An Introduction to Control Theory
From Classical to Quantum Applications

Course Lecture Notes

JM Geremia
Physics and Control & Dynamical Systems
California Institute of Technology

Heidelberger Graduiertenkurse Physik
6 – 10 October 2003
Ruprecht-Karls-Universität Heidelberg
Heidelberg, Deutschland

Abstract

Control theory is a vital component of all modern technology because it allows the construction of high performance devices despite large uncertainty in the individual components used to build them. This course provides a basic introduction to modern control theory in a series of five lectures, each intended to be roughly three hours long. It focusses on the development of fundamental concepts, such as stability, performance and robustness, that lie at the heart of the subject. However, it also addresses practical issues in controller design and presents techniques for constructing and improving control systems typically found in an experimental physics setting.

In the first part of the course (Lectures 1-3) we develop basic definitions in classical dynamical systems and focus on linear feedback control. Experience has shown that a vast majority of physical systems, even those that are nonlinear, can be well controlled using a feedback theory of this type. We explore quantitative descriptions of stability, performance and robustness and use them to develop design criteria for engineering feedback controllers. In Lecture 3, we work through a full example of this loop-shaping design procedure.

The second part of the course (Lectures 4-5) makes a transition to quantum systems. As modern devices grow ever smaller, and due to the realization that quantum mechanics opens avenues for computation and communication that exceed the limits of classical information theories, there has been substantial interest in extending the field of feedback control to include quantum dynamical systems. This process is already being realized in both theoretical and experimental quantum optics for precision measurement and fundamental tests of quantum measurement theories. Forefront research in micro- and nano-scale physics has produced mechanical systems that are rapidly approaching the quantum regime. We will look at the formulation of quantum feedback control theory for continuously observed open quantum systems in a manner that highlights both the similarities and differences between classical and quantum control theory. The final lecture will involve a discussion of special topics in the field and is meant to provide a casual overview of current experiments in quantum control.

Contents

List of Figures

Acknowledgements

These notes would have been completely impossible without countless discussions on the subjects of classical and quantum control that I have had with Andrew Doherty, John Stockton, Michael Armen, Andrew Berglund and Hideo Mabuchi over the course of the past several years while in the Quantum Optics group at Caltech.

Lecture 1

Introduction to Control and Dynamical Systems

The purpose of this course is to provide an introduction to the branch of science and mathematics known as *control theory*, a field that plays a major role in nearly every modern precision device. In the classical engineering world, everything from stereos and computers to chemical manufacturing and aircraft utilizes control theory. In a more natural setting, biological systems, even the smallest single-celled creatures, have evolved intricate, life-sustaining feedback mechanisms in the form of biochemical pathways. In the atomic physics laboratory (closer to home), classical control theory plays a crucial experimental role in stabilizing laser frequencies, temperature, and the length of cavities and interferometers. The new field of experimental quantum feedback control is beginning to improve certain precision measurements to their fundamental quantum noise limit.

Apparently, there are at least two reasons for studying control theory. First, a detailed knowledge of the subject will allow you to construct high precision devices that accurately perform their intended tasks despite large disturbances from the environment. Certainly, modern electronics and engineering systems would be impossible without control theory. In many physics experiments, a good control system can mean the difference between an experimental apparatus that works most of the time and one that hardly functions at all. The second reason is potentially even more exciting— it is to develop general mathematical tools for describing and understanding dynamical systems. For example, the mathematics that has been developed for engineering applications, such as flying airplanes, is now being used to quantitatively analyze the behavior of complex systems including metabolic biochemical pathways, the internet, chemical reactions, and even fundamental questions in quantum mechanics.

There is substantial interest in extending the concepts from classical control theory to settings governed by quantum mechanics. Quantum control is particularly exciting because it satisfies both of the reasons for studying control theory. First, there is a good motivation for controlling quantum systems. As electronic and mechanical technologies grow ever smaller, we are rapidly approaching the point where quantum effects will need to be addressed; quantum control will probably be necessary in ordrer to implement quantum computation and communication in a robust manner. On the other hand, control theory provides a new perspective for testing modern theories of quantum measurement, particularly in the setting of continuous observation. Many of

the mathematical tools used for analyzing classical dynamical systems can be applied to quantum mechanics in a manner that is complementary to the traditional tools available from physics. The mathematical analysis in control theory has tended to be more rigorous than typical approaches in physics, and quantum control provides a fantastic setting for demonstrating the power of these tools.

The defining characteristic of a quantum system, from the perspective of control theory, is that measurement causes its state to change. This property is referred to as *backaction*, an inevitable consequence of the Heisenberg uncertainty principle. In the classical world, it is (in principle) possible to perform a non-invasive measurement of the system as part of the control process. The controller continuously adjusts the system dynamics based on the outcome of these non-invasive measurements. However, in quantum mechanics, the very act of observation causes the state of the system to wander even farther away from the control target. In this sense, the effect of backaction is to introduce a type of quantum noise into the measurement of the system's state. Fortuately, it is possible to think of quantum control in a context that is similar to classical control, but with an additional noise source due to the measurement backaction. Therefore, it is theoretically possible to adapt existing control knowledge to make it apply in quantum settings, although there are some mathematical surprises since (unlike classical dynamics) continuous observation of quantum systems prevents them from achieving steady-state.

For this reason, our course begins with a description of classical systems, particularly that of linear feedback control. This is not only practical, since it will allow you to construct servos and controllers for your common laboratory experiments (like locking cavity mirrors and laser frequencies), it is also pedagogical. Understanding the particular details which make quantum systems different from classical systems requires that we first look at classical control theory. Additionally, in my experience, there is a reasonable amount of misconception regarding servos and simple feedback systems in the experimental physics community. In part, most of this is likely due to the fact that one has more interesting problems to deal with when constructing a complicated experimental apparatus. Stabilizing some boring classical detail such as a cavity length in an experiment to do something grand, like generate single photons, is often viewed as a necessary evil. Often, this type of feedback is implemented in a mediocre way rather than with a custom tailored, robust feedback loop. I assure you, the latter will always perform better!

1.1 State Space Models and Fundamental Concepts

In order to proceed any farther, we require a general mathematical description of dynamics.[1] This section develops a quantitative model of physical systems and provides the cornerstone upon which the rest of these lectures are constructed. The most important abstraction we make is that of an *input–output system model*, a concept that is depicted by Fig. 1.1. The entire dynamical system is schematically reduced to a single box, labelled by $\mathbf{x}(t)$. There is a control input, $u(t)$, and an output, $y(t)$, that reflects the response of the system's internal dynamics to the input.

1.1.1 State-Space Models

Every dynamical system is described by a collection of time-dependent variables, $\{x_1(t), x_2(t), \ldots, x_n(t)\}$, that we will group into a vector, $\mathbf{x}(t) \in \mathbb{R}^n$. Physically speaking, $\mathbf{x}(t)$ describes the *state* of the system, and it reflects all the information that the system has available to itself in order to propagate forward in time. In other words, the future state of the system for, $t > 0$, can be determined from the initial conditions, $\mathbf{x}_0 = \mathbf{x}(t = 0)$, and an equation of motion,

$$\frac{d}{dt}\mathbf{x}(t) = \mathbf{a}\left(\mathbf{x}_{[0,t)}, t\right), \quad \mathbf{a} : \mathbb{R}^n_{[0,t)} \times t \to \mathbb{R}^n . \tag{1.1}$$

This model says that the rate of change of the system variables at time, t, can be determined by some function, $\mathbf{a}$, of the system's evolution, $\mathbf{x}_{[0,t)}$, and the current value of t. Here, the symbol, $\mathbf{x}_{[0,t)}$, is shorthand notation for $\{\mathbf{x}(t'),\ 0 \leq t' < t\}$, i.e. it is the system history starting at time, $t = 0$, and ending at t. The output of the vector function, $\mathbf{a}$, necessarily has dimension n so that the number of variables in $\mathbf{x}(t)$ is constant in time.

The state-space model, given by Eq. (1.1), reflects information that in principle is only known to the system; it reflects all the physics. In general, we as an external observer can only gain information about the system by making measurements. In the general case, it is impossible to acquire complete information about the system from the measurement, meaning that we are not always privy to the values of $\mathbf{x}(t)$. Instead, we have a *partially observed system*,[2,3]

$$\mathbf{y}(t) = \mathbf{c}(\mathbf{x}(t), t), \quad \mathbf{c} : \mathbb{R}^n \times t \to \mathbb{R}^k . \tag{1.2}$$

$\mathbf{c}(\mathbf{x}(t), t)$ is some function of the state of the system at time, t, and the dimension of $\mathbf{y}(t)$ is k, which may be different from n. Physically speaking, $\mathbf{y}$, is the information about the system that is available to an observer who has performed some experiment to probe the state of the system at time, t.

At this point, we have developed, via Eqs. (1.1) and (1.2), an extremely general mathematical description of our dynamical system. The evolution and the observation processes are free to be completely nonlinear if necessary. The restriction to real-valued functions is not at all confining because complex-valued functions, $\mathbf{x}(t) \in \mathbb{C}^n$, can be redefined in terms of a real state vector, $\mathbf{x}'(t) = \{\mathrm{Re}[\mathbf{x}(t)], \mathrm{Im}[\mathbf{x}(t)]\} \in \mathbb{R}^{2n}$.

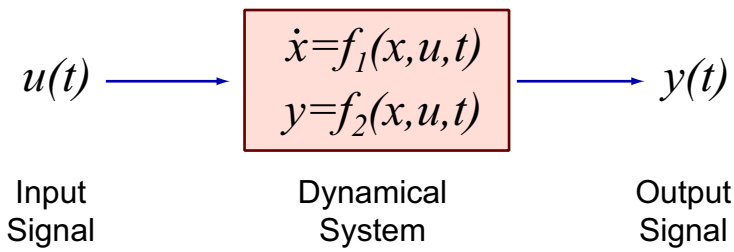

Figure 1.1: *Input-output model of the control system, which acts as a transducer that takes an input signal, $u(t)$, and imprints its internal dynamics onto an output signal, $y(t)$.*

There are, however, two important restrictions imposed by the structure of Eqs. (1.1) and (1.2). The first is that we have assumed a continuous time model (what one would expect for analog systems). In discrete-time (digital) systems, the only difference is that the evolution equations take the form of delay-update equations,

$$\mathbf{x}(t_i) = \mathbf{a}\left(\mathbf{x}_{[t_0,\ldots,t_{i-1}]}, t_i\right)\Delta t + \mathbf{x}(t_{i-1}) \tag{1.3}$$

$$\mathbf{y}(t_i) = \mathbf{c}\left(\mathbf{x}(t_i), t_i\right) \tag{1.4}$$

where $t_i = t_{i-1} + \Delta t$, rather than differentials. In principle, there is little qualitative difference between the continuous and discrete evolution equations. For the purpose of this course, I will tend to focus on continuous time models because even when designing a digital controller, it is common to perform the design in continuous time and then discretize the process as a second step. Occasionally, I will discuss digital systems when there is a particular point that is not obvious by simply replacing the continuous system dynamics with a discrete evolution model. There will be more on this Lecture 3.

The second simplification is that the dynamics are deterministic, in the sense that there are no terms reflecting random processes, such as environmental disturbances or measurement noise. For the time being, we will neglect these things. However, we will certainly return to them a little later in the course when we address quantum control systems. As can be expected, quantum control *must* adopt a stochastic, probabilistic model of the dynamics since $\mathbf{x}$ will be replaced with a representation of state, such as the wavefunction or density operator, and these are probabilistic quantities. Additionally, the measurement process will inherently introduce quantum fluctuations into the dynamics in the form of backaction projection noise.

1.1.2 Controlled Dynamics

The evolution generated by Eq. (1.1) represents the free evolution of the system; there is no external forcing. We need to add a term in the system evolution that will allow us to affect the dynamics by applying an input signal, $u(t)$. Following a general approach, as in the above section, this gives us a new equation,

$$\frac{d}{dt}\mathbf{x}(t) = \mathbf{a}\left(\mathbf{x}_{[0,t)}, t\right) + \mathbf{b}\left(\mathbf{x}_{[0,t)}, \mathbf{u}_{[0,t)}, t\right) \tag{1.5}$$

where $\mathbf{b}\left(\mathbf{x}_{[0,t)}, \mathbf{u}_{[0,t)}, t\right)$ can be thought of as the coupling term that describes the interaction between the input signal, $\mathbf{u}$, and the system. The dimension of $\mathbf{u}$ is free to be anything, which we will call k, but the output of the function $\mathbf{b}$ must have dimension, n in order to be consistent with the system dynamics. Similarly, the observation process, $\mathbf{y}(t)$, may be affected by the input signal, so that generally speaking,

$$\mathbf{y}(t) = \mathbf{c}\left(\mathbf{x}(t), \mathbf{u}(t), t\right) . \tag{1.6}$$

1.1.3 Linear Time-Invariant Systems

The equations above are extremely general. In some sense that is good because it allows for a theory that includes most possible input signals as well as the majority of physically realistic dynamical systems. The drawback is that with such generality, it is extremely difficult to perform any significant analytical analysis. Nonlinear functions cannot generally be solved exactly, and we would then have to resort to treating every new control system as a unique problem. Instead, it is beneficial to restrict the form of the functions, $\mathbf{a}$, $\mathbf{b}$, and $\mathbf{c}$ such that they still (approximately) reflect a large number of physical systems, but can be exactly solved for the entire class.

Such a restriction is provided by the class of dynamical equations referred to as *linear time-invariant* (LTI) models.[4–6] This approximation corresponds to treating the functions, $\mathbf{a}$, $\mathbf{b}$, and $\mathbf{c}$, as linear and time-stationary,

$$\frac{d}{dt}\mathbf{x}(t) = \mathbf{A}\mathbf{x}(t) + \mathbf{B}\mathbf{u}(t) \tag{1.7}$$

$$\mathbf{y}(t) = \mathbf{C}\mathbf{x}(t) + \mathbf{D}\mathbf{u}(t) \tag{1.8}$$

where $\mathbf{A} \in \mathbb{R}^{n\times n}$, $\mathbf{B} \in \mathbb{R}^{k\times n}$, $\mathbf{C} \in \mathbb{R}^{n\times m}$ and $\mathbf{D} \in \mathbb{R}^{k\times m}$. In other words, the functions have been replaced by time-invariant matrices.

The best thing about Eqs. (1.7) and (1.8) is that they can be solved exactly for all reasonable choices of $\mathbf{A}$ and $\mathbf{B}$. In fact, doing so only requires basic methods from the field of ordinary differential equations,

$$\mathbf{x}(t) = e^{\mathbf{A}t}\mathbf{x}_0 + e^{\mathbf{A}t}\int_0^t e^{-\mathbf{A}\tau}\mathbf{B}\mathbf{u}(\tau)\,d\tau\,. \tag{1.9}$$

We see that the system state is a superposition of a complimentary and a particular solution. The former is given by $\exp(\mathbf{A}t)\mathbf{x}_0$ and depends upon the initial condition, $\mathbf{x}_0$, while the particular solution is given by a convolution over the input signal. It contains all the physics of the interaction between the input signal and the internal system dynamics. Eq. (1.9) can be substituted in Eq. (1.8) in order to obtain an analytic expression for $\mathbf{y}(t)$.

At this point, we have to pause for a minute in order to justify the huge approximation that we made in reducing the system dynamics to a linear model. One argument for the approximation in Eqs. (1.7) and (1.8) is that these equations are locally true even for nonlinear dynamics. This can be seen from the fact that they reflect a first order Taylor expansion of the full equations in Eqs. (1.5) and (1.6). So, even if the full system dynamics are nonlinear, these linear equations will be true in a neighborhood surrounding each point, $\mathbf{x} \in \mathbb{R}^n$. Of course, this raises the question of how effective a controller design can be if the model use to construct it is only locally applicable. There are two possible responses to this criticism. The first one is simply to state the fact that engineering control theory survived for 80 years before it even attempted to go beyond LTI models. In that time, it was able to successfully produce controllers for highly nonlinear devices such as airplanes and space missions that made it to the moon and back. In that sense, the justification comes in the form of a proof by demonstration. More generally, the answer to this problem is that once the controller is doing its job, it will generally act to keep the system in a region of its phase-space where the linear approximation is good. Many examples where this is true can be

found in standard control theory texts.[7] You will have to take my word for this at the moment, but before long, the point will become clear from the mathematics.

1.1.4 Definition of the Control Problem

By applying a forcing term, $\mathbf{u}(t)$, it is possible to drive the system such that it evolves according to trajectories that depart from the natural dynamics. In this sense, the function, $\mathbf{u}(t)$, steers the outcome of the evolution. Everything in control theory involves selecting an appropriate signal, $\mathbf{u}(t)$, such that the dynamical system evolves in a specified manner. For instance, a control goal might be to choose a time, $t = T$, and a target value, $y * t = y(T)$, and identify a $\mathbf{u}^*(t)$ that achieves this outcome. Another objective is for the system output, $y(t)$, to track a reference signal, $\mathbf{r}(t)$, by minimizing

$$\|\mathbf{y}(t) - \mathbf{r}(t)\| \,. \tag{1.10}$$

The signal norm, $\|\cdots\|$ is some appropriate measure of the difference between the two signals (refer to the Lecture 1 Appendix for a review of norms).

There are generally a large number of signals, $\{\mathrm{u}_1^*(t), \mathbf{u}_2^*(t), \ldots\}$, that achieve the target objective, and each of these is referred to as a *permissable* control. However, applying $\mathbf{u}(t)$ involves expending energy in some sense or another, and it is desirable to select the $\mathbf{u}_i^*(t)$ that requires the least effort to reach the objective. This can be accomplished by adding an optimality component to the control problem by minimizing over a signal norm,

$$\mathbf{u}^\circ(t) = \min_{\mathbf{u}^*(t)} \|\mathbf{u}^*(t)\|_2 = \min_{\mathbf{u}^*(t)} \mathcal{J}[u(t)] \,. \tag{1.11}$$

The 2-norm is most commonly used in this cost functional because it is closely associated with the energy expenditure.[5] All control problems are intrinsically an optimization.

1.2 Classification of Control Systems

We have already seen that control systems can be classified according to their dynamics, as either continuous or discrete and linear or nonlinear. Another important classification involves the method that is used to determine the optimal control signal, $\mathbf{u}^\circ(t)$. Roughly speaking, all controllers can be grouped into one of three classes.

