Feedback and Stability in Control Systems

Leandra Vicci
Department of Computer Science
University of North Carolina at Chapel Hill
9 November 2009

Introduction

Many different forms of control systems exist, some being human-engineered artifacts, but many are natural systems. This report focuses on the control of a single variable in a *plant* (system) that can be represented in terms of linear ordinary differential equations (ODEs). ODEs can describe the behavior of *lumped constant* components. In Newtonian mechanics these consist of springs, masses, and friction; in electrical circuits they are resistors, capacitors and inductors, and there are counterparts in other domains such as Thermodymanics, Chemistry and Biology.

The purpose of any control system is to evoke a predictable response of a controlled variable to a reference input variable. In many cases the goal is a linear response, which is what we will be considering here. Most plants do not behave linearly. For example, the force needed to stand a ladder upright is large to start with but decreases to zero as the (controlled) angle approaches vertical. Moreover, there are often disturbing influences affecting the controlled variable. An example of this would be the effect of varying wind velocity on the speed of a bicycle.

Feedback is widely used to linearize the response of a controlled variable to a reference input. If you can measure the difference between the reference input and the controlled variable to obtain an error, you can apply a correction in opposition to the error (in the jargon, apply *negative feedback*) to reduce it. You might think that the more vigorously you oppose it (increase the *control gain*), the smaller the residual error should be. Conversely, *positive feedback* seems like it should increase the error. These intuitions are correct up to a point. But a more accurate understanding of feedback relies on (a relatively simple) mathematical model, which however yields some non-intuitive results that are borne out in practice.

This report will present a simple static model and its behavior, then will generalize it to the lumped constant dynamic case of a single-input-single-output (SISO) plant with linear time invariant (LTI) components. In particular, feedback stability will be treated.

Model of a feedback control system

A mathematically workable treatment of feedback must account for the time dependent dynamics of the plant and the control gain components. For LTI systems this is most tractably couched in terms of a complex frequency variable $s = \sigma + i\omega$. The time domain ODEs of the sytem are transformed into frequency domain where they become ordinary algebraic equations. The system variables take the form $v(s) = a\exp(st) = a\exp(\sigma t)\exp(i\omega t)$. This reveals that the real part σ represents exponential growth or decay of v with time, according to its sign, while the imaginary part ω represents a periodic variation in time, with a frequency $f = 2\pi\omega$.

Referring to Figure 1, let us model a control loop consisting of three basic components:

- A *plant*, which is represented by a dynamical *state variable* $C(s)$, which responds to a *plant control* signal $B(s)$ and various *disturbing influences* $D(s)$ according to $C(s) = B(s)F(s) + D(s)$.
- A *subtractor* which compares $C(s)$ to a *reference input* signal $R(s)$ to produce a difference Δ, or *error* signal $E(s)$, which is to be minimized.
- A *control gain* component G which drives the plant control input with $B(s) = E(s)G$.

