Control Theory

Reading and Communication Club

Cambridge, 12th January 2012

Roger Frigola

Agenda

- Essentials
- Design Objectives of a Control System
- Feedback
- State Estimation
- Identification and Adaptive Control
- Optimal Control
- A General Controller Structure
- Bibliography

What is Control?

- Sensing + Computation + Actuation
- Goals: Performance, Stability, Robustness.
- Demo

Canonical Control Problems

- Regulator problem: we want to keep the system at some nominal operating point regardless of external disturbances.

- End-point control problem: we want to reach a target operating point from our current one.

- Servomechanism problem: we want to control the system to follow some path.

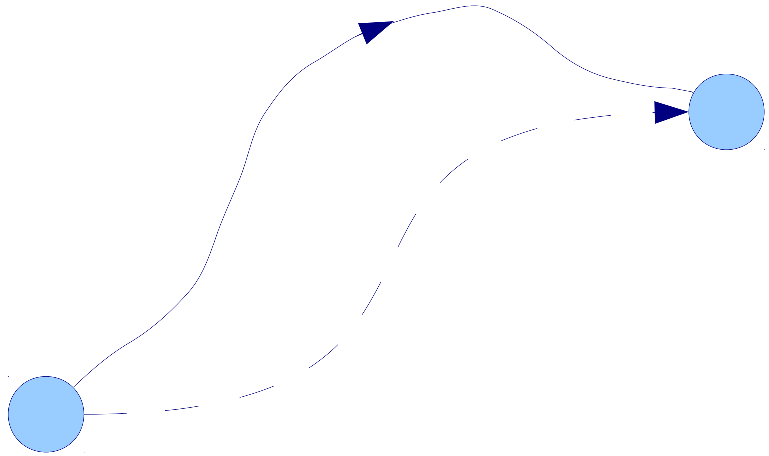

Control vs. Planning

- The boundary between planning and control is not always clear.
- Roughly, planning is strategic whereas control is tactical.

Dynamical Systems

- Dynamic(al) system: system whose behaviour changes over time, often in response to external stimulation or forcing.
- State: collection of variables that summarise past information of a system for the purpose of predicting the future.
- Deterministic systems defined via ODEs or difference equations

$$\dot{x} = f(x, u) \qquad x_{i+1} = \tilde{f}(x_i, u_i)$$
$$y = g(x, u) \qquad y_i = \tilde{g}(x_i, u_i)$$

State Space and Phase Portraits

- State space: set of all possible states.
- Phase portrait: vector field showing the dynamics of a non-actuated system:

$$\dot{x} = f(x, u_{fixed})$$

- Predicting the future equates to solving an initial value problem

$$x(t) = x(0) + \int_0^t \dot{x}dt = x(0) + \int_0^t f(x,u)dt$$

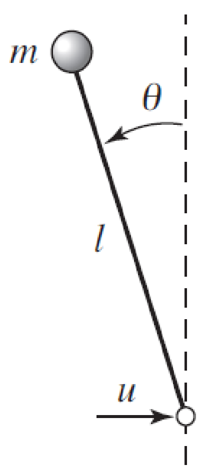

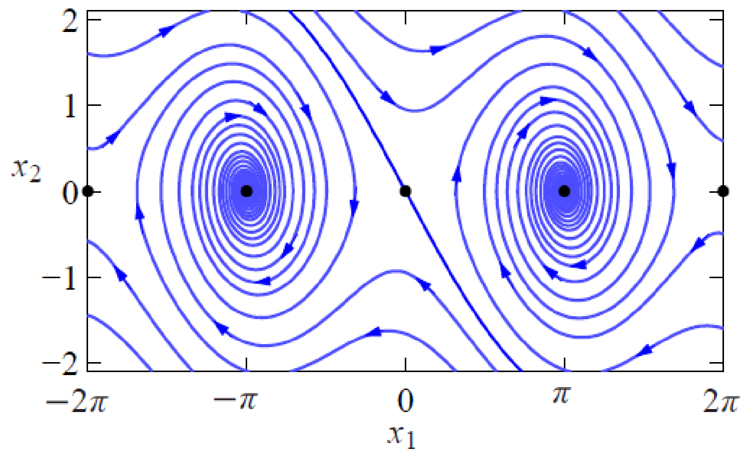

Linear Systems

- Enormous amounts of literature.
- Very useful, e.g. aircraft dynamics (no aerobatics)
- Linear time-invariant (LTI) system:

$$\dot{x} = Ax + Bu$$
$$y = Cx + Du$$

- Laplace and Fourier transforms of the system transport the analysis to very useful spaces.