- **Open-loop Control**: In open loop control, it is assumed that the dynamical model of the system is well known, that there is little or no environmental noise and that the control signal can be applied with high precision. This approach is generally utilized when there is a target value, $y^*(T)$, to achieve at a particular final time, T. Under these assumptions, it is possible to solve the dynamical equations of motion and directly perform the optimization in Eq. (1.11). The disadvantage of open-loop control is that the performance of the controller is highly susceptible to any unanticipated disturbances. Additionally, if there is any error in

the model used to describe the system dynamics, the designed control signal is unlikely to perform as expected.

- **Feedback Control**: In feedback control, continuous or discrete time measurements of the system output, $\mathbf{y}(t)$, are used to adjust the control signal in real time. At each instant, the observed process, $\mathbf{y}$ is compared to a tracking reference, $\mathbf{r}(t)$, and used to generate an error signal. The feedback controller interprets this error signal and uses it to generate a $\mathbf{u}^{\circ}(t)$ that maintains tracking. Feedback controllers have the potential to be adaptive and highly effective even if there is large uncertainty in the system model or if there are environmental disturbances or measurement noise. Feedback therefore provides the backbone of most modern control applications.
- **Learning Control**: In learning control, a measurement of the system, $\mathbf{y}(t)$, is also used to design the optimal feedback signal; however, it is not done in real time. Instead, a large number of trial control signals are tested in advance, and the one that performs best is selected to be $\mathbf{u}^{\circ}(t)$. Learning control has a number of disadvantages compared to feedback, because it does not permit as detailed a quantitative analysis of the controller performance and stability. However, it is useful for applications where the dynamics occur too fast for the system to be actuated in a feedback configuration.

1.3 Control Performance and Stability

Regardless of the controller that is used, there are three major objectives that must be considered when assessing the quality of the feedback system.

1. **Performance** is the ability of the controller to realize its dynamical objective, such as to achieve the target value, $y^*(T)$, or to follow a reference signal, $r(t)$.
2. **Stability** is a measure of how well-behaved the control system is. When feedback is utilized, there is the potential for instability. If caution is not taken when designing the feedback system, parasitic signals or environmental disturbances might cause the entire dynamical system to oscillate out of control and potentially experience physical damage.
3. **Robustness** is the ability of the controller to remain effective despite unanticipated changes in either the system being controlled or the environment.

In general, realistic controllers must achieve a balance between performance, stability and robustness. The particular objective that is most important is application dependent. For example, in an airplane, robustness and stability take higher priority than performance. This is because it is preferable for a plane to go slightly off course than it is for it to become unstable and potentially crash. However, in precision measurement, the opposite is true, and it is better for a laser or an interferometer to achieve high performance at the expense of having to re-lock it every few days, or when someone drops a screwdriver on the optical table.

1.3.1 Example 1: Open-Loop Control

In order to demonstrate these points, it is best to consider several examples of controller design. In this first case, we utilize open-loop control to illustrate the concepts of performance and robustness (this following example, and similar variations of it, can be found in most standard discussions of open-loop control theory). Consider the linear control system governed by the following one-dimensional state-space dynamical model,

$$\dot{x}(t) = ax(t) + bu(t) \tag{1.12}$$
$$y(t) = cx(t) \tag{1.13}$$

where $a, b, c \in \mathbb{R}$ and the objective is to achieve the value, $y(T) = y^*$ at time, $t = T$. This system is a specialization of the linear time invariant (LTI) model described in Section 1.1.3. Its solution is given by Eq. (1.9),

$$y(t) = c\left(e^{at}x_0 + be^{at}\int_0^t e^{-a\tau}u(\tau)\,d\tau\right). \tag{1.14}$$

By requiring that $u(t)$ is an admissible control function, we can express the control objective as follows,

$$\frac{y^*}{c} = e^{aT}x_0 + be^{aT}\int_0^T e^{-a\tau}u(\tau)\,d\tau \tag{1.15}$$

which can be rearranged to give,

$$\int_0^T e^{-a\tau}u(\tau)\,d\tau = \frac{c^{-1}y^* - e^{aT}x_0}{be^{aT}}. \tag{1.16}$$

Now we need to solve the minimization problem,

$$u^\circ(t) = \min_{u(t)} \mathcal{J}[u(t)] = \min_{u(t)} \int u^2(t)\,dt \tag{1.17}$$

which can be accomplished by squaring Eq. (1.16) and utilizing the Cauchy-Schwarz inequality,

$$\left[\int_0^T e^{-a\tau}u(\tau)\,d\tau\right]^2 \leq \int_0^T u^2(\tau)\,d\tau \int_0^T e^{-2a\tau'}\,d\tau' \rightarrow \mathcal{J}[u(t)]\int_0^T e^{-2a\tau}\,d\tau \tag{1.18}$$

and therefore,

$$\mathcal{J}[u(t)] \geq \frac{\left(c^{-1}y^* - e^{aT}x_0\right)^2}{b^2e^{2aT}\int_0^T e^{-2a\tau}\,d\tau}. \tag{1.19}$$

In order for the minimum to be achieved, we need to find a $u^*(t)$ that satisfies the equality in Eq. (1.19),

$$\mathcal{J}[u(t)] = 2a\frac{\left(c^{-1}y^* - e^{aT}x_0\right)^2}{b^2(e^{2aT}-1)} \tag{1.20}$$

which (this is the tricky part of the problem) happens if $u(t)$ is proportional to $\exp(-at)$,

$$u^*(t) = Ce^{-at} \tag{1.21}$$

to give,

$$\left[\int_0^T e^{-at} u*(t)\,dt\right]^2 = C^2\left[\int_0^T e^{-2at}\,dt\right]^2 . \tag{1.22}$$

Using Eq. (1.16), it can be shown that,

$$C = \frac{y^* c^{-1} - e^{aT}x_0}{be^{aT}\int_0^T e^{-2at}\,dt} \tag{1.23}$$

and finally that,

$$u^\circ(t) = 2a\frac{y^* c^{-1} - e^{aT}x_0}{be^{aT}(1-e^{-2aT})}e^{-at} \tag{1.24}$$

So it seems that everything worked out fine. We were able to find a solution that minimizes $\mathcal{J}[u(t)]$ and that achieves the control objective. However, what if the value of a, b or c was slightly wrong? In order to address this question, we should ideally evaluate the sensitivity of the control outcome. This could be accomplished by taking the derivatives, $\partial y(t)/\partial a|_{t=T}$ and so on. It would be found that the control outcome is exponentially sensitive to some of the control parameters, meaning that errors in the system model could produce an exponentially large tracking error! This method of controller design is exceptionally non-robust.

1.3.2 Example 2: Feedback and Stability

For the second example, we will try to design a feedback controller to stabilize the following system,

$$\frac{d}{dt}\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} 0 & \omega \\ -\omega & 0 \end{pmatrix}\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} + u(t)\begin{pmatrix} 0 \\ 1 \end{pmatrix} \tag{1.25}$$

$$y(t) = \begin{pmatrix} 1 & 0 \end{pmatrix}\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} \tag{1.26}$$

subject to the initial conditions that $\mathbf{x}_0 = \begin{pmatrix} 1 & 0 \end{pmatrix}^T$.

Our goal is to make $y(t)$ track the reference signal, $r(t) = 0$. Granted, we do not yet know anything about feedback control, but let us proceed in the dark and assume a controller that applies a proportional feedback signal opposite the measured displacement,

$$u(t) = -C(y(t) - r(t)) = -Cy(t) = -Cx_1(t)\,. \tag{1.27}$$

This seems like a very reasonable feedback control law. At each time, we measure how far the system is from the desired state and then apply a push in the right direction that is stronger if we are farther away from the target and gets smaller as we approach it. The dynamical system, using the choice, $u(t) = -Cy(t)$ can be solved analytically (it is an LTI model) to give the following equation of motion,

$$y(t) = \cosh\left(\sqrt{-\omega(C+\omega)}t\right) \tag{1.28}$$

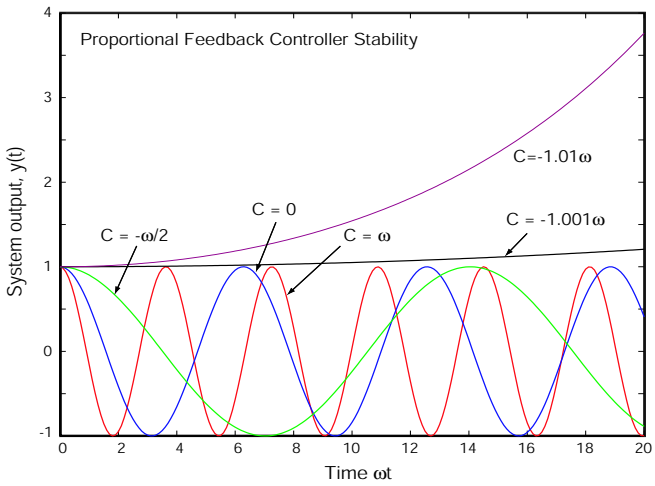

FIGURE 1.2: *Solution to the controlled dynamical system in Example 2 for different values of C. As can be seen, the feedback system provides poor tracking of the desired reference signal, $r(t) = 0$.*

which shows that $y(t)$ depends very strongly on the choice of C. If $C > 0$, then the system oscillates with a frequency that is somewhat larger than ω. However, if $C < -\omega$, then the system blows up, as $t \to \infty$, y exponentially increases. Fig. 1.2 shows a plot of $y(t)$ for several different choices of C, and highlights the undamped behavior. Here, C acts as a *gain*, and the hope would be that some judicious choice of C would result in stable feedback. However, this is clearly not the case. Making $C = 0$ corresponds to no feedback, and increasing $C > 0$ only makes matters worse by increasing the oscillation frequency. Certainly, $C < 0$ is

Well, our first attempt at feedback control was not very successful despite the fact that we did something entirely reasonable— something that a perfectly well-educated physicist might have tried (in fact, something I mistakenly did myself) in the laboratory. This exercise demonstrates that stability is not a guarantee when working with feedback, it must be specifically engineered, and that will be the subject of tomorrow's lecture.

1.4 LECTURE 1 APPENDICES

1.4.1 SIGNAL NORMS

A norm is an mathematical measure that provides a generalized measure of the magnitude of its argument, and therefore provides a way of ranking the relative "size" of different functions. Mathematically, in order for a measure, $\|\cdots\|$, to be considered a norm, it must satisfy the following requirements:

(i) $\|u\| \geq 0, \forall u(t)$ which is the requirement that norms be non-negative

(ii) $\|u\| = 0$ iff $u(t) = 0, \forall t$ which is the minimum bound requirement

(iii) $\|au\| = a\|u\|, \forall a \in \mathbb{R}$ which is the distributive property for scalars

(iv) $\|u + v\| \leq \|u\| + \|v\|$ which is the so-called triangle inequality

A set of measures that satisfy these requirements and that are regularly encountered in the mathematical analysis used in control theory are the n-norms, defined by

$$\|u\|_n \equiv \left(\int_{-\infty}^{\infty} |u(t)|^n \, dt \right)^{\frac{1}{n}}, \tag{1.29}$$

where particulary important n-norms include the 2-norm (or energy norm),

$$\|u\|_2 = \left(\int |u(t)|^2 \, dt \right)^{\frac{1}{2}} \tag{1.30}$$

and the ∞-norm (or maximum value norm),

$$\|u\|_\infty = \sup_t |u(t)| \, . \tag{1.31}$$

When the argument is a vector function, one can apply the norm to each component of the vector and then take any combination of those individual norms that satisfies the 4 requirements listed above. For example,

$$\|\mathbf{f}(t)\|_n \equiv \left(\sum_i \|f_i(t)\|_n^n \right)^{\frac{1}{n}} \tag{1.32}$$

1.4.2 A Note on References

The material covered in this lecture reflects topics that are addressed in a number of standard texts on the subject of control theory. Anyone looking to learn more about the field is is immediately confronted with a number of possible sources and references. Texts that I have found to be particularly useful are listed here,

1. Feedback Control Systems[8] by Phillips and Harbor is a very good general introduction the field of feedback control. It has a particularly accessible description of the Nyquist stability criterion.

2. Introduction to Control Theory[4] by Jacobs is a good introduction to the field of control theory. It discusses a large number of topics without going heavily into detail and provides a good starting point.

3. Feedback Control Theory[5] by Doyle, Francis and Tannenbaum is a fairly readable introduction to the field of robust control, but not a general introduction to feedback (as the title implies). This text requires only moderate background in mathematical analysis.

4. Essentials of Robust Control[9] by Zhou and Doyle is a somewhat more involved discussion of robust control theory and assumes a stronger background in mathematical analysis.

5. Modern Control Systems[7] by Dorf and Bishop is a good general introduction that discusses digital feedback controller design in reasonable detail.

6. Digital Signal Processing[10] by Proakis and Manolakis is the standard text on the subject of designing filters in discrete time systems and is an invaluable resource to anyone interested in digital control.

References

[1] L. Perko, *Differential Equations and Dynamical Systems* (Springer-Verlag, New York, 2001), 3rd ed.

[2] G. Chen, G. Chen, and S.-H. Hsu, *Linear Stochastic Control Theory* (CRC Press, London, 1995).

[3] B. Oksendal, *Stochastic Differential Equations* (Springer Verlag, 1998), 5th ed.

[4] O. L. R. Jacobs, *Introduction to Control Theory* (Oxford University Press, New York, 1996), 2nd ed.

[5] J. Doyle, B. Francis, and A. Tannenbaum, *Feedback Control Theory* (Macmillan Publishing Co., 1990).

[6] G. E. Dullerud and F. Paganini, *A Course in Robust Control Theory: A Convex Approach* (Springer Verlag, New York, 2000), 1st ed.

[7] R. C. Dorf and R. H. Bishop, *Modern Control Systems* (Prentice Hall, Upper Saddle River, NJ, 1991), 9th ed.

[8] C. L. Phillips and R. D. Harbor, *Feedback Control Systems* (Prentice Hall, Englewood Cliffs, 1996), 3rd ed.

[9] K. Zhou and J. C. Doyle, *Essentials of Robust Control* (Prentice-Hall, Inc., New Jersey, 1997), 1st ed.

[10] J. G. Proakis and D. G. Manolakis, *Digital Signal Processing: Principles, Algorithms and Applications* (Prentice Hall, Upper Saddle River, 1996), 3rd ed.

Lecture 2

Feedback and Stability

Yesterday, we attempted to design a feedback controller and were, unfortunately, not particularly successful. The problem we encountered was instability— depending on the controller gain, C, the output, $y(t)$, either oscillated or diverged to infinity. For one special case, $C = -\omega$, the value of $y(t)$ remained constant at $y(t) = 1$ for all $t > 0$. This value for the controller gain was (in some sense) more stable, but it failed to achieve tracking since the goal was $y(t) = 0$.

What went wrong? In order to uncover the problem, we can rewrite the dynamical equations from Eqs. (1.25) and (1.26) in a more convenient form. Making the change of variables, $z = x_2$, we see that the dynamics are effectively a second order differential equation (there is only a little algebra involved here)

$$\ddot{z}(t) = -\omega(C+\omega)z(t), \quad z_(0) = 1 \tag{2.1}$$

$$y(t) = z(t)\,. \tag{2.2}$$

Suddenly, the reason for the instability is clear. We have a perfect harmonic oscillator with no damping,

$$y(t) = \cos(\sqrt{\omega(C+\omega)}t) \tag{2.3}$$

when $C + \omega \geq 0$. Otherwise, we have exponential growth of the signal,

$$y(t) = \frac{1}{2}\left(e^{\sqrt{\omega(C+\omega)}t} + e^{-\sqrt{\omega(C+\omega)}t}\right) \tag{2.4}$$

for $C + \omega < 0$.

Physically speaking, this feedback controller operates in the following manner: it looks at the displacement, $y(t) - r(t)$, and then applies a torque to counteract the motion— it tries to actuate $\ddot{y}(t)$. As a result, feedback provides no damping of the system's internal dynamics. With no feedback, $C = 0$, the system is an oscillator, $y(t) = \cos(\omega t)$, and the addition of the controller only changes the oscillation frequency, not the qualitative behavior. In retrospect, we should have designed a controller that was proportional the *derivative* of the displacement,

$$u(t) = C\frac{d}{dt}\left[y(t) - r(t)\right] \tag{2.5}$$

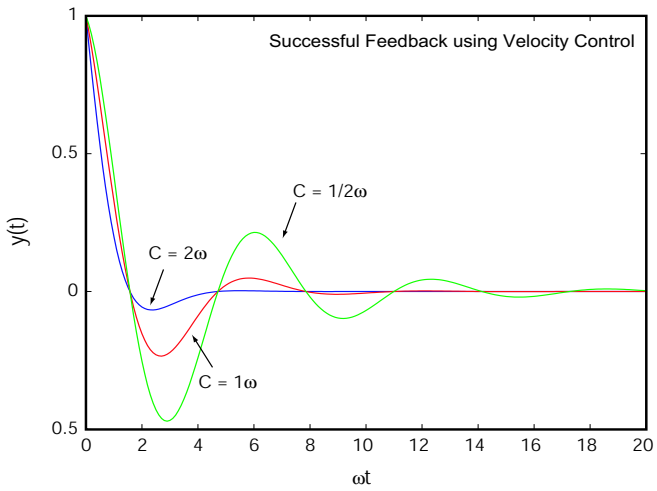

FIGURE 2.1: *Solution to the controlled dynamical system using velocity control. $y(t)$ is plotted for several difference choices for the controller gain.*

because this would produce an equation of motion for the feedback system,

$$\ddot{y}(t) = -\omega^2 y(t) - C\dot{y}(t) \tag{2.6}$$

which is solved by a damped sinusoid when $C > 0$,

$$y(t) = e^{-Ct/2\omega} \cos(\omega t)\,. \tag{2.7}$$

Now, $y(t)$ decays exponentially with time, and therefore rapidly approaches the tracking objective, $r(t) = 0$. Increasing the controller gain, C, increases the damping coefficient, $\Gamma = C/2\omega$, and helps the feedback achieve its objective even faster. This is illustrated in Fig. 2.1, where $y(t)$ is plotted as a function of time for several different choices of the controller gain.

The essential point is that the feedback controller needed to incorporate information about $y(t)$ from earlier times, not just t, in forming $u(t)$ (Eq. (2.6) does this through the derivative term which involves $y(t)$ at multiple times). In general, the controller signal should involve a convolution over the past values of $y(\tau)$, from $\tau = 0, \ldots, t$,

$$u(t) = \int_0^t F(t-\tau) y(\tau)\, d\tau\,. \tag{2.8}$$

Everything in feedback controller design involves engineering $F(t)$. By the end of this lecture, we will have a better idea of how to do this by looking at the requirements that $F(t)$ must satisfy in order to achieve both stability and performance.