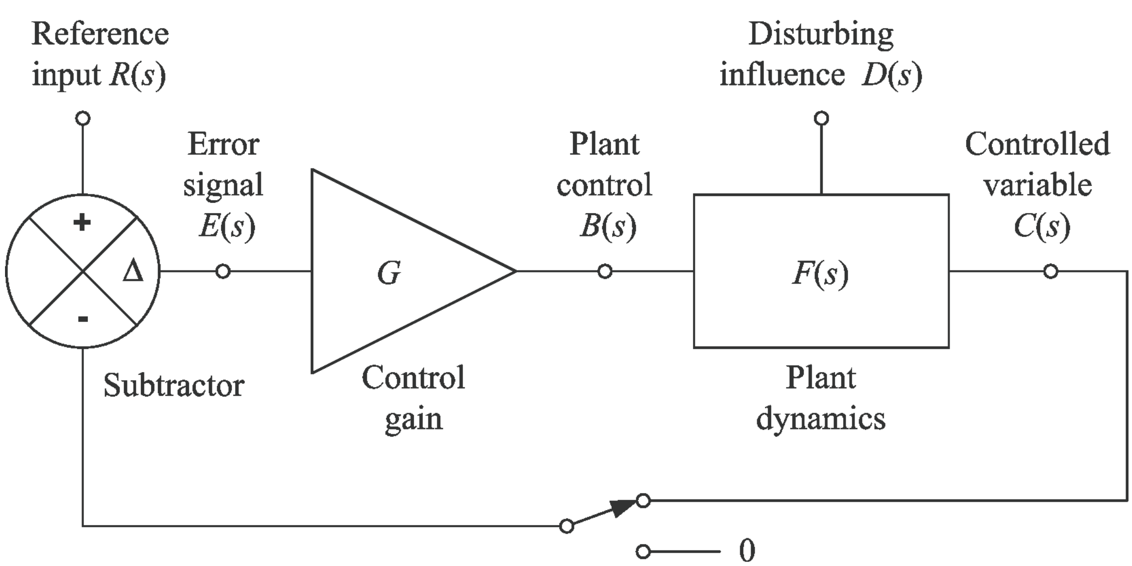

Figure 1: Model of a SISO control loop

Benefits and properties of feedback

Let's first consider the static case $s = 0$, where we can drop all references to complex frequency s. We have said above that $C = BF + D$ and $B = EG$, so $C = EGF + D$. Let us define the *open loop gain* $A_o(s) = G(s)F(s)$, i.e., the product of the transfer functions $G(s)$ and $F(s)$ traversing the loop. We then have,

$$C = A_o E + D. \tag{1}$$

With feedback disabled (switch set to the "0" position), $E = R$, and the controlled variable response to the reference input is simply $C = A_o R + D$, the open loop response. In this condition, the controlled variable is subject to both disturbances D and non-linearities in A_o. In some cases, both D and A_o can vary with time, and A_o can be seriously non-linear, all of which makes precise control of C problematical.

With feedback enabled however (switch set to close the feedback loop), $E = R - C$. Substituting into eq(1) we write, $C = A_o(R - C) + D$, and solve for C to get,

$$C = R\Big(\frac{A_o}{1 + A_o}\Big) + \Big(\frac{D}{1 + A_o}\Big). \tag{2}$$

It is immediately apparent that the effect of D on C can be made arbitrarily small by picking a sufficiently large value of A_o. This is one of the principal benefits of feedback: the suppression of noise and other spurious disturbances.

Now let's explore another principal benefit, linearizing the response of C. Let us define the *closed loop gain* as,

$$A_c = \frac{A_o}{1 + A_o}. \tag{3}$$

Substituting this into eq(2), and for the moment ignoring disturbances by setting $D = 0$, we get

$$C = A_c R. \tag{4}$$

By definition, for C to respond linearly to R, A_c must be constant. We have already conceded that A_o may behave non-linearly, varying significantly with R. Nevertheless, if the minimum value taken by $|A_o|$ is sufficiently large, the value of A_c will remain arbitrarily close to unity, and therefore constant.

Thus for better response linearity and disurbance suppression we wish to provide a large value of A_o. Given some arbitrary plant response F, we are free to increase the control gain G such that the lowest value of A_o is sufficiently large to achieve the desired feedback benefits. There are limits to this of course, which we will revisit later.

Example plots

To illustrate, consider an open loop gain $0 \leq A_o \leq 100$. Figure 2 shows A_c vs. A_o. For large values of A_o, the value of A_c approaches its asymptotic value of unity. For $A_o > 10$ we see that any variation in A_c must be less than 10%.

