Control Theory Is Not For Control Freaks

Mitch Gooding

November 16, 2012

Control theory is a field of mathematics and engineering concerned with manipulating the output of dynamical systems. This might sound like the perfect field of study for a control freak, however, the reader from a pure mathematical background, who is familiar with calculating exact solutions and rigorous proofs, should beware. It is only very rarely that the systems studied in Control Theory can be controlled perfectly. One of the most important tasks in this subject is deciding what can't be controlled, or what is impractical to control, and focusing on what can. This article aims to highlight these kinds of necessary compromises in the level of control a control theorist must make. Before we do this though, some background must be covered.

Background

A dynamical system is the representation of some physical system which changes the *state* it is in over time. An example of such a dynamical system is a car driving under cruise control. The physical object in the system is referred to as the plant, and the state could be any measurable property of the plant, such as its position, orientation, or temperature. Usually, the state is a list (or vector) of several important properties, rather than a single measurement. If there are n important properties, the state is represented as a function of time, $x(t) \in \mathbb{R}^n$. In the cruise control example, the plant is the car, and depending on the detail of the representation, the state could be as simple as the number of revolutions per second the engine is performing, or a vector of values describing the velocity and orientation of each of the car's wheels. The properties of the plant that we want to control are referred to as the output, denoted $y(t) \in \mathbb{R}^p$ for some $p \in \mathbb{N}$. In the case of this example, the output is the net velocity of the car.

For the dynamical system to be relevant to control theory, there must be some way for the system to affect the state it is in; this process is known as the control. In the car example, the cruise control of the car can affect the car's velocity by controlling the acceleration of the engine. The affect the control has on the plant is referred to as the input of the system, denoted $u(t) \in \mathbb{R}^m$. All that is needed to model this dynamical system is an equation or system of equations which describe the evolution of the state with time, and an equation or system of equations which maps the state to the output [1, Chapter 2.2]. These dynamic systems can be represented visually with a diagram known as a block diagram, such as in Figure 1.

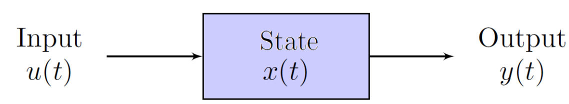

Figure 1: A block diagram illustrating an uncontrolled system.

SISO Systems

Even in the very simplest of systems, a control theorist must make compromises over the aspects of control that can be taken over the system. An example of such a system is the linear, time-invariant (LTI) single-input single-output (SISO) system. In these systems, the the dimension of the input and output is 1, and the only state variable of interest is the output variable. Such systems can usually be described by a single linear ODE relating $y(t)$ and $u(t)$. Taking the Laplace transform of this ODE and rearranging for the $Y(s)$ yields a function of the form

$$Y(s) = H(s)U(s)$$

for some function $H(s)$. This function is known as the transfer function of the system. In Figure 1, as there is no control, the input is just the default effect the environment has on the plant. This uncontrolled environmental signal is referred to as the *reference*, denoted $r(t)$. One way to control such a system is for a control (G) to be placed before the plant, which converts the reference into the input of the plant. For the control to have information about the output of the plant, part of the output is fed back into the control. A second control (F) can be placed on the feedback signal, to yield the block diagram in Figure 2. Taking the Laplace transform of this system now gives the equation

$$Y(s) = H_c(s)R(s) = R(s)\frac{G(s)H(s)}{1+F(s)G(s)H(s)}.$$

The function $H_c(s)$ is called the closed loop transfer function [1].

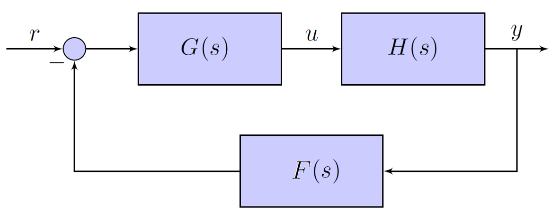

Figure 2: A plant with closed loop control.

To control this system, the controls G and F must be chosen so that certain reference signals produce some desired output. A sudden change in reference value is a signal which can have a acute effect on a plant, so a common reference signal that needs to be controlled is a step function. Two desirable aspects of the step function response is that the output adjusts to the new value quickly, but without overshooting the new value drastically. The time taken to adjust to the new value is called the rise time, t_r, and the percent that the response overshoots the final value is denoted M_p.

Unfortunately, even in a system as simple as this, it is rarely that both these goals can be satisfied. A low t_r usually implies a high M_p, and vice versa. The control theorist must tune the control so that an appropriate balance is achieved between these two goals. This is a common theme of control theory; a compromise must be made between different desired outcomes of the system being controlled.