Yesterday, we learned that the control problem should also satisfy a minimization over the cost of the driving term, $\mathbf{u}(t)$. This optimization is necessary to prevent the controller from utilizing signals that expend excess energy, or that are infinite. Feedback control provides no exception to this requirement. Although we will not consider the control signal minimization problem in detail, feedback control of the form in Eq. (2.8) satisfies a cost functional minimization analogous to Eq.

(1.11),

$$\mathcal{J}[\mathbf{u}(t)] \equiv \int_0^T \left(\mathbf{x}^T(t)\mathbf{P}\mathbf{x}(t) + \mathbf{u}^T(t)\mathbf{Q}\mathbf{u}(t)\right)\, dt \tag{2.9}$$

for appropriate choices of the matrices, $\mathbf{P}$ and $\mathbf{Q}$ that we will not derive here. For our purposes, it is sufficient to state that feedback control does satisfy all the necessary requirements for signal energy minimization. More information on deriving feedback control in terms of a cost functional optimization can be found in standard texts on the subject.[1]

2.1 Stability

While the harmonic motion in our first attempt at a feedback controller was better than exponential growth, it certainly failed to achieve tracking because it oscillated around the target indefinitely. In control theory, the term *unstable* refers to a system that diverges for long time. This includes both exponential growth and undamped harmonic oscillation, although the latter is typically given the distinction of *marginal stability* because $y(t \to \infty)$ remains finite, but not zero.

When discussing stability in a feedback control system, there are two questions to address. The first has to do with the intrinsic dynamics of the system (without any feedback), and the second involves the effect of closing the loop.

- **Internal Stability** refers to the dynamical system, $\dot{\mathbf{x}}(t) = \mathbf{A}\mathbf{x}(t)$, without feedback. To assess internal stability, we ask whether the output, $\mathbf{y}(t)$, experiences oscillations, or exponential growth.
- **Closed-loop Stability** refers to the behavior of the system when feedback is applied. To assess this question, we look for properties of the feedback loop that lead to oscillation or exponential growth in closed-loop. Feedback loops can be unstable even when its components are internally stable, as well as the opposite.

A quantitative analysis of these types of stability is an essential aspect of control theory, and it requires the development of several mathematical concepts.

2.1.1 Internal Stability

Let us return to our feedback system but focus only on the internal dynamics (those involving $\mathbf{x}$ and not $\mathbf{u}$),

$$\dot{\mathbf{x}}(t) = \mathbf{A}\mathbf{x}(t) = \begin{pmatrix} 0 & \omega \\ -\omega & 0 \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \end{pmatrix}. \tag{2.10}$$

Since this system is linear, we know that the solution is given by the form, $\mathbf{x}(t) = \exp(\mathbf{A}t)\mathbf{x}_0$. For all reasonable system matrices, we can diagonalize $\mathbf{A}$,

$$\mathbf{A} = \mathbf{O}\Lambda\mathbf{O}^\dagger = \mathbf{O}\begin{pmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_n \end{pmatrix}\mathbf{O}^\dagger \tag{2.11}$$

where $\mathbf{O}$ is normally orthogonal or unitary. $\mathbf{O}$ provides a similarity transform that reduces $\mathbf{A}$ to a diagonal matrix, Λ, whose elements are the eigenvalues, $\{\lambda_1, \ldots, \lambda_n\}$ of $\mathbf{A}$. With this transformation, it is possible to evaluate the matrix exponential,

$$e^{\mathbf{A}t} = \exp(\mathbf{O}\Lambda\mathbf{O}^\dagger t) = \mathbf{O}\begin{pmatrix} e^{\lambda_1 t} & & \\ & \ddots & \\ & & e^{\lambda_n t} \end{pmatrix}\mathbf{O}^\dagger \tag{2.12}$$

which shows that the internal dynamics are governed by a superposition of the eigenvalues exponentials (which are in general complex), $\lambda_i = \sigma_i + j\omega_i$, $j = \sqrt{-1}$. The imaginary components, ω_i of the eigenvectors describe the *fundamental frequencies* of the linear dynamical system, and the real components, σ_i relate to the damping coefficients. In order for the dynamics to be internally stable, it is necessary for *all* the system eigenvalues to have negative real parts,

Definition: A system is internally stable if and only if, $\mathrm{Re}\lambda_i < 0, \forall\lambda_i$, where λ_i is an eigenvalue of $\mathbf{A}$. This leads to exponential decay of all the internal dynamics, and the open-loop dynamical system rapidly achieves a steady state, $\mathbf{x}_{\mathrm{ss}} \to 0$, for $t > 0$.

We can immediately see that the eigenvalues of $\mathbf{A}$ for our example system are $\lambda_i = \pm j\omega$, and $\mathrm{Re}\lambda_i = 0$. A dynamical system with zero real eigenvalues is marginally stable— it is an undamped harmonic oscillator and will not converge to a steady-state for long time. A system matrix whose eigenvalues all lie in the left half-plane is referred to as *Hurwitz*, and internally stability is achieved when the Hurwitz criterion is satisfied.

2.2 Transfer Functions

In order to address the question of closed-loop stability, we need a tool for analyzing the system's response to a particular input. This can be represented in the following manner,

$$y(t) = G[u(t)] \tag{2.13}$$

where the notation, $[\cdots]$, indicates that G takes a function, $u(t)$, as its input. Also note that G outputs a function, $y(t)$, and therefore serves as a liaison between the input signal used to drive the system, and its response.

For linear systems, it is particularly convenient to analyze G for the case where the inputs are sinusoids, $u(t) = a\cos(\omega t) + b\sin(\omega t) = r\exp(j\omega t + \phi)$. This is because driving a linear system with a sinusoid,

$$x(t) = x_0 e^{at} + e^{at}\int_0^t e^{-a\tau} r e^{j\omega\tau}\, d\tau = \frac{e^{aT}(1 + ax_0 - j\omega x_0)}{a - j\omega} - \frac{re^{j\omega t}}{a - j\omega} \tag{2.14}$$

produces an output with a term that oscillates at the same frequency, ω, as the driving term plus a complimentary solution due to the the internal dynamics. Assuming that the system is internally stable, the later component exponentially decays on a timescale that is often faster that $1/\omega$. The system output, $y(t)$, is therefore,

$$y(t) = \frac{C}{a - j\omega} r e^{j\omega t} = Z(\omega)u(t) \tag{2.15}$$

meaning that the output differs from the input only by a change in its magnitude (due to the real part of a) and a phase shift (due to the imaginary part). The quantity, $Z(\omega)$, is generally referred to as the *complex frequency response* of the linear system at the frequency, ω. This point is made even clearer by taking the Laplace transform of the dynamical equations, including all the signals as well as the differential equations to give,

$$\mathcal{L}[\dot{\mathbf{x}}(t)] \quad \rightarrow s\mathbf{x}(s) \quad = \mathbf{A}\mathbf{x}(s) + \mathbf{B}\mathbf{u}(s) \tag{2.16}$$

$$\mathcal{L}[\mathbf{y}(t)] \quad \rightarrow \mathbf{y}(s) \quad = \mathbf{C}\mathbf{s}(s) \tag{2.17}$$

where the Laplace variable, $s = \sigma + j\omega$, and we have ignored the terms that correspond to the complementary solution to the differential equation.[a] For internally stable systems, these neglected terms depend upon the initial conditions and are proportional to the exponentially decaying dynamics. Therefore, the description given by Eqs. (2.16) and (2.17) are valid in steady state. Under the approximation that the internal dynamics decay quickly compared to the time-dependence of $\mathbf{u}(t)$, it is possible to algebraically eliminate $\mathbf{x}(s)$ from the above equations to give,

$$\mathbf{y}(s) = \mathbf{C}\left(s\mathbf{I} - \mathbf{A}\right)^{-1}\mathbf{B}\mathbf{u}(s) \tag{2.18}$$

or,

$$\mathbf{G}(s) \equiv \frac{\mathbf{y}(s)}{\mathbf{u}(s)} = \mathbf{C}\left(s\mathbf{I} - \mathbf{A}\right)^{-1}\mathbf{B} \tag{2.19}$$

where $\mathbf{G}(s)$ is referred to as the system *Transfer Function.*

2.2.1 Properties of Transfer Functions

All of the system dynamics are captured by the transfer function, and even though they are defined in terms of sinusoidal inputs, this imposes no restriction since for linear systems,

$$\mathbf{y}(s) = \mathbf{G}\left(\sum_k \mathbf{u}_k(s)\right) = \sum_k \mathbf{G}\mathbf{u}_k \tag{2.20}$$

[a] *I will use the convention that a function, $f(s)$, always corresponds to the Laplace transform of an associated time-domain function, $f(s) = \mathcal{L}[f(t)]$.*

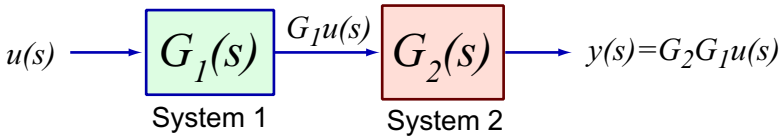

FIGURE 2.2: *When two systems are coupled by a common signal, their transfer functions can be multiplied to give the transfer function of the composite system, from initial input to final output.*

and any input signal can be decomposed according to its Fourier expansion. As can be seen from Eq. (2.19), the transfer function is time-independent (due to the fact that we are using an LTI model), and it represents the complex frequency response of the system as a function of frequency.

Transfer functions are convenient due to their algebraic simplicity— convolution integrals in the time-domain correspond to products of functions in the Laplace domain. The product of $\mathbf{G}(s)$ and a signal, $\mathbf{u}(s)$, is another signal. And, coupled systems, as depicted in Fig. 2.2, with individual transfer functions, $\mathbf{G}_k(s)$, have a total transfer function is given by the product,

$$\mathbf{G}(s) = \frac{\mathbf{y}(s)}{\mathbf{u}(s)} = \prod_{k=1} \mathbf{G}_k(s) \tag{2.21}$$

As a matter of convention, system transfer functions are written using uppercase letters, such as $\mathbf{G}$, whereas signals are expressed using lowercase letters, such as $\mathbf{u}(s)$.

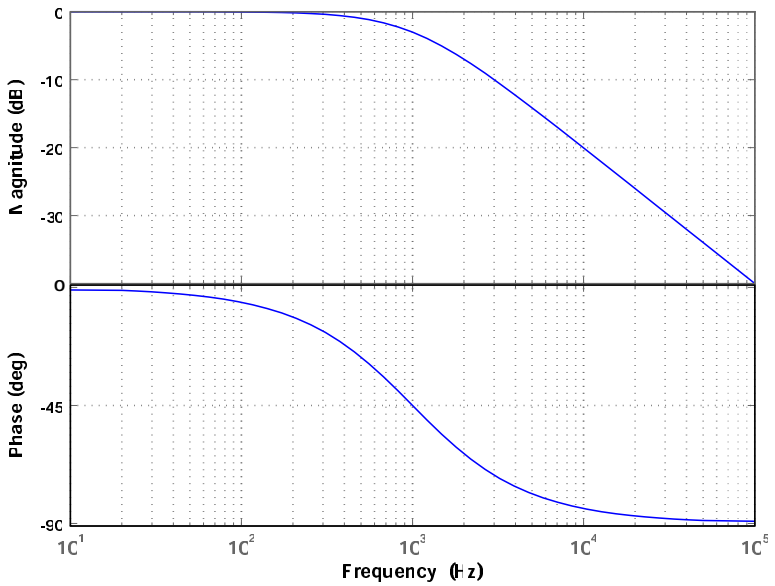

FIGURE 2.3: *An example of a Bode plot for a low pass filter with a corner frequency of 1 kHz. The utility of the Bode plot is that provides an easy to interpret graphical representation of a system transfer function by extracting the frequency dependent gain and phase.*

2.2.2 Bode Plots

The transfer function, $G(s) \equiv G(j\omega)$ gives the complex response of the system as a function of frequency.[b] For linear dynamics, we have already seen that driving a system with a sinusoidal signal produces a response that oscillates at the same frequency, but with a complex coefficient. In the laboratory, it is usually more convenient to think in terms of real numbers that corresponds to physical quantities rather than complex quantities. Since the frequency of the driving, $u(t)$, and response, $y(t)$, signals are the same, the only differences between them can be in the amplitude and phase,

$$G(j\omega) = g(\omega)e^{j\omega t + \phi(\omega)} \tag{2.22}$$

where $g(\omega)$ is the real-valued *gain* or change in amplitude, and $\phi(\omega)$ is the real-valued phase-shift between the input and the output. A convenient method for viewing transfer functions is to plot their frequency-dependent gain and phase, referred to as a Bode plot. As a matter of convention, Bode plots are generated using a log frequency scale, and separate axes for the gain and phase functions, as in Fig. 2.3. The gain is plotted in units of dB, defined for signals in relative energy units as,

$$G_{\mathbf{dB}} = 20\log_{10}|g(\omega)| \tag{2.23}$$

such that 0 dB is equivalent to unity gain, and every 6 dB corresponds to doubling the signal amplitude. There is a related definition for dB in terms of relative power units given by, $10\log_{10}|g_{\mathrm{Pow}}(\omega)|$.

The transfer function depicted by Fig. 2.3 is a Butterworth low pass filter with a corner frequency of $f_c = 1$ kHz. The effect of the filter is to allow low frequency components of the input signal, $u(t)$ to pass with no attenuation and little phase shift; however, higher frequency components are heavily attenuated and also experience a 90° phase shift.

2.2.3 Transfer Function Poles and Zeros

It is often convenient to express transfer functions using ratios of polynomials in $s = j\omega$. This is because the Laplace transform turns differential equations, such as those found in Eq. (1.7), into algebraic equations (refer to the appendix for a review of Laplace transforms). Therefore, the transfer function, $G(s)$, can be represented,

$$G(s) = \frac{N(s)}{M(s)} \equiv \frac{v_n s^n + v_{n-1}s^{n-1} + \ldots + v_1 s + v_0}{w_m s^m + w_{m-1}w^{n-1} + \ldots + w_1 s + w_0} \tag{2.24}$$

where the $\{v_0, \ldots, v_n\}$ are coefficients for the numerator polynomial, $N(s)$, and $\{w_0, \ldots, w_m\}$ are coefficients for the denominator polynomial, $M(s)$ where n and m are the orders of the polynomials.[c]. Since $N(s)$ and $M(s)$ are polynomials, they can be factored into their n and m roots,

[b] *For the remainder of this lecture, I will only consider scalar rather than vector functions because the algebra is a little easier and there is nothing essential to be gained from retaining the vector quantities.*

[c] *I apologize for re-using variables, but the n and m correspond to the order of the polynomials and are not necessarily related to the n an m which provided the dimensions of the matrices,* **A** *and* **B**

respectively,

$$N(s) = k_n (s - z_n)(s - z_{n-1}) \cdots (s - z_0) \tag{2.25}$$

$$M(s) = k_m (s - p_m)(s - p_{m-1}) \cdots (s - p_0) \tag{2.26}$$

where the polynomial zeros are complex, $p_i, z_i \in \mathbb{C}$. The transfer function can be represented,

$$G(s) = K \frac{(s - z_n)(s - z_{n-1}) \cdots (s - z_1)}{(s - p_m)(s - p_{m-1}) \cdots (s - p_1)} \tag{2.27}$$

where K is a constant gain factor, the $\{z_i\}$ are the *zeros* of the transfer function, and the $\{p_i\}$ are the *poles*. Like the magnitude, $g(\omega)$ and phase, $\phi(\omega)$, the quantities, $\{K, p_i, z_i\}$, are sufficient to completely characterize the transfer function.

The usefulness of the pole-zero-gain (PZK) representation is that it is closely connected to the issue of stability. This can be seen by comparing Eq. (2.19) to Eq. (2.27), and noting that the eigenvalues of the dynamical system matrix $\mathbf{A}$ are associated with the poles of the transfer function,

$$(\mathbf{I}s - \mathbf{A})^{-1} = \frac{1}{\det |\mathbf{I}s - \mathbf{A}|} \text{adj} (\mathbf{I}s - \mathbf{A}) \tag{2.28}$$

according to the definition of the matrix inverse. Since the transfer function poles include the eigenvalues of $\mathbf{A}$, stability requires that the transfer function, $G(s)$, contains no right-half plane (RHP) poles. If the real parts of the poles of $G(s)$ are all negative, then system is internally stable, otherwise if $\exists\, p_i \,|\, \text{Re}(p_i) \geq 0$, there will be exponential growth of the system dynamics, or marginal stability (harmonic motion) if $\text{Re}(p_i)$=0.

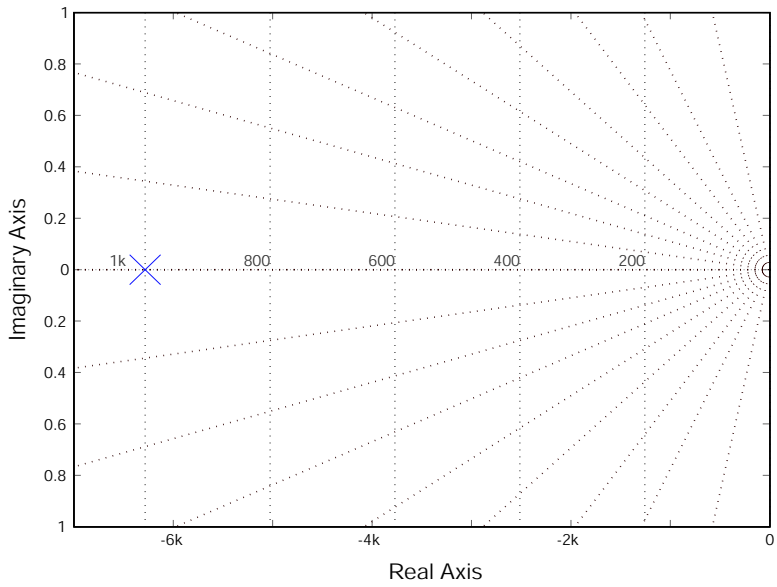

FIGURE 2.4: *An example of a pole zero plot for the same low-pass filter transfer function depicted in Fig. 2.3 and described in Section 2.2.3.*

This stability criterion is elegantly expressed by a pole-zero map, which is simply a plot of the locations of the poles and zeros of $G(s)$ on the complex s-plane. It provides a convenient visual

tool for assessing the internal stability of the dynamical system. Typically, the symbol, $\times$, is used to denote the location of a pole, and $\circ$ is used to denote a zero, as is illustrated in Fig. 2.4. This plot corresponds to the Butterworth filter which we saw in Section 2.2.2 with transfer function,

$$G(s) = \frac{2\pi \times 10^3}{s + 2\pi \times 10^3} \tag{2.29}$$

expressed according to a pole-zero-gain model.

