Bifurcations in second-order nonlinear systems

(from Strogatz, Nonlinear dynamics and chaos, Perseus Books, Cambridge, Ma, 1994, Ch. 8)

The bifurcations of fixed points discussed in Chapter 3 have analogs in two dimensions (and indeed, in *all* dimensions). Yet it turns out that nothing really new happens when more dimensions are added—all the action is confined to a one-dimensional subspace along which the bifurcations occur, while in the extra dimensions the flow is either simple attraction or repulsion from that subspace, as we'll see below.

Saddle-Node Bifurcation

The saddle-node bifurcation is the basic mechanism for the creation and destruction of fixed points. Here's the prototypical example in two dimensions:

$$\begin{aligned}\dot{x} &= \mu - x^2 \\ \dot{y} &= -y\,. \end{aligned} \qquad (1)$$

In the x-direction we see the bifurcation behavior discussed in Section 3.1, while in the y-direction the motion is exponentially damped.

Consider the phase portrait as μ varies. For $\mu > 0$, Figure 8.1.1 shows that there are two fixed points, a stable node at $(x^*, y^*) = (\sqrt{\mu}, 0)$ and a saddle at $(-\sqrt{\mu}, 0)$. As μ decreases, the saddle and node approach each other, then collide when $\mu = 0$, and finally disappear when $\mu < 0$.

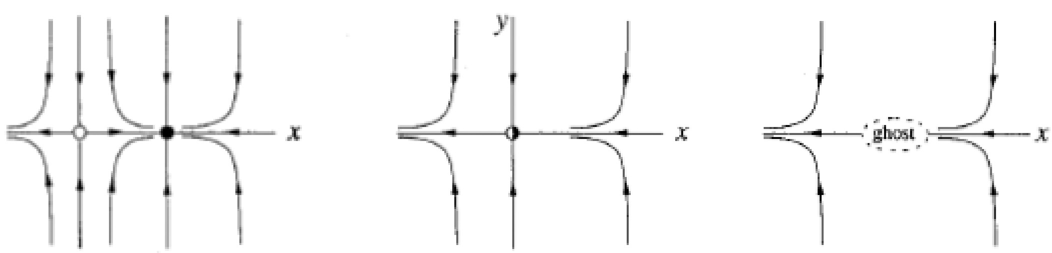

$\mu > 0$ $\qquad$ $\mu = 0$ $\qquad$ $\mu < 0$

Figure 8.1.1

Even after the fixed points have annihilated each other, they continue to influence the flow—as in Section 4.3, they leave a *ghost*, a bottleneck region that sucks trajectories in and delays them before allowing passage out the other side. For the same reasons as in Section 4.3, the time spent in the bottleneck generically increases as $(\mu - \mu_c)^{-1/2}$, where μ_c is the value at which the saddle-node bifurcation occurs.

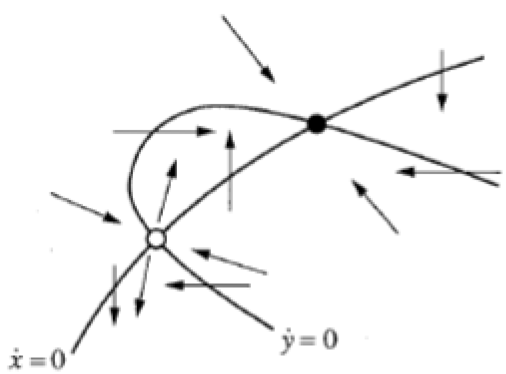

Figure 8.1.2

Some applications of this scaling law in condensed-matter physics are discussed by Strogatz and Westervelt (1989).

Figure 8.1.1 is representative of the following more general situation. Consider a two-dimensional system $\dot{x}=f(x,y)$, $\dot{y}=g(x,y)$ that depends on a parameter μ. Suppose that for some value of μ the nullclines intersect as shown in Figure 8.1.2. Notice that each intersection corresponds to a fixed point since $\dot{x}=0$ and $\dot{y}=0$ simultaneously. Thus, to see how the fixed points move as μ changes, we just have to watch the intersections. Now suppose that the nullclines pull away from each other as μ varies, becoming *tangent* at $\mu=\mu_c$. Then the fixed points approach each other and collide when $\mu=\mu_c$; after the nullclines pull apart, there are no intersections and the fixed points disappear with a bang. The point is that *all* saddle-node bifurcations have this character locally.