MIMO systems

Another lesson which must be learnt in control theory is that some aspects of some systems cannot be controlled. An example of a system that suffers from this setback is the multi-input multi-output (MIMO) system described by

$$\begin{cases} \dot{x}(t) = Ax(t) + Bu(t) \\ y(t) = Cx(t) + Du(t) \end{cases} \tag{1}$$

where $A \in \mathbb{R}^{n\times n}$, $B \in \mathbb{R}^{n\times m}$, $C \in \mathbb{R}^{p\times n}$ and $D \in \mathbb{R}^{p\times m}$. To have complete control over the plant, there needs to be an input, which takes arbitrary state $x_s \in \mathbb{R}^n$ to arbitrary state $x_d \in \mathbb{R}^n$. If for any state $x_s \in \mathbb{R}^n$, there is an input which takes x_s to the origin in finite time, the state x_s is called controllable, and if for any state $x_d \in \mathbb{R}^n$, there is an input which takes the origin to x_d in finite time, the state x_d is called reachable. A system is controllable/reachable if every state is controllable/reachable. It can be shown that if a state is reachable, it is controllable. Unfortunately, not all systems are controllable. A system is controllable if and only if [1, Chapter 6.2]

$$rank\left(\left[A, AB, \ldots, A^{n-1}B\right]\right) = n.$$

While feedback control will result in quite sophisticated control for controllable systems, the same level of control cannot be achieved in an uncontrollable system. The control theorist must make another compromise, either the system must be modified in some way to make it controllable, or the goal of complete control over the system must be abandoned, and replaced with a less ambitions goal, such as keeping the system bounded. Such a goal involves designing the system so that the each state does not diverge to infinity, but the systems stays within some finite subspace of the total state space.

A similar problem is the Linear Quadratic Regulator (LQR). In this problem, we have the same ODE describing the state variable, however, our objective is no longer to control the output y, but to minimise the cost functional

$$\int_0^T (x(t)'Qx(t) + u(t)'Ru(t)) \ \mathrm{d}t \ + x(T)'Q_f x(T), \tag{2}$$

where T is the *time horizon* (the length of time the process will run for), Q is a symmetric nonnegative definite matrix, and Q_f and R are symmetric positive definite matrices. Problems like this arise often. For example, the plant may be a vehicle, which needs to travel to a specific location by time T. Then $x(t)$ could represent its position, and Q maps the position to fuel consumption. The cost (in terms of fuel consumption) of performing different maneuvers in this vehicle could be represented by R, and the input u is the sequence of maneuvers which the vehicle is instructed to do in order to reach the required location. The control which results in the lowest fuel consumption but still moves the vehicle to the required location is desired.

Like the previous problem, there is some very powerful theory which shows that the optimal solution to the LQR problem is the input

$$u(t) = -R^{-1}B'P(t)x(t), \tag{3}$$

where the matrix $P(t)$ can be found by solving the Riccati equation

$$-\dot{P}(t) = A'P(t) + P(t)A - P(t)BR^{-1}B'P(t) + Q$$

with boundary condition

$$P(T) = Q_f.$$

This Riccati equation is fairly simple to solve approximately through numerical approximation, which makes this is an extremely useful theory. Once again, however, the control theorist must abandon hopes for calculating the exact optimal solution u, as the Riccati equations in general do not have an explicit solution [2, Chapter 3.12].

These examples of problems from control theory should illustrate that the goal of achieving perfect control of a system is unnecessary and impractical. Once this goal is abandoned, and a compromise is made on the level of the control over the system, the theory for controlling such systems to achieve these more practical goals is quite powerful.

Physical limitations

The theory discussed so far can control the many idealised problems mentioned so far quite accurately, albeit with some limitations. However, in the real world, measurements have errors, and these errors may cause otherwise stable controls to diverge. This is the ultimate reason for why a control theorist cannot be a control freak; despite the most comprehensive theory, the limitations of physical problems would render perfect control useless. However, there are a number of ideas from control theory developed to deal with these limitations.

So far we have assumed that any mathematically valid control is possible. However, all possible states and controls may not be practical. In the cruise control example, the car has an upper bound on its acceleration, so an input prescribing great rates of change is not possible. To deal with situations like this, a method known as model predictive control (MPC) can be used. Before describing this method, it should be noted that unlike the previous control techniques discussed, which calculate the control 'before' the dynamical system is run (offline), this method of control updates the control rule while the process is running (online). Consequently, this method is used in petro-chemical industries, where reactions are 'slow enough' for these online calculations can be made in time to be implemented [4].