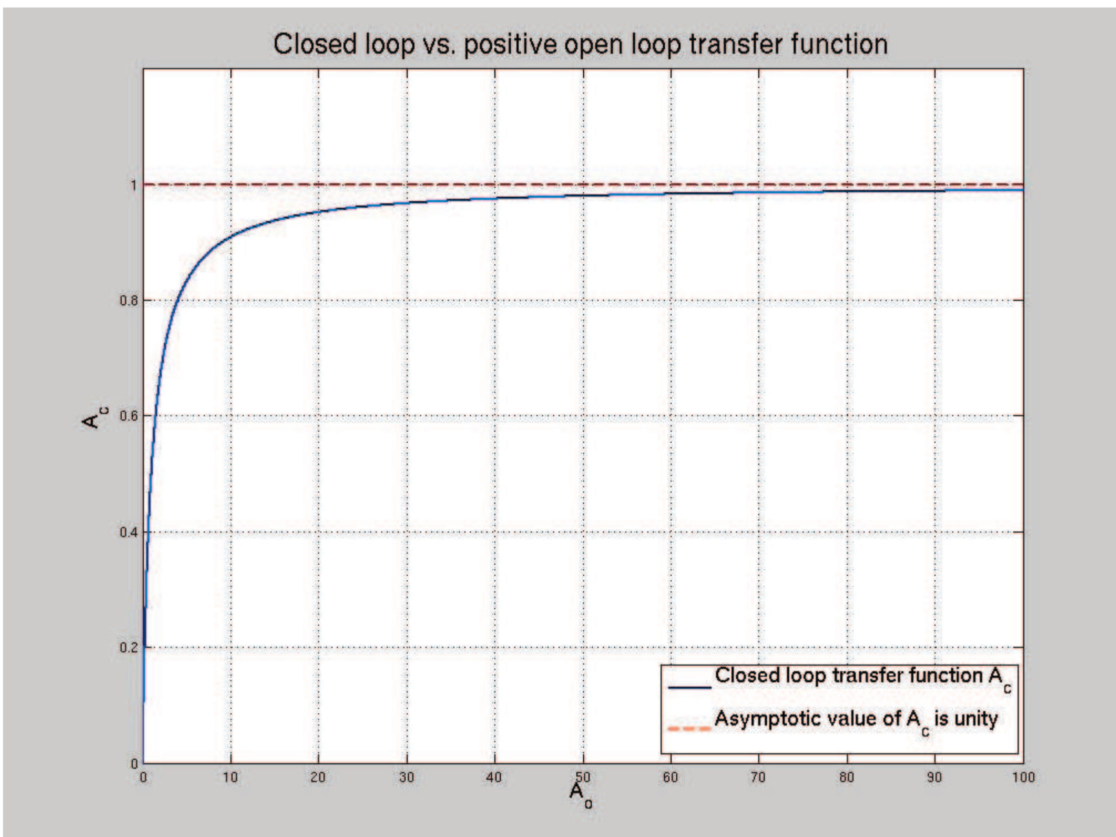

Figure 2: Closed and open loop gains vs reference input

This quite neatly fits our intuition as expressed in the introduction. The greater the open loop gain, the more vigorously the feedback opposes any difference between R and S, leading to the asymptotic approach of A_c to unity.

Consider however, if A_o goes negative. This is tantamount in Figure 1 to swapping the signs of the subtractor, resulting in positive feedback. In this case, one would expect the error E to be magnified rather than reduced. In fact this does occur, but to a significant degree only for small range of negative values. Figure 3 extends the plot of Figure 2 to include negative values of A_o to illustrate the behavior of this positive feedback.