Frequency Domain Analysis

- The frequency response measures the way in which a system responds to a sinusoidal excitation.
- Offers a very powerful tool for linear system analysis and controller synthesis.

$$u = \sin \omega t$$
$$y = g(\omega) \sin(\omega t + \varphi(\omega))$$

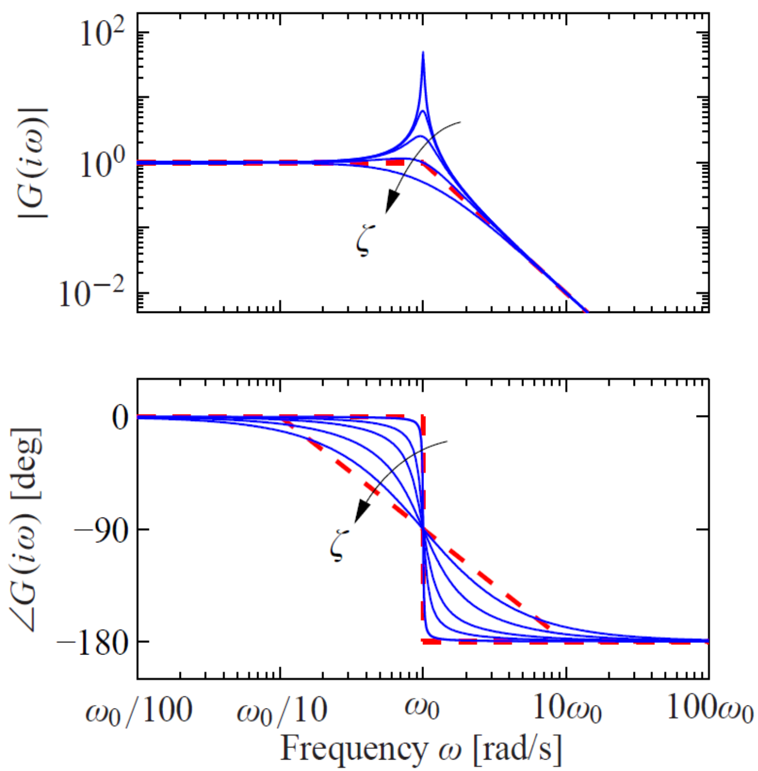

Time Domain Analysis

- Simulate the behaviour of the system to a given input signal.
- Even for very complex non-linear models we can simulate the time response by solving an initial value problem numerically, e.g.

$$x(t + \Delta t) \approx x(t) + \Delta t \; \dot{x}(t) = x(t) + \Delta t \;\; f(x(t), u(t))$$

- There is a trade-off between computation time, accuracy and numerical stability when solving the initial value problem. This leads to more advanced integration methods:
  - Use of high-order derivatives.
  - Adaptive step size.
  - Implicit methods.

Observability

- A system is observable if the current state $x(t_0)$ can be determined using the input and output signals of the system during a finite time interval $[t_0,t]$.

- No test for observability for a general nonlinear system. However, the linear time invariant case is neat and illustrative:

$$
\begin{aligned}
y_0 &= Cx_0 \\
y_1 &= Cx_1 = CAx_0 \\
y_2 &= Cx_2 = CA^2x_0 \\
&\vdots
\end{aligned}
$$

$$
\begin{bmatrix} y_0 \\ y_1 \\ \vdots \\ y_{n-1} \end{bmatrix} = \underbrace{\begin{bmatrix} C \\ CA \\ \vdots \\ CA^{n-1} \end{bmatrix}}_{\mathcal{O}} x_0
$$

observability $\Leftrightarrow$ uniqueness of solution $\Leftrightarrow$ $\mathcal{O}$ has full rank

Controllability/Reachability

- A system is controllable if it is possible to steer it from any state x to any state z in finite time.