EXAMPLE 8.1.1:

Consider the following example, derived from a model of a genetic control system.

In dimensionless form, the equations are

$$\dot{x}=-ax+y$$

$$\dot{y}=\frac{x^2}{1+x^2}-by$$

The constants a and b are supposed to be positive; while a may vary, b is kept fixed.

Show that the system has three fixed points when $a<a_c$, where a_c is to be determined. Show that two of these fixed points coalesce in a saddle-node bifurcation when $a=a_c$. Then sketch the phase portrait for $a<a_c$.

The problem can be faced in a geometric way, by drawing the two nullclines $y = ax$ and $by = x^2 / (1 + x^2)$

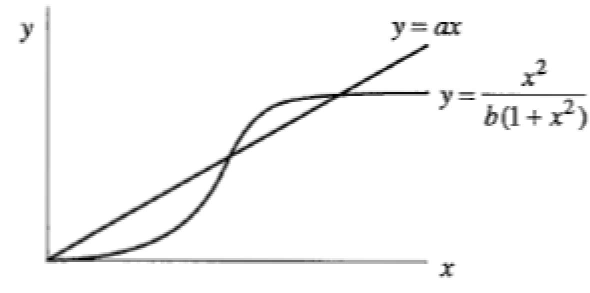

Figure 8.1.3

For small a there are three intersections, as in Figure 8.1.3. As a increases, the top two intersections approach each other and collide when the line intersects the curve tangentially. For larger values of a, those fixed points disappear, leaving the origin as the only fixed point.

To find a_c, we compute the fixed points directly and find where they coalesce. The nullclines intersect when

$$ax = \frac{x^2}{b(1+x^2)}.$$

One solution is $x^* = 0$, in which case $y^* = 0$. The other intersections satisfy the quadratic equation

$$ab(1+x^2) = x \qquad (2)$$

which has two solutions

$$x^* = \frac{1 \pm \sqrt{1-4a^2b^2}}{2ab}$$

if $1-4a^2b^2 > 0$, i.e., $2ab < 1$. These solutions coalesce when $2ab = 1$. Hence

$$\boxed{a_c = 1/2b.}$$

Note that, with a at the bifurcation value $1/2b$, the fixed point is $x^* = 1$.

The nullclines (Figure 8.1.4) provide a lot of information about the phase portrait for $a < a_c$. The vector field is vertical on the line $y = ax$ and horizontal on the sigmoidal curve. Other arrows can be sketched by noting the signs of $\dot{x}$ and $\dot{y}$. It appears that the middle fixed point is a saddle and the other two are sinks. To confirm this, we turn now to the classification of the fixed points.

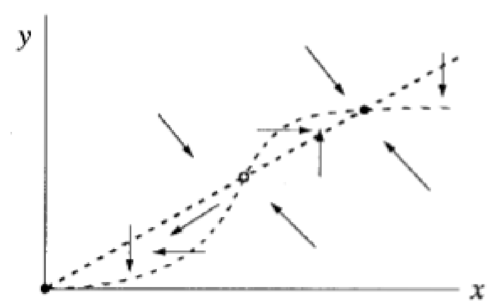

Figure 8.1.4

The Jacobian matrix at (x, y) is

$$A = \begin{pmatrix} -a & 1 \\ \frac{2x}{(1+x^2)^2} & -b \end{pmatrix}.$$

At $x^* = 0$, the eigenvalues are both negative, hence **the origin is always a stable node**. In the two other fixed points ($a < a_c$), the eigenvalues are given by (using Eq. 2):

$$\lambda^2 + (a+b)\lambda + ab\frac{\left(x^*\right)^2 - 1}{\left(x^*\right)^2 + 1} = 0 \qquad (1)$$

Hence, at the middle fixed point, which has $0 < x^* < 1$, we have two real eigenvalues of opposite sign. Therefore, this fixed point is a **saddle** point.
The other fixed point has always $x^* > 1$ and hence is always a **stable node.**

The phase portrait is plotted in Figure 8.1.5. By looking back at Figure 8.1.4, we can see that the unstable manifold of the saddle is necessarily trapped in the narrow channel between the two nullclines. More importantly, the *stable* manifold separates the plane into two regions, each a basin of attraction for a sink.

