Stability Analysis for ODEs

Marc R. Roussel

May 13, 2004

1 Linear stability analysis

Equilibria are not always stable. Since stable and unstable equilibria play quite different roles in the dynamics of a system, it is useful to be able to classify equilibrium points based on their stability.

Suppose that we have a set of autonomous ordinary differential equations, written in vector form:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}). \tag{1}$$

Suppose that $\mathbf{x}^*$ is an equilibrium point. By definition, $\mathbf{f}(\mathbf{x}^*) = \mathbf{0}$. Now suppose that we take a multivariate Taylor expansion of the right-hand side of our differential equation:

$$\begin{aligned} \dot{\mathbf{x}} &= \mathbf{f}(\mathbf{x}^*) + \left.\frac{\partial \mathbf{f}}{\partial \mathbf{x}}\right|_{\mathbf{x}^*} (\mathbf{x} - \mathbf{x}^*) + \ldots \\ &= \left.\frac{\partial \mathbf{f}}{\partial \mathbf{x}}\right|_{\mathbf{x}^*} (\mathbf{x} - \mathbf{x}^*) + \ldots \end{aligned} \tag{2}$$

The partial derivative in the above equation is to be interpreted as the **Jacobian matrix**. If the components of the state vector $\mathbf{x}$ are $(x_1, x_2, \ldots, x_n)$ and the components of the rate vector $\mathbf{f}$ are $(f_1, f_2, \ldots, f_n)$, then the Jacobian is

$$\mathbf{J} = \begin{bmatrix} \frac{\partial f_1}{\partial x_1} & \frac{\partial f_1}{\partial x_2} & \cdots & \frac{\partial f_1}{\partial x_n} \\ \frac{\partial f_2}{\partial x_1} & \frac{\partial f_2}{\partial x_2} & \cdots & \frac{\partial f_2}{\partial x_n} \\ \vdots & \vdots & & \vdots \\ \frac{\partial f_n}{\partial x_1} & \frac{\partial f_n}{\partial x_2} & \cdots & \frac{\partial f_n}{\partial x_n} \end{bmatrix}.$$

Now define $\delta\mathbf{x} = \mathbf{x} - \mathbf{x}^*$. Taking a derivative of this definition, we get $\dot{\delta\mathbf{x}} = \dot{\mathbf{x}}$. If $\delta\mathbf{x}$ is small, then only the first term in equation 2 is significant, since the higher terms involve powers of our small displacement from equilibrium. If we want to know how trajectories behave *near* the equilibrium point, e.g. whether they move toward or away from the equilibrium point, it should therefore be good enough to keep just this term.[1] Then we have

$$\dot{\delta\mathbf{x}} = \mathbf{J}^* \, \delta\mathbf{x},$$

where $\mathbf{J}^*$ is the Jacobian evaluated at the equilibrium point. The matrix $\mathbf{J}^*$ is a constant, so this is just a linear differential equation. According to the theory of linear differential equations, the solution can be written as a superposition of terms of the form $e^{\lambda_j t}$ where $\{\lambda_j\}$ is the set of eigenvalues of the Jacobian.

The eigenvalues of the Jacobian are, in general, complex numbers. Let $\lambda_j = \mu_j + i\nu_j$, where μ_j and ν_j are, respectively, the real and imaginary parts of the eigenvalue. Each of the exponential terms in the expansion can therefore be written

$$e^{\lambda_j t} = e^{\mu_j t} e^{i\nu_j t}.$$

The complex exponential in turn can be written

$$e^{i\nu_j t} = \cos(\nu_j t) + i \sin(\nu_j t).$$

The complex part of the eigenvalue therefore only contributes an oscillatory component to the solution. It's the real part that matters: If any one of the $\mu_j > 0$, $e^{\mu_j t}$ grows with time, which means that systems will tend to move away from the equilibrium point. This leads us to a very important theorem:

Theorem 1 *An equilibrium point* $\mathbf{x}^*$ *of the differential equation 1 is stable if all the eigenvalues of* $\mathbf{J}^*$*, the Jacobian evaluated at* $\mathbf{x}^*$*, have negative real parts. The equilibrium point is unstable if at least one of the eigenvalues has a positive real part.*

Because we are only keeping an approximate, local, linear approximation to the original differential equation, an analysis based on this theorem is called a **linear stability analysis**.

Note that the theorem is silent on the issue of what happens if some of the eigenvalues have zero real parts while the others are all negative. This case can't

[1] This assumes that the Jacobian evaluated at the equilibrium point isn't, in a sense to be defined later, zero.

be decided based on linear stability analysis. The nonlinear terms we left out of equation 2 in fact determine the stability in this case. Unfortunately, there isn't a general technique to deal with this case. We will discuss systems with zero eigenvalues from time to time and discuss their treatment as appropriate.

