

Control theory at a glance

Tinos Nitsopoulos

November 13, 2012

Abstract

Control theory covers many concepts and tools widely used in engineering to design robust, optimal and stable systems. However the theory, being inspired by our interaction with nature, results in many intriguing theoretical results. In this paper I survey the field by covering topics such as system linearisation, frequency space representations, cost minimisation (Euler and Hamilton equations) and stochastic conditional expectation.

1 State Space

In general we deal with a multi-valued system x , and multi-valued output, y . We use the latter to deduce information about the former, so that we can effectively control it with a multi-valued input, u . The system changes in accordance with its current state, the input, time and stochastic noise. At an introductory level, we assume noise is negligible, and that all relations are linear with time being irrelevant. That is, we assume local linearisation is possible using the Jacobian (derivative). We start with vector equations

$$\begin{aligned}\dot{x} &= f(x, u, \xi_x, t) \\ y &= g(x, u, \xi_y, t)\end{aligned}$$

where ξ is the noise, and this becomes, after a shift,

$$\begin{aligned}\dot{x} &= Ax + Bu \\ y &= Cx + Du\end{aligned}$$

with the matrices given by

$$\begin{aligned}A &= \partial_x f \\ B &= \partial_u f \\ C &= \partial_x g \\ D &= \partial_u g \\ \xi_x &= \xi_y = 0\end{aligned}$$

A is square though the rest need not be. Of course, systems in which the attempted linearisation gives zero matrices are suspicious, since the non-linearity dominates in that case. Note we have also assumed no memory; i.e. system depends only on current state. This standardisation is useful as a standard set of algorithms exist for solving this (A, B, C, D) system.

Note higher order differential equations are reducable to first order by using $\begin{pmatrix} x \\ \dot{x} \end{pmatrix}$ as x , etc.

All equations can be written analogously in the discrete case, too. In that case we write equations for (x_{n+1}, y_{n+1}) rather than the derivatives. That is, conventionally we **don't** use difference equations.

As is well known in pure maths, the solution to the linearised problem with initial conditions is guaranteed to exist via Picard iterations. i.e.

$$x_{n+1}(t) = x_0(t) + \int^t d\tau (A(\tau)x_n(\tau) + B(\tau)u(\tau))$$

By recursion with time invariance we obtain

$$x = e^{At}x_0 + \int_0^t d\tau e^{A(t-\tau)}Bu(\tau)$$

Exponentiating a matrix is defined by the usual exponential Taylor series. The second term is a convolution, $e^{At} * Bu$. This hints that frequency space may be a useful representation of the system.

But first a few notes regarding the matrix exponentiation. We have the following equalities:

$$e^{At} = Pe^{Jt}P^{-1} = \sum_{i=0}^{\infty} \frac{(tA)^i}{i!} = \sum_{i=0}^n A_i e^{\lambda_i t} = \mathcal{L}^{-1} \left((s - A)^{-1} \right)$$

The first equality uses the Jordan normal form (JNF) J , and transition matrix P (which takes JNF space to real space). Exponentiating the JNF is easy (block diagonal), and converting to JNF is also easy, versus calculating the second equality.

The third equality means a matrix can be split up into its modes A_i . There are n such modes, which is the the state space dimension. λ_i are the eigenvalues of A . This representation follows from the Cayley-Hamilton theorem; that the eigenvalue equation is true for the matrix itself.

The fourth equality shows that Laplace transforms (and the inverses) can be performed on matrix expressions. This is sometimes the easiest route for calculating the exponential if a concrete prediction is required. For scalars, $\mathcal{L}^{-1} \left(\frac{1}{s-\alpha} \right) = e^{\alpha t}$ is a well known result.

