Engineering Tripos Part IIB Module 4F3
Optimal and Predictive Control
2014 Crib

E. N. Hartley

12 May, 2014

1. (a) i. The definitions of the state and input increments are: [20%]

$$\Delta x = x(k) - x(k-1) \qquad \Delta u = u(k) - u(k-1)$$

The original system is of the form:

$$x(k+1) = Ax(k) + Bu(k) \qquad z(k) = Cx(k).$$

For the state update equation:

$$\begin{aligned} x(k+1) - x(k) &= Ax(k) + Bu(k) - x(k) \\ &= Ax(k) + Bu(k) - (Ax(k-1) + Bu(k-1)) \\ &= A(x(k) - x(k-1)) + B(u(k) - u(k-1)) \\ \Delta x(k+1) &= A\Delta x(k) + B\Delta u(k). \end{aligned}$$

For the output:

$$\begin{aligned} z(k+1) &= C(x(k+1)) \\ &= C(x(k) + \Delta x(k+1)) \\ &= Cx(k) + C\Delta x(k+1) \\ &= z(k) + CA\Delta x(k) + CB\Delta u(k). \end{aligned}$$

Therefore:

$$F = \begin{bmatrix} A & 0 \\ CA & I \end{bmatrix} \qquad G = \begin{bmatrix} B \\ CB \end{bmatrix}.$$

ii. Sufficient conditions for convexity: $P \geq 0$, $Q_1 \geq 0$, $Q_2 \geq 0$, $R \geq 0$. ($R \geq 0$ sufficient for convexity, but $R > 0$ needed for uniqueness). [15%]

iii. Convexity is important because it means that a local optimum is also a global optimum, and that the optimisation problem can be solved efficiently. [10%]

(b) i. At equilibrium: $x(k+1) = x(k)$. [20%]

$$\implies Ax_\infty + Bu_\infty = x_\infty$$
$$\implies (A-I)x_\infty + Bu_\infty = 0.$$

For compatibility with the tracked output reference: $Cx_\infty = r$. So, combined:

$$\begin{bmatrix} (A-I) & B \\ C & 0 \end{bmatrix} \begin{bmatrix} x_\infty \\ u_\infty \end{bmatrix} = \begin{bmatrix} 0 \\ r \end{bmatrix}$$

The matrix

$$\begin{bmatrix} (A-I) & B \\ C & 0 \end{bmatrix}$$

should be of full row rank for a solution to exist, and it should be invertible for the solution to be unique (i.e. also of full column rank).
It is *not sufficient* simply for the number of outputs to be equal to the number of inputs.

ii. $\Delta x_\infty = 0$, $\Delta u_\infty = 0$ (steady state), $z_\infty = r$ (tracking reference). [10%]

iii. Modify the cost function by substituting $z_i \leftarrow z_i - r$, $z_N \leftarrow z_N - r$. No further changes needed since $\Delta x_\infty = 0$ and $\Delta u_\infty = 0$! [10%]

iv. Advantage: [15%]

- No need for target calculator for reference tracking.
- Automatic integral action (zero offset to constant disturbances)
- Can constrain slew rates easily
- Damping controlled by R because it suppresses input moves

Disadvantage:

- Constraints may be defined on x and u rather than Δx and Δu
- Computing differences Δx might amplify sensor noise
- Response to setpoint changes and disturbances may be sluggish
- Can be more difficult to tune satisfactorily (oscillatory responses easy to obtain).

Differences in computational load and memory requirements with respect to MPC with a "standard" model are problem specific, and likely to be insignificant. The computational bottleneck is solution of the (contrained) quadratic program. If the state is eliminated using the prediction matrices, then the number of decision variables in the optimisation problem is the same as for "standard" MPC.