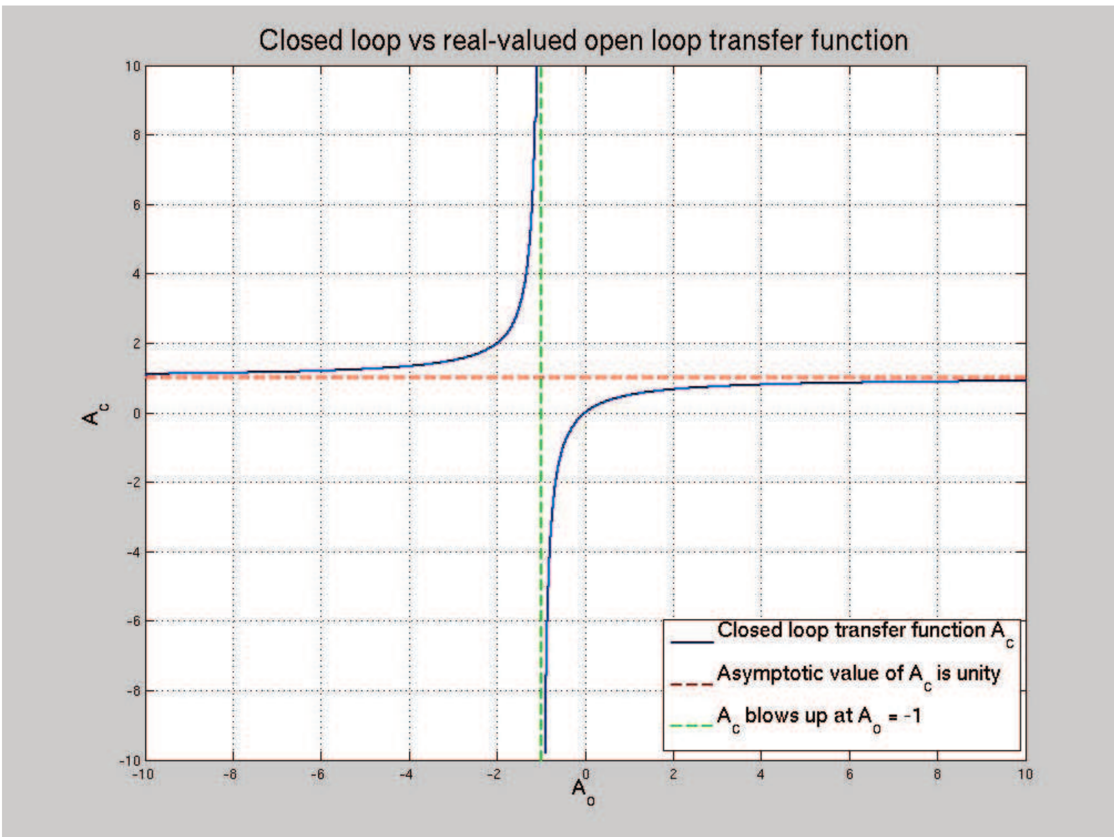

Figure 3: Four quadrant plot of closed and open loop gains

Following the intuition valid for negative feedback, one might expect positive feedback to cause A_c to become increasingly exaggerated (and negative) as the magnitude of (the negative) A_o is increased. Instead, we see this behavior only in the range of $0 \geq A_o > -1$, where it blows up. For $A_o < -1$, A_c flips sign, then asymptotically approaches (positive) unity as $A_o \to -\infty$. Thus, for large values of $|A_o|$, feedback will cause A_c to approach positive unity, regardless of the positive or negative sense of A_o. Most people find this to be counterintuitive.

In preparation for returning to the dynamical case, let us re-introduce the complex frequency variable $s = \sigma + i\omega$, which results in complex values of the plant dynamics function $F(s)$. Consequently the open loop gain $A_o(s) = GF(s)$ is also in general complex. Evaluating eq(3) for A_o extended into the complex domain fills out the picture suggested by Figures 2 and 3. Figure 4 plots the magnitude (height) and phase (color) of $A_c(A_o)$ for $A_o \in (\pm 10 \pm 10i)$.

