|\hat{y}_n\| + 2,$$

which implies a bound on $\|\hat{y}_n\|$, which in turn—using (‡) again—implies a bound on $\|f^n(\hat{z})\|$, which contradicts the assumed boundedness of $f^n(z)$ and the exponential separation of $f^n(\hat{z})$ from $f^n(z)$.

Thus, the proof in the case $r_+ > 1$ is complete. If $r_+ \leq 1$, and if the orbit of z remains in the unit ball for all n, then

$$\|f^n(z) - f^n(\hat{z})\| \geq (r_+)^n \|z - z_1\| \quad \text{for all } n,$$

whereas, by Proposition ??,

$$\|f^n(\hat{z})\| \leq (r_-)^n \|z_1\| \quad \text{for all } n.$$

Since $r_- < r_+$,

$$\liminf_{n\to\infty} \frac{1}{n} \log \|f^n(z)\| \leq \liminf_{n\to\infty} \frac{1}{n} \log \left((r_+)^n \|z - z_1\| - (r_-)^n \|z_1\| \right) = \log r_+.$$

□

As in the preceding section, we can sharpen the above: If z is not in the graph of w, and if $f^n(z)$ remains in the unit ball for all n, then

$$\liminf_{n\to\infty} \frac{1}{n} \log \|f^n(z)\| \geq \log \rho_+.$$

We omit the proof. As a consequence: If $\rho_- < 1$, then the graph of w consists exactly of all those z in the unit ball whose orbits remain in the unit ball for all n and for which

$$\limsup_{n\to\infty} \frac{1}{n} \log \|f^n(z)\| \leq \log \rho_-;$$

we can replace the limit superior on the left by limit inferior, and we can replace ρ_- by any number between ρ_- and ρ_+. Any one of these variants on the characterization of the points of w gives a very satisfactory *uniqueness assertion* for w.

7.3 Reformulation of results

We now want to re-express the main results obtained so far in a way which is less closely tied to specific choices of coordinates, norms, and so on. We consider a mapping f, of class $\mathcal{C}^r$ with $r \geq 1$, defined on an open set $\mathcal{U}$ in a Banach space $\mathcal{E}$, and we let z_0 denote a fixed point of f. We let ρ be a positive real number such that the spectrum of $Df(z_0)$ does not intersect the circle $\{\lambda : |\lambda| = \rho\}$ but such that it does intersect both the inside and the outside of this circle. We define

$$\begin{aligned} \rho_- &:= \sup\{|\lambda| : \lambda \in \sigma(Df(z_0)), |\lambda| < \rho\} \\ \rho_+ &:= \inf\{|\lambda| : \lambda \in \sigma(Df(z_0)), |\lambda| > \rho\}. \end{aligned}$$

In the usual way, the splitting of the spectrum of $Df(z_0)$ into the parts inside and outside the circle $\{|\lambda| = \rho\}$ gives a direct sum splitting of $\mathcal{E}$ as $\mathcal{E}_- \oplus \mathcal{E}_+$. We will correspondingly represent z_0 as (x_0, y_0). We will use the phrase *local $\mathcal{C}^r$ manifold tangent to $\mathcal{E}_-$ at z_0* to refer to a set which can be represented as the graph of a $\mathcal{C}^r$ function w, defined on an open neighborhood $\mathcal{V}_w$ of x_0 in $\mathcal{E}_-$, taking values in $\mathcal{E}_+$, with $w(x_0) = y_0$ and $Dw(x_0) = 0$. We will also use the following terminology: A set X is said to be *locally invariant for f at z_0* if there is a neighborhood $\mathcal{W}$ of z_0 such that

$$z \in X \cap \mathcal{W} \quad \text{implies} \quad f(z) \in X.$$

i.e., such that $f(X \cap \mathcal{W}) \subset X$, and we say that two sets X_1 and X_2 *coincide near z_0* if there is a neighborhood $\mathcal{W}$ of z_0 such that

$$\mathcal{W} \cap X_1 = \mathcal{W} \cap X_2.$$

Theorem 7.3.1 (Generalized stable manifold theorem) *Let the notation be as above. If $\rho_- \geq 1$ assume*

i. $r < \infty$

ii. $(\rho_-)^r < \rho_+$

iii. There is a real-valued $\mathcal{C}^r$ function ψ on $\mathcal{E}_-$ identically equal to one on a neighborhood of 0 and vanishing outside the unit ball.

Then:

1. There exists a $\mathcal{C}^r$ manifold $W^{(-)}_{loc}$ tangent at z_0 to $\mathcal{E}_-$ and locally invariant there for f.

2. If, further, $\rho_- < 1$, then:

- *$W^{(-)}_{loc}$ can be taken—by judicious choice of $\mathcal{V}_w$—to be invariant in the literal sense, i.e., mapped into itself by f. (Such a choice is however not assumed in the following assertions.)*
- *$W^{(-)}_{loc}$ is "locally unique" in the sense that any two locally invariant $\mathcal{C}^1$ manifolds tangent to $\mathcal{E}_-$ at z_0 coincide near z_0.*
- *if $\mathcal{W}$ is any sufficiently small open neighborhood of z_0, then $W^{(-)}_{loc}$ coincides near z_0 with the set of all $z \in \mathcal{W}$ whose forward orbits remain in $\mathcal{W}$ for all n and converge to z_0 sufficiently rapidly that*

$$\limsup \frac{1}{n} \log(\|f^n(z) - z_0\|) \leq \log(\rho_-).$$

3. If $\rho_+ > 1$ (whether or not $\rho_- < 1$), we have the following partial uniqueness result: If $W^{(-)}_{loc}$ is any $\mathcal{C}^1$ locally invariant manifold tangent to $\mathcal{E}_-$ at z_0, there is an open neighborhood $\mathcal{W}$ of z_0 such that any z such that $f^n(z) \in \mathcal{W}$ for all $n = 0, 1, \ldots$ must be in $W^{(-)}_{loc}$.

4. If $\rho_- < 1 < \rho_+$ (hyperbolic case), then for any sufficiently small neighborhood $\mathcal{W}$ of z_0, $W^{(-)}_{loc}$ coincides near z_0 with the set of z such that $f^n(z)$ is in $\mathcal{W}$ for all $n = 0, 1, \ldots$.

Chapter 8

Generalized unstable manifolds

8.1 Introduction

Given a splitting $\mathcal{E}_- \oplus \mathcal{E}_+$ as in the preceding sections, we want to investigate invariant manifolds tangent to $\mathcal{E}_+$. We will call these *generalized unstable manifolds*. The most important case is that where $\rho_- \leq 1 < \rho_+$; a corresponding invariant manifolds are called an *unstable* manifold. Similarly, if $\rho_- > 1$, we speak of a *strong unstable* manifold, and, if $\rho_+ = 1$, of a *center-unstable* manifold. If f is invertible, we can deduce the existence and properties of generalized unstable manifolds by applying our results about stable manifolds to f^{-1}. It turns out, however, that a theory of unstable manifolds completely parallel to that for stable manifolds can be developed *without assuming the invertibility of* f. This is a useful thing to do, since in applications to partial differential equations one expects that derivatives of the mappings encountered will often be compact operators and hence not invertible but with finite-dimensional $\mathcal{E}_+$'s, so restricting to a generalized unstable manifold may give a means of extracting the essential finite-dimensional part from such infinite-dimensional dynamical systems.

As before, we want to study mappings of the form

$$f(x, y) = (\Lambda_- x + f_-(x, y), \Lambda_+ y + f_+(x, y)),$$

with $f_\pm$ small in Lipschitz norm and vanishing at the origin, and we seek invariant sets which are graphs of mappings w from $\mathcal{E}_+$ to $\mathcal{E}_-$ with Lipschitz norm no greater than one.

The first question we have to address is: In what sense do we want to require that a small piece of manifold W be invariant? There are two plausible candidates:

$$fW \supset W \quad \text{and} \quad f^{-1}W \subset W,$$

which are *not* equivalent in the non-invertible case. For orientation, consider the case of a linear f with non-trivial null space. In this case, it is clear that $f\mathcal{E}_+ \supset \mathcal{E}_+$, whereas any W containing 0 and invariant in the second sense must contain the null space of f. The lesson we draw from this is that we should look for a function w whose graph W satisfies

$$fW \supset W.$$

The general plan is now much as before: We are looking for a function whose graph is invariant in this sense. We convert the condition of invariance of the graph of w to a functional equation of the form

$$w = \mathcal{F}w.$$

We then prove that $\mathcal{F}$ is contractive in an appropriate functions space (provided that $\mathrm{Lip}(f_\pm)$ are small enough.) With the unique candidate for an invariant manifold thus in hand, we proceed to analyze its smoothness and—in nice cases—to give it a dynamical characterization. In spite of the similarity of the general plan to that used for investigating the generalized stable manifold, there are significant differences in detail, and, unfortunately, the analysis in the present case turns out to be noticeably more complicated than for generalized stable manifolds.

8.2 The Lipschitz invariant manifold

Our first task is to convert invariance of the graph into a functional equation for w. That is: We want to express in terms of w the condition that every point $z = (w(y), y)$ of the graph of w can be written as $f(w(\bar{y}), \bar{y})$. Using the detailed form for f we find the pair of conditions

$$y = \Lambda_+\bar{y} + f_+(w(\bar{y}), \bar{y}) \tag{8.1}$$

and

$$w(y) = \Lambda_- w(\bar{y}) + f_-(w(\bar{y}), \bar{y}). \tag{8.2}$$

We will handle these equations by showing first that, if f_- is small enough, (8.1) can be solved uniquely for $\bar{y}$. Let us accept for the moment that this can be done, and let us denote the solution by

$$\bar{y} = v(y, w).$$

In contrast to the analogous function encountered in the proof of the existence of generalized stable manifolds, v here depends of the *whole function* w and not just on its value at one point.