This transfer function has a single pole at $z_1 = -1$ kHz. Since the pole lies in the left-half of the complex s-plane, it corresponds to a stable transfer function. Evidently, this is why a Butterworth low-pass filter is often referred to as a "single-pole low-pass." There are some useful comparisons to make between the Bode plot in Fig. 2.3 and Fig. 2.4. First, the single pole in the PZ map corresponds to the corner frequency, f_c, or -3 dB point, on the Bode plot. Evidentally, the LHP pole results in a 6 dB per octave *roll-off*, or gradual decrease in the gain. It also appears that a single pole results in a gradual onset of phase delay that increases from 0° to -90°, with the inflection point located at the pole-frequency. Although not depicted by this transfer function, a *zero* has the opposite effect of a pole— it causes the gain to rise at 6 dB per octave and results in $+90^\circ$ of phase. Finally, it is necessary to define several common phases that are used when describing transfer functions,

1. **Proper:** A transfer function is proper if $G(s)$ is finite for $\omega \to \infty$.
2. **Strictly Proper:** A transfer function is strictly proper if $G(s) \to 0$ for $\omega \to \infty$.
3. **Minimum Phase:** A transfer function is called minimum-phase if it has no right-half plane zeros, $\mathrm{Re} z_i < 0, \forall z_i$.

2.3 Stability of the Feedback Control System

We finally have enough mathematical equipment to tackle the problem of stability in a closed-loop feedback control system, and we begin by defining the close-loop feedback transfer function using the flow diagram in Fig. 2.5 where the signals are defined as follows,

1. $r(s)$ is the tracking reference signal.
2. $y(s)$ is the system output.
3. $e(s)$ is referred to as the error signal and is given by the difference between the output and the reference, $e(s) = y(s) - r(s)$.
4. $u(s)$ is the controller signal used to drive the system in response to the error signal.

and the system transfer functions are given by,

1. $C(s)$ is the controller transfer function and represents the implementation of Eq. (2.8)

2. $P(s)$ is the *plant* transfer function, and represents the system that is being controlled. The term, "plant" is used generally in control theory, and reflects the fact that control theory was historically first utilized in optimizing large manufacturing facilities.

2.3.1 Closed Loop Transfer Functions

Using what we know about transfer functions and how they apply to signals, it is possible to make the following statements about Fig. 2.5:

$$u(s) = C(s)e(s) \equiv C(s)\left[r(s) - y(s)\right] \tag{2.30}$$

and

$$y(s) = P(s)u(s) \tag{2.31}$$

These two equations can be combined by eliminating $e(s)$ to solve for the closed loop transfer function from the tracking signal, $r(s)$, to the system output, $y(s)$,

$$T(s) \equiv \frac{y(s)}{r(s)} = \frac{P(s)C(s)}{1 + P(s)C(s)}\,. \tag{2.32}$$

$T(s)$ captures all the necessary response information about the behavior of the feedback loop, and it can be used to assess assess both the system stability and the tracking performance.

Definition: Feedback stability– The feedback system characterized the closed-loop transfer function, $T(s)$, is stable if and only if it contains no poles in the right half complex s-plane, $\mathrm{Re}(p_i) < 0, \forall p_i$, where p_i is a pole of $T(s)$.

The requirement that $T(s)$ has no RHP poles translates into specifications for the plant, $P(s)$, and controller, $C(s)$, transfer functions. This can be seen by considering the denominator of the feedback transfer function, $1+PC$, a quantity that is sometimes referred to as the *return difference*. Stability of the feedback system is assured provided that $1+PC$ contains no right-half plane zeros, which correspond to the poles of $(1+PC)^{-1}$.

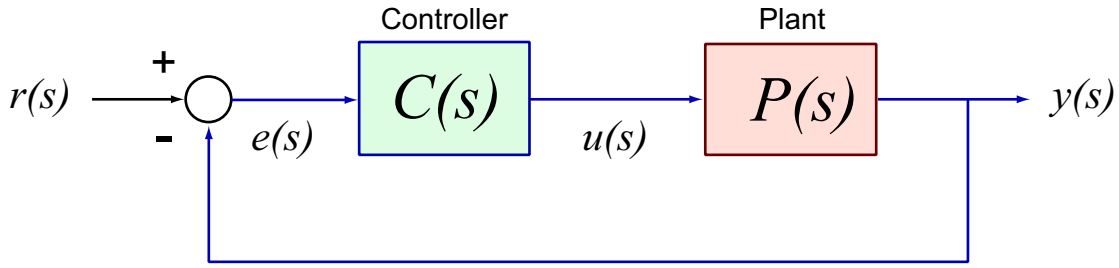

Figure 2.5: *Flow diagram for a simple feedback control system where the $r(s)$ is the reference signal, $y(s)$ is the system output, $e(s)$ is the error-signal, $u(s)$ is the control signal, and $C(s)$ and $P(s)$ are, respectively, the controller and plant transfer functions.*

Technically, there is second requirement for stability: that there are no pole-zero cancellations in the *loop transfer function*,

$$L(s) = P(s)C(s)\,. \tag{2.33}$$

A pole-zero cancellation occurs when two transfer functions with complimentary roots are multiplied together. For example, consider,

$$G_1(s) = \frac{1}{(s+2)(s-1)}, \quad \text{and} \quad G_2(s) = \frac{(s+1)(s-1)}{(s+2)}\,. \tag{2.34}$$

Clearly, G_1 has a RHP pole at $z = 1$ and is therefore internally unstable. However, G_2 has a RHP zero at exactly the same point, $s = 1$, and when multiplied, these roots cancel,

$$G(s) = G_1 G_2 = \frac{s+1}{(s+2)^2} \tag{2.35}$$

and it appears as if $G(s)$ is internally stable. Therefore, it is theoretically possible to make an unstable plant into a well-behaved device by engineering the feedback controller to contain a RHP zero to match each unstable pole in the plant. This technique is used in practice all the time despite the fact that it can be an unsound procedure. For instance, an inherently unstable system that is made stable by feedback in closed-loop configuration will run wildly out of control if something breaks the loop, such as a malfunction in a sensor or a wire. In control applications where instability could lead to a tragedy (like an airplane) pole-zero cancellation should generally be avoided.

2.3.2 The Nyquist Criterion

In practice, a pole zero plot and a sometimes a Bode plot are sufficient to assess the stability of a feedback control loop. However, there is another stability analysis technique that you may come across in the references and control theory literature referred to as the Nyquist stability criterion. It is somewhat more mathematically involved than the previous discussions and requires a background in complex analysis to fully appreciate. The Nyquist criterion is not a necessary technique for engineering a high performance feedback controller, but is instead most useful for its power in formal mathematical theorems involving feedback control stability. This section provides only a brief outline of the procedure because we will use it in passing to describe robust stability in a later section. A longer discussion that is extremely accessible can be found in the text by Phillips and Harbor.[2]

The general idea behind the Nyquist criterion involves generating a map in the s-plane of the loop transfer function, $L(s)$, (recall that $L = PC$ is the loop transfer function) which corresponds to an infinite contour in the s-plane. For anyone familiar with concepts such as contour integration and conformal mapping, this will sound normal; to everyone else, it may seem a bit bizarre. The map is constructed by tabulating the values of $L(s)$, beginning from the point $s = 0$, and proceeding up the imaginary axis to $j\omega \to +\infty$. The contour then traces a loop in the right half plane of infinite radius, σ, to the point, $j\omega = -\infty$. Finally, the contour returns to the $s = 0$ point by proceeding up the negative imaginary axis. While traversing the imaginary axis, any $\mathrm{Re}(s) = 0$ poles in $L(s)$

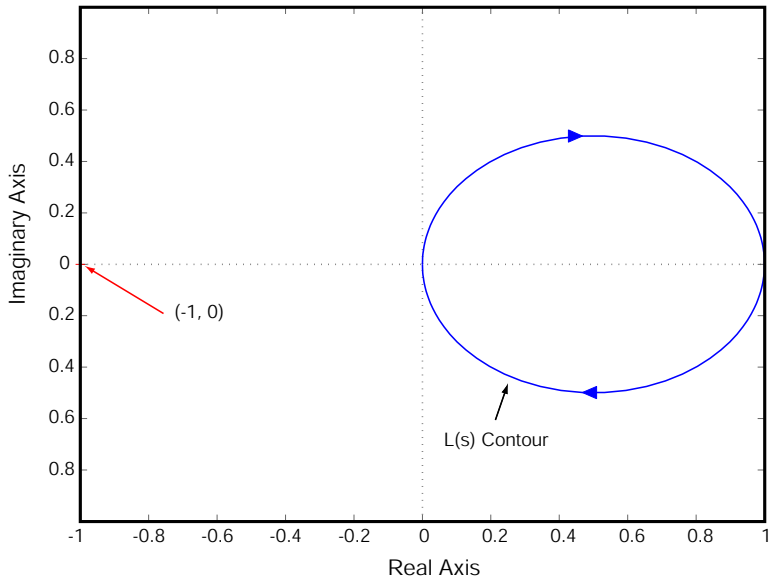

FIGURE 2.6: *An example of a Nyquist plot for the single-pole low-pass filter transfer function described in Eq. (2.29).*

should be traversed by encircling them to the left. This contour can be used to generate a Nyquist plot, such as the one in Fig. 2.6 generated from the low-pass filter transfer function in Eq. (2.29).

Nyquist Stability Criterion A feedback transfer function is stable if its $L(s)$ contour does not encircle the point the $(0, -1)$ in the s-plane if there are no right-half plane poles in $L(s)$. Otherwise, stability is only achieved if the number of counterclockwise encirclements of the $s = (-1, 0)$ point is equal the number of RHP poles in $L(s)$.

The grounds for the Nyquist criterion crome from the Cauchy principle in complex analysis which states that the $L(s)$ contour will encircle the (0,0) point for each RHP zero and pole for an s-plane contour that includes the full right-half s-plane. Since we are interested in the zeros of $1 + PC$ for stability, it is convenient to translate everything by subtracting off the 1 and counting the encirclements of $L = PC$ around the -1 point rather than the origin. If $L(s)$ has no RHP zeros (which would produce instability) then it should not encircle the -1 point, provided that it has no RHP poles, in which case it should encircle -1 exactly once for each pole. This point is demonstrated by Fig. 2.6 for the transfer function in Eq. (2.29), which we know to be stable. As is easily seen, its contour does not encircle the -1 point.

2.4 TRACKING PERFORMANCE

Until now, we have focussed nearly all our attention on developing quantitative measures of stability and have payed little attention to other issues such as performance and robustness. In part, this

is because stability is the most difficult property to achieve. However, in any good servo design, it is necessary to balance the three considerations of stability, tracking and robustness (this will be the subject of tomorrow's lecture). For the time being, it is only necessary to develop quantitative measures of these factors.

As can be expected, the feedback transfer function provides a measure of the tracking since it describes the relationship between the reference signal, $r(s)$, and the system output, $y(s)$. If perfect tracking is achieved, then these two signals will be identical, $y(s) = r(s)$, so that,

$$T(s) \equiv \frac{y(s)}{r(s)} = 1 = \frac{P(s)C(s)}{1 + P(s)C(s)}\,. \tag{2.36}$$

Nominally we can see that tracking is improved by increasing the magnitude of $C(s)$ since rearranging Eq. (2.36) yields,

$$C = \frac{1 + PC}{P} \Rightarrow \frac{1}{C} = 0 \tag{2.37}$$

and this leads to the common control mantra that *gain is good*. Of course, it is not generally possible to make $C(t)$ arbitrarily large because this may cause the stability to suffer.

2.4.1 The Sensitivity Function

The quantity,

$$1 - T(s) \equiv 1 - \frac{P(s)C(s)}{1 + P(s)C(s)} \tag{2.38}$$

provides a quantitative measure of the tracking error. This new function, $1 - T(s)$, has a special interpretation; it is the transfer function from the reference input, $r(s)$, to the error signal, $e(s)$, and is given the symbol, $S(s)$,

$$S(s) = \frac{e(s)}{r(s)} = \frac{1}{1 + P(s)C(s)}\,. \tag{2.39}$$

$S(s)$ is referred to as the sensitivity function because in addition to quantifying the tracking ability of the feedback system, it provides a measure of how changes in $P(s)$ affect the control system. In order to see this, we will take a generalized derivative of $T(s)$ with respect to P,

$$\lim_{\Delta P \to 0} \frac{\Delta T/T}{\Delta P/P} = \frac{dT}{dP}\frac{P}{T} = \frac{1}{1 + P(s)C(s)} = S(s) \tag{2.40}$$

and this is, in fact, the sensitivity function, $S(s)$, as defined in Eq. (2.39). Quantifying a requirement on the tracking performance is accomplished by employing a system norm (refer to the Lecture 2 Appendices for a review of system norms),

$$\|S(s)\|_\infty < \epsilon \tag{2.41}$$

to impose the design requirement that the ratio, $e(s)/r(s)$, must not exceed ϵ. In practice, it is common to express the tracking performance using a potentially frequency dependent weighting function, $W_1(s)$,

$$\|W_1(s)S(s)\|_\infty < 1 \tag{2.42}$$

to quantify the tracking performance requirements of the feedback controller.

2.5 Robustness

The basic idea behind robustness is that sometimes, knowledge of the plant transfer function, $P(s)$, is incomplete or slightly wrong. In other cases, environmental fluctuations such as variable temperature, air pressure, or similar exogenous factors, can be modelled by altering the plant transfer function. The danger with an uncertain plant is that these fluctuations might cause the feedback loop, which was stable for $P(s)$, to become unstable for $P'(s) = (1 + \Delta(s))P(s)$. In the language that we have developed, this would mean that although $P(s)$ contains no unstable poles ($\mathrm{Re}(p_i) < 0, \forall p_i$), it might be the case that the perturbation, $\Delta(s)$, causes a LHP pole to drift across the imaginary axis into the right half of the s-plane. Returning to state-space modes, the dynamical matrix, $\mathbf{A}$, corresponding to to stable plant, P, has no RHP eigenvalues, however, $\mathbf{A}'$ does.

In the field of robust control theory, robustness (particularly robust stability) is achieved by considering a family of different plant transfer functions,

$$\mathcal{P} = \{P(s)(1 + \Delta(s)W_2(s)),\ \|\Delta\|_\infty \leq 1\} \tag{2.43}$$

that are generated by considering a collection of different perturbations, $\Delta W_2(s)$. Here, W_2 is fixed and is referred to as the weight, while Δ generates $\mathcal{P}$. This family is referred to as an *unstructured multiplicative set* because of the fact that Δ and W_2 are allowed to take the form of any stable transfer function.[d] Physically, ΔW_2 is a multiplicative perturbation to $P(s)$ in the sense that their transfer functions multiply, which can be thought of as a fixed nominal plant, $P(s)$ in series with a fluctuating plant, ΔW_2.[e] The convenience of the multiplicative uncertainty in Eq. (2.43) is that is has a simple interpretation in the s-plane. At each frequency, the value of $\mathcal{P}(s)$ in the complex plane is given by a disk of radius $\|W_2\|$ centered at the point, $P(s)$, as is depicted in Fig. 2.7. Robust stability can be expressed using the Nyquist contour method described in Section 2.3.2 and the graphical description of the ΔW_2 uncertainty illustrated by Fig. 2.7. The entire family, $\mathcal{P}$ is stable provided that no member of it violates the Nyquist criterion, which graphically corresponds to the requirement that no uncertainty disk enclose the $(-1, 0)$ point in the s-plane.

Quantitatively, the requirement for robust stability of a feedback control system corresponds to the following inequality,

$$\|W_2 T(s)\|_\infty \leq 1\,. \tag{2.44}$$

Although the argument is somewhat mathematically involved, the argument is roughly the folowing: it can be shown from complex analysis that the ∞-norm is a measure of the minimum distance between the $(-1, 0)$ point and the Nyquist contour of $L(s)$. Using this fact, along with the disk radius measured by $\|W_2 L(s)\|_\infty$, it is possible to derive Eq. (2.44), although a formal proof is somewhat involved, and is worked through in greater detail in several robust control texts.[1,3,4]

[d] *As before, there is a technical requirement that Δ and W_2 not cancel any unstable poles in $P(s)$.*

[e] *In the field of robust control theory, it is also customary to consider other unstructured families of plant perturbations. For more information, a good source is the text by Doyle, Francis and Tannenbaum*[3]

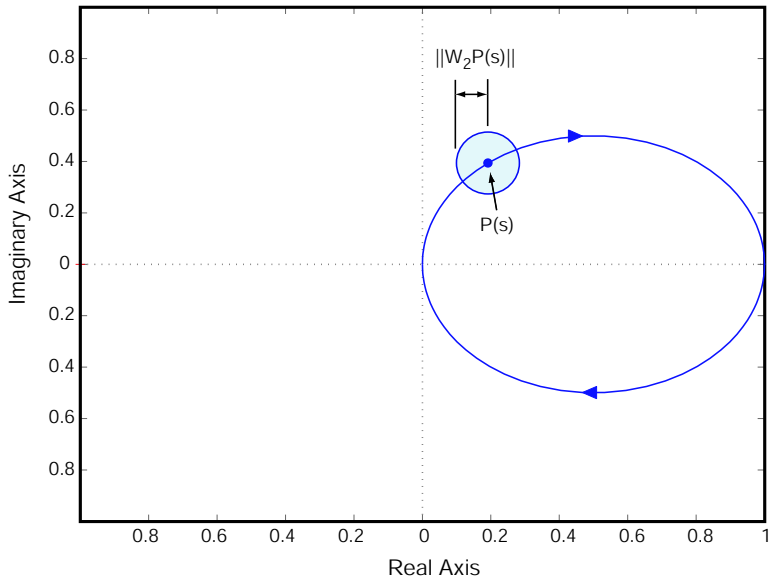

FIGURE 2.7: *Graphical interpretation of robust stability in the s-plane for a family of transfer functions generated from the unstructured model,* $(1 + \Delta W_2)P(s)$.