2. (a) i. The receding horizon principle in the context of model predictive control is as follows: [20%]

1) The current state is measured.
2) Using a model of the plant to make predictions of future states, a finite horizon (constrained) optimal control problem is solved with the state trajectory starting at the current measured state.
3) The first part of the optimal input trajectory is applied to the plant, and the rest is discarded

4) At the next time step a new optimal control problem is solved over a horizon of the same length (so now viewing one step further into the future, hence “receding”), the first part is applied to the plant, and the process is repeated.

ii. The question specifically asks for the advantages and disadvantages of receding horizon control with respect to infinite horizon control. It is not sufficient to simply state advantages/disadvantages of predictive control with respect to arbitrary other control schemes. [10%]

- Advantages
  - Receding horizon control enables a computationally tractable numerical computation of the control policy in cases where an infinite horizon control policy is not computationally tractable. For example, when no analytical solution exists and “gridding” is impractical. In these cases, the receding horizon control can be viewed as an “approximation” of the infinite horizon control policy if the latter is what is actually desired. Examples of such situations:
    * constraints on inputs and/or states;
    * nonlinear plant models;
    * discrete-valued/switched/hybrid systems;
  - When a system is time-varying an infinite horizon may not be appropriate
- Disadvantages
  - More care must be taken to ensure closed-loop stability and to analyse closed-loop performance.
  - Computational cost can still be prohibitive for a given application.
  - Numerical computation of a receding horizon control policy may be viewed as inappropriate when there is an analytical solution to the infinite horizon problem.

(b) i. [20%]

$$\begin{bmatrix} x_0 \\ x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} I \\ A \\ A^2 \\ A^3 \end{bmatrix} x(k) + \begin{bmatrix} 0 & & \\ B & & \\ AB & B & \\ A^2B & AB & B \end{bmatrix} \begin{bmatrix} u_0 \\ u_1 \\ u_2 \end{bmatrix}$$

$$\Phi = \begin{bmatrix} I \\ A \\ A^2 \\ A^3 \end{bmatrix}, \quad \Gamma = \begin{bmatrix} 0 & & \\ B & & \\ AB & B & \\ A^2B & AB & B \end{bmatrix}$$

ii. [15%]

$$F = \begin{bmatrix} C & 0 & 0 & 0 \\ 0 & C & 0 & 0 \\ 0 & 0 & C & 0 \\ 0 & 0 & 0 & T \end{bmatrix}, \quad G = \begin{bmatrix} D & 0 & 0 \\ 0 & D & 0 \\ 0 & 0 & D \\ 0 & 0 & 0 \end{bmatrix}, \quad g = \begin{bmatrix} e \\ e \\ e \\ t \end{bmatrix}$$

iii. [10%]

$$F(\Phi x(k) + \Gamma\underline{\mathbf{u}}) + G\underline{\mathbf{u}} \leq g$$
$$(F\Gamma + G)\underline{\mathbf{u}} \leq g - F\Phi x(k)$$

(c) i. The admisible set is the set of x such that $Cx + Du \leq e$, i.e. that $(C + DK)x \leq e$. A set $\mathbb{T}$ is positively invariant with respect to the the system $x(k+1) = (A + BK)x(k)$ if $x(k) \in \mathbb{T} \implies x(k+1) \in \mathbb{T}$. Therefore, a positively invariant admissible set is: [10%]

$$\mathbb{T} = \{x : (C + DK)x \leq e, \quad (A + BK)x \in \mathbb{T}\}$$

Alternatively, and equivalently:

$$\mathbb{T} = \left\{x : (C + DK)(A + BK)^i x \leq e, \quad \forall i = 0, \ldots, \infty\right\}.$$

ii. If the terminal constraint is admissible then the control law $u = Kx$ is always an admissible action after the end of the horizon for at least one step. If the terminal constraint is also positively invariant, applying the control law $u = Kx$ after the end of the prediction horizon will guarantee constraint satisfaction over an infinite horizon. [15%]

This means that if the feasible (optimal) control input is applied to the plant at the current time step, there is an N-step sequence of inputs that is feasible at the next time step, constructed by removing the first element u_0 from the beginning previous input sequence and appending $u_{N+1} = Kx_N$ at the end. This guarantees *recursive feasibility* of the *closed-loop system* with the *receding horizon control law* and is usually a key step towards proving closed-loop stability.

NB. On its own the terminal constraint is not sufficient to provide a theoretical guarantee of asymptotic stability and is usually used in conjunction with other constructs such as a specially designed terminal cost.