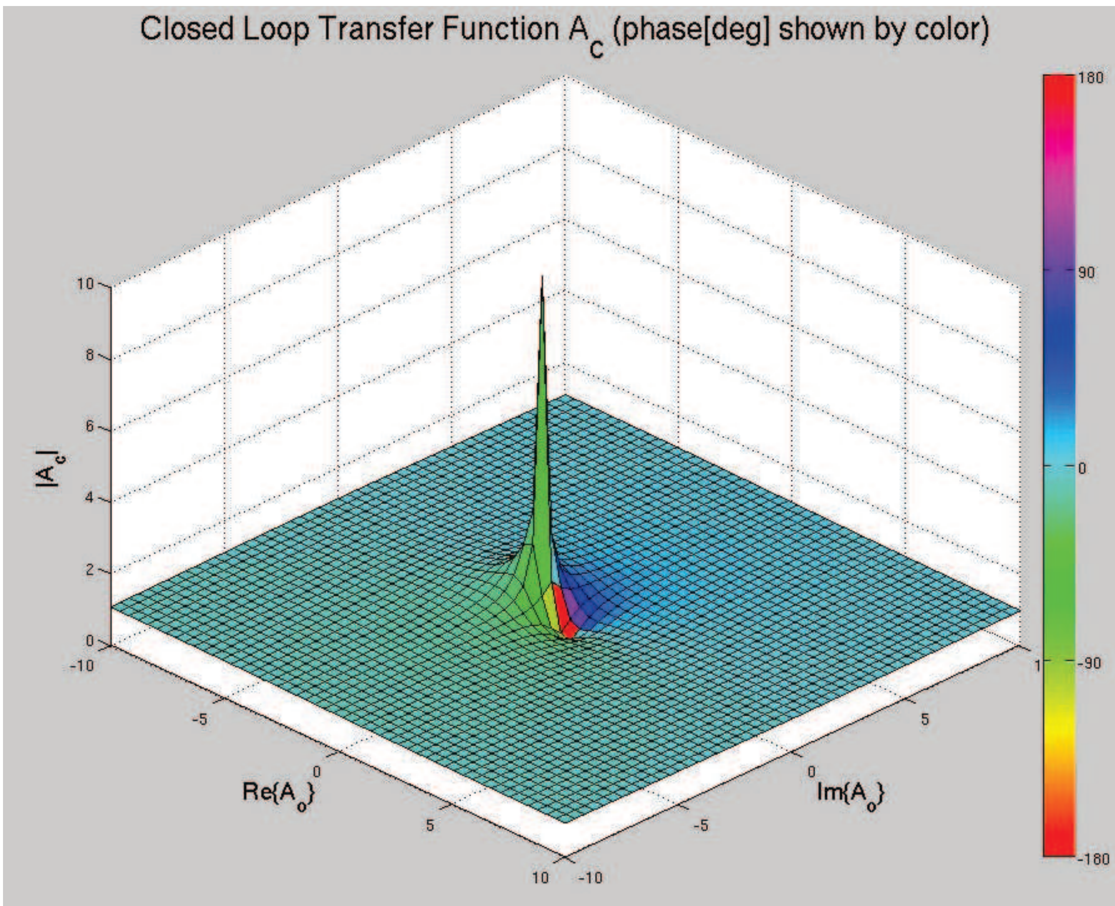

Figure 4: Complex closed loop gain vs. complex open loop gain

As before, we see $A_c = 0$ when $A_o = 0$. We also see that $A_c \to \infty$ as $A_o \to -1$. In the jargon, these are known as a *zero* and a *pole*, respectively. The most significant property to notice however is that $A_c \to 1$ as $|A_o| \to \infty$. This generalizes the notion observed in Figure 3 that large $|A_o|$ provide for $A_c \simeq 1$ irrespective of sign and variations in A_o. Here, the notion of sign generalizes to phase ϕ where the special cases of $\phi = 0$ and $\phi = \pi$[radians] correspond to signs of $+$ and $-$ respectively. It is an important property of the closed loop gain that its value is nearly +1 no matter what the open loop gain is, provided it's large.

Stability

Feedback systems are infamous for stability problems, and it is often regarded as near black magic to implement them with high performance while avoiding oscillation. This perception has developed because the underlying mechanisms are not well understood. Instead, stability criteria such as Routh's and Nyquist's are routinely taught and used which, while entirely valid methods, are not particularly enlightening. In the same spirit as in the previous section, I will attempt to make stability sensible in terms of ordinary algebraic equations.

The problem can be couched in the form of finding values of C that can exist in the loop when there are no disturbances or reference signals present. The differential equations describing this situation are called *homogeneous*, and represent the so-called *unforced* case. Any non-trivial (non zero) solutions represent signals that are supported by the loop in the absence of any other influences. It seems reasonable to regard any such signals as potential instability, and all other signals as dependent on external forcing factors such as R and D for their existence.

Accordingly, let us seek such solutions. We start with eq(4) which represents the loop behavior with $D = 0$. Substituting eq(3) into it and solving for R we have $C(1 + A_o)/A_o = R$. Then setting $R = 0$ we have the homogeneous case,

$$C\Big(\frac{1 + A_o}{A_o}\Big) = 0. \tag{5}$$

There are two possible solutions to eq(5): $C = 0$, and $1 + A_o = 0$. The non trivial solution is $A_o = -1$, where C can assume any finite value. It is only under this condition that the feedback loop can support a value of C independent of R and D. This mathematics is simple to follow, yet gives results that are magically deep. This magic emerges when we introduce explicit time dependence arising from the dynamics of the plant, and an equivalent frequency dependence of A_o,

$$A_o(s) = -1. \tag{6}$$

Understanding the plant model

As mentioned in the introduction, the components of the plant are described by ODEs. This report avoids the complexity of differential equations by invoking the *Laplace transform* which converts time domain ODEs into ordinary algebraic equations in frequency domain. I shall forgo exposition of the transform itself: you will just have to trust me that it works. It is worth summarizing some of its *properties* however, and how the physics of the plant components lead to a structural model in s domain.