Example 1.1 Let us return to the Lindemann mechanism, for which phase-plane analysis has already shown us that the equilibrium point is stable. The dimensionless ODEs are

$$\begin{aligned} \dot{a} &= -a^2 + \alpha ab, \\ \dot{b} &= a^2 - \alpha ab - b. \end{aligned}$$

The equilibrium point is $(0,0)$. The Jacobian matrix is

$$\mathbf{J} = \begin{bmatrix} \frac{d\dot{a}}{da} & \frac{d\dot{a}}{db} \\ \frac{d\dot{b}}{da} & \frac{d\dot{b}}{db} \end{bmatrix} = \begin{bmatrix} -2a + \alpha b & \alpha a \\ 2a - \alpha b & -\alpha a - 1 \end{bmatrix}.$$

Evaluating the Jacobian at the equilibrium point, we get

$$\mathbf{J}^* = \begin{bmatrix} 0 & 0 \\ 0 & -1 \end{bmatrix}.$$

The eigenvalues of a 2×2 matrix are easy to calculate by hand: They are the solutions of the determinant equation

$$|\lambda \mathbf{I} - \mathbf{J}| = 0.$$

In this case,

$$\begin{vmatrix} \lambda & 0 \\ 0 & \lambda + 1 \end{vmatrix} = \lambda(\lambda + 1) = 0.$$

The solutions of this equation can be read by inspection: $\lambda = 0$ or $\lambda = -1$. One of the eigenvalues is zero, so we can't tell from the linear stability analysis alone whether or not the equilibrium point is stable. Of course, we already know from the phase-plane analysis that it is.

Example 1.2 The law of microscopic reversibility says that we can't have truly irreversible elementary chemical reactions, although this might be a good approximation if the reaction is strongly product-favored. Consider therefore the fully reversible Lindemann mechanism:

$$\begin{aligned} \mathrm{A}+\mathrm{A} &\underset{k_{-1}}{\overset{k_1}{\rightleftharpoons}} \mathrm{A}+\mathrm{B}, \\ \mathrm{B} &\underset{k_{-2}}{\overset{k_2}{\rightleftharpoons}} \mathrm{P}. \end{aligned}$$

Defining

$$\begin{aligned} a &= k_1 A/k_2, & \alpha &= k_{-1}/k_1, \\ b &= k_1 B/k_2, & \beta &= k_{-2}/k_2, \\ \tau &= k_2 t, & a_0 &= k_1 A_0/k_2, \end{aligned}$$

the dimensionless ODEs become

$$\begin{aligned} \dot{a} &= -a^2 + \alpha a b, \\ \dot{b} &= a^2 - \alpha a b - b + \beta(a_0 - a - b). \end{aligned}$$

If we set $\dot{a} = \dot{b} = 0$, we discover that there are *two* equilibrium points:

$$\begin{aligned} (a^\dagger, b^\dagger) &= \left(0, \frac{\beta a_0}{1+\beta}\right), \\ (a^*, b^*) &= \left(\frac{\alpha\beta a_0}{1+\beta(1+\alpha)}, \frac{\beta a_0}{1+\beta(1+\alpha)}\right). \end{aligned}$$

The Jacobian of the differential equations is

$$\mathbf{J} = \begin{bmatrix} -2a + \alpha b & \alpha a \\ 2a - \alpha b - \beta & -\alpha a - 1 - \beta \end{bmatrix}.$$

The Jacobian evaluated at the first equilibrium point is

$$\mathbf{J}^\dagger = \begin{bmatrix} \alpha b^\dagger & 0 \\ \alpha b^\dagger - \beta & -1 - \beta \end{bmatrix}.$$

We could substitute the value of $b^\dagger$ directly into $\mathbf{J}^\dagger$, but I usually leave this for a later step, and then only if it's necessary to determine the sign of the eigenvalues.

The eigenvalues satisfy the equation

$$\begin{vmatrix} \lambda - \alpha b^\dagger & 0 \\ -(\alpha b^\dagger - \beta) & \lambda + 1 + \beta \end{vmatrix} = (\lambda - \alpha b^\dagger)(\lambda + 1 + \beta) = 0.$$

The eigenvalues associated with $(a^\dagger, b^\dagger)$ are therefore

$$\begin{aligned} \lambda_1 &= \alpha b^\dagger, \\ \text{and } \lambda_2 &= -(1+\beta). \end{aligned}$$

Since $b^\dagger$ is positive, λ_1 is clearly positive. It follows that $(a^\dagger, b^\dagger)$ is an *unstable* equilibrium point.

