Companion Notes for 3F2: Systems and Control

Andrew Guy

The thoughts that occur to me or things that have confused students, ordered by which question we were looking at when the words started falling out of my mouth. *Caveat lector*: this may contain errors, is mostly devoid of references and is not yet complete (particularly for EP3).

Version 1
April 12, 2023

Examples Paper 1

Question 1

Key points:

- Remember that blocks in the block diagram are transfer functions from one signal to another. Don't be afraid to write down equations relating every signal in the information given, then assemble the diagram from those — you can always go back later and lump blocks together.
- For equations of motion of order n with respect to time $\partial_t^n y = g(y, \partial_t y, \ldots, \partial_t^{n-1} y)$ we need to keep n state variables for the derivatives $x_i = \partial_t^{i-1} y$ for $1 \leq k \leq n$ in our state vector x, we get $n-1$ of the derivatives in $\dot{x}$ for free by just shifting them along ($\dot{x}_i = x_{i+1}$ for $1 \leq i < n$, and a block of A will be I_{n-1}) and $\dot{x}_n = g(x)$. We can start our block diagram with a chain of n integrators to get our signals $\partial_t^n y \to y$, then add connections from those signals to make up the dynamics.

$$x = \begin{bmatrix} y \\ \partial_t y \\ \vdots \\ \partial_t^{n-1} y \end{bmatrix} \quad \dot{x} = \begin{bmatrix} x_2 \\ \vdots \\ x_n \\ g(x) \end{bmatrix} \quad A = \begin{bmatrix} & 1 & & \\ & & \ddots & \\ & & & 1 \\ \frac{\partial g}{\partial y} & \frac{\partial g}{\partial(\partial_t y)} & \cdots & \frac{\partial g}{\partial(\partial_t^{n-1} y)} \end{bmatrix}$$

- When deriving a state space model, if a candidate state variable can be expressed in terms of the others then it is not a state variable.
- Remember to revise the implicit function theorem for $F(\dot{x}, x, u) = 0$: set $\frac{\partial F}{\partial x}\delta x + \frac{\partial F}{\partial \dot{x}}\delta \dot{x} + \frac{\partial F}{\partial u}\delta u = 0$ at your equilibrium point $(\dot{x}, x, u) = (0, x_0, u_0)$.

So, for $x = [\theta_o, \dot{\theta}_o] = [x_1, x_2]$ and $J\ddot{\theta}_o = \tau$ we will have a block $\frac{1}{Js}$ from output shaft torque to $\dot{\theta}_o = x_2$ and $\frac{1}{s}$ to $\theta_o = x_1$. From there, the connection diagram remains mostly intact with the exception of needing to measure motor speed $\dot{\theta}_m$. As we have the relationship $n\dot{\theta}_o = \dot{\theta}_m$, we can detach the link from the motor output to the feedback path $k_d k_t$ and reattach it at $\dot{\theta}_o$ by inserting a block to multiply it by n.

Question 2

Bet you thought you could leave electronics behind, huh?

- For real world systems, good candidates for state variables are things that store energy (and a good candidate for proving stability is showing that total system energy is dissipated). For the state variable, choose whatever gets squared in the energy term:

Component	Energy	State variable
Capacitor	$\frac{1}{2}Cv^2$	v
Inductor	$\frac{1}{2}Li^2$	i
Mass	$\frac{1}{2}m\dot{x}^2$	$\dot{x}$
Spring	$\frac{1}{2}kx^2$	x

  The reason will become clear once you take 4F2, but if you view the evolution of the state using a flow field analogy with $u = \dot{x}$ (where u is the fluid velocity), then by taking the total energy of the system in a state as $V(x)$ our material derivative (time-independent, so $u \cdot \nabla$) of system energy becomes $f(x) \cdot \nabla V(x)$. Try showing that the total energy of a mass-spring system without damping remains constant using this method, then add a damper and show that energy is always dissipated.

- The universe *hates* differentiators: suppose the spectrum of a signal has bounded tail $k\frac{1}{\omega^{0.5+\varepsilon}} < |F(\omega)| < K\frac{1}{\omega^{0.5+\varepsilon}}$ for $\varepsilon > 0$, $k < K$ and $\omega > \omega'$. Then the total signal energy associated with the tail $\int_{\omega'}^{\infty} |F(\omega)|^2 \, d\omega < \infty$, but putting it through the differentiator yields $\int_{\omega'}^{\infty} \omega^2 |F(\omega)|^2 \, d\omega$, which for suitably small ε or enough differentiators places unbounded energy in the high frequency parts of the filtered signal. *You can approximate differentiators in the low frequency range, but this behaviour must eventually break down.*

- Another pretty good reason that we don't like direct differentiators is that we then link $\dot{x} = f(\dot{u})$, and our model doesn't actually place any guarantees on the differentiability of our input signals at all (we only require enough integrability for $\int_0^t e^{A(t-\tau)} Bu(\tau) \, d\tau$ to be meaningful). *Differentiators don't break causality: they only provide a guess of future input at a point in time and can be estimated arbitrarily closely using backward differences, but do not actually pull any information from the future.*

  The modern control theorist may wish to exploit the fact that the theory based on $\dot{x} = f(x, u)$ does not require u to be differentiable, and control their system by doing some computation on a microcontroller and setting a constant value of $u_i(t) = u_k^i U_i(t - k\Delta t)$ for $k\Delta t < t \leq (k+1)\Delta t$, where u_k^i is the computed level on input i at sampling point k and $U(\tau)$ a shape function over the holding interval (which in the ideal case might be $H(\tau) - H(\tau - \Delta t)$ for unit step H). There's nothing special about $t = 0$, so we have:

$$x_{k+1} = e^{A\Delta t} x_k + \int_0^{\Delta t} e^{A(\Delta t - \tau)} BU(\tau) u_k \, d\tau.$$

  As long as Δt is constant, the state transition matrix doesn't change and the input contribution can be split into a set of linear transfer functions Γ_{ij} from $u_k^i \to x_{k+1}^i$ dependent on Δt, B, $U(\tau)$ and e^{At} to yield $x_{k+1} = \Phi x_k + \Gamma u_k$. More on this later.