Inserting the solution into (8.2) gives a functional equation, in the form of a fixed-point problem, for w:

$$w = \mathcal{F}w,$$

where

$$(\mathcal{F}w)(y) = F(w(v(y,w)), v(y,w)) \quad \text{with} \quad F(x,\bar{y}) = \Lambda_- x + f_-(x,\bar{y}). \qquad (8.3)$$

Thus: Assuming solvability of (8.1) for all relevant y, the graph W of w satisfies $fW \supset W$ if and only if $w = \mathcal{F}w$.

Before turning to the detailed estimates, let us anticipate the fact that, as for generalized stable manifolds, we will need to carry out two parallel versions of the argument. What distinguishes the good from not-so-good cases here is whether or not $\rho_+ > 1$. We will use the following notation: $\mathcal{B}_-$ will denote—in both cases—the closed unit ball in $\mathcal{E}_-$. Then either

- **Expansive case.** $\rho_+ > 1$, in which case we choose the norm so that $\|\Lambda_+^{-1}\| < 1$, and we let $\mathcal{B}_+$ denote the closed unit ball in $\mathcal{E}_+$, or

- **Non-expansive case.** $\rho_+ \leq 1$, in which case—from the general hypothesis $\rho_- < \rho_+$—we have $\rho_- < 1$, so we choose the norm so that $\|\Lambda_-\| < 1$, and we let $\mathcal{B}_+$ denote $\mathcal{E}_+$.

In both cases, $\mathcal{B}$ will denote $\mathcal{B}_- \times \mathcal{B}_+$. The functions $f_\pm$ are defined and Lipschitz continuous on $\mathcal{B}$; Lipschitz norms mean norms over $\mathcal{B}$. We always assume $f_\pm(0) = 0$ (and, when we assume differentiability, we also assume $Df_\pm(0) = 0$.) In the non-expansive case, we require that

$$f_\pm(x,y) = 0 \quad \text{for } \|y\| > 1.$$

It then follows that

$$\|f_\pm(x,y)\| \leq \mathrm{Lip}(f_\pm) \quad \text{for all relevant } x, y.$$

The functions w which we consider are mappings from $\mathcal{B}_+$ to the unit ball of $\mathcal{E}_-$ with

$$\mathrm{Lip}(w) \leq 1 \quad \text{and} w(0) = 0.$$

In the expansive case, the requirement that

$$\|w(y)\| \leq 1 \quad \text{for all } y$$

follows from the preceding assumptions; in the non-expansive case, it must be imposed as a separate condition, and we have to check that it is preserved by the operator $\mathcal{F}$. From the formula (8.3) for $\mathcal{F}$, it follows that $\|\mathcal{F}w(y)\| \leq 1$ for all y provided that $\|w(y)\| \leq 1$ for all y and

$$\|\Lambda_-\| + \mathrm{Lip}(f_-) \leq 1, \tag{8.4}$$

and we accordingly add this estimate to our list of standing hypotheses (in the non-expansive case.)

The solvability of (8.1) is an easy application of the contraction mapping principle: We rewrite the equation in question as

$$\bar{y} = \Lambda_+^{-1} y - \Lambda_+^{-1} f_+(w(\bar{y}), \bar{y}) \tag{8.5}$$

i.e., as a fixed-point problem for the unknown $\bar{y}$. The right-hand side of (8.5) is a Lipschitz-continuous function of $\bar{y}$ with Lipschitz constant not greater than $\|\Lambda_+^{-1}\|\mathrm{Lip}(f_+)$ (provided that $\mathrm{Lip}(w) \leq 1$.) Thus, if

$$\|\Lambda_+^{-1}\| \, \mathrm{Lip} f_+ < 1, \tag{8.6}$$

the right-hand side is contractive. To apply the contraction mapping principle, we need to specify a domain in which contractivity holds and which is mapped into itself. In the expansive case, we take as domain the closed unit ball of $\mathcal{E}_+$. A straightforward estimate shows that, if

$$\|\Lambda_+^{-1}\| \, [1 + \mathrm{Lip}(f_+)] \leq 1, \tag{A$_\mathrm{e}$}$$

the right-hand side of (8.5) maps the closed unit ball into itself for all relevant y and w. By good fortune, the condition ($\mathrm{A_e}$) already includes the contractivity condition (8.6) so:

If (A_e) holds, then, for any given w and y as above, there is exactly one $\bar{y}$ in the unit ball such that (8.1) holds. As already indicated, we denote this $\bar{y}$ by $v(y, w)$.

In the non-expansive case $\rho_+ \leq 1$, we need only assume (8.6) (with Lipschitz norm now taken over all of $\mathcal{B}_- \times \mathcal{E}_+$,) and we can then solve for arbitrary y; the solution, of course, need not lie in the unit ball. On the other hand, because we want to restrict the x variable to the unit ball, we need to impose (8.4). Thus, our first quantitative condition in the non-expansive case is

$$\|\Lambda_+^{-1}\| \operatorname{Lip}(f_-) < 1 \quad \text{and} \quad \|\Lambda_-\| + \operatorname{Lip}(f_-) \leq 1, \tag{A_{ne}}$$

Conditions (A_e) respectively (A_{ne}) suffice to guarantee the solvability of (8.1) and hence the definition of $v(y,w)$ for y in the closed unit ball respectively all of $\mathcal{E}_+$.

We next want to estimate the Lipschitz norm of v with respect to y. By definition, the mapping $y \mapsto v(y,w)$ is the inverse of the mapping

$$y \longmapsto \Lambda_+ y + f_+(w(y), y),$$

so an estimate on the Lipschitz norm of v means a lower bound on the extent to which this latter mapping expands distances. Thus, let y_1, y_2 be two points of $\mathcal{B}_+$, i.e., either the closed unit ball in $\mathcal{E}_+$ (expansive case) or all of $\mathcal{E}_+$ (non-expansive case), and let x_1, x_2 two points of the closed unit ball $\mathcal{B}_-$ of $\mathcal{E}_-$ such that

$$\|x_1 - x_2\| \leq \|y_1 - y_2\|.$$

Then

$$\begin{aligned} &\|\{\Lambda_+ y_1 + f_+(x_1, y_1)\} - \{\Lambda_+ y_2 + f_+(x_2, y_2)\}\| \\ &\quad \geq \|\Lambda_+ (y_1 - y_2)\| - \operatorname{Lip}(f_+)\|y_1 - y_2\| \\ &\quad \geq \left(\|\Lambda_+^{-1}\|^{-1} - \operatorname{Lip}(f_+)\right) \|y_1 - y_2\|. \end{aligned}$$

We now let $L(v)$ be the smallest constant such that

$$\|\{\Lambda_+ y_1 + f_+(x_1, y_1)\} - \{\Lambda_+ y_2 + f_+(x_2, y_2)\}\| \geq \frac{\|y_1 - y_2\|}{L(v)}$$

whenever the x_i and y_i are as above; the computation we have just done shows that

$$L(v) \leq \frac{\|\Lambda_+^{-1}\|}{1 - \|\Lambda_+^{-1}\| \operatorname{Lip}(f_+)} = \|\Lambda_+^{-1}\| + O(\operatorname{Lip}(f_+)).$$

The rule of thumb is thus that we can make $L(v)$ almost as small as $1/\rho_+$, by choice of norm and smallness conditions on $\operatorname{Lip}(f_+)$. In the interest of keeping the formulas simple, it is useful to note that—by an easy but not-quite-obvious calculation—*our standing assumption (A_e) in the expansive case implies that* $L(v) \leq 1$.

It is evident from the definitions that, for any w with $\operatorname{Lip}(w) \leq 1$,

$$\operatorname{Lip}(y \mapsto v(y,w)) \leq L(v).$$

We note for later use that, if the $f_\pm$ are continuously differentiable and if A is a linear operator from $\mathcal{E}_+$ to $\mathcal{E}_-$ with $\|A\| \le 1$, then

$$\left\| [\Lambda_+ + D_1 f_+ A + D_2 f_+]^{-1} \right\| \le L(v); \tag{8.7}$$

this can be seen by applying the definition of $L(v)$ with y_1 and y_2 very close together.

It is now an easy matter to estimate the Lipschitz norm of $\mathcal{F}w$ in terms of $L(v)$: The construction of $\mathcal{F}w(y)$ can be factorized as

$$y \mapsto \bar{y} := v(y, w) \mapsto z := (w(\bar{y}), \bar{y}) \mapsto F(z).$$

The first factor has Lipschitz norm $\le L(v)$, and the second Lipschitz norm unity. Hence

$$\mathrm{Lip}(\mathcal{F}w) \le L(v)\,\mathrm{Lip}(F).$$

We thus impose as a second quantitative condition

$$\mathrm{Lip}(F)\,L(v) \le 1, \tag{B}$$

and it then follows that $\mathrm{Lip}(\mathcal{F}w) \le 1$ for all w.

From the definition (8.3) of F,

$$\mathrm{Lip}(F) \le \|\Lambda_-\| + \mathrm{Lip}(f_-),$$

so we can make $\mathrm{Lip}(F)$ almost as small as ρ_-. Since, as already noted, we can make $L(v)$ almost as small as $1/\rho_+$, condition (B) can be satisfied by first choosing the norms properly and then taking $\mathrm{Lip}(f_\pm)$ small enough. It is straightforward to check that $\mathcal{F}$ preserves the condition $w(0) = 0$.