3. (a) i. Dynamic programming (in a discrete time setting) is a systematic procedure for transforming an optimisation over a sequence of h inputs into h minimisations over 1 inputs (but all states). [20%]

The minimisation "for all states" can be done by enumeration (if there are a finite number of possible states), or analytically if an analytical expression exists. This enables a "backwards recursion" to find a sequence of value functions.

The limitation is that it must be possible to perform the 1-step minimisation over "all" steps, and the procedure cannot always be implemented in practice. Problems can be approximated by "gridding" and enumeration, but complexity can grow exponentially ("curse of dimensionality").

ii. [20%]

$$
\begin{aligned}
V_d(x,k) &= \min_{u_k,\dots,u_{h-1}} \left(\sum_{i=k}^{h-1} c_d(x_i,u_i) + J_h(x_h) \right) \\
&= \min_{u_k,\dots,u_{h-1}} \left(c_d(x_k,u_k) + \sum_{i=k+1}^{h-1} c_d(x_i,u_i) + J_h(x_h) \right) \\
&= \min_{u_k} \min_{u_{k+1},\dots,u_{h-1}} \left(c_d(x_k,u_k) + \sum_{i=k+1}^{h-1} c_d(x_i,u_i) + J_h(x_h) \right) \\
&= \min_{u_k} c_d(x_k,u_k) + \min_{u_{k+1},\dots,u_{h-1}} \left(\sum_{i=k+1}^{h-1} c_d(x_i,u_i) + J_h(x_h) \right) \\
&= \min_{u_k} \left(c_d(x_k,u_k) + V_d(x_{k+1},k+1) \right)
\end{aligned}
$$

(b) i. [40%]

$$
\begin{aligned}
V(x,t) &= \min_u \left(c(x,u)\delta t + V(x + f(x,u)\delta t, t+\delta t) \right) + \mathcal{O}(\delta t^2) \\
V(x+\delta x, t+\delta t) &= V(x,t) + \frac{\partial V(x,t)}{\partial x}\delta x + \frac{\partial V(x,t)}{\partial t}\delta t + \text{ higher order terms} \\
V(x,t) &= \min_u \left(c(x,u)\delta t + V(x,t) + \frac{\partial V(x,t)}{\partial x}\delta x + \frac{\partial V(x,t)}{\partial t}\delta t \right) + \mathcal{O}(\delta t^2)
\end{aligned}
$$

Note that $V(x,t)$ does not depend on u:

$$
0 = \min_u \left(c(x,u)\delta t + 0 + \frac{\partial V(x,t)}{\partial x} f(x,u)\delta t + \frac{\partial V(x,t)}{\partial t}\delta t \right) + \mathcal{O}(\delta t^2)
$$

Divide by δt and rearrange:

$$
-\frac{\partial V(x,t)}{\partial t} = \min_u \left(c(x,u) + \frac{\partial V(x,t)}{\partial x} f(x,u) \right) + \frac{\mathcal{O}(\delta t^2)}{\delta t}
$$

Take limit as $\delta t \to 0$:

$$
-\frac{\partial V(x,t)}{\partial t} = \min_u \left(c(x,u) + \frac{\partial V(x,t)}{\partial x} f(x,u) \right)
$$

ii. [20%]

$$
\begin{aligned}
f(x,u) &= 2x + u \\
c(x,u) &= x^2 + u^2 \\
V(x,t) &= p(t)x(t)^2
\end{aligned}
$$

Noting that we are taking *partial derivatives* and **not** total derivatives:

$$\frac{\partial V(x,t)}{\partial t} = \dot{p}(t)x(t)^2$$
$$\frac{\partial V(x,t)}{\partial x} = 2p(t)x(t)$$

By substitution:

$$\begin{aligned}-\dot{p}(t)x(t)^2 &= \min_u(x(t)^2 + u(t)^2 + 2p(t)x(t)(2x(t) + u(t))) \\ &= \min_u(x(t)^2 + u(t)^2 + 4p(t)x(t)^2 + 2p(t)x(t)u(t))\end{aligned}$$

$$\frac{\mathrm{d}}{\mathrm{d}u}\left(x(t)^2 + u(t)^2 + 4p(t)x(t)^2 + 2p(t)x(t)u(t)\right) = 2u(t) + 2p(t)x(t)$$