We could try to work out the stability of the other point by hand, but it's messy. In this case, it's far better to use Maple. The steps in the analysis are much the same, although it takes a few tricks to get to the bottom of this exercise. We start by defining the differential equations:

```
> adot := (a,b) -> -a^2+alpha*a*b;
```

$$adot := (a, b) \rightarrow -a^2 + \alpha\, a\, b$$

```
> bdot := (a,b) -> a^2 - alpha*a*b - b + beta*(a0-a-b);
```

$$bdot := (a, b) \rightarrow a^2 - \alpha\, a\, b - b + \beta\,(a0 - a - b)$$

Define the steady state:

```
> bstar := beta*a0/(1+beta*(1+alpha));
```

$$bstar := \frac{\beta\, a0}{1 + \beta\,(1+\alpha)}$$

```
> astar := alpha*bstar;
```

$$astar := \frac{\alpha\,\beta\, a0}{1 + \beta\,(1+\alpha)}$$

Verify:

```
> simplify(adot(astar,bstar));
```

$$0$$

```
> simplify(bdot(astar,bstar));
```

$$0$$

Calculate the Jacobian:

```
> with(linalg):
Warning, the protected names norm and trace have been
redefined and unprotected
> J := jacobian([adot(a,b),bdot(a,b)],[a,b]);
```

$$J := \begin{bmatrix} -2a+\alpha b & \alpha a \\ 2a-\alpha b-\beta & -\alpha a-1-\beta \end{bmatrix}$$

Substitute the steady state into the Jacobian:

```
> Jstar := map(x->subs(a=astar,b=bstar,x),J);
```

$$Jstar := \begin{bmatrix} -\dfrac{\alpha\beta a0}{\%1} & \dfrac{\alpha^2\beta a0}{\%1} \\ \dfrac{\alpha\beta a0}{\%1}-\beta & -\dfrac{\alpha^2\beta a0}{\%1}-1-\beta \end{bmatrix}$$

$$\%1 := 1+\beta(1+\alpha)$$

Determine the characteristic polynomial (the polynomial whose roots are the eigenvalues). Collect the result in powers of the eigenvalue λ:

```
> collect(charpoly(Jstar,lambda),lambda);
```

$$\lambda^2 + \frac{(\alpha^2\beta a0+\alpha\beta a0+\beta\alpha+\beta^2\alpha+\beta^2+2\beta+1)\lambda}{1+\beta+\beta\alpha} + \frac{\alpha^2\beta^2 a0+\alpha\beta^2 a0+\alpha\beta a0}{1+\beta+\beta\alpha}$$

This polynomial is in the form

$$\lambda^2 + c_1\lambda + c_0.$$

Note that the coefficients c_0 and c_1 are both positive. The roots of this polynomial are

$$\lambda = \frac{1}{2}\left\{-c_1 \pm \sqrt{c_1^2 - 4c_0}\right\}.$$

Since c_0 is positive, the quantity under the square root is either smaller than c_1^2, or it is negative. If negative, the solutions are complex with real part $-c_1$, which is negative. Otherwise, the square root must be smaller in absolute value than c_1, so that the two eigenvalues must still be negative. Either way, we conclude that the steady state is stable since the real parts of both eigenvalues must be negative.

Note that the possibility of a zero eigenvalue disappeared in the last example when we considered the effect of the law of microscopic reversibility. The zero eigenvalue was therefore an artifact of having an incomplete chemical model since this feature disappears when we include a nonzero value for the rate constant k_{-2}, no matter how small this value is. This feature of the irreversible Lindemann mechanism is not **structurally stable** since an arbitrarily small change to the model makes it disappear. This is a minor structural instability since the qualitative behavior of the model (approach to the equilibrium point) is pretty much the same, whether the leading eigenvalue is zero or negative. However, there are models which have severe structural instabilities, i.e. models in which the whole behavior of the model is significantly different depending on whether one includes a small term or not. Structurally unstable models are generally considered to be bad models since we generally don't know the exact evolution equations for a model. If the model's behavior changes when we make small changes to the equations, then we probably can't trust the model's predictions very well.