Consider the LQR problem with the added constraint

$$F \begin{bmatrix} x(t) \\ u(t) \end{bmatrix} \leq b$$

for some matrix F and some vector b. The original control for an LQR problem can be found, but if the control reaches the boundary of the constraints at some time τ, this control may no longer be optimal after τ. Using MPC, the time scale is discretised in to time steps (which will be assumed to be size 1), and a time horizon, $N < T$, is chosen. At the first time step, the optimal control only until time N is calculated, subject to constraints. Because the time horizon has been restricted from T to N, the calculated control will not in general be the same as the optimal control over all of T. However, this control is only used for the first time step, and at the next time step the optimal control for the next N time steps is calculated. This method of control is sub-optimal, but can result in more sensible solutions in physical applications, as it avoids the overambitious task of predicting the optimal control over the whole duration of the process.

Plant uncertainty

Another major physical limitation is that measurements have errors, and the parameters that model these systems have uncertainties. These errors, however small, can build upon each other to result in significant inaccuracies. One way to avoid this problem is to add a noise term to the dynamic system model. For example, consider the LQR problem in discrete time, with the following modification to the dynamic system model (1):

$$\begin{cases} x(n+1) &= Ax(n) + Bu(n) + \xi_x(n) \\ \quad y(n) &= Cx(n) + Du(n) + \xi_y(n), \end{cases}$$

where $\xi_x(n)$ and $\xi_y(n)$ are random processes, which represent the random errors in the measurements. Naturally, this implies that the state and output variables $x(n)$ and $y(n)$ are random processes also. It is reasonable to assume that the $\xi_i(n)$ as sequences are independent, and by the assumption of time-invariance, identically distributed. Furthermore, the errors at each time step are assumed to be the sum of many tiny independent disturbances, so by the central limit theorem, it may be assumed that $\xi_i(n)$ are Gaussian random variables. If there is no net drift in the value of the errors, the mean of these Gaussian variables will be assumed to be 0, and the covariance of these variables are denoted Σ_x and Σ_y. This problem is known as the linear quadratic Gaussian regulator (LQG), and the objective is to find the $u(n)$ which minimises the expected value of the discrete version of (2),

$$J[u] = \mathbb{E}\left(\sum_{i=0}^{T} \left(x(i)'Qx(i) + u(i)'Ru(i)\right)\right).$$

In order to control this system, the true value of the state must be estimated. To do this, a technique known as Kalman filtering is used. This method uses minimum mean squared estimation of the state variables for time steps from 0 to N based on the known input and output variables $y(i)$ and $u(i)$, for $i \in \{0, \dots, N-1\}$. It turns out that the optimal control for a LQG problem is to use Kalman filtering to estimate the state, and then use same control (3) derived for the deterministic LQR problem on the estimated system [3, Chapter 4.3.2].

There are a number of other methods for controlling a plant with uncertainty, rather than model the dynamic system with noise. One of these methods is robust control. The goal of this method is to be able to control the plant even in the 'worst case' scenario. To do this, the system is modeled by

$$\begin{cases} \dot{x}(t) = (A + \delta G)x(t) + Bu(t) \\ y(t) = Cx(t) + Du(t), \end{cases} \tag{4}$$

for some arbitrary matrix G and some small δ. If the uncertainty in the plant is bounded by some magnitude ϵ, then if a control is found for (4), for any small $\delta < \epsilon$ and any G in some class of matrices, then the plant with uncertainty will also be controlled by this control. Thus, the plant will be controlled despite the uncertainty.

Another method for controlling a plant with uncertainty is adaptive control. This method is especially used for systems where the uncertainty is in the model itself, i.e. the parameters A, B, C, D in

(1) are unknown. Adaptive control methods can estimate these parameters while calculating the optimal control online, but in some cases these parameters do not need to be estimated. Finally, the uncertainty in the system may not just come from uncertain parameters, but the dynamic system being modeled may be inherently stochastic. Arrival processes, such as in telecommunications and birth and death processes are examples of such inherently stochastic systems. In these systems with uncertainty, the goal of exercising exact control over the system is a faint memory.

Conclusion

This has been a very brief outline of some of the problems encountered in control theory. Despite the name, one of the most important lessons in this theory is that perfect control can only rarely be achieved, and the vast majority of the time, a compromise in the detail of the control must be made. However, when the right compromises are made, very powerful theory can be derived, which result in the near optimal control of many real world systems.

References

[1] Panos J. Antsaklis and Anthony N. Michel. *A linear systems primer.* Birkhäuser Boston Inc., Boston, MA, 2007.

[2] Donald E. Kirk. *Optimal control theory: an introduction.* Prentice-Hall, Englewood Cliffs, N.J., 1970.

[3] Huibert Kwakernaak and Raphael Sivan. *Linear optimal control systems.* Wiley-Interscience [John Wiley & Sons], New York, 1972.

[4] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. M. Scokaert. Constrained model predictive control: stability and optimality. *Automatica J. IFAC*, 36(6):789–814, 2000.