Therefore minimum achieved when $u(t) = -p(t)x(t)$.

$$-\dot{p}(t)x(t)^2 = x(t)^2 + p(t)^2x(t)^2 + 4p(t)x(t)^2 - 2p(t)^2x(t)^2$$

$$-\dot{p}(t) = 1 - p(t)^2 + 4p(t)$$

Check from datasheet:

$$A = 2, \quad B = 1, \quad Q = 1, \quad R = 1, X = p.$$

$$\begin{aligned}-\dot{X} &= Q + XA + A^TX - XBR^{-1}B^TX \\ -\dot{p} &= 1 + (2p + 2p) - p^2.\end{aligned}$$

Boundary condition: $p(T) = 10$.

4. (a) Let [20%]

$$y = C(sI - A)^{-1}Bu$$

$$\|y\|_\infty \leq \frac{1}{2\pi}\|G(s)\|_2\|u\|_2.$$

where:

$$\|y\|_\infty = \sup_t \sqrt{y(t)^Ty(t)}$$

i.e. is the maximum magnitude of y over all time.

$$\|u\|_2 = \sqrt{\int_0^\infty u(t)^Tu(t)\,\mathrm{d}t}$$

is effectively the energy in the input signal.

(b) i. From the definition in the question [25%]

$$Q = \int_0^\infty e^{A^T\tau} C^T C e^{A\tau}\, \mathrm{d}\tau.$$

Substitute this into the Lyapunov equation:

$$\begin{aligned}
A^T Q + QA + C^T C &= 0\\
A^T \left(\int_0^\infty e^{A^T\tau} C^T C e^{A\tau}\, \mathrm{d}\tau \right) + \left(\int_0^\infty e^{A^T\tau} C^T C e^{A\tau}\, \mathrm{d}\tau \right) A + C^T C &= 0\\
\int_0^\infty A^T (e^{A^T\tau} C^T C e^{A\tau}) + (e^{A^T\tau} C^T C e^{A\tau}) A\, \mathrm{d}\tau + C^T C &= 0\\
\int_0^\infty \frac{\mathrm{d}}{\mathrm{d}\tau} (e^{A^T\tau} C^T C e^{A\tau})\, \mathrm{d}\tau + C^T C &= 0\\
\left[e^{A^T\tau} C^T C e^{A\tau} \right]_0^\infty + C^T C &= 0\\
0 - C^T C + C^T C &= 0
\end{aligned}$$

(Using stability of system to conclude that as $t \to \infty$, $e^{At} \to 0$). This confirms the required result.

ii. $Q > 0$ implies that for every $x \neq 0$ there exists a $t \in [0, \infty)$ such that $Ce^{At}x \neq 0$. If (A, C) were not observable, there would exist an $x \neq 0$ such that $Ce^{At}x = 0$ for all t. Therefore (A, C) must be observable. [20%]

iii. The impulse response $g(t) = Ce^{At}B$. [5%]

iv. The $\mathcal{H}_2$ norm of the plant $\|G(s)\|_2$, by Parseval's theorem is: [30%]

$$\sqrt{2\pi}\|g(t)\|_2.$$

Define B_i to be the i-th column of B and $g_i(t)$ to be the response to an input on the i-th input, i.e. $g_i(t) = Ce^{At}B_i$.

$$\begin{aligned}
\|g_i(t)\|_2 &= \sqrt{\int_0^\infty B_i^T e^{A^T t} C^T C e^{At} B_i\, \mathrm{d}t}\\
&= \sqrt{B_i^T \left(\int_0^\infty e^{A^T t} C^T C e^{At}\, \mathrm{d}t \right) B_i}\\
&= \sqrt{B_i^T Q B_i}
\end{aligned}$$

$$\begin{aligned}
\|g(t)\|_2^2 &= \sum_i \|g_i(t)\|_2^2\\
&= \sum_i B_i^T Q B_i\\
&= \mathrm{trace}(B^T Q B)\\
\|G(s)\|_2 &= \sqrt{2\pi}\sqrt{\mathrm{trace}(B^T Q B)}
\end{aligned}$$