Asymptotic stability is achieved iff all e-values have negative real part. This is defined as a guarantee that any initial condition will eventually lead to zero. A temporary input (say transient noise) can be seen as such an initial state. Stability alone requires just that the e-values have non-positive real part, and that any e-value with zero real part has the same geometric and algebraic multiplicity. Geometric multiplicity is the number of e-vectors corresponding to the e-value (what physicists sometimes call degeneracy). Algebraic multiplicity is how many times the e-value appears in the total e-value equation. The latter is always less than or equal to the former. If it is less then instability occurs due to too much geometric multiplicity terms (polynomial instability, rather than exponential).

Often an engineer wishes to make the system stable by moving the e-values to the LHP (left-half plane). This can be achieved in this linearised regime with the use of feedback controllers. In this case, we set

$$u = Fx + \text{other input}$$

where F is the linear feedback controller, which can be represented as a matrix. This means $A \rightarrow A + BF$, and so the F must be defined so as to make the e-values of this new matrix are in the LHP.


There is a theorem which states it is possible to design such an F iff $\text{con}(A, B)$, the controllability matrix, is full rank (i.e. rank n). The controllability matrix is defined by

$$\text{con}(A, B) = \begin{pmatrix} B & AB & A^2B & \dots & A^{n-1}B \end{pmatrix}$$

If it is not full rank, that means some subspace of state-space will be uncontrollable (we are unable to reduce it to zero with feedback). There is a systematic algorithm for generating F if the controllability matrix is full rank (see Ref. ?).

2 Frequency Domain

Rather than analysing a scalar system of high order as a vector system of single order, we can alternatively model just the scalar system. In this pursuit we have a plant H and control system (G_1, G_2) . The controlled plant is $\frac{G_1 H}{1 + G_1 G_2 H}$ (this is quickly derived from the figure). Negative feedback G_2 is usually set at a constant.



One simple application of this is to control the gain of a transistor. If the transistor is maxed out, the amplification will be very choppy. H is the transistor gain and the controls are voltage dividers. Assuming the $G_1 H$ is sufficiently high, $G_1 G_2 H \gg 1$, then the controlled transistor is just $\frac{1}{G_2}$. So we control the transistor gain with the smallness of G_2 .

What is this H ? It is known as the transfer function of the system and equals $\mathcal{L}(h)$, where h is the impulse response. That is, the response to a sharp input (δ input). The (A, B) system from before can be viewed simply as a convolution acting on the input. This is represented in the figure as a multiplication, because of the convolution theorem with frequency transforms (convolutions are equivalent to multiplication in frequency space). Hence $\mathcal{L}^{-1}\left(\frac{G_1 H}{1 + G_1 G_2 H}\right)$ gives us the total output. (G_1, G_2) need to be designed so that the poles of the controlled plant are in the LHP, because of the laplace inverse (in the simplest case) is $e^{\lambda_i t}$.

One of the oldest and most commonly used controllers is the PID (proportional, integral, derivative). This one uses constant G_2 (a subsystem which only changes an input δ by a constant), and a $G_1(s)$ such that

$$G_1 = P + Ds + \frac{I}{s}$$

Note that this can be easily made equivalent to a state space representation $(A + BF, B)$ with $n = 3$. B is G_1 and F is G_2 . However the frequency space representation is very useful because it allows us to conduct frequency analysis. This often involves Bode and Nyquist plots. Bode, for example, plots the magnitude and phase of the controlled plant versus the frequency of a sinusoidal input. Where this peaks we see which frequencies of noise (input) the system is most sensitive to. The phase tells us the time-lag between input and its affect on the system state. Space does not permit a longer discussion.

Another useful property is the final value theorem,

$$\begin{aligned} \lim_{s \rightarrow 0} sH &= h(\infty) \\ \lim_{s \rightarrow \infty} sH &= h(0^+) \end{aligned}$$

The first formula can be used to find the steady state error if we desire a certain input to be reduced to zero. For instance, if we have a constant input to have no affect, we need $sH_c \rightarrow 0$ as $s \rightarrow 0$, where H_c is the controlled plant transfer function. Note in time domain we can convolve h with a step input, which gives us important stability parameters such as the rise time, settling time, and overshoot.