- Again, no general results to test controllability in the nonlinear case but an illustrative solution for LTI systems:

$$
\begin{aligned}
x_1 &= Ax_0 + Bu_0 \\
x_2 &= A(Ax_0 + Bu_0) + Bu_1 = A^2x_0 + ABu_0 + Bu_1 \\
&\vdots \\
x_n &= A^n x_0 + \underbrace{\begin{bmatrix} A^{n-1}B & \cdots & AB & B \end{bmatrix}}_{\mathcal{C}} \begin{bmatrix} u_0 \\ \vdots \\ u_{n-2} \\ u_{n-1} \end{bmatrix}
\end{aligned}
$$

$$\text{controllability} \Leftrightarrow \text{solution } \exists \, \forall x_n \Leftrightarrow \mathcal{C} \text{ has full rank}$$

Performance

- Good dynamic properties:
  - Adequate damping and oscillation frequencies.
  - Low/no steady state error.
  - Quick response.
  - Low overshoot.
  - ...
- Good disturbance rejection.
- Low control effort to minimise control energy and actuator wear.
- ...

Stability

- Equilibrium point

$$\dot{x} = f(x, u)$$

$$0 = f(x_e, u_e) \Leftrightarrow x_e \text{ is an equilibrium state}$$

- x_e is a stable equilibrium point if for all $\varepsilon > 0$ there exists a $\delta > 0$ such that

$$\|x_{pert} - x_e\| < \delta \Rightarrow \|x(t; x_{pert}) - x_e\| < \varepsilon \quad \forall\, t > 0$$