The correct answer can also be achieved by defining:

$$V = x(t)^T Q x(t)$$

and letting $z = \dot{V} + y^T y$.

$$\begin{aligned}
\dot{V} + y(t)^T y(t) &= \frac{\mathrm{d}}{\mathrm{d}t}\left(x(t)^T Q x(t)\right) + y(t)^T y(t) \\
&= x(t)^T A^T Q x(t) + x(t) Q A x(t) + x(t)^T C^T C x(t) \\
&= x(t)^T (A^T Q + QA + C^T C) x(t)
\end{aligned}$$

$$\int_0^\infty z(t)\,\mathrm{d}t = \int_0^\infty \dot{V}(t) + y(t)^T y(t)\,\mathrm{d}t = x(t)^T (A^T Q + QA + C^T C) x(t)$$

$$\begin{aligned}
[V(t)]_0^\infty + \|y\|_2^2 &= x(t)^T (A^T Q + QA + C^T C) x(t) \\
V(\infty) - V(0) + \|y\|_2^2 &= 0 \\
\implies \|y\|_2^2 &= x(0)^T Q x(0)
\end{aligned}$$

Then using

$$\|g_i(t)\|_2^2 = B_i^T Q B_i$$

the same arguments as above can be followed.

Examiners report

Q1 MPC Convexity and Target Tracking

23 attempts, average mark 13.42/20, Maximum 20, Minimum 7.

A popular and straightforward question, well-answered by most candidates, although some candidates struggled with the derivation of the linear velocity form. Most appreciated the importance of convexity of the cost function. Some candidates erroneously thought that an equal number of inputs and outputs was sufficient for a unique equilibrium pair (ignoring the matrix invertibility), and a small minority erroneously thought that the terminal weighting matrix should be changed to enable reference tracking.

Q2 Receding Horizon and Constraints Formulation

29 attempts, average mark 14.97/20, Maximum 19, Minimum 10.

The most popular question, attempted by all candidates, and reasonably well answered by most. All candidates were able to satisfactorily describe the receding horizon principle, although some failed to mention that the predictions should be made starting from a measurement or estimate of the current plant state. Unfortunately a sizeable minority of candidates chose to describe the advantages/disadvantages of MPC in general (as has been often asked in previous years) rather than comparing receding horizon versus infinite horizon control. The prediction matrix formulation questions were generally well answered, contributing to the high average mark and low variation for this question, although some candidates did not correctly form the constraint matrices. Justifications for the invariant, admissible terminal constraint were mostly vague with very few candidates mentioning recursive feasibility in closed loop with the receding horizon control law, and a minority instead described conditions on terminal cost matrices!

Q3 Dynamic Programming and the HJB Equation

25 attempts, average mark 12.60/20, Maximum 18, Minimum 2.

The second most popular question. The description of dynamic programming was generally well answered, as was the derivation of the discrete-time dynamic programming equation. The derivation of the HJB equation was more divisive: in general, candidates were able to do it correctly or not, with little middle ground. Application to the LQR problem posed more difficulties. Many candidates evaluated a total derivative rather than a partial derivative of the value function with respect to time, and thus obtained an incorrect answer. A few candidates correctly identified this was a finite-horizon LQR problem and applied the Riccati equation from the data sheet, for which partial credit was given for the correct final answer, despite the question specifically requesting a derivation.

Q4 Signal Norms, Controllability and Infinite Horizons

9 attempts, average mark 12.00/20, Maximum 19, Minimum 1.

The least popular and most divisive question: either answered very well, or with just a token effort. Most candidates correctly identified how the $\mathcal{H}_2$ norm related the input signal to the output signal, although worryingly some confused the definition of the $\mathcal{L}_\infty$ norm of the output with that of the $\mathcal{H}_\infty$ norm. Most candidates were able to show that the infinite-horizon integral could be computed by solving the Lyapunov equation, but fewer were able to make a coherent argument regarding the observability of the system (although there were also some pleasing correct answers). Derivation of the $\mathcal{H}_2$ norm in terms of the solution to the Lyapunov equation was divisive: candidates were generally either able to (following one of two valid methods covered in the course), or not able to, with little in the way of middle ground.