The Laplace transform:

$$F(s) = \int_0^\infty f(t)e^{-st}dt \tag{7}$$

converts a time dependent function $f(t)$ in the domain $0 \le t \le \infty$ into a frequency dependent function $F(s)$ in the complex domain $\mathcal{C}$ (**n.b.** use of F here is *not* the same as the plant dynamics F above!). It has the following properties:

- $df(t)/dt \Rightarrow sF(s)$
- $\int f(t)dt \Rightarrow F(s)/s$
- $f(t) + g(t) \Rightarrow F(s) + G(s)$
- $af(t) \Rightarrow aF(s), \quad a \in \mathcal{C}$

Finally, $s = \sigma + i\omega$ defines a frequency having a corresponding time domain behavior $e^{st} = e^{(\sigma+i\omega)t} = e^{\sigma t}e^{i\omega t}$. The identity $e^{i\omega t} = \cos(\omega t) + i\sin(\omega t)$ shows that ω represents an oscillatory frequency, while σ represents an exponantial growth or decay rate, according to its sign.

The plant dynamics:

As a conceptual aid, let us imagine a plant comprising masses, springs, and "dashpots," the latter being an abstract linear frictional device one could think of as a piston in a cylinder which pumps an ideal viscous fluid through an orifice. This has the property that the harder you push, the faster the piston moves; its force law is $F_b = bv$ where b is the frictional coefficient and v is the velocity. Masses behave according to Newton's Second Law, $F_m = ma$, where m is the mass and a is acceleration. Ideal springs behave according to Hooke's Law, $F_k = kx$ where k is the spring constant and x is the length, or position of one end relative to the other.

Now velocity is the rate of change of position, $v = dx/dt$, and acceleration is the rate of change of velocity, $a = dv/dt = d^2x/dt^2$. We can thus express these behaviors in time domain and transform them into frequency domain as follows:

- $F_k(t) = k\,x(t) \Rightarrow f_k(s) = k\,x(s)$,
- $F_b(t) = b\,dx/dt \Rightarrow f_b(s) = b\,s\,x(s)$,
- $F_m(t) = m\,d^2x/dt^2 \Rightarrow f_m(s) = m\,s^2x(s)$.

The plant can be composed of an arbitrarily large number of such components connected such that each connection sums forces or sums positions. Without going into detail, the behavior of such a connected network of components can always be represented in terms of ratios of polynomials P and Q in s.

Applying a force (or motion) $B(s)$ at some place in the plant (cf. Figure 1) will result in a force (or motion) $C(s) = P(s)/Q(s)$ at some other place. It is necessary that $C(s) \to 0$ as $s \to \infty$, otherwise velocities $s\,x(s)$ and accelerations $s^2x(s)$ would also have to become infinite, a clearly unphysical situation. Consequently, the degree of $Q(s)$ must always exceed that of $P(s)$ for a physical system.

Stability of the physical loop

Thus we can write, $A_o(s) = GP(s)/Q(s)$ where G, $P(s)$ and $Q(s)$ are derived from the plant and control gain. Substituting from eq(6) we find potentially unstable conditions when $GP(s)/Q(s) = -1$. The roots s_k of this equation represent k discrete frequencies where $C(s)$ can occur independent of $R(s)$ and $D(s)$. Recall that the real parts σ_k of these roots represent exponentially growing or decaying signals, according to their signs.

All roots having $\sigma < 0$ represent components of $C(t)$ which decay with time, and are, while perhaps in some cases a bother, not considered unstable. Any root with a positive real part however, grows exponentially with time, oscillating at a frequency corresponding to its imaginary part. This represents an unstable condition where the oscillation increases to the maximum level the system can support. In this state, the system is entirely out of control.