Note also that we have our first example of a model with two equilibrium points, one of which is stable while the other is unstable. It is instructive to look at the phase portrait, which we'll draw using Maple this time since we have it running anyway. We first have to pick some values for our parameters. Since our analysis indicates that the stability properties of the two fixed points never change, it doesn't much matter what we pick.

```
> alpha := 1;
```

$$\alpha := 1$$

```
> beta := 1;
```

$$\beta := 1$$

```
> a0 := 1;
```

$$a0 := 1$$

The coordinates of the stable equilibrium point are

```
> astar; bstar;
```

$$\frac{1}{3}$$

$$\frac{1}{3}$$

The unstable equilibrium sits on the *b* axis at coordinate

```
> bdagger := beta*a0/(1+beta);
```

$$bdagger := \frac{1}{2}$$

We're now ready to create the phase portrait.

```
> with(DEtools):
Warning, the name adjoint has been redefined
> with(plots):
Warning, the name changecoords has been redefined
> pp := phaseportrait([diff(a(t),t)=adot(a(t),b(t)),
  diff(b(t),t)=bdot(a(t),b(t))], [a(t),b(t)], t=0..10,
  [[a(0)=i/5,b(0)=0]$i=0..5, [a(0)=i/5,b(0)=1]$i=0..5,
  [a(0)=0.02,b(0)=0], [a(0)=0.02,b(0)=1]], a(t)=0..a0,
  b(t)=0..a0, arrows=NONE, stepsize=0.01,
  linecolor=green):
```

We're storing the plot in the variable `pp`. We'll also plot the nullclines, and store them temporarily:

```
> a_null1 := plot(a/alpha,a=0..1,color=red):
> a_null2 := plot([0,b,b=0..1],color=red,thickness=3):
> b_null := plot((a^2+beta*(a0-a))/(alpha*a+1+beta),
  a=0..1):
```

The `display()` function puts all four pictures out on one graph:

```
> display({pp,a_null1,a_null2,b_null}, labels=[a,b],
  scaling=CONSTRAINED);
```

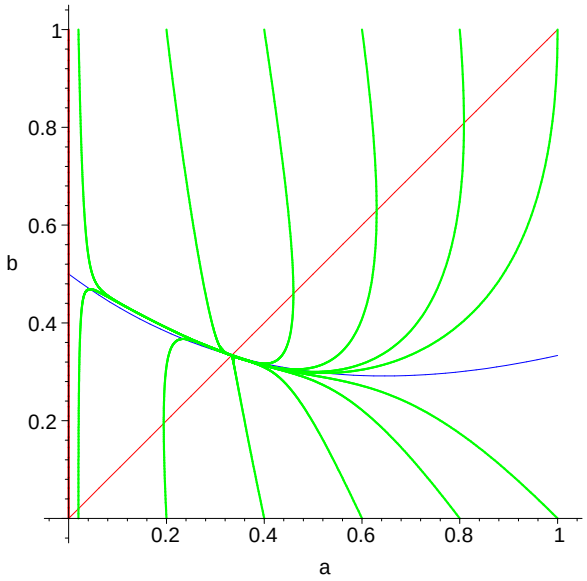

Note that $\dot{a} = 0$ on the b axis. Accordingly, if we start a system with $a = 0$, it moves along the b axis toward the unstable fixed point. However, if we start anywhere *off* the b axis, we eventually end up at the stable fixed point. We say that $(0, b^\dagger)$ is a **saddle point**, i.e. a point with one stable and one unstable direction. The stable direction corresponds to the negative eigenvalue while the unstable direction corresponds to the positive eigenvalue.

The stable equilibrium point (a^*, b^*) on the other hand is called a **stable node**. It has two real, negative eigenvalues.[2]

[2] We've shown that the two eigenvalues always have negative real parts. Proving that they are also real is a pain, but it is possible.

When the eigenvalues of a fixed point are complex, the point is called a **focus**. Fixed points of planar systems can be almost completely classified based on their eigenvalues:

Eigenvalues	**Fixed point**	**Flow**
complex with positive real parts	unstable focus	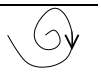
complex with negative real parts	stable focus	
real and positive	unstable node	
real and negative	stable node	
one positive and one negative	saddle point	

The only problem is with systems which have one or more eigenvalues with real part zero. In these cases, all of the above behaviors are possibilities, along with one other, namely a **centre**, which is a closed orbit (e.g. a circle) which is neutrally stable. The latter means that if we displace the system off its orbit, it will continue to circle the equilibrium point, but at a new radius determined by the the disturbance.

Linear stability analysis can be carried out for higher-dimensional systems although, predictably, it gets harder to do things analytically. The ideas remain intact however, and the classification of fixed points follows similar principles, although new possibilities emerge.