As in the investigation of generalized stable manifolds, we will prove contractivity of $\mathcal{F}$ with respect to the norm

$$\|w\| = \sup_y \frac{\|w(y)\|}{\|y\|}.$$

We consider a pair of mappings w_1, w_2, and we write

$$\begin{aligned}(\mathcal{F}w_1 - \mathcal{F}w_2)(y) &= F(w_1(v_1), v_1) - F(w_2(v_1), v_1) \\ &\quad + F(w_2(v_1), v_1) - F(w_2(v_2), v_2),\end{aligned}$$

where we have abbreviated $v(y, w_i)$ by v_i, $i = 1, 2$. Now

$$\begin{aligned}\|F(w_1(v_1), v_1) - F(w_2(v_1), v_1)\| &\le \mathrm{Lip}_1(F)\|w_1 - w_2\|\|v_1\| \\ &\le \mathrm{Lip}_1(F)\|w_1 - w_2\|L(v)\|y\|,\end{aligned}$$

and

$$\|F(w_2(v_1), v_1) - F(w_2(v_2), v_2)\| \leq \mathrm{Lip}(F)\|v_1 - v_2\|.$$

To make effective use of this last bound, we need to estimate $\|v_1 - v_2\|$ carefully. Recalling the definition of $v(y, w)$ as the solution of (8.1), we write

$$\Lambda_+ v_1 + f_+(w_1(v_1), v_1) = \Lambda_+ v_2 + f_+(w_2(v_2), v_2)$$

or

$$\begin{aligned}\{\Lambda_+ v_1 + f_+(w_1(v_1), v_1)\} - \{\Lambda_+ v_2 + f_+(w_1(v_2), v_2)\} \\ = f_+(w_1(v_2), v_2) - f_+(w_2(v_2), v_2).\end{aligned}$$

Taking norms and using the definition of $L(v)$ gives

$$\begin{aligned} \frac{\|v_1 - v_2\|}{L(v)} &\leq \| \{\Lambda_+ v_1 + f_+(w_1(v_1), v_1)\} - \{\Lambda_+ v_2 + f_+(w_1(v_2), v_2)\} \| \\ &= \|f_+(w_1(v_2), v_2) - f_+(w_2(v_2), v_2)\| \\ &\leq \mathrm{Lip}_1(f_+)\, \|w_1(v_2) - w_2(v_2)\| \\ &\leq \mathrm{Lip}_1(f_+)\, \|w_1 - w_2\| \|v_2\| \\ &\leq \mathrm{Lip}_1(f_+)\, \|w_1 - w_2\|\, L(v)\, \|y\|, \end{aligned}$$

(where, to get the last expression, we used the estimate

$$\mathrm{Lip}(y \mapsto v(y, w_2)) \leq L(v).)$$

Simplifying the above:

$$\|v_1 - v_2\| \leq \mathrm{Lip}_1(f_+)\, (L(v))^2\, \|w_1 - w_2\|\, \|y\|.$$

Combining the above estimates gives

$$\begin{aligned} \| (\mathcal{F} w_1 - \mathcal{F} w_2)(y)\| \quad \leq \quad & \|w_1 - w_2\|\, \|y\| \times \\ & \times \left(\mathrm{Lip}_1(F) L(v) + \mathrm{Lip}_1(f_+)\mathrm{Lip}(F)(L(v))^2\right) \end{aligned}$$

Thus, if we impose

$$\left(\mathrm{Lip}_1(F) + \mathrm{Lip}_1(f_+)\mathrm{Lip}(F)L(v)\right)L(v) < 1, \qquad (C)$$

it follows that $\mathcal{F}$ is contractive and hence that it has exactly one fixed point in the space of w's under consideration. Since we can make $\mathrm{Lip}_1(F)$ almost as small as ρ_-, $L(v)$ almost as small as $1/\rho_+$, and $\mathrm{Lip}_1(f_+)$ as small as we like, condition (C) can be satisfied.

Summarizing what we have shown so far:

If (A_e) (in the expansive case) or (A_{ne}) (in the non-expansive case), (B), and (C) hold, then there is a unique $w : \mathcal{B}_- \to \mathcal{B}_+$ with Lip(w) ≤ 1 and $w(0) = 0$, whose graph W satisfies $fW \supset W$.

From now on,

- We will assume ($\mathrm{A_e}$) respectively ($\mathrm{A_{ne}}$), (B), and (C).
- w will denote the unique fixed point of $\mathcal{F}$.

8.3 Differentiability

We now assume that the $f_\pm$ are of class $\mathcal{C}^1$ with derivatives vanishing at 0, and we want to show that w is also of class $\mathcal{C}^1$ with derivative vanishing at 0. Again, this goes roughly as before: We first find a functional equation which must be satisfied by the derivative of w if this derivative exists; we then show that this equation has a unique solution (subject to some bounds); and we then show that the unique solution is indeed the derivative of w.

If w is differentiable than so—by the inverse function theorem—is v (as a function of y, with w fixed), and

$$(Dv)(y) = \left[\Lambda_+ + D_1 f_+(w(v), v) Dw(v) + D_2 f_+(w(v), v)\right]^{-1};$$

then differentiating the equation for w gives

$$Dw(y) = \left[D_1 F Dw(v) + D_2 F\right] Dv(y),$$

where suppressed arguments are now to be understood to be $(w(v(y)), v(y))$, and where we are now regarding v as a function of the single variable y. Hence, if w is differentiable, then, writing σ for Dw, we have

$$\sigma(y) = \left[D_1 F + D_2 F \sigma\right]\left[\Lambda_+ + D_1 f_+ \sigma(v) + D_2 f_+\right]^{-1}.$$

We will denote the right-hand side of this equation as $(\mathcal{K}\sigma)(x)$, and we want to show that $\mathcal{K}$ is a contraction on the space of continuous mappings $y \mapsto \sigma(y)$

- from $\mathcal{B}_+$ to the closed unit ball in $\mathcal{L}(\mathcal{E}_+, \mathcal{E}_-)$
- with $\sigma(0) = 0$,

equipped with the supremum norm. From $D_1 F(0,0) = 0$ and $v(0) = 0$ it follows that $\mathcal{K}$ preserves the condition $\sigma(0) = 0$. From (8.7),

$$\left\| \left[\Lambda_+ + D_1 f_+ \sigma(v) + D_2 f_+\right]^{-1} \right\| \leq L(v). \qquad (*)$$

Also, by the argument in the fine-print section of §6.1,

$$\|D_1F + D_2F\sigma\| \le \mathrm{Lip}(F).$$

Hence $\mathcal{K}$ preserves the condition $\|\sigma(y)\| \le 1$ for all y provided

$$\mathrm{Lip}(F)\mathrm{Lip}(v) \le 1,$$

and this is condition (B) again.

To find conditions which make $\mathcal{K}$ contractive, we choose σ_1 and σ_2 and write

$$\begin{aligned}(\mathcal{K}\sigma_1)(y) - (\mathcal{K}\sigma_2)(y) = D_1F\,[\sigma_1(v) - \sigma_2(v)]\,[\Lambda_+ + D_1f_+\sigma_1(v) + D_2f_+]^{-1} \\ + [D_1F + D_2F\sigma_2(v)]\,[\Lambda_+ + D_1f_+\sigma_1(v) + D_2f_+]^{-1}\,D_1f_+\times \\ \times [\sigma_2(v) - \sigma_1(v)]\,[\Lambda_+ + D_1f_+\sigma_1(v) + D_2f_+]^{-1},\end{aligned}$$

where we have used the standard operator identity

$$U^{-1} - V^{-1} = U^{-1}(V - U)V^{-1}.$$

Using the bound $(*)$ repeatedly, we get

$$\|\mathcal{K}\sigma_1(y) - \mathcal{K}\sigma_2(y)\| \le \left[\mathrm{Lip}_1(F)L(v) + \mathrm{Lip}_1(f_+)\mathrm{Lip}(F)(L(v))^2\right]\|\sigma_1(v) - \sigma_2(v)\|,$$

which shows that $\mathcal{K}$ is contractive provided that condition (C) holds.

Once again we let σ denote the unique fixed point for $\mathcal{K}$; we write

$$M = \sup_y \limsup_{\delta y \to 0} \frac{\|w(y+\delta y) - w(y) - \sigma(y)\delta y\|}{\|\delta y\|};$$

we observe that $M \le 2$; and we want to use the functional equation to show that $M \le \kappa M$ for some $\kappa < 1$. We will use the abbreviations

$$\delta v := v(y+\delta y, w) - v(y,w);$$

$$\delta w := w(v(y,w) + \delta v) - w(v(y,w)).$$

(The latter, while a little unnatural, is convenient because $w(v(y,w))$ occurs frequently on the right-hand sides of the functional equations whereas $w(y)$ does not.) Note also that suppressed arguments in $\mathcal{E}_+$ are to be taken to be $v(y,w)$ and in $\mathcal{E}_-$ are to be taken to be $w(v(y,w))$.

Writing the equations satisfied by $v(y,w)$ and by $v(y+\delta y, w)$, subtracting, and using the differentiability of f_+, gives

$$\begin{aligned}\delta y &= \Lambda_+\delta v + D_1f_+\delta w + D_2f_+\delta v + o(\delta y) \\ &= [\Lambda_+ + D_1f_+\sigma + D_2f_+]\,\delta v + D_1f_+\,[\delta w - \sigma\delta v] + o(\delta y).\end{aligned}$$

Similarly, writing the functional equation for w and using the differentiability of the $f_\pm$, we get

$$\begin{aligned} w(y+\delta y) - w(y) &= D_1 F \delta w + D_2 F \delta v + o(\delta y) \\ &= [D_1 F\sigma + D_2 F]\, \delta v + D_1 F\, [\delta w - \sigma \delta v] + o(\delta y) \\ &= [D_1 F\sigma + D_2 F]\, [\Lambda_+ + D_1 f_+ \sigma + D_2 f_+]^{-1}\, \delta y + D_1 F\, [\delta w - \sigma\delta v] \\ &\quad - [D_1 F\sigma + D_2 F]\, [\Lambda_+ + D_1 f_+ \sigma + D_2 f_+]^{-1}\, D_1 f_+\, [\delta w - \sigma \delta v] + o(\delta y). \end{aligned}$$

The first term on the right of this last equation can be recognized as $\sigma(y)\delta y$, so we get

$$\begin{aligned} & w(y+\delta y) - w(y) - \sigma(y)\delta y = \\ & \quad \Big(D_1 F - (D_1 F \sigma + D_2 F)\, (\Lambda_+ + D_1 f_+ \sigma + D_2 f_+)^{-1}\, D_1 f_+ \Big)\, (\delta w - \sigma \delta v) \\ & \qquad + o(\delta y). \end{aligned}$$

Taking norms, dividing by $\|\delta y\|$, letting δy tend to zero, and using

$$\limsup_{\delta y \to 0} \frac{\|\delta w - \sigma \delta v\|}{\|\delta y\|} \le M\, \mathrm{Lip}(v),$$

we get

$$\begin{aligned} & \limsup_{\delta y \to 0} \frac{\|w(y+\delta y) - w(y) - \sigma(y)\delta y\|}{\|\delta y\|} \\ & \qquad \le \Big(\mathrm{Lip}_1(F) + \mathrm{Lip}(F) \mathrm{Lip}(v) \mathrm{Lip}_1(f_+) \Big)\, \mathrm{Lip}(v)\, M. \end{aligned}$$

The factor multiplying M on the right is < 1 by (C), so $M = 0$, so w is differentiable with derivative σ.