2.6 LECTURE 2 APPENDICES

2.6.1 REVIEW OF LAPLACE AND INVERSE LAPLACE TRANSFORMS

The Laplace transform of a time-dependent function is defined by,

$$f(s) = \mathcal{L}[f(t)] = \int_0^\infty e^{-st} f(t)\, dt, \quad s = \sigma + j\omega \tag{2.45}$$

where σ is a real number that is arbitrary provided that the integral converges, i.e. that σ is larger than a quantity, α, called the *radius of convergence*. Similarly, the inverse Laplace transform is given by the contour integral,

$$f(t) = \frac{1}{2\pi} \int_{\sigma - j\infty}^{\sigma + j\infty} e^{st} f(s)\, ds \tag{2.46}$$

where, again, σ must be larger than the radius of convergence. Mathematically, the Laplace transform is a generalization of the Fourier transform, and it is more general because it allows for a spectral (or weighting) decomposition over all exponential functions, e^{st}, rather than only sinusoids.

Of particular mathematical importance is the fact that the Laplace transform changes differential equations into algebraic equations,

$$\mathcal{L}\left[\frac{d}{dt} f(t)\right] = sf(s) + f(0) \tag{2.47}$$

and because of the final value theorem,

$$\lim_{s \to 0} [sf(s)] = f(0) + \int_0^\infty \frac{d}{dt} f(t)\, dt = \lim_{t \to \infty} [f(t)] \tag{2.48}$$

which relates the steady-state value of $f(t)$ to the Laplace transform, and serves as the justification for using Laplace transforms to construct transfer functions of linear dynamical systems in control theory.

TABLE 2.1: *Laplace transforms that are commonly encountered in control theory and the study of linear dynamical systems.*

$f(t)$, $t \geq 0$	$f(s) = \mathcal{L}[f(t)]$
1	s^{-1}
t	s^{-2}
t^n	$n!s^{-n-1}$
e^{-at}	$(s+a)^{-1}$
$\cos(at)$	$s(s^2+a^2)^{-1}$
$\sin(at)$	$a(s^2+a^2)^{-1}$

2.6.2 System Norms

The term *system norm* is generally used to describe a norm over a transfer function, $G(s)$, where $s = j\omega$ and $j = \sqrt{-1}$. These norms must satisfy a set of requirements analogous to those listed above in Section 1.4.1 for frequency arguments. Commonly encountered norms for transfer functions are the H_2-norm,

$$\|G(j\omega)\|_2 = \left(\int_{-\infty}^{\infty} G(j\omega)\, d\omega \right)^{\frac{1}{2}} \tag{2.49}$$

and the H_∞-norm,

$$\|G(j\omega)\|_\infty = \sup_{\omega} |G(j\omega)| \tag{2.50}$$

It can be shown using techniques from complex analysis that the 2-norm of $G(s)$ is finite iff G is strictly proper and has no poles on the imaginary axis. Similarly, the ∞-norm of $G(s)$ is finite iff G is proper and has no poles on the imaginary axis.

References

[1] G. E. Dullerud and F. Paganini, *A Course in Robust Control Theory: A Convex Approach* (Springer Verlag, New York, 2000), 1st ed.

[2] C. L. Phillips and R. D. Harbor, *Feedback Control Systems* (Prentice Hall, Englewood Cliffs, 1996), 3rd ed.

[3] J. Doyle, B. Francis, and A. Tannenbaum, *Feedback Control Theory* (Macmillan Publishing Co., 1990).

[4] K. Zhou and J. C. Doyle, *Essentials of Robust Control* (Prentice-Hall, Inc., New Jersey, 1997), 1st ed.

Lecture 3

Loop Shaping: Balancing Stability, Performance and Robustness

We left off yesterday with formal mathematical descriptions of stability, performance and robustness that were expressed in terms of the feedback transfer function,

$$T(s) \equiv \frac{P(s)C(s)}{1 + P(s)C(s)}\,, \tag{3.1}$$

the sensitivity function,

$$S(s) \equiv \frac{1}{1 + P(s)C(s)} = 1 - T(s) \tag{3.2}$$

and a set of frequency-dependent weighting functions, $W_1(s)$ and $W_2(s)$. We saw that the sensitivity function provides a measure of the tracking error, $e(s)$, since $S(s)$ is the transfer function from input, r, to error signal, e,

$$S(s) = \frac{e(s)}{r(s)}\,. \tag{3.3}$$

The requirements for tracking and robust stability where then summarized using ∞-norms, to provide two inequalities,

$$\|W_1 S(s)\|_\infty < 1 \quad \text{and} \quad \|W_2 T(s)\|_\infty < 1 \tag{3.4}$$

for specifying tracking and robust stability, respectively. For all of this to be valid, both T and S must be strictly proper transfer functions, $T(s), S(s) \to 0$, for $t \to \infty$. Of course, $S(s) = 1 - T(s)$, and therefore, both the feedback transfer function and sensitivity function cannot be simultaneously much smaller than one. This fact leads to a tradeoff between stability and performance— ideally we wish to satisfy,

$$\| \, |W_1 S(s)| + |W_2 T(s)| \, \| < 1 \tag{3.5}$$

and with a little algebra, the inequality can be simplified (by removing the sum), to give,[1]

$$\left\| \frac{W_1 S(s)}{1 + \Delta W_2 T(s)} \right\|_\infty < 1\,. \tag{3.6}$$

In practice, the plant is provided to you— it is the system that needs to be controlled. From the perspective of a control theorist, you have very little freedom to adjust $P(s)$, because this

would require physically altering the plant. Therefore, everything in control theory comes down to engineering $C(s)$ in so that Eq. (3.6) is satisfied. In many cases where the plant is sufficiently well behaved, it will be possible to implement an aggressive feedback controller that provides both high performance, $\|W_1 S\| \ll 1$ and good robustness, $\|W_2 T\| \ll 1$, without sacrificing stability. Unfortunately, there are occasions where the plant displays some parasitic behavior that makes it difficult to achieve the desired amount of simultaneous tracking, robustness and stability. It is always necessary to achieve a balance between these three competing design requirements.[1]

Today, we will explore the trade-offs between tracking, robustness and stability, and will develop a procedure for designing $C(s)$. In general, controller design is an immense and involved subject, especially when robustness is required. For the most part, I am going to introduce qualitative guidelines that are inspired by a mathematically rigorous treatment, but I will not formally prove anything. In nearly all cases typical of a physics laboratory, the techniques that I describe work very well.

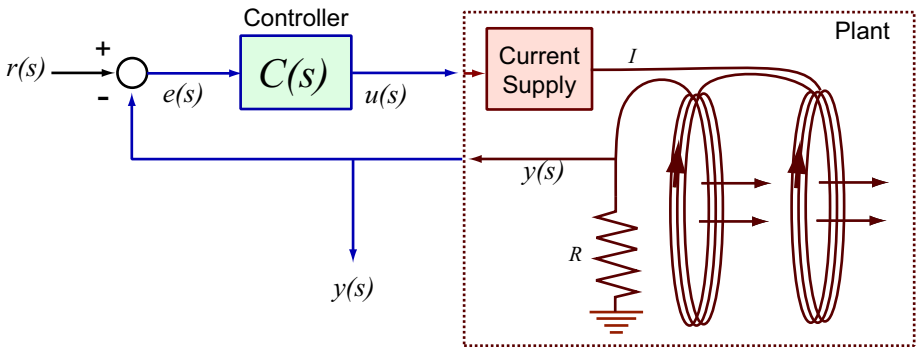

FIGURE 3.1: *Schematic for a feedback control system typical of one that might be found in a physics laboratory, where the objective is to cause the output sense voltage, $y = IR$, to track the programming signal, r.*

3.1 Example Feedback Controller Design

Our goal is to engineer a stabilizing feedback controller for the system depicted in Fig. 3.1. The schematic depicts a very common setup found in many physics laboratories, where the goal is to produce a time-dependent magnetic field by driving an electromagnet. The coil is powered by a current supply that accepts a programming input (in the form of a voltage), and a measurement of the magnetic field is performed by monitoring the current flowing through the magnet (we assume that a calibration of I to B via the inductance of the Helmholz coil is well-characterized). The system output, y is obtained from a current→voltage conversion performed by passing the magnet current through a sense resistor, R. And, the feedback objective is to force y to track the programming input r, thus resulting in a magnetic field that also tracks r. Here, the *plant* consists of the current supply, the coils and the sense resistor, and the controller will be designed such that it can be implemented using analog electronics.

3.1.1 Plant Transfer Function: Internal Stability

The first step in designing the feedback controller is to measure the plant transfer function, $P(s)$, in open loop. This can be accomplished in the laboratory using a network analyzer- wonderful instruments which perform a swept sine measurement of the complex response of the system being tested— they can be configured to output a Bode plot. A plant transfer function measurement for Fig. 3.1 would very likely look like the Bode plot in Fig. 3.2. This transfer function is characterized by a low-frequency gain of approximately 100, but rolls off to its unity gain point around 100 kHz. The phase shift between the current supply programming input and the measured coil current is a relatively complicated function that varies between $0°$ and $-270°$ over the full bandwidth. This phase profile is likely the result of stray phase shifts in the current supply and the effect of the coil inductance.

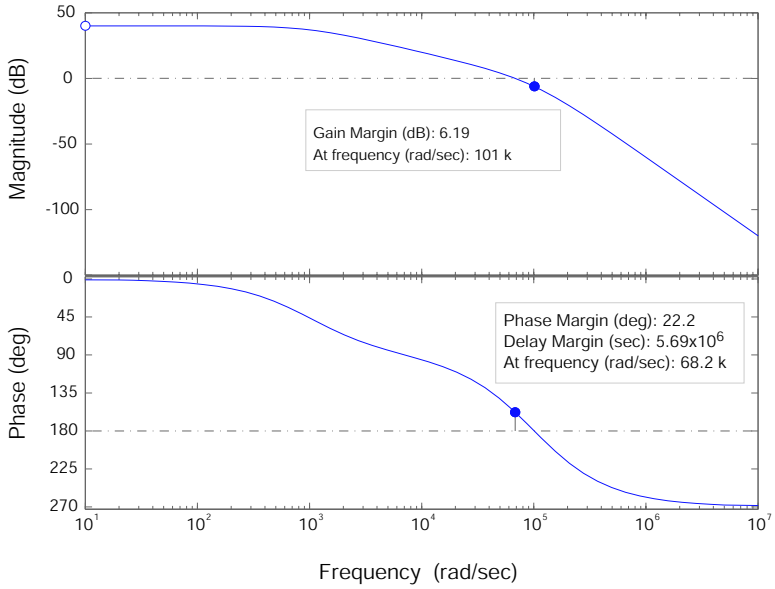

Figure 3.2: *Bode plot for $P(s)$ which is typical of one that you might encounter in the physics laboratory. It has relatively high low frequency gain, rolls off and higher frequency, and has a complicated phase profile.*

The Bode plot in Fig. 3.2 can be fit to a rational polynomial transfer function model, as in Eq. (2.24) to give,

$$P(s) = \frac{10^{15}}{s^3 + (2.01 \times 10^5)s^2 + (1.02 \times 10^{10})s + 10^{13}} \tag{3.7}$$

and the denominator can be factored to find the poles,

$$p_i = \{-10^3, -10^5, -10^5\}\,. \tag{3.8}$$

Clearly, $P(s)$ has no zeros since the numerator in Eq. (3.7) is a constant. An immediate inspection of the poles of $P(s)$ shows that they are all located in the left half of the complex plane: P has two

real poles at $-100/2\pi$ kHz and one at $-1/2\pi$ kHz. Therefore, the plant is stable, and similarly, the transfer function is strictly proper,

$$\lim_{\omega\to\infty} P(s) = 0 \tag{3.9}$$

since the order of the denominator polynomial in Eq. (3.7) is larger than that of the numerator. Therefore, the limit of $P(s)$ vanishes for $s \to \infty$. These properties should come as no surprise since most laboratory devices are internally stable on their own and stop responding at some high-frequency cut-off rendering them strictly proper.

3.1.2 Feedback Stability and Robustness

Now what happens when we add feedback? Since, the plant is internally stable, as a first guess, we might try proportional control by choosing $C(s)$ to be a constant gain. In feedback configuration this will give a closed-loop transfer function,

$$T(s) = \frac{P(s)}{C^{-1} + P(s)} \tag{3.10}$$

and we see that making C large will lead to good tracking,

$$\lim_{C\to\infty} T(s) = \frac{P(s)}{P(s)} = 1 \tag{3.11}$$

and therefore the sensitivity function, $S(s)$, will be zero and the tracking requirement,

$$\|W_1 S(s)\|_\infty < 1 \tag{3.12}$$

will be satisfied for all possible choices of W_1. So, initially it appears that we can meet our tracking objective using a simple amplifier with a large gain for our feedback controller.

However, before we consider our job done, we need to check the stability of our proportional control strategy. Fig. 3.3 provides Nyquist diagrams for two different choices of the controller gain. In plot (A), where $C = 1$, it clear that the feedback system is stable since the contour does not encircle the $(-1, 0)$ point. However, in plot (B), where $C = 3$, something went terribly wrong again— the Nyquist plot most definitively encircles the $(-1, 0)$ point. It appears that increasing the gain has caused the feedback system to go unstable. A quick check of the poles of the closed-loop transfer function, $T(s) = 3P(s)/(1+3P(s))$ shows that increasing the gain by a factor of three pushed the two plant poles at -10^5, into the right half-plane,

$$p_i = \{-2.1767 \times 10^5, (.0833 + 1.1743j) \times 10^5, (0.0833 - 1.1743j) \times 10^5\} \tag{3.13}$$

to produce the instability. Therefore, our plan to increase the gain for good tracking is inherently flawed because the system will go unstable for $C \gtrsim 5/2$.

The behavior we just observed is also closely related to the concept of robustness.[2,3] If we assume that we set the controller gain to unity, $C(s) = 1$, because we know that this procedure

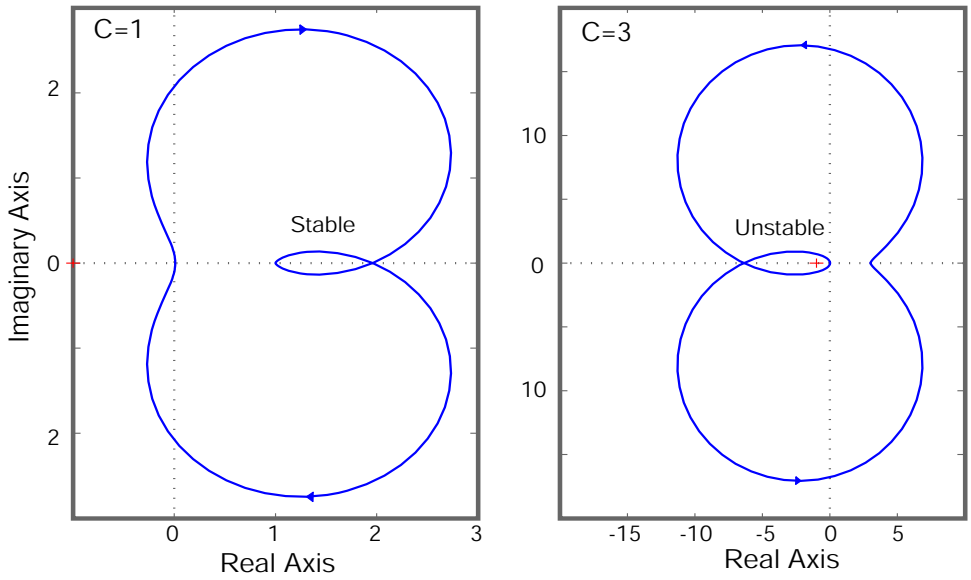

FIGURE 3.3: *Nyquist plot for $P(s)$ which shows that the plant is described by a stable transfer function.*

will be stable for the plant transfer function that we measured in Fig. 3.2, we ask the question: How much must the plant change before the system becomes unstable? Without performing any detailed mathematical analysis, such as by computing $\|W_2 T(s)\|_\infty$, we already know that increasing the plant gain by a factor of three will produce instability. We know this because in the feedback transfer function,

$$T(s) = \frac{P(s)C(s)}{1 + P(s)C(s)}\,, \tag{3.14}$$

an overall change in the gain of $C(s)$ is indistinguishable from an identical change in the gain of $P(s)$. Therefore, using unity gain proportional feedback to control our system is not particularly robust since the plant gain need only drift a few dB before the servo stops working. There are probably some low-quality, high-power laboratory current supplies whose open-loop gain changes by a factor of 2-3 between the time it is turned on, and the time it has warmed up.

3.1.3 Gain Margin, Phase Margin and Bandwidth

Although proportional feedback is stable for $C = 1$, the system is (in some sense) very close to being unstable since only small changes in the plant or controller cause it to fail. What we need is a measure of how far our feedback system is from the stable/unstable border in order to help us improve the controller design. These requirements lead to two quantities, known as the *gain margin* and *phase margin* which can be easily identified from a Bode plot of the closed-loop feedback system.

The Bode plots depicted in Fig. 3.4 illustrate the concepts of gain and phase margin, and these points can be used to identify qualitative reasons why proportional control with $C = 1$ is

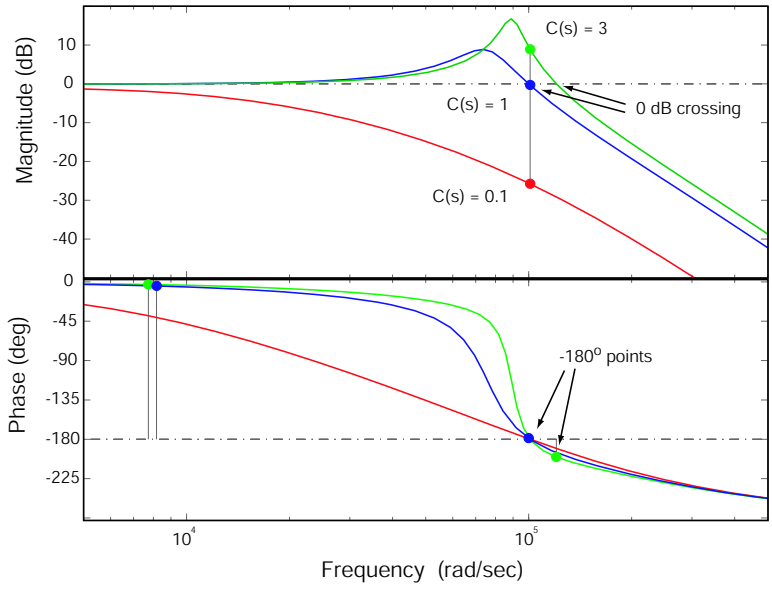

FIGURE 3.4: *Bode plot of $T(s)$ for proportional controllers with three different gains.*

stable while $C = 3$ is not. The blue curve (labelled with $C = 1$) depicts the closed-loop transfer function, $T(s)$. At low frequency, the tracking is good since the gain is very close to 0 dB [meaning no amplification or attenuation of $r(s)$] and there is little phase shift. However, as the frequency increases up to the point $\omega = 900$ kRad/s, the closed-loop gain rises above 0 dB, achieves a maximum and then drops back below the unity gain level (0 dB). Similarly the phase drops to below $-180°$ near the same frequency, $\omega = 900$ kRad/s.