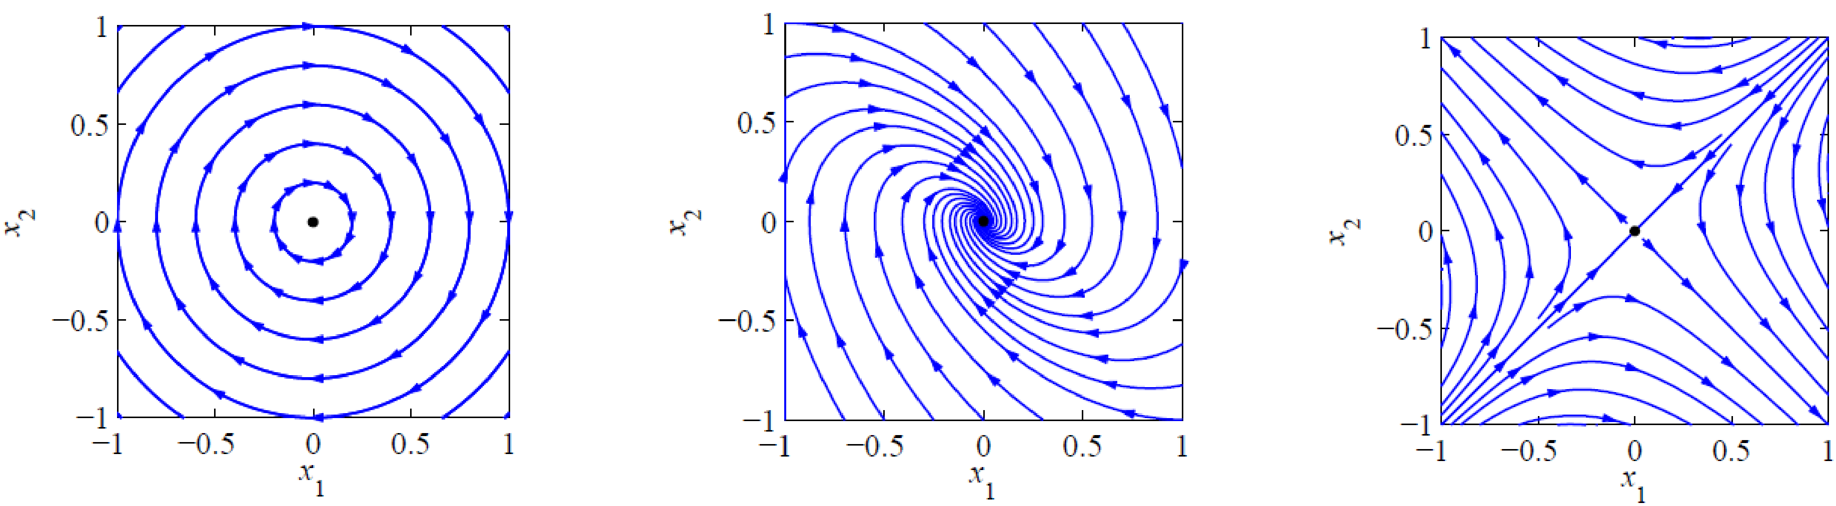

Robustness

- Robustness to uncertainty: performance and stability need to be achieved even when the dynamical system is not exactly the way we thought it was.

- Parameter uncertainty and unmodelled dynamics.

- Robust stability: a controller provides robust stability if it stabilises all dynamical systems belonging to a certain set constructed around the nominal system.

- Robust performance: same idea.

Feedback

- Feedback: usage of a system output to compute its input.

$$\dot{x} = f(x, u) \qquad\qquad u = \pi(y)$$
$$y = g(x) \qquad\qquad \dot{x} = f(x, \pi(y))$$

- Used in engineering systems but it is ubiquitous in nature.

Feedback

Feedback provides two key benefits:

- Robustness to uncertainty
  - High performance can be achieved even when our model of reality is far from perfect.
  - Sensing allows comparison between actual and desired and feedback provides a correction.

- Modification of the natural dynamics
  - Feedback allows the closed-loop system to have the desired dynamics regardless of the natural dynamics of the original system (within limits).

State Feedback

- Measure the full state and use it to determine the input to the system:

$$u = \pi(x)$$
$$\dot{x} = f(x, \pi(x))$$

- We have all possible information about the system. However, it may not be controllable...

- Controller design for linear systems: pole placement, LQR...

Output Feedback

The full state is not measured, we have two main alternatives:

- If we can define our performance objectives as a function of the measured output.

$$u = \pi(y)$$
$$\dot{x} = f(x, \pi(y))$$

- In a more general case we can use the measured outputs to estimate the value of the states and then use state feedback.

$$u = \pi(\hat{x})$$
$$\dot{x} = f(x, \pi(\hat{x}))$$

State Estimation

- If the system is observable , an observer solves the state estimation inference problem

$$p(x_k | y_k, y_{k-1}, \cdots, u_k, u_{k-1}, \cdots)$$

- In general, the observer has to deal with noisy measurements.

- Several approaches to design an observer:
  - Traditionally the observer is another dynamical system where the error between estimated and real states tends to zero with time, e.g.

$$\frac{d\hat{x}}{dt} = A\hat{x} + Bu + L(y - \hat{y})$$

  - Currently one can sometimes afford to use particle filters (i.e. Monte Carlo) to solve nonlinear/non-Gaussian state estimation problems.

Kalman Filter

- Bayes filter for linear systems with Gaussian noise in the measurements and the dynamics.

- The Kalman filter uses a gain L that minimises the mean square error of the state prediction given the noise covariance matrices and the system matrices.

- A graphical model may help.

System Identification

- Learning problem: given a collection of measured inputs and outputs, make a model of the dynamical system operating behind the scenes.
- Typically uses some form of regression.
- Identification in the time domain and in the frequency domain.

Adaptive Control

- Needs the identification step? AC is reacting to changes, or at least to more information gathered from the system.

- Dual control: simultaneous identification of the system dynamics and satisfaction of the controller goals. Exploration-exploitation trade-off. Several strategies available:
  - Certainty equivalence: ignore uncertainty.
  - Probing: active learning.
  - Cautious: use a conservative loss function.