As long as $|A_o(s)|$ is large, of course, this condition can be avoided. However because the degree of $Q(s)$ must exceed that of $P(s)$, $A_o(s) \to 0$ as $s \to \infty$. As we have seen, in its useful working range of frequencies $|A_o(s)|$ must be large to be effective. Accordingly we pick a control gain providing $|A_o(s)| \gg 1$ over some prescribed domain in s. Now holding $A_o(s)$ fixed, let us expand the boundaries of the domain. As each point on the boundary is displaced outward, if $A_o(\infty) = 0$, it must encounter at least one value of s where $|A_o(s)| = 1$. Thus, there is (at least one) expanded domain in s having a contour along which $|A_o(s)| = 1$. This cannot be avoided. The phase of $A_o(s)$ of course may vary along this contour, and any location where the phase $\phi = \pi$ radians, we have the condition eq(6) where instability may occur. If this occurs in the right (positive) half the s plane, $\sigma > 0$ and the system is indeed unstable.

So the trick is to select a control gain which provides a sufficient open loop gain over the desired frequency domain, while ensuring the roots of eq(6) remain in the left half s plane by monkeying around with the number and locations of the poles and zeros of eq(6). The synthesis of $A_o(s)$ from physical components is beyond the scope of this report, but I can point out a few useful abstract properties. First, the poles and zeros of A_o are due to zeros of $Q(s)$ and $P(s)$ respectively. The locations of the poles and zeros in the s plane uniquely determine $A_o(s)$ to within a scale factor. Thus, the phase ϕ of $A_o(s)$ depends only on the pole and zero locations. Traversing any closed contour around a zero in s accumulates a phase shift of $360°$; around a pole, it's $-360°$. Thus, a contour encircling an equal number of poles and zeros accumulates no phase shift, although the phase may vary substantially at various locations on the contour.

To help visualize this, two plots of typical open loop gains are shown, both having four poles and no zeros. These are nearly identical in that they have the same open loop gain $A_o(0) = 80$ ("DC gain" in electrical engineering lingo), and the same three poles at $s = -10$ and $s = -800 \pm 800i$. In Figure 5, a fourth pole is located at a relatively fast damping rate of $s = -2000$, while in Figure 6, this pole is shifted to a slower rate of $s = -500$. In these plots, height represents the magnitude, while color represents the phase of $A_o(s)$. A black contour of $x = 0$ is overlaid to distinguish the left and right halves of the s plane. The contour where $|A_o| = 1$ is shown in white, and the contour where $\phi = 180$ degrees is shown in blue. The potentially unstable points occur at intersections of the white and blue contours. As previously mentioned, any of these that occur in the right half s plane are unstable.