The analysis of Hölder continuity of the first derivative and of higher differentiability goes in almost exactly the same way as for generalized stable manifolds. We can write the equation satisfied by $\sigma = Dw(y)$ in the form

$$\sigma(y) = G(y, \sigma(v(y))),$$

where

$$G(y, s') = \Big(D_1 F + D_2 F s' \Big) \Big(\Lambda_+ + D_1 f_+ s' + D_2 f_+ \Big)^{-1}.$$

Lemmas 6.2.2, 6.3.2, and 6.3.3 apply here—with a G depending only on two of the original three variables—so the analysis given for the stable manifold can be applied directly. The outcome can be summed up as follows:

Proposition 8.3.1 *Let the $f_\pm$ be of class $\mathcal{C}^r$, $1 \le r \le \infty$, with $Df_\pm(0) = 0$. Assume (A), (B), (C) and also, if $Lip(v) > 1$,*

$$(Lip_1(F) + Lip(F)Lip_1(f_+)Lip(v))\, Lip(v)^r < 1. \qquad (C_r)$$

(from which it follows that $r < \infty$.) Then $\mathcal{F}$ has a unique fixed point w with $w(0) = 0$ and $Lip(w) \le 1$. This w is of class $\mathcal{C}^r$, and its first derivative vanishes at the origin.

8.4 Expanding and contracting wedges

We now want to apply the expanding and contracting wedges idea of §7.2. We have defined the separation between two points $z_i = (x_i, y_i)$ to be predominantly expansive if $\|x_1 - x_2\| < \|y_1 - y_2\|$. If the separation is not predominantly expansive, we will say that it is *predominantly contractive* (in spite of the fact that the strict inequality in the definition of "predominantly expansive" makes the two definitions slightly unsymmetric.) Also, by a *backward orbit* for a point z under a (not necessarily invertible) mapping f, we will mean a sequence z_{-1}, z_{-2},...,z_{-n} with $f(z_{-1}) = z$ and $f(z_{-j}) = z_{-j+1}$ for $j = 2, 3, \ldots, n$. Since we are not assuming invertibility, there may be no backward orbit at all, or there may be more than one.

Proposition 8.4.1 *Let the hypotheses and notation be as in §7.2, but assume in addition that*

- $r_+ > 1$
- *There is a function w from the unit ball in $\mathcal{E}_+$ into $\mathcal{E}_-$ with $w(0) = 0$ and $Lip(w) \le 1$ whose graph is mapped onto itself by f.*
- *$f_\pm$ are continuously differentiable with derivatives vanishing at $(0, 0)$.*

Then any z in the graph of w admits a unique infinite backward orbit (z_{-n}) in the graph of w. This backward orbit converges to $(0, 0)$ fast enough so that

$$\limsup_n \frac{1}{n} log\|z_n\| \le -\log(\rho_+)$$

Conversely, a point z in the unit ball but not in the graph of w does not admit an infinite backward orbit in the unit ball such that

$$\limsup_n \frac{1}{n} \log \|z_n\| \le -\log(r_-).$$

We now drop the assumption that $r_+ > 1$ and assume instead that

$$r_- < 1.$$

Then if z is a point of the unit ball such that the $f^j(z)$ up to $j = n$ are all in the unit ball, the distance from $f^n(z)$ to the graph of w is not greater than $2r_-^n$.

From the last assertion it follows in particular that any orbit which stays in the unit ball for all n converges to the graph of w, and any orbit which starts very near to fixed point—and so stays in the unit ball for a long time—but which nevertheless eventually escapes from the unit ball must escape "along" the graph of w.

Proof. Under our hypotheses, f is injective on the graph W of w and $fW \supset W$. Hence, there is a uniquely determined sequence of successive preimages z_{-n} in W, and these are easily seen to converge to the origin at least as fast as r_+^{-n}. Improving this rate of convergence to the one indicated is done by the same arguments as those used to prove proposition 7.1.2.

Suppose that z is not in the graph of w but that it admits successive preimages z_{-1}, z_{-2}, ...,z_{-n}, all in the unit ball. Write $z = (x, y)$; put $z' = (w(y), y)$; and write z'_{-j} for the successive preimages of z' in the graph. By construction, the separation between z and z' is predominantly contractive, so – applying repeatedly the last assertion of lemma 7.2.1 – the separations between all the z_j and z'_j are predominantly contractive and

$$\|z - z'\| \leq (r_-)^n \|z_n - z'_n\|.$$

Since the left-hand side of this inequality is non-zero and independent of n, $\|z_n - z'_n\|$ cannot go to zero faster than $(r_-)^{-n}$. But we already know that $\|z'_n\|$ goes to zero at least as fast as $(r_+)^{-n}$, so $\|z_n\|$ cannot go to zero faster than $(r_-)^{-n}$.

We now turn to the last assertion. If $f^n(z)$ is in the graph of w, there is nothing to prove. Otherwise, the argument of the preceding paragraph with z of that paragraph replaced by $f^n(z)$ shows that there is a z' on the graph of w whose separation from z is predominantly contractive and such that the distance from $f^n(z)$ to the graph of w is not greater than $(r_-)^n\|z - z'\|$. Since $z' = (x', y')$ is on the graph of w, $\|x'\| \leq 1$, and since z is in the unit ball and has predominantly contractive separation from z', it follows that

$$\|z - z'\| = \|x - x'\| \leq 2$$

Hence the distance from $f^n(z)$ to the graph of w is not greater than $2(r_-)^n$, as asserted. □

8.5 Reformulation of results

Theorem 8.5.1 (Generalized unstable manifold theorem) *The general hypotheses are as for theorem 7.3.1. If $\rho_+ \leq 1$, we assume further that*

i. $r < \infty$

ii. $\rho_- < (\rho_+)^r$

iii. There is a real-valued $\mathcal{C}^r$ function ψ on $\mathcal{E}_+$ identically equal to one on a neighborhood of 0 and vanishing outside the unit ball.

Then:

1. There exists a $\mathcal{C}^r$ manifold $W^{(+)}_{loc}$ tangent at z_0 to $\mathcal{E}_+$ and locally invariant there for f.

2. If, further, $\rho_+ > 1$, then:

- $W^{(+)}_{loc}$ *can be taken to satisfy* $fW^{(+)}_{loc} \supset W^{(+)}_{loc}$
- $W^{(+)}_{loc}$ *is locally unique in the sense that any two locally invariant $\mathcal{C}^1$ manifolds tangent to $\mathcal{E}_+$ at z_0 coincide near z_0.*
- *if $\mathcal{W}$ is any sufficiently small open neighborhood of z_0, then $W^{(+)}_{loc}$ coincides near z_0 with the set of all $z \in \mathcal{W}$ which admit backward orbits (z_{-n}) defined for all $n = 1, 2, \ldots$, remaining in $\mathcal{W}$, and converging to z_0 rapidly enough so that*

$$\limsup \frac{1}{n} \log(\|z_{-n} - z_0\|) \leq -\log(\rho_+).$$

3. If $\rho_- < 1$ (whether or not $\rho_+ > 1$), and if $W^{(+)}_{loc}$ is any $\mathcal{C}^1$ locally invariant manifold tangent to $\mathcal{E}_+$ at z_0, there is an open neighborhood $\mathcal{W}$ of z_0 such that if $z_n := f^n(z)$ stays in $\mathcal{W}$ for all $n = 0, 1, \ldots$ then $z_n \to W^{(+)}_{loc}$.

4. If $\rho_- < 1 < \rho_+$ (hyperbolic case), then for any sufficiently small neighborhood $\mathcal{W}$ of z_0, $W^{(+)}_{loc}$ coincides near z_0 with the set of z which admit arbitrarily long backward orbits in $\mathcal{W}$.

Chapter 9

Invariant manifolds: Miscellany

9.1 Real-analytic mappings

If f is real-analytic, *and if* $\rho_- < 1$, it is easy to show that the invariant manifold tangent to $\mathcal{E}_-$ is real-analytic: Instead of giving the fundamental proof of existence in a space of Lipschitz continuous functions on a real neighborhood of zero, one can work in a space of complex-analytic functions on a complex neighborhood of zero. The proof of contractivity is just the same, and gives immediately an analytic w, thus bypassing completely the complicated analysis of smoothness given above. Similarly for the invariant manifold tangent to $\mathcal{E}_+$ when $\rho_+ > 1$. In the cases $\rho_- \geq 1$ and $\rho_+ \leq 1$, however, it is necessary to cut off the non-linear terms before applying the contraction argument, and the cut-off necessarily spoils analyticity. In fact, we saw in §6.4 an example of an analytic mappings with no analytic center manifold.