3 Observations

Going back to the state space representation; there is another issue. We cannot measure x directly, but only y . How, then, do we estimate x ? We call this estimation the Luenberger observer, $\hat{x}$. We generate it by guessing an initial value, and then setting

$$\begin{aligned} \dot{\hat{x}} &= A\hat{x} + Bu + K(y - \hat{y}) \\ \hat{y} &= C\hat{x} + Du \end{aligned}$$

$\hat{y}$ is the predicted output were we correct in our initial guess. K tells us how to alter the guess in order to determine the correct state. Is a state always observable? That is, can we always generate a K that will always take any output and produce an asymptotically accurate Luenberger observer? The answer is yes iff the observability matrix is full rank. The $\text{obs}(A, C)$ matrix is

$$\text{obs}(A, C) = \begin{pmatrix} C \\ CA \\ CA^2 \\ \vdots \\ CA^{n-1} \end{pmatrix}$$

where this matrix has the submatrices embedded into it. Again there is are systematic algorithms for generating the K in the discrete and continuous cases when the obs matrix is full rank.

An important point is that there is an independence between controlling the system and observing it. When both F & K are in our system, they do not affect each other's success, and so may be designed separately. K becomes the Kalman filter used in stochastic modelling when the ξ are introduced.

4 Optimality

Thus far we have considered only predictability, stability, controllability and observability. But what if there is a cost to be minimised in a given system or problem? There are many examples, such as that the electrons in a piece of metal move so as to minimise the amount of heat produced. One way of representing the problem is to write

$$J(x(t), t) = J_T + \int_t^T d\tau g(x(\tau), u(\tau), \tau)$$

where J is the total optimum cost, J_T is the cost at the end-point, and g is the cost per unit time, and u is the optimum input. g may depend on the trajectory the system is taking, and the input at that time.

There is an important principle, the Bellman principle of optimality. It states that the minimum cost at a future time is the minimum cost now less the cost to get to that time with the optimum input. The goal is to use that principle to obtain an equation for what u should be.

One approach is to derive the HJB equation (Hilbert-Jacobi-Bellman). A simple derivation is to simply use the chain rule in differentiating the cost w.r.t. (with respect to) time:

$$-g = J_t + J_x \dot{x}$$

Both sides are the total time derivative of J . The minus is just due to how we've represented the cost. One common way to write it is as

$$\begin{aligned} 0 &= J_t + \min_u H(x, u, t) \\ H &= g + J_x \dot{x} \end{aligned}$$

The H here is the Hamiltonian, and corresponds exactly to the energy function known to physicists.

The HJB formulation is complicated by two things; it does not do well with higher dimensional trajectories, and it involves the overall cost J . The calculus of variations (Euler's equation, or Lagrange's equation to physicists) does not suffer from these deficits. In that formulation we do not minimise J directly, because there are constraints that must be satisfied in higher dimensions (e.g. speed and position are related to each other). The constraint equation we write as

$$\dot{x} = a(x, t)$$

Then we minimise the integral over $g + p \cdot (a - \dot{x})$, where $p(t)$ is the momentum vector (can be set exactly to the quantity in physics). That is, momentum is just a set of Lagrange multipliers. Minimising the composite quantity is a necessary condition. If it were not minimised, one could vary the trajectory without violating the constraint yet still minimise J further. By "varying" trajectory x by δx (a small change at each time that leaves the path continuous), one can easily derive the Euler equation via integration by parts of δJ :

$$\begin{aligned} \partial_x g_c &= d_t \partial_{\dot{x}} g_c \\ g_c &= g + p \cdot (a - \dot{x}) = \text{constrained marginal cost} \end{aligned}$$

d_t is the scalar total time derivative, and ∂_x & $\partial_{\dot{x}}$ are vector derivatives. If there are n dimensions, and m constraints, we now have $n + m$ equations (Euler gives n and the constraints the other m) which yield the m unknown momentum components, as well as the n unknowns in the optimum trajectory.