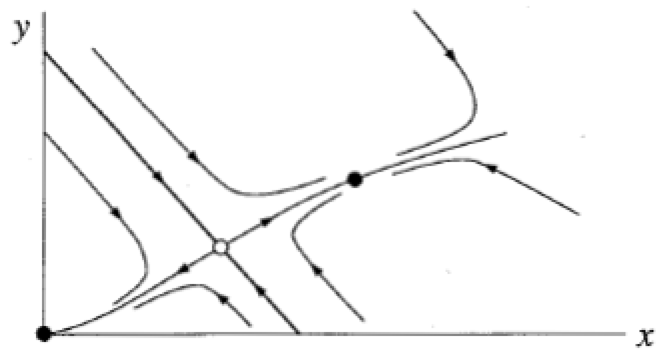

Figure 8.1.5

Transcritical and Pitchfork Bifurcations

Using the same idea as above, we can also construct prototypical examples of transcritical and pitchfork bifurcations at a stable fixed point. In the x-direction the dynamics are given by the normal forms discussed in Chapter 3, and in the y-direction the motion is exponentially damped. This yields the following examples:

$\dot{x} = \mu x - x^2, \quad \dot{y} = -y$ (transcritical)

$\dot{x} = \mu x - x^3, \quad \dot{y} = -y$ (supercritical pitchfork)

$\dot{x} = \mu x + x^3, \quad \dot{y} = -y$ (subcritical pitchfork)

The analysis in each case follows the same pattern, so we'll discuss only the supercritical pitchfork, and leave the other two cases as exercises.

EXAMPLE 8.1.2:

Plot the phase portraits for the supercritical pitchfork system $\dot{x} = \mu x - x^3$, $\dot{y} = -y$, for $\mu < 0$, $\mu = 0$, and $\mu > 0$.

Solution: For $\mu < 0$, the only fixed point is a stable node at the origin. For $\mu = 0$, the origin is still stable, but now we have very slow (algebraic) decay along the x-direction instead of exponential decay; this is the phenomenon of "critical slowing down" discussed in Section 3.4 and Exercise 2.4.9. For $\mu > 0$, the origin loses stability and gives birth to two new stable fixed points symmetrically located at $(x^*, y^*) = (\pm\sqrt{\mu}, 0)$. By computing the Jacobian at each point, you can check that the origin is a saddle and the other two fixed points are stable nodes. The phase portraits are shown in Figure 8.1.6. ■

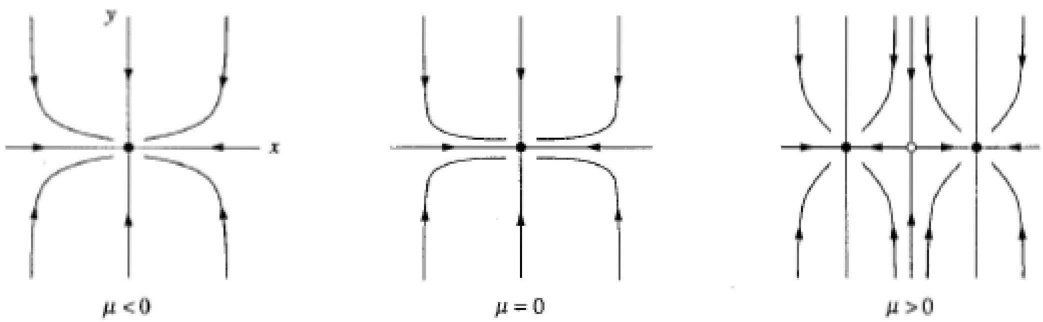

Figure 8.1.6

8.2 Hopf Bifurcations

Suppose a two-dimensional system has a stable fixed point. What are all the possible ways it could lose stability as a parameter μ varies? The eigenvalues of the Jacobian are the key. If the fixed point is stable, the eigenvalues λ_1, λ_2 must both lie in the left half-plane $\text{Re}\,\lambda < 0$. Since the λ's satisfy a quadratic equation with real coefficients, there are two possible pictures: either the eigenvalues are both real and negative (Figure 8.2.1a) or they are complex conjugates (Figure 8.2.1b). To destabilize the fixed point, we need one or both of the eigenvalues to cross into the right half-plane as μ varies.