9.2 Flows

Consider a differential equation

$$\frac{dz}{dt} = X(z)$$

with solution flow f^t, and let 0 be a stationary solution of this differential equation. Let σ be a real number such that the the spectrum of $DX(0)$ does not intersect the vertical line Re $\lambda = \sigma$ but does intersect both half-planes bounded by this line. We let

$$\sigma_- = \sup\{\text{Re }\lambda : \lambda \in \sigma(DX(0)) \quad \text{and} \quad \text{Re }\lambda < \sigma\}$$

and

$$\sigma_+ = \inf\{\text{Re }\lambda : \lambda \in \sigma(DX(0)) \quad \text{and} \quad \text{Re }\lambda > \sigma\}$$

There is then a splitting of the state space into the direct sum of the spectral subspaces $\mathcal{E}_-$ associated with the part of the spectrum of $DX(0)$ in the left half-plane $\{\text{Re }\lambda \leq \sigma_-\}$ and $\mathcal{E}_+$ associated with the part of the spectrum in the right half-plane $\{\text{Re }\lambda \geq \sigma_+\}$. As we saw earlier,

$$Df^t(0) = \exp(tDX(0)),$$

so, for each $t > 0$, $\mathcal{E}_- \oplus \mathcal{E}_+$ is the direct sum decomposition corresponding to an annulus

$$\exp(t\sigma_-) < |\lambda| < \exp(t\sigma_+)$$

disjoint from the spectrum of $Df^t(0)$, and there are therefore invariant manifolds for f^t tangent to $\mathcal{E}_-$ and $\mathcal{E}_+$. These invariant manifolds might, in principle, depend on t. What we want to show here is that there is a single manifold tangent to $\mathcal{E}_-$ locally invariant for *all the f^t simultaneously*, and similarly for $\mathcal{E}_+$. We will restrict our discussion to invariant manifolds tangent to $\mathcal{E}_-$, i.e., generalized stable manifolds.

In the case $\sigma_- < 0$ (stable and strong stable manifolds), this is quite easy: Let W be a manifold locally invariant for f^1 tangent to $\mathcal{E}_-$. From the stable manifold theorem for maps, we know that such a manifold exists and that it is locally unique. *Since the f^t commute,* f^tW is another such locally invariant manifold, and so must coincide near 0 with W. Thus, for any t, W is locally invariant for f^t at 0.

The case $\sigma_- \geq 0$ requires a little more care. In the analogous case for maps f we constructed a locally invariant manifold by constructing a globally invariant manifold for a cut-off version of f, and this globally invariant manifold is unique. In the case at hand, if we cut off the non-linear terms in the solution mappings f^t, it is hard to see how we will preserve the commutativity needed to apply the argument of the preceding paragraph. What we do instead is to cut off the non-linear terms in the differential equation. That is, we write

$$X(x, y) = (A_-x + X_-(x, y), A_+y + X_+(x, y)),$$

where the $X_\pm$ vanish, together with their first derivatives, at the origin. As in the case of mappings, we can make the $X_\pm$ small on the unit ball by magnification, and we can then multiply them by a smooth function which is identically one on a neighborhood of the origin and identically zero outside the unit ball to get a globally defined differential equation with globally small non-linear terms. For any t, the time-t solution mappings for the cut-off differential equation agrees with that for the original equation on a neighborhood of the origin (but this neighborhood may become very small as $|t|$ becomes large.)

We now consider only the cut-off equation, which we denote by the same symbols as before. It is not difficult to see that the solution mappings have globally small non-linear terms. Hence, the time-one mapping, for example, admits a unique global invariant manifold W which is the graph of a mapping with Lipschitz norm not greater than one. We would like to argue that W is mapped into itself by all the f^t. By the group property of the solution flow, it is enough to prove this for small t. By commutativity, $f^t W$ is mapped into itself by f^1. Hence, if we can show that $f^t W$ is the graph of a mapping with Lipschitz norm not greater than one, we are done, by the uniqueness of such a globally invariant graph. One way to proceed here is as follows: By imposing further smallness condition on the non-linear terms in f^1, we can ensure that W is in fact the graph of a mapping with Lipschitz norm not greater than $1/2$, say. Uniqueness still holds, of course, in the class of maps with Lipschitz norm not greater than one. By making t small enough, we can make f^t globally as near as we like to the identity in the $\mathcal{C}^1$ sense. By making it near enough, we can guarantee that that image under f^t of the graph of any mapping with Lipschitz norm not greater than $1/2$ is again the graph of a mapping, this time with Lipschitz norm not greater than one. Thus, the desired invariance follows by uniqueness.

There are a number of respect in which the arguments we have given are less than optimal:

1. In the case $\sigma_- < 0$, we would like to show how to construct an invariant manifold which is mapped into itself by all f^t with $t > 0$. Such a manifold would be realized as the graph of a mapping of the unit ball in $\mathcal{E}_-$ into $\mathcal{E}_+$, and to make it invariant for all positive t it is necessary to choose the norm on $\mathcal{E}_-$ with more care than we have used. The remark we need is that, for any $s > \sigma_-$, it is possible to choose a norm on $\mathcal{E}_-$ which is equivalent to the original norm and for which

$$\|\exp(tA_-)\| \leq \exp(ts) \quad \text{for all } t \geq 0.$$

If we use such a norm, with $s < 0$, it is not difficult to see that, by making the non-linear terms in X sufficiently small on the unit ball, we can arrange that the invariant manifold for f^1 defined as the graph of a mapping defined on the unit ball in $\mathcal{E}_-$ is indeed mapped into itself by all f^t with $t > 0$.

2. The argument we have given for the case $\sigma_- \geq 0$, involving cutting off the non-linear terms in the differential equation, does not seem to be very well adapted to the case of partial differential evolution equations. In this latter case, the f^t for small t need not be $\mathcal{C}^1$-near the identity. It is more likely, however, that the cutting off can be arranged to make them $\mathcal{C}^1$-near to $\exp(tDX(0))$. If this can be done it is not difficult to see that the *preimage* under f^t, t small and positive, of the graph of a mapping with Lipschitz norm not greater than $1/2$ is the graph of a mapping with Lipschitz norm not greater than one, so uniqueness will imply that $(f^t)^{-1} W = W$ for small positive t and hence for all positive t. A different approach which may work better in this case is suggested on p. 48 of [Marsden-McCracken].

9.3 Center manifolds.

Suppose we have a fixed point z_0 for a smooth mapping f such that spectrum of $Df(z_0)$ does intersect the unit circle non-trivially but also such that the part of the spectrum on the unit circle is separated from the rest. There is then a splitting of the state space into a direct sum of three subspaces:

$$\mathcal{E} = \mathcal{E}_s \oplus \mathcal{E}_c \oplus \mathcal{E}_u,$$

where $\mathcal{E}_c$, the *center subspace*, is the spectral subspace associated with the part of the spectrum on the unit circle and $\mathcal{E}_s$, $\mathcal{E}_u$ are, as before, the spectral subspaces associated with the parts of the spectrum inside and outside the unit circle respectively. We want to argue that there exists a locally invariant manifold passing through z_0 and tangent there to $\mathcal{E}_c$. Such an invariant manifold is called a *center manifold*. We do this in two steps. We first apply the generalized unstable manifold theorem to see that there is a locally invariant manifold W_{cu} (a *center-unstable* manifold) tangent to $\mathcal{E}_c \oplus \mathcal{E}_u$. We know that such an invariant manifold exists; we don't know that it is unique, but we know that it can be chosen to be of class $\mathcal{C}^r$, r finite, if f is of class $\mathcal{C}^r$. The idea now is to apply the generalized *stable* manifold theorem to the restriction of f to the center-unstable manifold. The spectrum of the derivative of the restriction of f is just the spectrum of the restriction of $Df(z_0)$ to $\mathcal{E}_c \oplus \mathcal{E}_u$, i.e., the part of the spectrum of $Df(z_0)$ on and outside the unit circle. The generalized stable manifold theorem applies to show that there is a locally invariant manifold for f, contained in the center-unstable manifold and tangent to $\mathcal{E}_c$ at z_0. Again, this center manifold can be taken to be of class $\mathcal{C}^r$ for any finite r if f is. Evidently, a similar construction can be done for a separated part of the spectrum contained in an annulus centered at the origin, except that smoothness need not be preserved in this more general situation.

The center manifold, although useful for many things, turns out to be a somewhat unnatural object which only barely manages to exist and which is subject to a number of pathologies. We already saw in §6.4 examples which show

- the center manifold may be badly non-unique
- an analytic mapping, although it has a $\mathcal{C}^r$ center manifold for every finite r, need not have *any* $\mathcal{C}^\infty$ center manifold.

9.4 Derivatives of invariant manifolds.

We have shown how to reduce the problem of finding generalized stable and unstable manifolds to solving functional equations, and have proved existence theorems for those functional equations. We have not, however, been able to give "formulas" for the invariant manifolds. We want to note here that at least the derivatives at the fixed point of the function whose graph is the invariant manifold can be computed algebraically. Consider for definiteness a generalized stable manifold which is the graph of a function w satisfying

$$w(x) = \Lambda_+^{-1} \left\{ w(\Lambda_- x + f_-(x, w(x))) - f_+(x, w(x)) \right\}.$$

We have already seen what the equation given by differentiating this one once looks like. Differentiating a second time gives a discouragingly complicated equation which reduces to

$$D^2 w(0) = \Lambda_+^{-1} \left\{ D^2 w(0) \Lambda_-^2 - D_1^2 f_+(0,0) \right\}$$

at $x = 0$. This is a simple inhomogeneous linear equation for $D^2 w(0)$ which can be solved provided

$$\|\Lambda_+^{-1}\| \, \|\Lambda_-\|^2 < 1.$$

(It may also be solvable otherwise, but more delicate considerations are involved.)

It is not difficult to see that this procedure can be repeated: If f is of class $\mathcal{C}^n$ and if

$$\|\Lambda_+^{-1}\| \, \|\Lambda_-\|^n < 1,$$

(which guarantees that w is of class $\mathcal{C}^n$), then

$$D^n w(0) = \Lambda_+^{-1} D^n w(0) \Lambda_-^n + \Psi_n,$$

where Ψ_n is an expression involving derivatives of order up to n of the $f_\pm$ and derivatives up to order $n-1$ of w, all at the origin. Thus, the $D^n w(0)$ can in principle be computed explicitly.