Note it is necessary in field theory (e.g. the gravitational field) to have an optimum scalar field (in the physics sense). That is, one can have multiple "time" variables if you want to optimise a surface or volume manifold, rather than just a linear trajectory.

How does one obtain the Hamiltonian? This is just

$$\begin{aligned} H &= g + p \cdot a \\ \dot{x} &= \partial_p H \\ \dot{p} &= -\partial_x H \end{aligned}$$

The above are Hamilton's equations, and fully encapsulate Euler's equation and the constraints (this is quickly derived). This abstract but direct connection between Hamiltonian and Lagrangian mechanics is often not taught to physicists. To give a couple examples, the most famous Lagrangian (g) is that for motion of a beam of light, 1. A more complex example

is that to derive Newton's laws one uses the difference between kinetic and potential energy as the Lagrangian. This is easily derived from the first of the three equations just listed. Note that the input is obtained via

$$\partial_u H = 0$$

which is obtained via HJB or direct variation w.r.t. δu . The formalism is often used with the form

$$g = x^T Q x + u^T R u$$

with Q non-negative (all e-values non-negative) and R strictly positive. This problem is known as LQR (linear quadratic regulator). The regulation is via the cost, which is quadratic. The solution for u is linear in x . The feedback matrix is given by the solution of a Riccati equation, which can be directly derived from Hamilton's equations. A Riccati equation is a first order ODE which depends quadratically on the dependent function.

Physicists may find this form useful if we ignore the R because then it corresponds exactly to a simple harmonic oscillator (quadratic energy and kinetic energy). Alternatively x may be viewed as a finite dimensional wavefunction (i.e. one defined over finite discrete states). The fact that momentum seems to correspond to the J_x from above (in HJB) hints at the correspondence between momentum and differentiation (in quantum mechanics, the momentum operator in position space corresponds to $-i\partial_x$).

5 Stochastic Noise

If we introduce (ξ_x, ξ_y) into our deterministic (A, B, C, D) system we must solve an optimality problem to get the conditional expectation of the Luenberger observer given all the previous output, which is itself stochastic. The "cost" is defined as the MSE (mean square estimation). That is, $\mathbb{E}(|x - \hat{x}|^2)$. In the discrete case, the solution for $\hat{x}_n$ is just $\mathbb{E}(x_n, Y(n))$, where $Y(n)$ is all the output up to step n . In the simplest case we assume all the noise is independent of all the other noise. i.e. $\Sigma_{xy} = 0$ where Σ_{xy} is the covariance of the x & y noise.

The solution for the Luenberger K is obtained from a Riccati equation (see Ref. 3 for derivation). This is called the Kalman filter in the LMMSE case (linear minimum MSE). Linear because the ODEs involved are linear.

6 Hybrid Systems

Suppose we have a continuous system, such as the temperature in a room. But there is a discrete component to the system, such as the digital temperature setting on an air-conditioner. In this case we have a discrete and continuous state space. i.e. x has two components, temperature and temperature setting. The former is in $\mathbb{R}$, while the latter only in $\mathbb{N}$, and so a matrix ODE makes no sense. Optimising, controlling and predicting such hybrid systems forms an area of active research.

7 Conclusion

It is clear that control theory covers many topics which are of fundamental importance in engineering and science. The subject requires an understanding of many areas of mathematics; such as statistics, calculus and linear algebra. Many of the proofs involved are ingenious, and control theory provides a vital link between many applied fields and the pure mathematics from which they borrow ideas.

8 References

Antsaklis, Panos J., Michel, Anthony N., *A Linear Systems Primer*, Birkhauser, 2007. Covers many control topics in great depth.

Kwakernaak, Sivan R, *Linear Optimal Control Systems*, Wiley Interscience, 1972. Contains inverted pendulum example of feedback design.

Nazarathy, Yoni, *MATH4406 Course Notes*, UQ, 2012. Covers all mathematical topics with examples.