- Other strategies:
  - Gain scheduling.
  - Self-tuning regulators.

Optimal Control

- Design of controllers that optimise a criterion.
  - Regulator problems.
  - End-point problems, e.g. spacecraft rendezvous.

LQR: Linear Quadratic Regulator

- A practical optimal control method for linear time-invariant systems.

$$J = \int_0^\infty x^\top Q_x x + u^\top Q_u u \; dt$$

$$u = -Q_u^{-1} B^\top P \; x$$

$$PA + A^\top P - PBQ_u^{-1} B^\top P + Q_x = 0$$

- Duality control – state estimation.
- LQG: Linear Gaussian Regulator = Kalman filter + LQR
- Certainty equivalence.

Relationship to Reinforcement Learning

- RL is a combination of identification, adaptive control and optimal control.

- Control theory focuses more on actually implementing controllers than on trying to solve the big RL problem.

- RL uses feedback, not only open loop planning. Doesn't care about things like stability because it's normally aimed at longer time scale planning, e.g. traditionally we wouldn't use RL to stabilise an aircraft.

A General Controller Structure

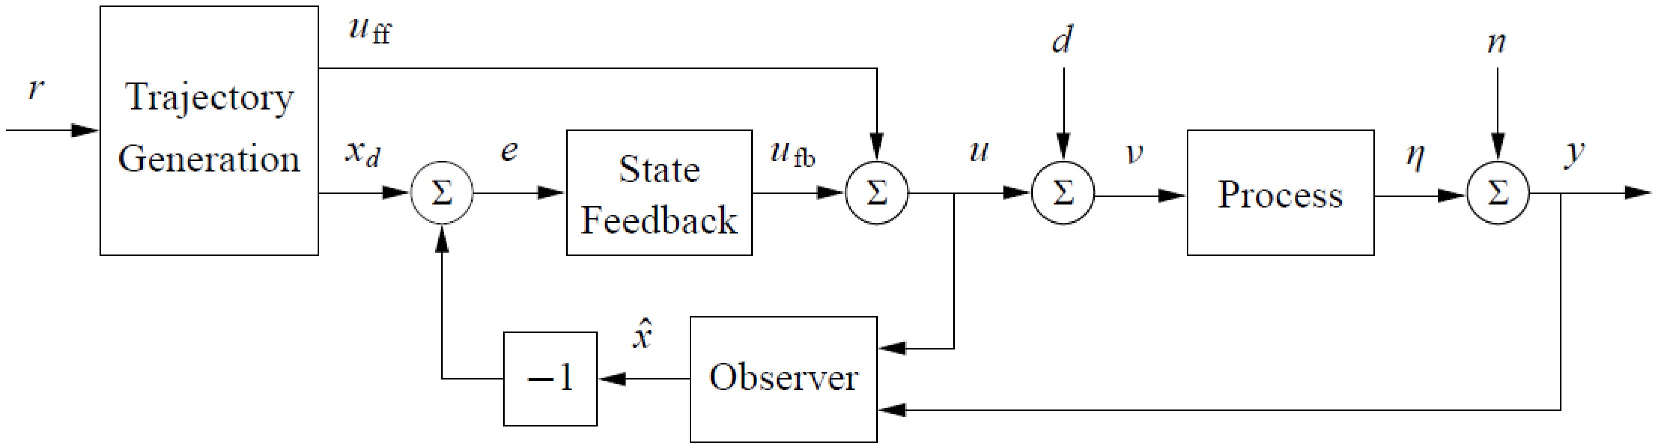

- The trajectory generation block has two functions:
  - Generate the reference state x_d.
  - Generate a feedforward control signal u_{ff}.

Bibliography

- Åström and Murray. *Feedback Systems*. Princeton University Press, 2008.
- Franklin et al. *Feedback Control of Dynamic Systems*. Prentice Hall, 2006.
- Murray (ed.). *Control in an Information Rich World*. Report of the Panel on Future Directions in Control, Dynamics, and Systems, 2002.