The frequency at which the phase passes through $-180°$ is an extremely important point on the Bode plot. With $-180°$ of phase, $y(s) = -r(s)$ and negative feedback is replaced with positive feedback, which certainly leads to oscillations instability since positive feedback produces right-half plane poles in the system transfer function. If there is any amplification (indicated by $|T(s)| > 0$ db) in regions of the Bode plot where there is more than 180° of phase, the feedback loop will be unstable because of positive feedback amplification. Stability requires that regions of the Bode plot with more than 180° of phase correspond to frequency components that are attenuated, $|T(s)| < 0$ dB. This concept leads to the definition of the gain and phase margins:

- **Phase margin:** If ω_0 corresponds to a point where the magnitude of $T(s)$ crosses 0 dB, $|T(j\omega_0)| = 0$ dB, then the phase margin is $\phi(\omega_0 - 180°$. It corresponds to the additional amount of phase that could be tolerated before the system becomes unstable. Feedback systems with positive phase margin are stable while those with negative phase margin are unstable.

- **Gain margin:** If ω_{-180} corresponds to a point where the phase of $T(s)$ crosses $-180°$, then the gain margin is defined as $-|T(s)|$ in dB. It corresponds to the amount of attenuation that positive feedback frequency components experience. Systems with positive gain margin are stable, while those with negative gain margin are unstable.

With these definitions, it is possible to classify the proportional feedback transfer functions in Fig. 3.4 according to their phase and gain margins,

1. The $C = 0.1$ feedback system has a large positive phase margin (nearly 90°) and a gain margin of nearly 25 dB. It is clearly stable.

2. The $C = 1$ feedback system has a phase margin of approximately 2° and a gain margin of 1 dB. It is barely stable.

3. The $C = 3$ feedback system has a phase margin of -20° and a gain margin of -9 dB. It is unstable.

We see that stability can be engineered into $T(s)$ by maximizing both the gain and phase margins. However, with the improved stability we also expect a decrease in tracking performance. This too can be seen from the Bode plot in Fig. 3.4 where the most stable controller, with $C = 1$, loses its tracking ability at a lower frequency than the other two. By 10^4 kRad/s, $|T(s)|$ for the $C = 1$ controller has already dropped to -3 dB, i.e. the bandwidth of the controller is 10^4 kRad/s, whereas the bandwidths for $C = 1$ and $C = 3$ exceed 10^5 kRad/s. The tradeoff for stability is manifest here as reduced effective feedback bandwidth.

3.1.4 Loop Shaping: Improving upon Proportional Control

The control design approach used until now did not allow for any frequency dependence in $C(s)$, and was therefore heavily constrained. We should expect much better bandwidth, stability and robustness from a frequency dependent controller transfer function. Of course, it is also possible to do worse by making a poor choice for $C(s)$. We want to "shape" the feedback loop by tailoring the frequency response of the controller transfer function to optimally improve the gain and phase margin without substantially reducing the bandwidth.

Loop-shaping techniques form a large fraction of the entire field of control theory. Therefore, you can imagine that it is not something that we can thoroughly explore in only a three hour lecture period. In most cases, it is possible to perform a relatively straightforward controller design that involves adjusting the frequency response of the loop transfer function, which we recall is given by,

$$L(s) = P(s)C(s)\,. \tag{3.15}$$

The argument for engineering the controller by adjusting $L(s)$ is purely a matter of convenience— it is a quantity that appears in both the closed-loop transfer function, $T(s)$, and sensitivity function, $S(s)$. Therefore it is often possible to satisfy design requirements on both S and T by adjusting a single function, L. Once the design is completed, the controller is found by factoring out the plant transfer function, $P(s)$, to yield $C(s)$.

We expect that for the same qualitative design objectives, posed in terms as requirements on $T(s)$ and $S(s)$, every loop-shaping application should qualitatively involve the same target loop

transfer function, $L(s)$. In order to see this from a more mathematical perspective, our objective will be to make $T = 1$ and $S = 0$ at low frequencies, and then reverse their roles at high frequencies. The rational behind this qualitative objective is that $T = 1$, $S = 0$ corresponds to good tracking. And, we expect to be less concerned about the tracking ability above a a certain maximum frequency, $\omega_{\max}$ so it is desirable to turn off the feedback at high frequency, $T(\omega > \omega_{\max}) \to 0$. Therefore, we qualitatively want $T(s)$ to look like a low-pass transfer function,

$$T(s) \equiv \frac{L(s)}{1 + L(s)} = \frac{\omega_{\max}}{s + \omega_{\max}} \tag{3.16}$$

and conversely, we hope to make $S(s)$ into a high-pass transfer function,

$$S(s) \equiv \frac{1}{1 + L(s)} = \frac{s}{s + \omega_{\max}} \,. \tag{3.17}$$

In the end, we will choose $\omega_{\max}$ to ensure a sufficient degree of phase and gain margin, thus ensuring adequate stability and robustness.

Straightforward algebra involving Eqs. (3.16) and (3.17) shows that the optimal family of loop transfer functions, $\mathcal{L}(s)$, that satisfy our design objectives is given by,

$$\mathcal{L}(s) = \omega_{\max} \frac{1}{s} \tag{3.18}$$

which has a single pole at the origin, $p_1 = 0$. This is the transfer function of integrator, and it is characterized by constant $-90°$ phase, infinite gain at $s = 0$ (infinite DC gain) and a constant 6 dB per octave roll-off. $\mathcal{L}(s)$ has a unity gain point at $\omega_{\max}$ and is strictly proper since $s^{-1} \to 0$ as $s \to \infty$. Most importantly an integrator transfer function has *infinite* phase and gain margins.

Extracting the optimized controller transfer function from $\mathcal{L}(s)$, is accomplished by factoring out the plant,

$$C^*(s) = \frac{\mathcal{L}(s)}{P(s)} = \frac{\omega_{\max}}{sP(s)} \,. \tag{3.19}$$

The only remaining issue is that $C^*(s)$ will not be strictly proper if $P(s)$ rolls-off faster than 6 db per octave, and in most practical applications, it is desirable to have all strictly proper transfer functions due to non-ideal behavior and modelling uncertainty at high frequency. Fortunately, $C^*(s)$ can be made proper by adding additional poles,

$$C^*_{\rm p} = \frac{\omega_{\max}}{sP(s)} \left(\frac{\omega_{\rm p}}{s + \omega_{\rm p}} \right)^n \tag{3.20}$$

where $\omega_{\rm p}$ must be chosen sufficiently large that it does not affect $T(s)$ at frequencies where good tracking performance is required. Here, $n > 1$ is an integer that should be chosen just large enough to make $C(s)$ strictly proper. Some caution must always be exercised when selecting $\omega_{\max}$ and $\omega_{\rm p}$ to ensure that there is sufficient phase and gain margin to achieve the required degree of stability and robustness. In practice, this most conveniently found from inspecting Bode plots of $T(s)$ for the closed-loop system for different values of the parameters, $\omega_{\max}$ and $\omega_{\rm p}$. A good rule of thumb for high precision laboratory control systems is that there should be $60°$ of phase margin and > 20 dB of gain margin to provide a comfortable window of stability with a good deal of robustness.

3.1.5 Finalized Controller Design

Using the plant transfer function defined in Eq. (3.7), and selecting the values $\omega_{\max} = 10^5$ rad/s and $\omega_{\rm p} = 10^7$ rad/s with $n = 3$, we can finalize our controller design,

$$C(s) = \frac{10^{29}s^3 + (2 \times 10^{34})s^2 + 10^{39}s + 10^42}{10^{15}s^4 + 3(\times 10^{23})s^3 + (3 \times 10^{31})s^2 + 10^{39}s}\,. \tag{3.21}$$

You should not be concerned with the large exponents in the polynomial coefficients. Remember that these numbers correspond to powers of ω, which itself has characteristic magnitudes of $10-10^8$ in common laboratory control applications.

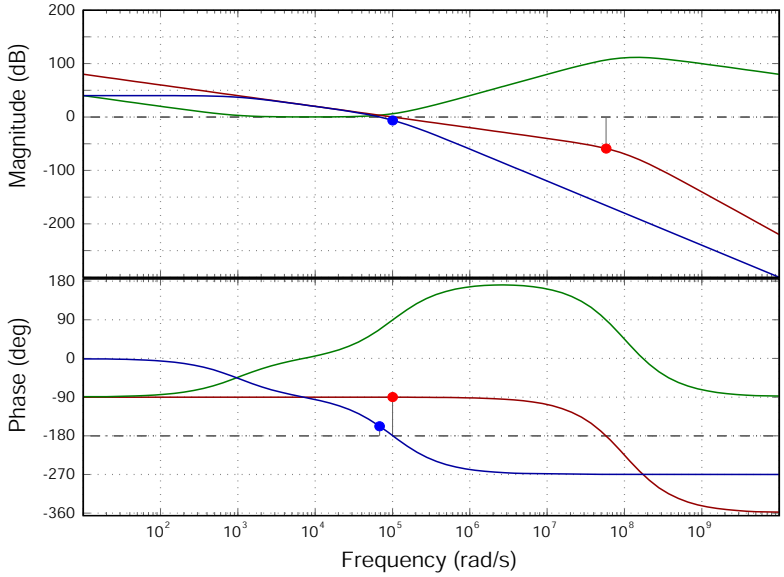

Figure 3.5: *A comparison of the Bode plots for the plant, $P(s)$, optimized controller, $C^*(s)$, and the loop transfer function, $L(s) = P(s)C(s)$ for our feedback controller designed using loop-shaping techniques.*

Fig. 3.5 shows Bode plots for the loop-shaped controller, $C^*(t)$, and compares it to the original plant transfer function, $P(s)$, and the target loop transfer function, $L(s)\,(\omega_{\rm p}/1 + \omega_{\rm p})^n$. The red curve is the target loop transfer function, and it is essentially an integrator with a -90° phase shift and a 6 dB per octave up to the cutoff-frequency, $\omega_{\rm p}$, where the $n = 3$ additional poles increase the roll-off rate and the phase. The green plot corresponds to $C(s)$ and it is clear that is is (in some sense) the inverse of $P(s)$ with respect to L. In regions of the Bode plot where $P(s)$ has a smaller magnitude or less phase that L, $C(s)$ compensates by increasing its magnitude and phase above that of an integrator. The opposite is true for regions where $P(s)$ is too large compared to the integrator transfer function, $L(s)$.

Finally, Fig. 3.6 shows a plot of of the closed-loop transfer function for our feedback system. The tracking bandwidth, which is set by $\omega_{\max}$, is correct. The -3 dB point in Fig. 3.6 is almost exactly located at $\omega_{\max} = 10^5$ rad/s. At slightly higher frequencies, $T(s)$ rolls-off at 6 dB per octave and displays an associated -90° of phase. There are no points where $\phi < -180^\circ$ with greater than

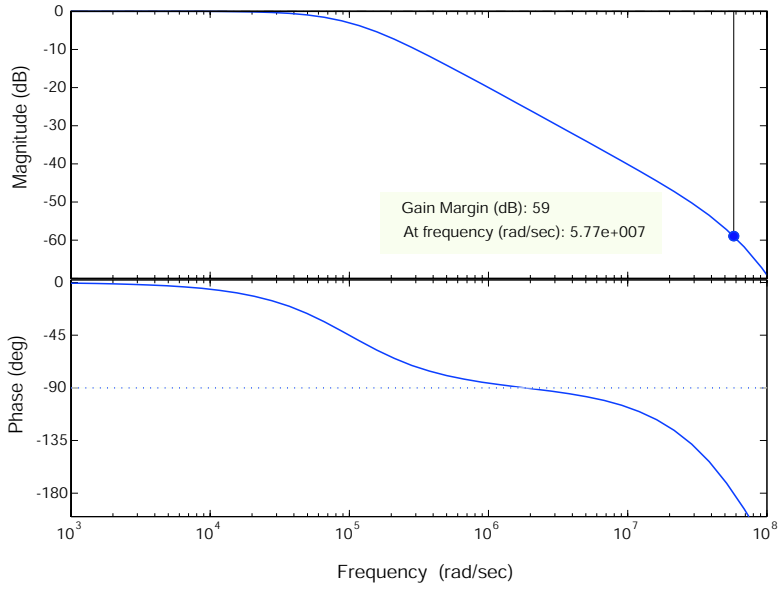

FIGURE 3.6: *Closed-loop feedback transfer function, $T(s)$ for the finalized control system designed using loop-shaping techniques.*

unity gain. In order to verify the stability, we can check the poles of $T(s)$,

$$p_i = \{-10^5, -9 \times 10^7, -(1.05 + 0.08j) \times 10^8, -(1.05 - 0.08j) \times 10^8\} \tag{3.22}$$

and see that all have negative real parts. Finally, Fig. 3.7 shows a Nyquist plot of the finalized feedback system, and it clear that the $(-1, 0)$ is not encircled.

3.2 Digital Controller Design

As a conclusion to our discussion on classical control problems, we will take a quick look at how to repeat the analog controller design in the above section using digital techniques. When working with a digital controller, the considerations for stability, robustness and tracking, are extremely similar to the analog case. Just as before, there will be a trade-off between stability, tracking bandwidth and robustness, and there are analogies to the Nyquist plot and to the gain and phase margin that will apply in a digital system. A full discussion of digital filter design can be found in the very complete text by Proakis and Manolakis[4] although this book is not about control theory. It focusses on the the theory of digital systems and how to faithfully implement filters using digital techniques.

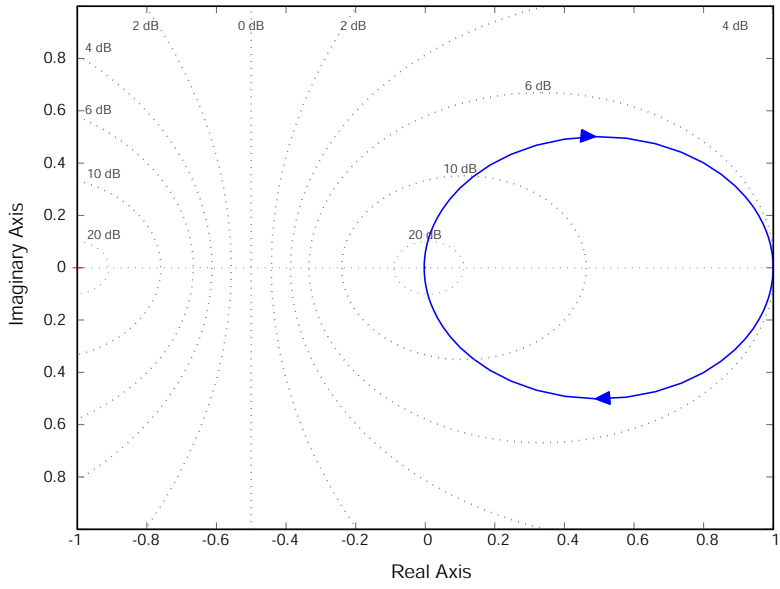

FIGURE 3.7: *Nyquist plot for the finalized feedback controller showing that is stable.*

3.2.1 DIGITAL FILTERS AND THE Z-TRANSFORM

Unlike continuous time dynamical systems whose evolution is described using differential equations, digital systems are described by delay-update equations,

$$\mathbf{x}_{t+1} = \mathbf{A}\mathbf{x}_t + \mathbf{B}\mathbf{u}_t, \quad t+1 = t + \Delta t\,. \tag{3.23}$$

In order to implement all of the controller design techniques that have been developed for analog systems in a relatively seamless manner, we require a digital version of the continuous-time transfer function. This is possible using a mathematical construction known as the z-transform,

$$X(z) = \sum_{i=0}^{\infty} x(t_i) z^{-i} \tag{3.24}$$

where $z = \sigma + j\omega$ is a discrete complex variable analogous to the Laplace variable, s. The z-transform converts a function, $x(t_i)$, that is defined on the discrete time-line to a function in the discrete complex plane, referred to as the z-plane. Just like continuous time transfer functions, z-transforms can be represented as ratios of polynomials in z and it is possible to compute their poles and zeros, and for example, the Hurwitz criterion applies to the z-plane. In fact, there are several procedures which can be used to convert directly from a Laplace transform to a z-transform such as the Tustin procedure.[4]

Therefore, designing a digital controller can be accomplished in one of two ways. If the dynamics are provided in the form of delay-update equations, then the z-transform can be used to compute digital transfer functions and all the same guidelines for feedback stability, performance and robustness can be repeated using discrete instead of continuous transfer functions. However, the usual case is that the system dynamics actually correspond to a continuous time system, in

which case the controller design can be performed using the analog techniques introduced in this Lecture. Once the analog controller has been designed, it can be converted into a digital filter. Fortunately, there is a substantial amount of digital design software, including functions in the Matlab control toolbox, that will perform the conversion.

3.2.2 FIR and IIR Filters

When building a digital controller, z-transforms can be implemented using two standard types of filters. These are the so-called *finite impulse response* (FIR) and *infinite impulse response* (IIR) circuits. FIR filters operate by constructing an output that is a linear combination of inputs at previous times,

$$y_{\mathrm{FIR}}(t_k) = \sum_{p=1}^{N} c_p t_{k-p} \tag{3.25}$$

where N is the number of these previous inputs, and the c_j are weighting coefficients. This circuit gets its name from the fact that it cannot produce an infinite output given a finite input signal because its output only depends on the input signal. In contrast, IIRs, also permit previous outputs to contribute to the filter output,

$$y_{\mathrm{IIR}}(t_k) = \sum_{p=0}^{N-1} c_p t_{k-p} + \sum_{q=0}^{N-1} d_q y_{k-q} \tag{3.26}$$

and can therefore increase without bound due to positive feedback. The requirement for stability of an IIR filter is that its coefficients, d_q, must be less than unity in order to avoid an output that increases without bound.