2 Liapunov functions

Linear stability analysis tells us how a system behaves near an equilibrium point. It does not however tell us anything about what happens farther away from equilibrium. Phase-plane analysis combined with linear stability analysis can generally give us a full picture of the dynamics, but things become much more difficult in higher-dimensional spaces. In this section, we consider a technique due to Liapunov which can be used to determine the stability of an equilibrium point "in the large", i.e. both near and far from the equilibrium point.

Liapunov's method is based on a simple idea. Suppose that $V(\mathbf{x})$ is a function of our state variables which has a minimum at an equilibrium point and which has no local minima. (Think of a paraboloid.) Now suppose that we can show that the dynamics of our system results in a steady decrease in V in some (possibly large) neighborhood of the equilibrium point. This necessarily means that we are tending toward the minimum of V, which is just the equilibrium point. Having shown this, we can conclude that the equilibrium point is stable over the entire neighborhood of $\mathbf{x}^*$ over which V decreases. A function V with these properties is called a **Liapunov function**.

Let us formalize this idea:

Definition 1 *Let U be a region of phase space containing the equilibrium point $\mathbf{x}^*$. Let $V : U \to \mathbb{R}$ be a continuous and differentiable function. V is a **positive definite function** for the point $\mathbf{x}^*$ if it satisfies the following two conditions:*

1. $V(\mathbf{x}^*) = 0$, *and*
2. $V(\mathbf{x}) > 0$ *for* $\mathbf{x} \in U - \{\mathbf{x}^*\}$.

Theorem 2 *Let $\mathbf{x}^*$ be an equilibrium point of the differential equation 1, and let V be a positive definite function for this point. The equilibrium point is **asymptotically stable** (the solutions tend to this point[3]) for initial conditions in the neighborhood U of $\mathbf{x}^*$ if $\dot{V}(\mathbf{x}) < 0$ for all $\mathbf{x} \in U - \{\mathbf{x}^*\}$.*

In order to use this theorem, we have to obtain a Liapunov function. Unfortunately, it's often really difficult to come up with a Liapunov function for a given system, except in some special cases where the physics of the problem suggests a particular choice.

Example 2.1 A Hooke's law spring is subject to a restoring force $-kx$, where x is the displacement from equilibrium. If the spring is pulling around a mass m on a lubricated surface, the damping (frictional) force is typically proportional to the velocity and opposite in direction. Thus,

$$F = ma = -kx - \mu v.$$

[3]The technical distinction being made here is between stability, which just means that the solutions don't wander away from the point, and asymptotic stability, which means that they actually tend toward the point with time.

By definition,

$$\begin{aligned} v &= \frac{dx}{dt}, \\ \text{and } a &= \frac{dv}{dt}. \end{aligned}$$

We therefore obtain the following planar system:[4]

$$\begin{aligned} \frac{dx}{dt} &= v, \\ \frac{dv}{dt} &= \frac{1}{m}\left(-kx-\mu v\right). \end{aligned}$$

The equilibrium point of this system is clearly $(x^*, v^*) = (0,0)$. In mechanical systems with friction, the total energy is often a Liapunov function. Since the force and potential energy[5] V are related by $F = -dV/dx$, the potential energy associated with the Hooke's law force is $\frac{1}{2}kx^2$. The kinetic energy is $\frac{1}{2}mv^2$, so the total energy is

$$E = \frac{1}{2}kx^2 + \frac{1}{2}mv^2.$$

The total energy is a positive definite function (zero at the equilibrium point, and increasing away fromt this point). Now consider

$$\begin{aligned} \dot{E} &= \frac{\partial E}{\partial x}\frac{dx}{dt} + \frac{\partial E}{\partial v}\frac{dv}{dt} \\ &= kxv + mv\frac{1}{m}\left(-kx-\mu v\right) \\ &= -\mu v^2 < 0. \end{aligned}$$

[4]Noting that $a = d^2x/dt^2$, this system can also be written as a single second-order ODE

$$m\frac{d^2x}{dt^2} = -kx - \mu\frac{dx}{dt},$$

and this is in fact the approach taken in many mechanics courses. What we are seeing here is that second-order ODEs (and higher-order equations, for that matter) are equivalent to sets of first-order ODEs, provided we use some of the derivatives (e.g. $v = dx/dt$) as variables.

[5]The potential energy V should not be confused with the Liapunov function V. The traditional notations of mechanics and of Liapunov functions are in conflict here.

Since E decreases everywhere except at the equilibrium point itself, the equilibrium point is globally asymptotically stable. In order words, starting from any initial conditions, we will eventually reach the equilibrium point.