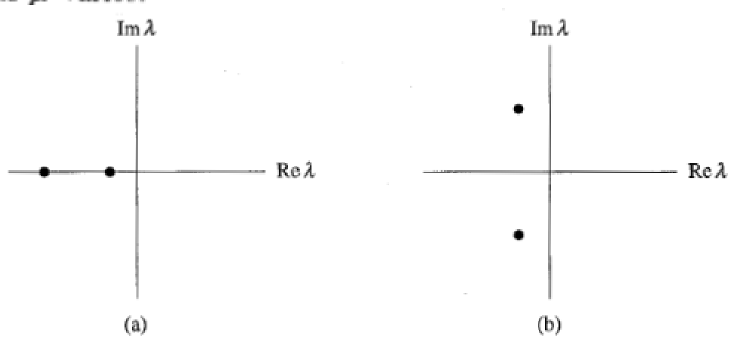

Figure 8.2.1

In Section 8.1 we explored the cases in which a real eigenvalue passes through $\lambda = 0$. These were just our old friends from Chapter 3, namely the saddle-node, transcritical, and pitchfork bifurcations. Now we consider the other possible scenario, in which two complex conjugate eigenvalues simultaneously cross the imaginary axis into the right half-plane.

Supercritical Hopf Bifurcation

Suppose we have a physical system that settles down to equilibrium through exponentially damped oscillations. In other words, small disturbances decay after "ringing" for a while (Figure 8.2.2a). Now suppose that the decay rate depends on a control parameter μ. If the decay becomes slower and slower and finally changes to *growth* at a critical value μ_c, the equilibrium state will lose stability. In many cases the resulting motion is a small-amplitude, sinusoidal, limit cycle oscillation about the former steady state (Figure 8.2.2b). Then we say that the system has undergone a ***supercritical Hopf bifurcation***.

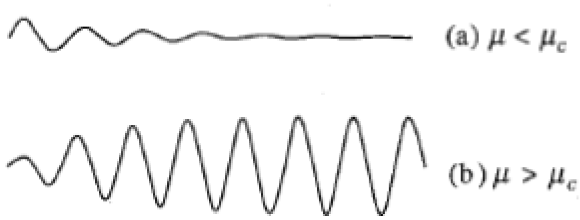

Figure 8.2.2

In terms of the flow in phase space, a supercritical Hopf bifurcation occurs when a stable spiral changes into an unstable spiral surrounded by a small, nearly elliptical limit cycle. Hopf bifurcations can occur in phase spaces of any dimension $n \geq 2$, but as in the rest of this chapter, we'll restrict ourselves to two dimensions.

A simple example of a supercritical Hopf bifurcation is given by the following system:

$$\dot{r} = \mu r - r^3$$
$$\dot{\theta} = \omega + br^2.$$

There are three parameters: μ controls the stability of the fixed point at the origin, ω gives the frequency of infinitesimal oscillations, and b determines the dependence of frequency on amplitude for larger amplitude oscillations.

Figure 8.2.3 plots the phase portraits for μ above and below the bifurcation. When $\mu < 0$ the origin $r = 0$ is a stable spiral whose sense of rotation depends on the sign of ω. For $\mu = 0$ the origin is still a stable spiral, though a very weak one: the decay is only algebraically fast. (This case was shown in Figure 6.3.2. Recall that the linearization wrongly predicts a center at the origin.) Finally, for $\mu > 0$ there is an unstable spiral at the origin and a stable circular limit cycle at $r = \sqrt{\mu}$.

$\mu < 0$ $\qquad\qquad$ $\mu > 0$

Figure 8.2.3

To see how the eigenvalues behave during the bifurcation, we rewrite the system in Cartesian coordinates; this makes it easier to find the Jacobian. We write $x = r\cos\theta$, $y = r\sin\theta$. Then

$$\begin{aligned}\dot{x} &= \dot{r}\cos\theta - r\dot{\theta}\sin\theta \\ &= (\mu r - r^3)\cos\theta - r(\omega + br^2)\sin\theta \\ &= \left(\mu - [x^2 + y^2]\right)x \;-\; \left(\omega + b[x^2 + y^2]\right)y \\ &= \mu x - \omega y + \text{cubic terms}\end{aligned}$$

and similarly

$$\dot{y} = \omega x + \mu y + \text{cubic terms.}$$

So the Jacobian at the origin is

$$A = \begin{pmatrix} \mu & -\omega \\ \omega & \mu \end{pmatrix},$$

which has eigenvalues

$$\lambda = \mu \pm i\omega.$$

As expected, the eigenvalues cross the imaginary axis from left to right as μ increases from negative to positive values.