3.3 Lecture 3 Appendix: Controller Design with Matlab

The Matlab programming environment is an extremely powerful tool for feedback controller design because it has a large number of functions (supplied as part of the Control Theory Toolbox). These functions include routines for forming and manipulating transfer functions, finding poles and zeros, and generating bode plots, Nyquist plots, and pole-zero plots. I used Matlab to generate all of the plots presented in these Lectures, and the Matlab script that follows is the code that I used to do so.

```
% %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
% controllerDesign( )
%
% JM Geremia, 25 September - 2 October, 2003
% California Institute of Technology
%
% This Matlab script generates most of the plots used in Lecture 3 for
% the feedback controller design
% %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

clear all;

    StartFrequency = 1E1;      % Starting Frequency Point for plots
    StopFrequency  = 1E7;      % Ending Frequency Point for plots

    % %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
    % CONTINUOUS TIME (ANALOG) CONTROLLER DESIGN
    % %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

    % Measured Plant Transfer Function in continous time
    %   The Matlab tf( ) function generates a transfer function
    %   Type "help tf" for more information
    PlantTF  = tf( [1E15], [1 201000 1.02E10 1E13] );

    pole( TF )                  % Compute the poles of the plant TF
    zero( TF )                  % Compute the zeros of the plant TF

    % Target Loop transfer function is an integrator that is rolled
    OmegaMax = 1E5;             % Closing frequency (rad/s)
    OmegaRollOff = 1E7;         % Roll-off Frequency (rad/s)
    n = 3;                      % Roll-off order for proper controller TF

    RollOffTF = tf( [OmegaRollOff], [1 OmegaRollOff] );
    TargetLoopTF = tf( [OmegaMax],[1 0] );

    % Compute the designed controller from the target loop
    % transfer function by dividing out the plant component
    ControllerTF = TargetLoopTF / PlantTF * RollOffTF^n;

    % Compute the closed-loop transfer function for the feedback system
    ClosedLoopTF = PlantTF * ControllerTF / ...
        ( 1 + PlantTF * ControllerTF );
```

```
pole( ClosedLoopTF )       % Compute the poles of T(s)
zero( ClosedLoopTF )       % Compute the zeros of T(s)

% Plot the plant, controller and closed-loop transfer functions
%   The bode( ) command generates a Bode plot
%   Type "help bode" for more information
figure(1);
bode( PlantTF , 'b', { StartFrequency, StopFrequency } );
hold on;
bode( ControllerTF, 'g', { StartFrequency, StopFrequency } );
bode( ClosedLoopTF, 'r', { StartFrequency, StopFrequency } );
hold off;

% %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
% DISCRETE TIME (DIGITAL) CONTROLLER DESIGN
% %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

DigitalClockFrequency = 2*pi*100E6;     % Clock Speed = 100 MHz

% FIR and IIR filters use a specified number of coefficients in order
% to define the filter.  This number depends on the hardware
% implementation of the digital filters
NumberOfCoefficients = 14;

% Digital filters are constructed using a certain number of bits
% to digitize the signals.  This number must be large enough to give
% the required precision.  As a general rule it should be at least
% twice the number of digital memory bits.
NumberOfBits = 28;

% The digital sampling frequency is the clock rate divided by the
% number of memory bits
SampleFrequency = DigitalClockFrequency / NumberOfBits;
SampleTime      = 1 / SampleFrequency;

% Here we convert the continuous time transfer functions into digital
% filters using standard techniques such as the "Tustin" approach to
% compute the Z-transform from the Laplace transfer function
%   Type "help c2d" for more information

PlantDigitalFilter = c2d( PlantTF, SampleTime, 'tustin');
```

```
    ControllerDigitalFilter = c2d( ControllerTF, SampleTime, 'tustin');

    pole( ControllerDigitalFilter )
    zero( ControllerDigitalFilter )
    LoopDigitalFilter = PlantDigitalFilter * ControllerDigitalFilter;

    % Generate Bode Plots for the digital filters
    figure(2);
    hold on;
    bode( ControllerDigitalFilter, 'g', { StartFrequency, StopFrequency } );
    bode( PlantDigitalFilter, 'b', { StartFrequency, StopFrequency } );
    hold off;

% %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
```

References

[1] J. Doyle, B. Francis, and A. Tannenbaum, *Feedback Control Theory* (Macmillan Publishing Co., 1990).

[2] C. L. Phillips and R. D. Harbor, *Feedback Control Systems* (Prentice Hall, Englewood Cliffs, 1996), 3rd ed.

[3] R. C. Dorf and R. H. Bishop, *Modern Control Systems* (Prentice Hall, Upper Saddle River, NJ, 1991), 9th ed.

[4] J. G. Proakis and D. G. Manolakis, *Digital Signal Processing: Principles, Algorithms and Applications* (Prentice Hall, Upper Saddle River, 1996), 3rd ed.

Lecture 4

Feedback Control in Open Quantum Systems

Today, we begin looking at quantum mechanical systems and we will attempt to adapt our knowledge of classical feedback control theory into something that applies when the dynamics are governed by quantum mechanics. Of course, quantum mechanics is a very large subject and I will need to assume some familiarity with it in order to even get started. Then again, many physicists who study quantum systems, particularly outside the field of quantum optics, are not necessarily familiar with theories on continuous measurement and conditional evolution. Therefore, I will spend some time *qualitatively* developing the equations of motion for a canonical quantum system in continuous contact with an external, classical environment. In all practical quantum control applications, the person building the controller lives in this external classical world. Therefore, the majority of the subject involves understanding how information about the state of the system propagates from its little quantum setting to our external classical world.

Observation plays a central role in feedback control— there can be no feedback without checking up on the state of the system to ensure that it is doing the right thing. However, everything we have been taught about measurement in quantum mechanics tells us that observation induces unavoidable randomness due to the projection postulate. Fortunately, quantum systems can be measured in many different ways, some of which are less intrusive than others. In the worst case scenario, a projective measurement instantaneously and completely collapses the quantum system into a eigenstate of the measurement operator. This drastic, sudden change in the quantum state appears to be a poor candidate for for feedback control. Instead, "weak measurements" of the quantum system result in less drastic fluctuations and only gradually reduce the system to a measurement eigenstate. Most importantly, weak measurements can be performed continuously and provides the backbone of quantum feedback control. There is a very elegant theory for describing the evolution of a quantum system subject to continuous observation known as quantum trajectory theory.[1–3] The approach I take is similar to those of Gardiner and Doherty; however, I will not completely derive quantum trajectory equations because this would require too much of a background in stochastic calculus. Instead, we will qualitatively construct the dynamical arguments, state the mathematical results, and then identify the physical significance of the different terms in the mathematics.

4.1 Similarities Between Classical and Quantum Dynamical Systems

In the first lecture we developed a concept of dynamical systems grounded upon an input/ouptut model. Therefore, if we hope to adapt any of the earlier material in the course to the quantum regime, it will be necessary to develop a similar intput/output formalism for quantum dynamics. Back in classical control, we described the system dynamics via state-space variables, $\mathbf{x}(t)$, that encode all the essential system information. In linear models, the matrix, $\mathbf{A}$ provides the necessary physical information to propagate $\mathbf{x}(t)$ forward in time starting from an initial condition, $\mathbf{x}_0$. Similarly $\mathbf{B}$ describes the effect of the control input, $\mathbf{u}(t)$ on the system dynamics, and $\mathbf{C}$ provides the relationship between the system state and the observed output, $\mathbf{y}(t)$.

We begin looking at the field of quantum control by making associations between the classical dynamical quantities, $\mathbf{x}(t), \mathbf{u}(t), \mathbf{y}(t), \mathbf{A}, \mathbf{B}$ and $\mathbf{C}$, and quantities from quantum mechanics. One immediate similarity with Lectures 1-3 is linearity. Quantum mechanics, in most cases, is described using linear operators, and so we expect that there will be dynamical analogies (if not equivalences) between classical and quantum dynamical systems. However, at the same time, we expect some new complications to arise due to the fact that quantum mechanics is a probabilistic theory while classical dynamics can be formulated deterministically.

- **State-Space → Hilbert-Space**: In classical systems, our state-space is given by $\mathbb{R}^n$, an n-dimensional vector space over the field of real numbers, i.e., $\mathbf{x}(t) \in \mathbb{R}^n, \forall t$. In quantum dynamical systems, the state-space corresponds to the system Hilbert-space, $\mathcal{H}_n$, where n denotes the dimension. In both quantum and classical dynamical systems, it is possible for $n \to \infty$, such as when describing fields.

- **State Vector → Density Operator**: In classical systems, $\mathbf{x}(t)$, represents all of the necessary information for describing the system. Instead, quantum mechanics utilizes a wavefunction, $|\psi(t)\rangle$, or more generally a density operator, $\hat{\rho}(t)$, to accomplish this task. The essential difference is that the classical state-vector, $\mathbf{x}(t)$, has definite values, while the quantum analog requires a statistical interpretation.

- **System Matrix, A, → Liouvillian, $\mathcal{L}_0$**: In classical mechanics, the internal dynamics were generated by the differential equation,

  $$\dot{\mathbf{x}}(t) = \mathbf{A}\mathbf{x},$$

  whereas in the quantum system, the Liovillian, $\mathcal{L}_0$ propagates the state,[a]

  $$d\,\hat{\rho}(t) = \mathcal{L}_0\hat{\rho}(t)\,dt\,.$$

  In the special case where $\mathcal{L}$ is a Hamiltonian, $\hat{H}_0(t)$, the dynamics are described by the von Neumann equation,[b]

  $$\frac{d}{dt}\hat{\rho} = -i\hbar[\hat{H}_0, \hat{\rho}]\,,$$

[a] *Certainly the classical system is also described by a Liovillian, even though we did not explicitly follow this route.*

[b] *In the quantum Lectures, I will use $i = \sqrt{-1}$, rather than j, because of tradition.*

but, in general, interactions between the quantum system and its surrounding environment and due to any measurement process, will lead to irreversible dynamics such as decoherence.

- **Forcing Term, $\mathbf{Bu}(t)$, $\rightarrow$ Interaction Hamiltonian,** $u(t)\hat{H}_{\rm fb}$ In classical control, the matrix, $\mathbf{B}$, described the interaction between the driving term, $\mathbf{u}(t)$ and the state-variables. In quantum systems, the control will be implemented through an additional term, $\hat{H}_{\rm fb}$, in the system Liovillian,

$$d\hat{\rho}(t) = \mathcal{L}_0\hat{\rho}(t)\,dt - i\hbar u(t)[\hat{H}_{\rm fb}, \hat{\rho}]$$

that describes the interaction between some external field, such as a laser or magnetic, and the system, where $u(t)$ is the time-dependent coupling strength.

- **Observation Process, $\mathbf{y} = \mathbf{Cx}$, $\rightarrow$ Conditional Evolution**: In classical dynamical systems, it was assumed that observing the system through $\mathbf{y}(t)$ did not affect the evolution, i.e. there was no term for $\mathbf{y}$ in the expression for $\dot{\mathbf{x}}(t)$. However, in quantum mechanics, we know that this cannot be the case since measurements result in backaction.

4.2 Continuously Observed Open Quantum Systems

The largest contrast between the classical and quantum control systems (according to the above comparisons) is measurement. As described above, direct projective measurements on the quantum system are poor candidates for the observation process since they induce discontinuous evolution of the quantum state due to quantum jumps. If we wish to develop a dynamical system model analogous to the classical control case, we require continuous (in some sense of the word) evolution.

4.2.1 Weak Measurements

A less intrusive observation process is provided by weak measurements. These operate by implementing *indirect* observation of the actual quantum system— they first entangle the system of interest with a disposable meter (such as an optical mode) and then perform a projective measurement on only the meter. This process still involves backaction because of the entanglement; however, for only moderate coupling between the system and meter, the backaction is much less severe and leads to continuous, stochastic evolution. As such, the field of continuous measurement is based solidly on the foundations of stochastic differential equations.[4] This field allows for rigorous derivations of the system evolution equations. We will use some stochastic calculus, but will not rigorously derive all the equations— it would take too long. This section provides a qualitative development of continuous measurement theory, hopefully, in a manner that will allow us to use stochastic methods later with minimal confusion.

The process of weak measurement is described schematically in Fig. 4.1. The system, $\hat{\rho}_{\rm S}$, belongs to a Hilbert space, $\mathcal{H}_{\rm s}$, and the meters, R_i, to a Hilbert space for the reservoir (the collection of many meters), $\mathcal{H}_{\rm R}$. Initially, the system and reservoir modes are separable because they have

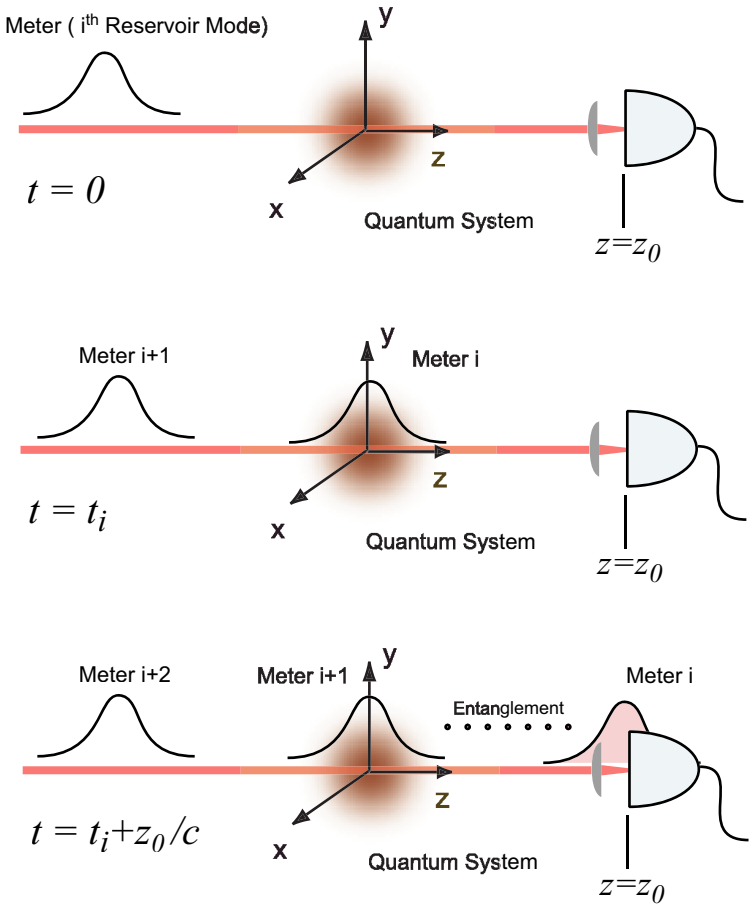

FIGURE 4.1: *A schematic of weak measurement schemes. A string of meter systems interact with the quantum system of interest, become entangled with it, and then are detected. By construction, the meters are independent do not interact with one another.*

never interacted,

$$\mathcal{H}_{\mathrm{S+R}} = \mathcal{H}_{\mathrm{S}} \otimes \mathcal{H}_{\mathrm{R}} \,. \tag{4.1}$$

The initial state of the composite system,

$$\hat{\chi}(0) = \hat{\rho}_S(0) \otimes \hat{\sigma}_i(0) \tag{4.2}$$

where $\hat{\chi}$ is the density operator for both the system and the meter, $\hat{\rho}_{\mathrm{S}}$ is the density operator for the system alone, and $\hat{\sigma}_i$ is the density operator for the i^{th} reservoir mode.[c]

At time, $t = t_i$, the i^{th} meter interacts with the system as it travels past it. Typically, the meters are provided by modes of an optical field and therefore travel at the speed of light, so the interaction time, Δt is very short. During this period, the system and i^{th} meter evolve together according to an interaction Hamiltonian, $\hat{H}_{\mathrm{SR}}$, and therefore become entangled, such that

[c] *I apologize for the fact that $\hat{\rho}$ is the Greek letter for "R" and that $\hat{\sigma}$ corresponds to "S", possibly a poor choice of notation, but $\hat{\rho}$ is traditionally the system density operator.*

$\hat{\chi} \neq \hat{\rho} \otimes \hat{\sigma}_i$ is no longer a product state. At time $t = t_i + z_0/c$, the mode is measured by projecting the combined system-reservoir state onto a measurement operator, $\hat{C} = \mathbb{1}_{\mathrm{S}} \otimes \hat{C}_{\mathrm{R}}$. Taking a partial trace over the reservoir leads to a new density operator for the system,

$$\hat{\rho}_c(t_i + z_0/c) = \mathrm{tr}_{\mathrm{R}}[\hat{C}\chi(t_i + z_0/c)], \tag{4.3}$$

where the subscript, c, on $\hat{\rho}_c$ indicates that it is conditioned on the outcome of the measurement. In other words, detection provides information about the particular random measurement outcome that resulted from the observation process, and therefore, our knowledge of the state is a probability conditioned on the measurement outcome. At each point in time, there are many possible random outcomes that could occur due to the measurement, and conditioning corresponds to a particular evolution path known as a *quantum trajectory.*[1–3]

Our goal is to develop an effective equation of motion for $\hat{\rho}_c(t)$ by eliminating the reservoir modes and relegating their influence on the system to terms in the evolution equations that resemble noise. Since we are propagating a density operator, and due to the measurement randomness, the evolution equation is referred to as a *stochastic master equation.* In order to derive it, we need to make several simplifying assumptions about the meters and their interaction with the system:

1. *Time-Independent Interaction*: It is assumed that the reservoir modes only interact with the quantum system of interest for a length of time, Δt, that is short compared to all of the internal dynamics in the system. Therefore, during that time, the Hamiltonian that couples them is effectively constant,

   $$\hat{H}_{\mathrm{SR}} = \sqrt{M}(sr^\dagger + s^\dagger r) \tag{4.4}$$

   where M is a coupling or *Measurement* strength with units of frequency,[d] s is an operator on the system Hilbert space, $\mathcal{H}_{\mathrm{S}}$ and r acts on the meter Hilbert space, $\mathcal{H}_{\mathrm{R}}$. After the Δt period has elapsed and the i^{th} mode has moved past the system, the interaction ceases.

2. *Independent and Non-Interacting Meters*: The reservoir consists of a long string of meters that interact sequentially with the system (refer to Fig. 4.1),

   $$R(t) = \sum_i R_i(t)\,. \tag{4.5}$$

   Each of the R_i are defined to be essentially orthogonal,

   $$[R_i(t), R_j^\dagger(t)] = \delta_{ij} \tag{4.6}$$

   and therefore,

   $$[R(t), R^\dagger(t')] = \delta(t - t') \tag{4.7}$$

   This assumption corresponds to a Markov approximation of the dynamics in the sense that each mode interacts with the system only for a small window surrounding t_i, and therefore the evolution of $\hat{\rho}_c$ at time t_i only depends on $\hat{\rho}_c(t_i)$.

[d] *I have chosen to work in units where* $\hbar = 1$.

3. *Insignificant Delay*: Since the modes are travelling at the speed of light, we will assume that it is unnecessary to worry about the retardation time, z_0/c.