One consequence of the above analysis, which we will have occasion to use later, is that, if the $f_\pm$ vanish to order n at zero, so does w.

9.5 Global stable and unstable manifolds.

If f is invertible, it is possible to define *global* stable and unstable manifolds. We need only consider the stable manifold, since the unstable manifold is simply the stable manifold for f^{-1}. If z_0 is a hyperbolic fixed point for f, the *(global) stable manifold* for f means the set W^s of all z such that

$$f^n(z) \to z_0 \quad \text{as } n \to \infty.$$

This definition of course makes sense whether or not f is invertible; the reason for restricting to invertible f is to ensure that this set has a nice structure. To see this, let W^s_{loc} be a "nice" local stable manifold, i.e., the graph of a smooth mapping from a small ball in $\mathcal{E}_s$ to $\mathcal{E}_u$, where the norm on $\mathcal{E}_s$ is chosen in such a way as to make W^s_{loc} be mapped into itself by f. Then W^s_{loc} is contained in W^s, and, furthermore, any orbit which stays in a small enough neighborhood of z_0 for all time must be in W^s_{loc}. Thus, for any $z \in W^s$, $f^n(z) \in W^s_{loc}$ for all sufficiently large n. Thus, W^s is a increasing union of a sequence of pieces $f^{-n}W^s_{loc}$, each diffeomorphic to a ball in $\mathcal{E}_s$. The global manifold W^s can nevertheless have a complicated structure; loosely speaking, it can double back on itself (but cannot cross itself). Technically, it is a *connected injectively immersed submanifold* of the state space, locally diffeomorphic to $\mathcal{E}_s$.

In fact, W^s is—under mild technical assumptions—the image under an injective $\mathcal{C}^r$ immersion of $\mathcal{E}^s$ itself. To see this: Let ϕ_0 be the mapping $x \mapsto (x, w(x))$ imbedding an ϵ-ball in $\mathcal{E}^s$ onto a small local stable manifold W^s_{loc}. We assume that the norm is chosen so that W^s_{loc} is mapped into itself by f. Pulling back the restriction of f under ϕ_0 gives a $\mathcal{C}^r$ mapping of the ϵ-ball in $\mathcal{E}_s$ into itself. The technical assumption referred to above is that there exists a globally defined and globally contractive $\mathcal{C}^r$ diffeomorphism $\tilde{f}_s$ of $\mathcal{E}^s$ which agrees with the pullback in some neighborhood of 0. If there is a $\mathcal{C}^r$ cut-off function on $\mathcal{E}^s$, we can construct such an $\tilde{f}_s$ by multiplying the non-linear part of the induced mapping by an appropriate cut-off. It is then easy to see that

$$f^{-n}\phi_0(\tilde{f}_s)^n(x)$$

is, for any given x, defined for all sufficiently large n and eventually independent of n. The point is that

$$f^{-(n+1)}\phi_0(\tilde{f}_s)^{n+1}(x) = f^{-n}\left(f^{-1}\phi_0\tilde{f}_s\right)(\tilde{f}_s)^n(x),$$

and, by the assumption that $\tilde{f}_s$ agrees with the pullback of f on a neighborhood of 0,

$$\phi_0 \tilde{f}_s = f \phi_0 \quad \text{near } 0.$$

Finally, because $\tilde{f}_s$ is by assumption globally contractive, $(\tilde{f}_s)^n(x)$ is near zero for every sufficiently large n. It then follows easily that

$$\phi(x) := \lim_{n \to \infty} f^{-n} \phi_0 (\tilde{f}_s)^n(x)$$

is an injective $\mathcal{C}^r$ immersion of $\mathcal{E}_s$ whose image is exactly W^s.

Strong stable and unstable manifolds can similarly be globalized. If, for example, z_0 is a fixed point of the differentiably invertible mapping f and if the spectrum of $Df(z_0)$ does not intersect the circle $\{|\lambda| = r\}$ with $r < 1$, then the set of z such that

$$\limsup_n \frac{1}{n} \log(\|f^n(z) - z_0\|) < \log(r)$$

is again an injectively immersed submanifold of the state space. Here, the state space can be allowed to be an arbitrary manifold with $\|f^n(z) - z_0\|$ understood to mean the distance between $f^n(z)$ and z_0 with respect to any Riemannian metric.

We mention, without going fully into the details, that there are reasonable definitions of global stable and unstable manifolds even without the assumption that f is invertible. Let us use the term *stable set* of the fixed point to refer to $\{z : f^n z \to z_0\}$ in the general case. The idea is roughly that the stable manifold is the maximal invariant submanifold of the stable set. The precise definition is that the global stable manifold is the set of all z such that, for some n,

- $f^n(z) \in W^s_{loc}$
- The image of $Df^n(z)$ and the tangent space to W^s_{loc} at $f^n(z)$ together span $\mathcal{E}$ (i.e., f^n is *transverse to* W^s_{loc} at z.)

It is not difficult to show[1] that this set is a submanifold in the same sense as is the global stable manifold for an invertible mapping, and it is manifestly mapped into itself by f and contained in the stable set of z_0.

[1]At least in the finite-dimensional case. The obvious proof requires that $\ker(Df^n(z))$ be a direct summand of $\mathcal{E}$, i.e., that there exist a closed subspace $\mathcal{F}$ of $\mathcal{E}$ such that $\mathcal{E} = \ker(Df^n(z)) \oplus \mathcal{F}$. I have no idea whether this additional assumption is really needed or not. In any event, it could simply be added to the conditions z must satisfy.

9.6 Example. Hyperbolic linear torus automorphisms.

Let $\mathbb{T}^2$ denote the quotient of $\mathbb{R}^2$ by the additive subgroup $\mathbb{Z}^2$ of points with integer coordinates, i.e., the standard realization of the two-dimensional torus. Let F denote the 2×2 matrix

$$F := \begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix}$$

and also the corresponding linear mapping of $\mathbb{R}^2$ to itself:

$$F : (x, y) \longmapsto (2x + y, x + y).$$

Since F is linear and maps $\mathbb{Z}^2$ to itself, it induces a mapping which we will call f from $\mathbb{T}^2$ to itself. Because the matrix F has unit determinant, its inverse also has integer entries, maps $\mathbb{Z}^2$ to itself, and hence induces a map from $\mathbb{T}^2$ to itself. It is easy to see that the mapping induced by F^{-1} inverts f, and hence that f is a diffeomorphism of $\mathbb{T}^2$

The image of $(0, 0)$ under the quotient mapping of $\mathbb{R}^2$ to $\mathbb{T}^2$ is evidently a fixed point for f, and is in fact easily seen to be its only fixed point. In the obvious coordinates, the derivative of f at this fixed point (and everywhere else) is the matrix F. The eigenvalues of F are easily computed to be

$$\lambda_\pm = \frac{3 \pm \sqrt{5}}{2}.$$

Since neither eigenvalue is on the unit circle, F is hyperbolic, i.e., the fixed point of f is hyperbolic. If we let $\phi_\pm$ denote (real) eigenvectors with eigenvalues $\lambda_\pm$, then it follows from the "linearity" of f that the global stable manifold of the fixed point is the image of the line

$$\{t\phi_- : -\infty < t < \infty\}$$

under the quotient mapping of $\mathbb{R}^2$ onto $\mathbb{Z}^2$. This line has irrational slope, so its image wraps densely around the torus without ever closing or crossing itself. The global unstable manifold is similarly the image under the quotient mapping of the line generated by ϕ_+ which again wraps densely around the torus.

This example gives some feeling for how complicated the global structure of stable and unstable manifolds can be even in the simplest cases. It is worth noting, however, that the situation is even more complicated than the preceding analysis shows: f has a great many periodic points, and the stable and unstable manifolds for each periodic point have the same complicated global structure as those for the fixed point. It comes as something of a surprise that one can easily identify all the periodic points of f (although the determination of the period of a given periodic point is not so obvious.) We will say that a point of $\mathbb{T}^2$ is *rational* if it has a representation as $\pi(x,y)$ with $x,y \in \mathbb{Q}$. Then

A point of $\mathbb{T}^2$ *is periodic for* f *if and only if it is rational.*

In one direction: For any positive integer N, the set X_N of points with representatives of the form $(j/N, k/N)$ with $j,k \in \mathbb{Z}$ is finite—it has exactly N^2 members—and is mapped into itself by f. Hence every point of this set must be periodic. Since every rational point is in a set of this form for some N, every rational point is periodic. In the other direction: If the point of $\mathbb{T}^2$ with representative $(x,y) \in \mathbb{R}^2$ is periodic with period p, then $F^p(x,y) = (x,y) + (j,k)$ for some $j,k \in \mathbb{Z}$. We write this equation as

$$(F^p - 1)(x,y) = (j,k).$$

The matrix $F^p - 1$ is invertible and has integer elements; hence, the matrix elements of its inverse are rational, so

$$(x,y) = (F^p - 1)^{-1}(j,k)$$

is rational.

If $z_0 = \pi(\tilde{z}_0)$ is a periodic point of period p for f, then the derivative of f^p at z_0, again in the obvious co-ordinates, is just F^p. The eigenvectors of F^p are the same as those of F, and from this it easily follows that the unstable and stable manifolds for f^p at z_0 are the images under the quotient mapping of the lines

$$\{\tilde{z}_0 + t\phi_\pm : -\infty < t < \infty\},$$

i.e., the stable manifolds (and the unstable manifolds) of periodic points are all "parallel" to each other and each of them is a line wrapping densely around the torus. This example can be generalized. Everything remains essentially unchanged if we replace $\mathbb{T}^2$ by $\mathbb{T}^m$ and F by any $m \times m$ matrix with integer entries and determinant ± 1 all of whose eigenvalues have modulus different from 1.