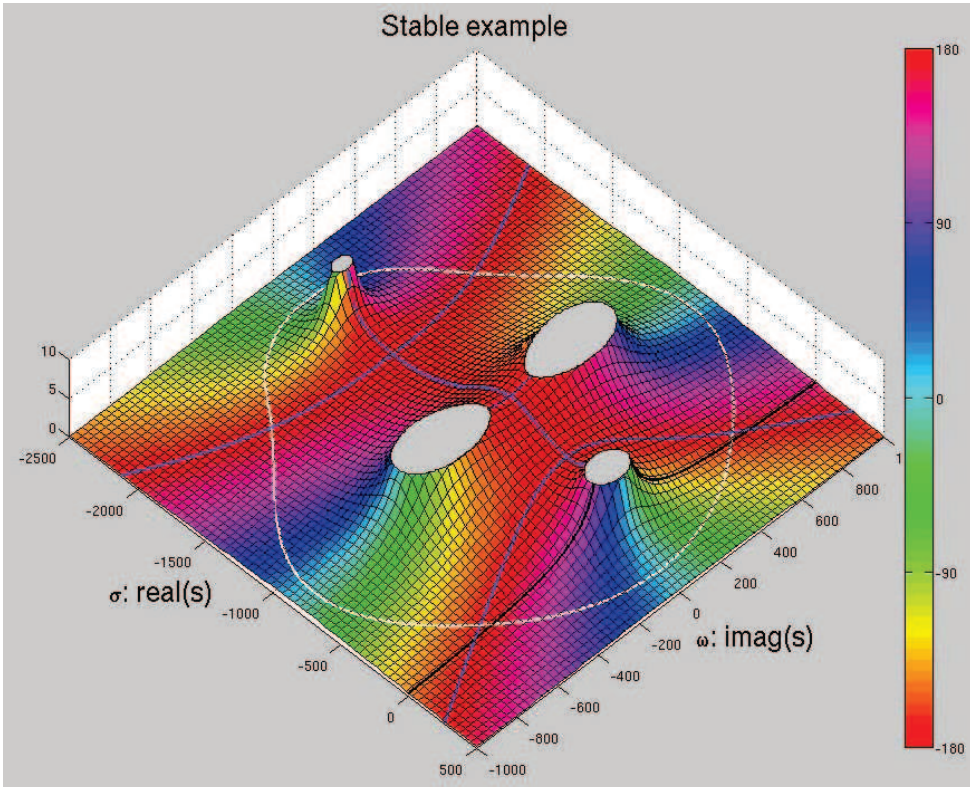

Figure 5: A stable four pole open loop gain function

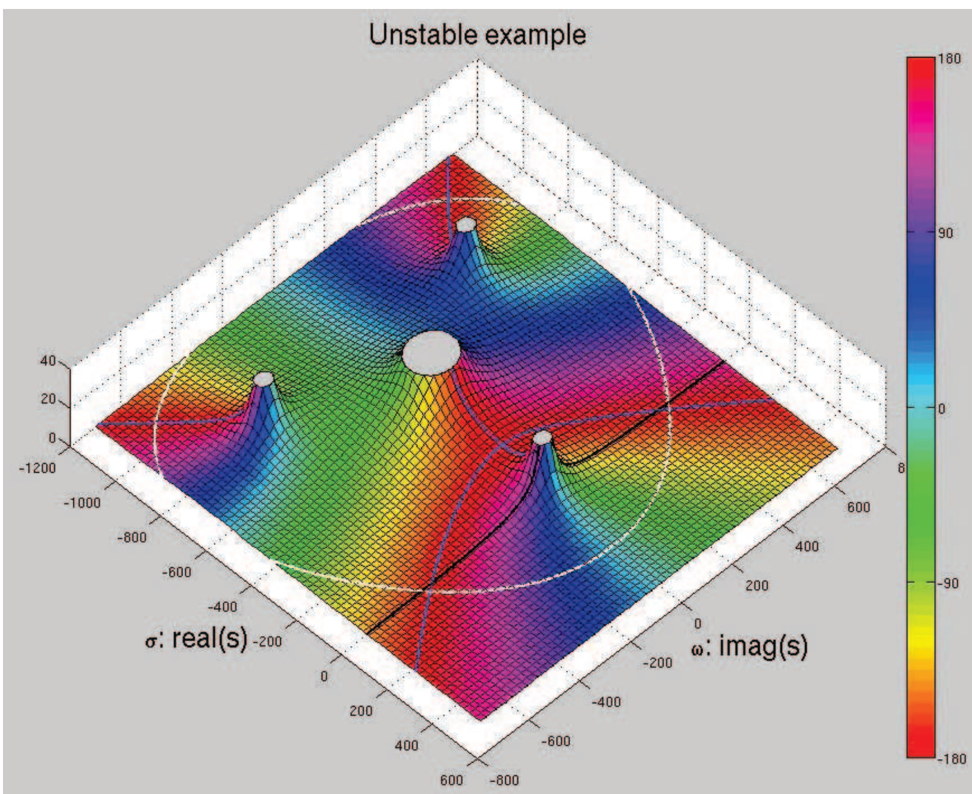

Figure 6: An unstable four pole open loop gain function

As a final remark, let me observe that the mechanical example is not the only physical domain where these principals apply. Lumped constant LTI electrical circuits behave in exactly the same way. In fact, chemical and biological processes that can be expressed in terms of ODEs and composition rules that result in rational polynomial transfer functions will all behave according to these rules. For example, diffusion transport and reaction rates may well be so modeled. If economic and population dynamics could be validly modeled in this fashion, they would also work. The trick in any domain is to find a model that composes in terms of lumped constant LTI structures that can in fact accurately describe the system in which feedback occurs. Cell Biology appears to be a domain where these techniques are not yet widely used. I hope this tutorial may promote the use of these well developed and powerful techniques by Cell Biologists.