As a matter of convention, we will also choose to adopt an interaction picture representation with respect to the system Hamiltonian,

$$\tilde{\chi}(t) = e^{-i\hat{H}_S}\hat{\chi}e^{i\hat{H}_S} \tag{4.8}$$

and adopt the convention that operators with an explicit time argument are in the interaction picture, while those without a time argument are Schrödinger picture operators.

4.3 Quantum Input-Output Formalism

In the interaction picture, the propagator for the quantum system is formally given by,

$$\hat{U}(t_i + \Delta t, t_i) = \vec{T}\exp\left[\sqrt{M}\hat{s}\int_{t_i}^{t_i+\Delta t}\hat{R}_{\rm in}^\dagger(t')\,dt' - \sqrt{M}\hat{s}^\dagger\int_{t_i}^{t_i+\Delta t}\hat{R}_{\rm in}(t')\,dt'\right] \tag{4.9}$$

where $\vec{T}$ is the time-ordering operator on the exponential. $\hat{R}_{\rm in}$ refers to the reservoir operator prior to the interaction, and has been renamed this in order to develop an input/output perspective. The idea is to express the reservoir operator at detection time in terms of $\hat{R}_{\rm in}$ and the interaction evolution by defining an output operator,

$$\hat{R}_{\rm out}(t) \simeq \hat{U}^\dagger(t,0)\hat{R}_{\rm in}(t)\hat{U}(t,0) \tag{4.10}$$

where the interaction time, t_i, lies between $t = 0$ and t, i.e., $0 < t_i < t$. Since there is no coupling between R and S before t_i or after $t_i + \Delta t$, the evolution operators can be factored,

$$\hat{U}(t,0) = \hat{U}(t, t_i + \Delta t)\hat{U}(t_i + \Delta t, t_i)\hat{U}(t_i, 0) \tag{4.11}$$

in accordance with the Markov approximation that we made above. For very small Δt (eventually we will take the limit $\Delta t \to dt$) the evolution operators can be expanded to first order,

$$\begin{aligned}
\hat{R}_{\rm out}(t) &= \hat{U}^\dagger(t,0)\hat{R}_{\rm in}(t)\hat{U}(t,0) &(4.12)\\
&= \hat{U}^\dagger(t,0)\left(\hat{R}_{\rm in}(t) + \sqrt{M}\hat{s}\right)\hat{U}(t,0) &(4.13)\\
&= \hat{R}_{\rm in}(t) + \sqrt{M}\hat{U}^\dagger(t,0)\hat{s}\hat{U}(t,0) &(4.14)
\end{aligned}$$

and separated into meter and system components using the commutation relations for $\hat{R}(t)$ given by $[\hat{R}(t), \hat{R}^\dagger(t')] = \delta(t - t')$. This result shows that the reservoir meter, at the time of detection, is given by a superposition of the input meter and a source term (this is effectively a scattering theory) determined by the system operator, $\hat{s}$.

The detector performs a measurement on the meter Hilbert space at $t \simeq t_i + z_0/c$ by acting on $\hat{R}_{\rm out}(t)$. Since the output operator should be interpreted in the Heisenberg picture, the measurement corresponds to an expectation value,

$$\mathrm{tr}\left[\mathbb{1}_s \otimes \hat{R}_{\rm out}(t)\hat{\chi}(0)\right] = \mathrm{tr}\left[\sqrt{M}\hat{U}^\dagger(t,0)\hat{s}\hat{U}(t,0)\hat{\chi}(0)\right] = \mathrm{tr}\left[\sqrt{M}\hat{s}\tilde{\chi}(t)\right] = \sqrt{M}\langle\hat{s}(t)\rangle \tag{4.15}$$

and the important final result is that the measurement outcome is proportional to the system operator, $\hat{s}$. Here, $\tilde{\chi}(t)$ indicates that the density operator has been expressed in an interaction picture, and therefore the system expectation should also be interpreted in that manner. Note that this result assumes that expectation values over $\hat{R}_{\text{in}}$ vanish, which is an assumption motivated by the fact that this is true in the laboratory for reservoirs in their ground state, or a coherent displacement from the vacuum.

Since it is most common in the laboratory for the meters to be modes (local in time) of an optical field such as a probe laser, the measurement result is often referred to as the *photocurrent*,

$$y(t) = \sqrt{M}\langle \hat{s}(t) \rangle \tag{4.16}$$

since it is acquired using a photoreceiver. Of course, $y(t)$ is the system output that we will use for control.

4.3.1 Conditional Evolution

Eq. (4.15) shows us that the measurements made on the meters provide information about the system through $\langle \hat{s} \rangle$. However, while this tells us the value of the system operator at the time of the interaction, we still need a method for determining the state of the system immediately after the measurement has been performed. To accomplish this, we return to the fact that the evolution is dissected into a sequence of steps, each Δt in length, during which the evolution equations are constant in time, $\hat{H}_{\text{int}} = \hat{s}\hat{R}_{\text{in}}^{\dagger} + \hat{s}^{\dagger}\hat{R}_{\text{in}}$. Therefore, we can integrate over this interaction period to define a set of course-grained operators,

$$\hat{c}_i = \frac{1}{\sqrt{\Delta t}} \int_{t_i}^{t_i+\Delta t} \hat{R}_{\text{in}}(t)\, dt\,. \tag{4.17}$$

A quick check of the commutation relations for these new input operators yields,

$$\left[\hat{c}_i, \hat{c}_i^{\dagger}\right] = \frac{1}{\Delta t} \int_{t_i}^{t_i+\Delta t} \int_{t_i}^{t_i+\Delta t} \left[\hat{R}_{\text{in}}(t'), \hat{R}_{\text{in}}^{\dagger}(t'')\, dt'\, dt''\right] = 1 \tag{4.18}$$

and justifies the $1/\sqrt{\Delta t}$ normalization, as it preserves the correct commutators. Propagating the conditional density operator, $\hat{\rho}_c(t)$, forward in time is accomplished using time-dependent perturbation theory for the full density operator, $\hat{\chi}(t)$ followed by a partial trace over the reservoir Hilbert space. The Born-Dyson expansion of $\hat{\chi}(t_i + \Delta t)$ under a Markov approximation is given by,

$$\hat{\chi}(t_i + \Delta t) = \hat{\chi}(t) + \frac{i}{\hbar}\left[\hat{H}(t_i), \hat{\chi}(t_i)\right]\Delta t - \frac{1}{\hbar^2}\left[\hat{H}(t_i), [\hat{H}(t_i), \hat{\chi}(t_i)]\right]\Delta t^2 + \cdots\,. \tag{4.19}$$

In this expression, it is not yet clear how far we must retain terms in the perturbation theory, however, expanding the commutators,

$$\hat{\chi}(t_i + \Delta t) = \hat{\chi}(t) - i\sqrt{M}\left[\hat{c}_i^{\dagger}\hat{s} + \hat{s}^{\dagger}\hat{c}_i, \hat{\chi}(t)\right]\sqrt{\Delta t} - 2M\left[\hat{c}_i^{\dagger}\hat{s} + \hat{s}^{\dagger}\hat{c}_i, \left[\hat{c}_i^{\dagger}\hat{s} + \hat{s}^{\dagger}\hat{c}_i, \hat{\chi}(t)\right]\right]\Delta t \tag{4.20}$$

shows that the second order terms are of order Δt and are therefore important. Computing third order terms would show that they are of order, $\Delta t^{3/2}$, and since we will ultimately take a formal

limit where, $\Delta t \to 0$, these can be safely ignored. We see that several of the terms here have what seem to be unconventional units, such as the square root of time. Although this arises naturally out of the commutation relations given our definition of the course-grained operators, it is characteristic of stochastic calculus, and rigorous justification for both Eqs. (4.19) and (4.20) requires a theory of stochastic differential equations.

Extracting an expression for the system alone, by taking a partial trace over the reservoir modes, requires that we properly order all the operators acting on the full Hilbert space. This is made much simpler if we assume that each of the meters is in its ground state, such that applying annihilation operators, $\hat{c}_i$, produces zero. Although we will not justify it here, it turns out that this restriction is not confining because the photon energy $\hbar\omega$ is often large. As a result, the electromagnetic field is usually very close to zero temperature. Using, the simplification, $\mathbb{1}_{\mathrm{S}} \otimes \hat{c}_i \hat{\chi} = 0$, leads to,

$$\hat{\chi}(t_i + \Delta t) = \hat{\rho} \otimes \hat{\sigma}_i - i\sqrt{M\Delta t}\left(\hat{s}\hat{\rho}(t) \otimes \hat{c}^\dagger \hat{\sigma}_i - \hat{\rho}(t)\hat{s}^\dagger \otimes \hat{\sigma}_i \hat{c}_i\right) \tag{4.21}$$

$$+M\Delta t\left(\hat{s}\hat{\rho}(t)\hat{s}^\dagger \otimes \hat{c}_i^\dagger \hat{\sigma}_i \hat{c}_i - \frac{1}{2}\hat{s}^\dagger \hat{s}\hat{\rho}(t) \otimes \hat{c}_i \hat{c}_i^\dagger \hat{\sigma}_i - \frac{1}{2}\hat{\rho}(t)\hat{s}^\dagger \hat{s} \otimes \hat{\sigma}_i \hat{c}_i \hat{c}_i^\dagger\right) \tag{4.22}$$

where, as a reminder, σ_i is the density operator for the i^{th} reservoir mode. This master equation is fairly typical of ones that we regularly encounter in quantum optics— it has terms which look like Linblad decay superoperators,

$$\mathcal{D}[\hat{s}]\hat{\rho}(t) = \hat{s}\hat{\rho}(t)\hat{s}^\dagger - \frac{1}{2}\left(\hat{s}^\dagger \hat{s}\hat{\rho}(t) + \hat{\rho}(t)\hat{s}^\dagger \hat{s}\right) \tag{4.23}$$

meaning that we should expect the measurement to (if nothing else) cause decoherence.

Taking the partial trace of this expression and evaluating the formal limit $\Delta t \to dt$ requires that we specify the particular measurement, such as photon-counting or quadrature detection in order to proceed. Since most proposed quantum feedback control experiments are expected to use a coherent laser field in order to perform the weak measurement, we will assume a quadrature detection model such as proposed by Wiseman and Milburn.[2] We now need to take a small leap of faith— mathematically this involves applying stochastic calculus to take the $\Delta t \to dt$ limit

$$d\hat{\rho}_c(t) = M\mathcal{D}[\hat{s}]\hat{\rho}_c(t)dt + \sqrt{M}\left(\hat{s}\hat{\rho}_c(t) + \hat{\rho}_c(t)\hat{s}^\dagger - \mathrm{tr}[(\hat{s} + \hat{s}^\dagger)\hat{\rho}_c(t)]\hat{\rho}_c(t)\right)dW \tag{4.24}$$

$$= M\mathcal{D}[\hat{s}]\hat{\rho}_c(t)dt + \sqrt{M}\mathcal{H}[\hat{s}]\hat{\rho}_c(t)dW(t) \tag{4.25}$$

to give the conditional master equation. We see that there is a prominent Linblad term that corresponds to deterministic decay of the system's coherence with rate, M. There is also an additional stochastic term. The $dW(t)$ is referred to as an Itó increment, and it is a random process that is Gaussian distributed with mean zero, and variance dt. In the language of stochastic calculus, the symbol $E[\cdots]$, refers to an average over many different stochastic trajectories. The stochastic increments have the property that $E[dW(t)] = 0$ and that $E[dW^2] = dt$. The important distinction between normal differential and Itó stochastic calculus is that the chain rules are different. We will not pursue stochastic differential equations in any depth in this course, only to say that Eq. (4.24) provides a recipe for propagating the conditional density matrix. In a simulation, one would

numerically integrate the stochastic master equation by randomly selecting values of $dW(t)$ from a Gaussian distribution with variance δt (the size of the integration step) during each integration step.

4.3.2 The Conditional Innovation Process

We will simply state the quantum trajectory result for the photocurrent,

$$y(t)dt = 2\sqrt{M}\langle\hat{s}\rangle_c(t)dt + dW \tag{4.26}$$

where the $dW(t)$ term now has a physical interpretation. It is the optical shotnoise on the probe laser used to perform the weak measurement, and is defined by,

$$dW(t) \equiv y(t)dt - 2\sqrt{M}\langle\hat{s}\rangle_c(t)\,. \tag{4.27}$$

Although this definition appears to be circular given Eq. (4.26), Eq. (4.27) should be taken as more fundamental. It has the following physical interpretation: Given knowledge of the initial system state, $\hat{\rho}(0)$, one can propagate the system density operator forward in time from $t = 0$ to $t = dt$ according to Eq. (4.24). At this point, a measurement of $\langle\hat{s}\rangle$ is performed; however, due to backaction and shotnoise, it will not exactly agree with the computed expectation value. Of course, we should trust the measurement over our stochastic differential equation, and therefore we need to correct the evolution equation by adding in the effect of the residual, or the difference between the measurement and computed value, $y(t)dt - 2\sqrt{M}\langle\hat{s}\rangle_c(t)dt$. This evolution is known as the *innovation* process, because at each time when we compare the measurement to our predicted value, we are a little bit surprised that the two do not agree. This surprise, of course, results from the inherent randomness of the quantum measurement process.

4.4 Quantum Feedback Control

Most of the hard work has already been done in developing the quantum trajectory equations of motion for the system under continuous evolution. We see, that as expected, the measurement produces a certain randomness to the evolution in the form of a stochastic differential equation and a noisy photocurrent. However, this process does provide a continuous stream of information about the system operator, $\langle\hat{s}\rangle_c(t)$, making it reasonable to consider a feedback scheme that controls $\langle\hat{s}\rangle$ by driving the quantum system with a control signal computed from the photocurrent. The basic theory of feedback control using quantum trajectory theory was introduced in a set of landmark papers by Wiseman and Milburn that are summarized in Howard Wiseman's dissertation.[5]

This additional driving term appears in the stochastic master equation as a feedback Hamiltonian, $\hat{H}_{\rm Fb}$,

$$d\hat{\rho}_c(t) = -i\left[u(t)\hat{H}_{\rm Fb}, \hat{\rho}_c(t)\right]dt + M\mathcal{D}[\hat{s}]\hat{\rho}_c(t)dt + \sqrt{M}\mathcal{H}[\hat{s}]\hat{\rho}_c(t)\,dW(t) \tag{4.28}$$

where the feedback strength, $u(t)$, is computed from the past history of the photocurrent,

$$u(t) = \int_0^t F(t-\tau) I_c(\tau)\, d\tau\,. \tag{4.29}$$

Remembering all the way back to Lecture 2, this equation looks very similar to Eq. (2.8). The difficulty, however, is that Eq. (4.29) involves a stochastic quantity, and computing the convolution integral cannot be relegated to a steady-state transfer-function analysis, as was done in classical control. The conditional evolution of the quantum system is not a steady-state process, and instead, a quantum feedback controller would need to integrate Eq. (4.29) in real-time. The difficulty is that performing the feedback convolution integral is computationally difficult because it requires propagating the stochastic master equation— the evolution reflects a different measurement history on each trajectory because of inherent randomness. Current proposals for quantum feedback control attempt to circumvent this problem in two ways: one is to use high-speed digital processing [6] and the other is to employ quantum filtering techniques [7–10] to simplify the real-time convolution integral. These more advanced quantum control topics are something we may consider discussing tomorrow in the final "special topics" lecture.

References

[1] H. Carmichael, *An open-syemtems approach to quantum optics* (Springer Verlag, Berlin, 1993).

[2] H. M. Wiseman and G. J. Milburn, Phys. Rev. A **47**, 642 (1993).

[3] C. Gardiner, *Quantum Noise* (Springer Verlag, New York, 1991).

[4] C. W. Gardiner, *Handbook of Stochastic Methods* (Springer, New York, 1985), 2nd ed.

[5] H. Wiseman, Ph.D. thesis, University of Queensland (1994).

[6] M. A. Armen, J. K. Au, J. K. Stockton, A. C. Doherty, and H. Mabuchi, Phys. Rev. Lett **89**, 133602 (2002).

[7] H. Mabuchi, Quantum Semiclass. Opt. **8**, 1103 (1996).

[8] V. Belavkin, Rep. on Math. Phys. **43**, 405 (1999).

[9] F. Verstraete, A. C. Doherty, and H. Mabuchi, Phys. Rev. A **64**, 032111 (2001).

[10] J. Geremia, J. K. Stockton, A. C. Doherty, and H. Mabuchi, quant-ph/0306192 (2003).

Lecture 5

Special Topics in Quantum Control

This lecture is intended to focus on special topics in quantum control. As such, the exact material that we will cover is something that we will decide during the course of the Lecture series. I am certainly open to suggestions for the exact format of this lecture; however, I though we might consider reviewing a paper or two

1. **Heisenberg Limited Magnetometry** Thomsen, Mancini and Wiseman[1] have shown that it is possible to produce spin-squeezing using feedback. This technique can be used, in conjunction with quantum Kalman filtering, to achieve the quantum noise limit for measuring the magnitude of a magnetic field.[2]

2. **Adaptive Measurements** It was shown by the group of Wiseman that it is possible to reach the quantum noise limit for phase measurement of a weak optical pulse with unknown initial phase by performing an adaptive measurement.[3] This procedure was then demonstrated at Caltech.[4]

3. **Classical Control of atomic Motion** There have been several recent experiments in which control theory has been used to affect the motion of atoms in an optical lattice. These include the recent papers from the Raithel Group, *Feedback Control of Atomic Motion in an Optical Lattice*,[5] and the Rempe group, *Feedback on the Motion of a Single Atom in an Optical Cavity.*[6]

In particular, it may be best to focus on the differences between these papers, focusing on what it actually means for the experimental demonstration of *quantum control*, a feet which has still not yet been conclusively demonstrated.

References

[1] L. K. Thomsen, S. Mancini, and H. M. Wiseman, Phys. Rev. A **65**, 061801 (2002).

[2] J. Geremia, J. K. Stockton, A. C. Doherty, and H. Mabuchi, quant-ph/0306192 (2003).

[3] H. Wiseman and R. Killip, Phys. Rev. A **57**, 2169 (1998).

[4] M. A. Armen, J. K. Au, J. K. Stockton, A. C. Doherty, and H. Mabuchi, Phys. Rev. Lett **89**, 133602 (2002).

[5] N. V. Morrow, S. K. Dutta, and G. Raithel, Phys. Rev. Lett **88**, 093003 (2002).

[6] T. Fischer, P. Maunz, P. W. H. Pinkse, T. Puppe, and G. Rempe, Phys. Rev. Lett **88**, 163002 (2002).