9.7 Example. The Hénon mapping.

The Hénon mapping is a mapping from the plane $\mathbb{R}^2$ to itself given by the formula

$$f(x,y) = (1 - ax^2 + by, x), \tag{9.1}$$

with a and b constants (parameters)[2]. We are going to be interested primarily in small b and, roughly, a between 1 and 2.

It is illuminating to split the Hénon mapping into a sequence of simple steps:

$$(x,y) \longmapsto (x,by) \longmapsto (x, 1-ax^2+by) \longmapsto (1-ax^2+by, x).$$

The first step is a vertical compression. The second is a *vertical shear:* Each vertical line $x = x_0$ is shifted up or down along itself by an amount $1 - ax_0^2$ depending on x_0. The last step is simply a reflection through the line $y = x$. The first step multiplies areas by $|b|$, while the second and third steps preserve areas, so the composite mapping (9.1) multiplies areas by a factor of $|b|$. This can also be seen algebraically by checking that the Jacobian of (9.1) is constant and equal to $-b$. Thus, small b corresponds to strong contraction of areas. (At the other extreme, for $b = \pm 1$ the mapping is area preserving. This is a very interesting case, but not one we will discuss here.) Note also that f is *orientation preserving* for b negative, and orientation reversing for b positive. This gives the cases of positive and negative b slightly different flavors. For example: The eigenvalues of Df^p at a periodic point of period p are either real or complex conjugates of each other. Since their product is $(-b)^p$, the second alternative is not possible for $b > 0$ and p odd. Thus: In the vicinity of a periodic point of odd period p, f^p can look like a rotation followed by a constant contraction (by $|b|^{p/2}$) if b is negative but not if b is positive. We are going to concentrate, for definiteness, on the case of positive b.

Somewhat surprisingly, in view of the x^2 term it contains, f is *invertible* as a mapping of the whole plane to itself (provided that $b \neq 0$.) This can be seen either by noting that each of the three elementary steps into which f is decomposed is separately invertible, or by computing its inverse in a straightforward way:

$$f^{-1}(x,y) = \left(y, \frac{1}{b}[x - (1 - ay^2)] \right).$$

Numerical studies first performed by Hénon, and since repeated and extended by many others, strongly suggest that the above mapping f, with $a = 1.4$ and $b = .3$, admits a "fractal strange attractor", i.e., that almost all orbits starting in a reasonably large region of the plane converge to a set with a self-similar structure which is Cantor-set like in some directions. This is best seen by looking a pictures showing the results of numerical experiments.

[2]The mapping is often defined instead as $(1 - ax^2 + y, bx)$, and it was in this form that Hénon originally wrote it. The two forms differ only by a rescaling of y by a factor of b, and the one we have chosen is better adapted to looking at what happens for small b.

This situation, simple as it seems, resisted mathematical investigation for a long time. In very important—and strikingly difficult—recent work, Benedicks and Carleson have shown that, for sufficiently small non-zero b there is a set of a's of positive linear Lebesgue measure for which the Hénon mapping really does have such a strange attractor.

Note added, September 1997: *The above paragraph is by now seriously out of date. The work of Benedicks and Carleson appeared as* The dynamics of the Hénon map, *Ann. of Math.* **133** *(1991) 73–576 and served as the starting point for a great deal of further work. For an accessible survey of this area, see Lai-Sang Young,* Ergodic theory of attractors *in Proceedings of the 1994 International Congress of Mathematicians, Birkhäuser, (1995) 1230–1237.*

We will not say any more about the work of Benedicks and Carleson here, but instead describe some relatively simple theory which has been developed for the Hénon mapping. Suppose we fix a at some value between 1 and 2. If we put $b = 0$, the mapping is of course no longer invertible; the whole plane is squeezed down onto the parabola $x = 1 - ay^2$. The square

$$\Delta_0 = \{(x, y) : -1 \le x \le 1, -1 \le y \le 1\}$$

is mapped onto the arc

$$\left\{(1 - ay^2, y) : -1 \le y \le 1\right\}.$$

This arc is contained in Δ_0 but touches its boundary at three points. If we open up Δ_0 slightly, say to the rectangle

$$\Delta_\epsilon = \{(x, y) : -(1 + \epsilon) \le x \le 1 + \epsilon, -(1 + 2\epsilon) \le y \le 1 + 2\epsilon\},$$

with ϵ small, then Δ_ϵ is mapped into its own interior by f (with $b = 0$.) By continuity, then, Δ_ϵ will be mapped into itself by all f with sufficiently small b.

Suppose we pick and fix such a b which is not exactly zero. From the fact that f maps Δ_ϵ into itself it follows that

$$\Delta_\epsilon \supset f\Delta_\epsilon \supset f^2\Delta_\epsilon \supset \ldots \supset f^n\Delta_\epsilon \ldots$$

On the other hand, the area (Lebesgue measure) of $f^n\Delta_\epsilon$ is smaller than that of Δ_ϵ by a factor of $|b|^n$. Hence

$$X = \cap_{n=0}^{\infty} f^n\Delta_\epsilon$$

is a compact set of zero Lebesgue measure, mapped onto itself by f, to which all orbits starting in Δ_ϵ converge. Although the above simple construction of an invariant rectangle doesn't work for Hénon's parameter values $a = 1.4$, $b = .3$, Hénon was nevertheless able, by looking at the mapping in slightly more detail, to construct an invariant quadrilateral in his case as well, so there is also a set X in that case.

It might be thought that the set X, since at attracts all orbits which start in the much larger set Δ_ϵ, would qualify to be called an *attractor.* On closer examination, this does not seem to be a good idea. To see why, suppose for example that f has an attracting periodic orbit in Δ_ϵ. (This certainly happens for some parameter pairs to which our analysis applies.) This orbit—indeed, any periodic orbit, attracting or not—must be in X. Around each point of the periodic cycle there will be a neighborhood of points whose forward orbits are asymptotic to the cycle and which are therefore, in a natural sense, *non-recurrent.* It is easy to see that X, as the decreasing intersection of a sequence of sets each homeomorphic to a rectangle, is connected, so some of these transient points have to belong to X. It seems natural to make attractors as small as possible by leaving out all transient points, and X does not meet this requirement. Accordingly, it is customary to distinguish between *attractors* and *attracting sets.* An *attracting set* simply means a compact set X mapped onto itself by f which has an open neighborhood $\mathcal{U} \supset X$ whose forward images $f^n\mathcal{U}$ shrink down to X in the sense that, for any open set $\mathcal{V} \supset X$, $f^n\mathcal{U} \subset \mathcal{V}$ for all sufficiently large n. What we earlier defined as an attracting fixed point is simply an attracting set with only one element. Just as in the fixed point case, one can add the requirement that $f\mathcal{U} \subset \mathcal{U}$ without making the definition more restrictive. An *attractor* then means an attracting set with some further properties. It is still controversial exactly what these properties should be, but two things one certainly wants are that every point of an attractor should be (in some sense) recurrent and that the motion on the attractor should be in some sense indecomposable. The set X constructed above is certainly an attracting set, but need not be an attractor.

It is immediate from the formula (9.1) for f that any fixed point must lie on the line $y = x$. The point (x_0, x_0) is a fixed point if and only if

$$1 - ax_0^2 + bx_0 = x_0.$$

For $a > 0$, whatever the value of b, this equation has one positive solution and one negative solution, so f has exactly two fixed points. As b goes to zero, these fixed points converge to the two points of intersection of the parabola $x = 1 - ay^2$ with the line $y = x$. We are going to study the solution with $x_0 > 0$, for small b. It is easy to check from the equation that this fixed point lies above and to the right of the parabola $x = 1 - ay^2$ for small positive b. Since the fixed point and the derivative of f at it can be computed explicitly, it is straightforward to determine the spectrum of the derivative of f at the fixed point. In this way it can be shown that the fixed point is hyperbolic, with one-dimensional unstable manifold, if and only if $a > 3(1-b)^2/4$, and that the expanding eigenvalue is negative. For $b = 0$, the contracting eigenvalue is zero, and it is easy to guess what the local stable and unstable manifolds are: The local stable manifold is a vertical line segment through the fixed point, and the local stable manifold is a small segment of the parabola $x = 1 - ay^2$.

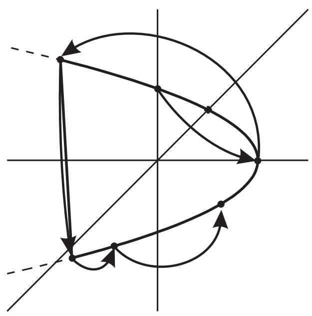

Figure 9.1: . Mapping of the parabola $x = 1 - ay^2$ by the Hénon mapping with $b = 0$.

The Hénon mapping with $b = 0$ provides a tame example of the problems with global stable and unstable manifolds for non-invertible mappings. It is nearly obvious that the parabolic arc

$$\left\{(1 - ay^2, y) : 1 - a \leq y \leq 1\right\}$$

is mapped into itself by f and it can be proved that—provided a is near enough to 2—the successive images of an arbitrarily small segment of this arc containing the fixed point expand to fill up the whole arc. Thus, if W^u_{loc} is a local unstable manifold, i.e., a small segment of the parabola, then

$$\bigcup_{n=0}^{\infty} f^n W^u_{loc} = \left\{(1 - ay^2, y) : 1 - a \leq y \leq 1\right\},$$

and the segment on the right-hand side of this equation is not a submanifold because it contains its end points.

As soon as $b \neq 0$, the pathologies of the preceding paragraph have to disappear. What happens for a slightly less than 2 and b very small is that the global unstable manifold winds back and forth infinitely often, always staying close to the parabolic arc, never crossing itself, with sharper and sharper bends.