Rules of Thumb

Our idealized case illustrates two rules that hold *generically* for supercritical Hopf bifurcations:

1. The size of the limit cycle grows continuously from zero, and increases proportional to $\sqrt{\mu - \mu_c}$, for μ close to μ_c.

2. The frequency of the limit cycle is given approximately by $\omega = \text{Im}\,\lambda$, evaluated at $\mu = \mu_c$. This formula is exact at the birth of the limit cycle, and correct within $O(\mu - \mu_c)$ for μ close to μ_c. The period is therefore $T = (2\pi/\text{Im}\,\lambda) + O(\mu - \mu_c)$.

Subcritical Hopf Bifurcation

Like pitchfork bifurcations, Hopf bifurcations come in both super- and subcritical varieties. The subcritical case is always much more dramatic, and potentially dangerous in engineering applications. After the bifurcation, the trajectories must *jump* to a *distant* attractor, which may be a fixed point, another limit cycle, infinity, or—in three and higher dimensions—a chaotic attractor. We'll see a concrete example of this last, most interesting case when we study the Lorenz equations (Chapter 9).

But for now, consider the two-dimensional example

$$\dot{r} = \mu r + r^3 - r^5$$
$$\dot{\theta} = \omega + br^2.$$

The important difference from the earlier supercritical case is that the cubic term r^3 is now *destabilizing*; it helps to drive trajectories away from the origin.

The phase portraits are shown in Figure 8.2.5. For $\mu < 0$ there are two attractors, a stable limit cycle and a stable fixed point at the origin. Between them lies an unstable cycle, shown as a dashed curve in Figure 8.2.5; it's the player to watch in this scenario. As μ increases, the unstable cycle tightens like a noose around the fixed point.

A ***subcritical Hopf bifurcation*** occurs at $\mu = 0$, where the unstable cycle shrinks to zero amplitude and engulfs the origin, rendering it unstable. For $\mu > 0$, the large-amplitude limit cycle is suddenly the only attractor in town. Solutions that used to remain near the origin are now forced to grow into large-amplitude oscillations.

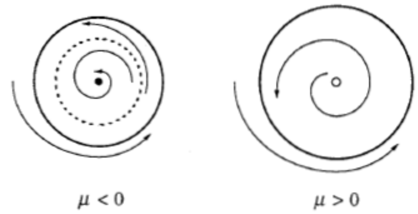

Figure 8.2.5

Note that the system exhibits *hysteresis*: once large-amplitude oscillations have begun, they cannot be turned off by bringing μ back to zero. In fact, the large oscillations will persist until $\mu = -1/4$ where the stable and unstable cycles collide and annihilate. This destruction of the large-amplitude cycle occurs via another type of bifurcation, to be discussed in Section 8.4.

Subcritical Hopf bifurcations occur in the dynamics of nerve cells (Rinzel and Ermentrout 1989), in aeroelastic flutter and other vibrations of airplane wings (Dowell and Ilgamova 1988, Thompson and Stewart 1986), and in instabilities of fluid flows (Drazin and Reid 1981).

Subcritical, Supercritical, or Degenerate Bifurcation?

Given that a Hopf bifurcation occurs, how can we tell if it's sub- or supercritical? The linearization doesn't provide a distinction: in both cases, a pair of eigenvalues moves from the left to the right half-plane.

An analytical criterion exists, but it can be difficult to use (see Exercises 8.2.12–15 for some tractable cases). A quick and dirty approach is to use the computer. If a small, attracting limit cycle appears immediately after the fixed point goes unstable, and if its amplitude shrinks back to zero as the parameter is reversed, the bifurcation is supercritical; otherwise, it's probably subcritical, in which case the nearest attractor might be far from the fixed point, and the system may exhibit hysteresis as the parameter is reversed. Of course, computer experiments are not proofs and you should check the numerics carefully before making any firm conclusions.