Although the detailed behavior of the unstable manifold for small non-zero b is complicated, it is worth taking at least a few preliminary steps in the direction of analyzing it. To get first-order estimates of the effect of a small but non-zero b, it is helpful to note that

$$f(x,y) = f_0(x,y) + b(y,0) \quad \text{where} \quad f_0(x,y) = (1-ax^2, x).$$

Thus, the effect of a small positive b is to displace the image of (x,y) to the right if $y > 0$ and to the left if $y < 0$. We will, for definiteness, consider in what follows only $b \geq 0$.

Since the unstable manifold is one-dimensional, removing the fixed point splits it into two "branches". The expanding eigenvalue at the fixed point is *negative*, so f interchanges these two branches. The first part of each branch will, for small b, lie along the corresponding part of the parabolic arc which is the local unstable manifold at $b = 0$. At $b = 0$ there is a piece of unstable manifold running from the fixed point to the point $(1,0)$ where the parabola crosses the x axis. Correspondingly, for small $b > 0$, there is a piece of stable manifoldn $w_0^{(b)}$ running from the fixed point to the x axis. By the remark of the preceding paragraph, it lies to the right of the parabola. At $b = 0$, the image of $(1,0)$ is the upper left-hand end of the piece of parabola which lies in the unstable set of the fixed point. Similarly, the image $w_1^{(b)}$ of $w_0^{(b)}$ runs from the fixed point to somewhere near this upper end, and the second image $w_2^{(b)}$—which contains $w_0^{(b)}$—runs from the fixed point to somewhere near the lower left-hand end of the parabolic arc. Now $w_2^{(b)}$ runs past the "tip" of the parabola. Applying f compresses the gentle bend in $w_2^{(b)}$ to a much sharper bend, which is mapped to the upper left-hand end. The image of the part of $w_2^{(b)}$ past the tip is shifted to the left of the image of $w_0^{(b)}$. The outer end of $w_2^{(b)}$ lies in $x < 0$ (for b small) and so maps to a point in $y < 0$. Since the image point must also lie near the parabolic arc, the image of the outer part of $w_2^{(b)}$ must run along the parabolic arc from its upper end down past the x axis. Thus: $w_3^{(b)} := fw_2^{(b)}$ starts at the fixed point, runs up to the upper-left end, turns sharply, then runs back along the parabola past the x axis. Applying f again produces a curve $w_4^{(b)}$ with a very sharp turn at the lower left end which then follows back along the parabola to another sharp turn at the upper left end. Beyond this point, the picture gets too complicated to follow in detail, but it is clear that the unstable manifold

– has a sequence of sharp turns along—roughly—the orbit of the place where $w_0^{(b)}$ crosses the x axis.

– crosses the x axis many more times, and has a corresponding sequence of sharp turns along the forward orbit of each of these crossings.

That this is not the whole story is indicated by the fact that we must also expect some of these sharp turns to lie back on the x axis; the next application of f must then be expected to make the turn *less* sharp. It is perhaps also necessary to explain that, for small finite b, it is not *exactly* the x axis where the bending takes place ...

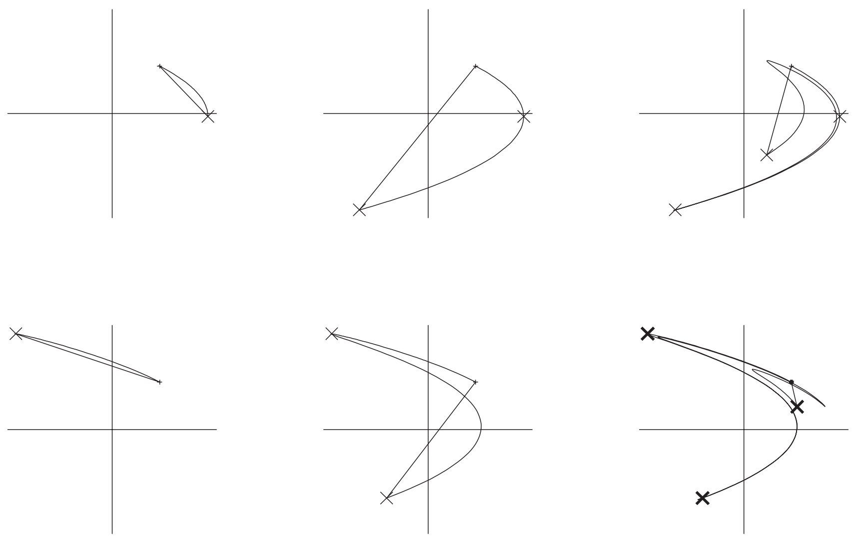

Figure 9.2: . Successive pieces of the unstable manifold of the Hénon mapping.

We will next give an argument indicating that the Hénon attractor, if it exists, cannot be too different from the closure of the global unstable manifold of the fixed point with $x_0 > 0$. For small b, the unstable manifold has a "first branch" lying roughly along

$$\left\{(1 - ay^2, y) : 1 - a \leq y \leq 1\right\},$$

and there is a large piece of the stable manifold lying roughly along the vertical line through the fixed point. The stable and unstable manifolds therefore cross at a point near $(x_0, -x_0)$. (A point of intersection of stable and unstable manifolds for a single hyperbolic fixed point is called a *homoclinic point;* the fixed point itself, of course, doesn't count. If the crossing takes place at non-zero angle, as will be the case here, one speaks of a *transverse* homoclinic point. We will see later on that the existence of a transverse homoclinic point implies that the mapping has quite a lot of complicated structure, and in particular has infinitely many periodic orbits.) The above argument for the existence of a homoclinic point only works for b small, but for the rest of our argument all we need is the existence of a homoclinc point and $|b| < 1$.

Let Δ denote the D-shaped region bounded by the segments of the stable and unstable manifolds joining the fixed point to the homoclinic point. We want to argue that all orbits starting in Δ converge to the global unstable manifold. The proof is very simple: Consider $\Delta_n = f^n\Delta$ with n large. Like Δ, Δ_n is homeomorphic to a disk and its boundary is made up a a piece of stable manifold and a piece of unstable manifold. The piece of stable manifold is however the image under f^n of the piece of stable arc joining the fixed point to the homoclinic point and so is small; most of the boundary is unstable manifold. On the other hand, the area of Δ_n is $|b|^n$ times the area of Δ, and so is small. This means that, for any fixed $\epsilon > 0$, Δ_n for n sufficiently large does not contain any disk of radius ϵ, i.e., every point of Δ_n is within a distance ϵ of the boundary. If, furthermore, n is big enough so that the part of the boundary consisting of stable manifold has length no larger than ϵ, it follows that every point of Δ_n is at distance no greater than 2ϵ from the unstable manifold, and thus every orbit starting in Δ converges to the unstable manifold.

Finally, we present a clever and convenient procedure, due to Franceschini and Russo, for numerical computation of the stable and unstable manifolds of the fixed points of the Hénon mapping. Franceschini and Russo pose the problem as follows: They look for analytic functions $x(t)$, $y(t)$, and a number λ, such that

$$f(x(t), y(t)) = (x(\lambda t), y(\lambda t)) \tag{9.2}$$

That is: They look for a representation of the invariant manifolds as parametrized curves, with parametrization chosen so that the action of f on the curve itself, using the parameter as a coordinate, is simply multiplication by λ. It is not difficult to show from general theory that the unstable manifold does admit a parametrization of this form with x and y analytic. It is furthermore apparent that if such a representation exists, it is non-unique in a trivial sense: The parameter t can be multiplied by an arbitrary non-zero constant. We can thus expect to have to impose a normalization condition to remove this arbitrariness.

The first observation to be made about (9.2) is that any solution must have $(x(0), y(0))$ a fixed point for f, and we should expect to be able to choose here either of the fixed points. From the explicit formula for f, the equation (9.2) is equivalent to the pair of equations

$$x(t) = y(\lambda t) \quad \text{and} \quad x(\lambda t) = 1 - ax(t)^2 + by(t).$$

We can use the first equation to eliminate $y(t)$ from the second, so what we really have to solve is

$$x(\lambda t) = 1 - ax(t)^2 + bx(t/\lambda) \tag{$*$}$$

Now write the Taylor series expansion for $x(t)$ as

$$x(t) = \sum_{n=0}^{\infty} x_n t^n.$$

Inserting into $(*)$ gives

$$\sum_{n=0}^{\infty} \left(x_n \lambda^n + a \sum_{j=0}^{n} x_j x_{n-j} - b\lambda^{-n} \right) t^n = 1.$$

Identifying the constant terms gives

$$x_0 = 1 - ax_0^2 + bx_0,$$

which is the fixed point condition. Identifying the linear terms gives

$$\left(\lambda^2 + 2ax_0\lambda - b\right) x_1 = 0. \tag{$\dagger$}$$

If the factor $(\lambda^2 + 2ax_0\lambda - b)$ does not vanish, then x_1 must vanish, and it is not hard to see that all higher x_n's then also vanish. This does give a solution of $(*)$ but not a very interesting one. Thus, we require

$$\lambda^2 + 2ax_0\lambda - b = 0.$$

Examination of the formulas shows that the λ's satisfying this equation are exactly the eigenvalues of the derivative of f at the fixed point (x_0, x_0). We can choose either of the two eigenvalues; which one we choose determines whether the calculation produces the stable or the unstable manifold.

The equation $(\dagger)$ now imposes no constraint on x_1, so we can set $x_1 = 1$; this fixes the normalization of t. For $n \geq 2$ we get

$$x_n \lambda^n + 2ax_0 x_n + a \sum_{j=1}^{n-1} x_j x_{n-j} - b\lambda^{-n} x_n = 0,$$

or

$$\left(\lambda^n + 2ax_0 - b\lambda^{-n}\right) x_n = -a \sum_{j=1}^{n-1} x_j x_{n-j}.$$

From this last equation it is straightforward to compute as many x_n's as desired numerically (subject to the limitations imposed by round-off error.) It is furthermore not difficult to see, directly from $(*)$, that any solution analytic in a neighborhood of 0 must be entire, so the series converges for arbitrary t, and hence gives the full stable and unstable manifolds.