(ii) Adding a series resistance decouples the derivative of voltage across the capacitor from the input derivative, as we now have $i = C\dot{x} = (u - x)/R$ instead of $\dot{x} = \dot{u}$. Your real world circuits have this tiny resistance in them anyway, maybe you can try modelling the effect at higher frequencies.

(iv) When faced with a scary looking circuit like this, the best bet is to ignore anything they've not given component values for and charge ahead as though it has absolutely no impact on the behaviour, because it probably doesn't.

Question 3

- Remember to differentiate the equations of motion at the equilibrium point, even though it might feel wrong! Our expansion around x_0, u_0 is given by

$$f(x_0 + \delta x, u_0 + \delta u) = f(x_0, u_0) + \frac{\partial f}{\partial x}(x_0, u_0)\delta x + \frac{\partial f}{\partial u}(x_0, u_0)\delta u + O(\delta^2),$$

  and the matrix element $A_{ij} = \frac{\partial f_i}{\partial x_j}(x_0, u_0)$.

- Block matrix inversion: maybe you get a matrix that can be nicely decomposed into blocks. Then if A, D are invertible:

$$\begin{bmatrix} A & B \\ C & D \end{bmatrix}^{-1} = \begin{bmatrix} (A - BD^{-1}C)^{-1} & 0 \\ 0 & (D - CA^{-1}B)^{-1} \end{bmatrix} \begin{bmatrix} I & -BD^{-1} \\ -CA^{-1} & I \end{bmatrix}.$$

  In the general case we only need either A or D invertible as long as $D - CA^{-1}B$ or $A - BD^{-1}C$ is also invertible (the Schur complement, if you draw the block matrix as mappings between two pairs of spaces you'll see that there's the direct path $x_2 \to y_2$ via D and $x_2 \to y_1 \to x_1 \to y_2$ via $CA^{-1}B$, allowing us to precompensate the component of $x_1 \to y_2$ in the right hand matrix, leaving a block diagonal matrix to invert), but the expression becomes too tedious to memorise. *I expect you'll be given any matrix inversion tricks required in the exam, it's not a test of memorising obscure linear algebra identities.*

In this case, $x_3 = \dot{\theta}$, $\dot{x}_3 = \ddot{\theta} = f_3(x, u)$ is given in the question preamble and $x_2 = z - a$ so we have $A_{3,2} = \frac{\partial(\ddot{\theta})}{\partial(z-a)}|_{x_0,u_0}$. We also have here

$$sI - A = \begin{bmatrix} sI & -I \\ -P & sI \end{bmatrix},$$

thus

$$(sI - A)^{-1} = \begin{bmatrix} (sI - s^{-1}P)^{-1} & 0 \\ 0 & (sI - s^{-1}P)^{-1} \end{bmatrix} \begin{bmatrix} I & s^{-1}I \\ s^{-1}P & I \end{bmatrix}$$

(I'm not solving this though, it's still far too much effort.)

Question 4

No comments, so let's go over $G(s)$:

- The Neumann series: let T be an operator and $S = \sum_{k=0}^{\infty} T^k$. Using the exact same trick from regular algebra, multiply by T and take the difference: $(I - T)S = I - \lim_{k\to\infty} T^k$. Provided that $\|T\| < 1$, we have $(I - T)S = I$, or that $(I - T)^{-1} = \sum_{k=0}^{\infty} T^k$.
- That $\mathcal{L}\{e^{At}\} = (sI - A)^{-1}$. Bash the expression into the Neumann series by pulling out an s: $s^{-1}(I - \frac{A}{s})^{-1}$, then expand to get:

  $$(sI - A)^{-1} = \sum_{k=0}^{\infty} \frac{A^k}{s^{k+1}}.$$

  Note that $\mathcal{L}\{t^k\} = \frac{k!}{s^{k+1}}$, and inverse transform each term of the series.

  $$\mathcal{L}^{-1}\{(sI - A)^{-1}\} = \sum_{k=0}^{\infty} \frac{A^k t^k}{k!} = \sum_{k=0}^{\infty} \frac{(At)^k}{k!},$$

  which by the definition of $e^{(\cdot)} = \sum_{k=0}^{\infty} \frac{(\cdot)^k}{k!}$ yields e^{At}.
- Once you're happy with $e^{At} \leftrightarrows (sI - A)^{-1}$, you can just use the fact that multiplication in s is convolution in t to justify

  $$x(t) = e^{At}x(0) + \int_0^t e^{A(t-\tau)}Bu(\tau)\, d\tau,$$

  or you can use the real-space justification of the system evolving on its own and the effect of inputs being weighted by the impulse response of the system at the time since they occurred.

Question 5

- Wherever possible, diagonalise your system. There are many equally crap ways to compute e^{At}, but if A has an eigenbasis then nothing beats transforming it to the diagonal form, letting the eigenstates evolve independently of each other and recombining them when we need to do things in real space again.

Problem-specific notes:

(a) Remember that for the system $\dot{x}(t) = Ax(t)$, $x(0) = x_0$ where $Ax_0 = \lambda x_0$, we can instead express $x(t)$ as $\alpha(t)x_0$ with $\alpha(0) = 1$. Now our equation of motion becomes $\dot{x}(t) = \dot{\alpha}(t)x_0 = \lambda\alpha(t)x_0$, giving our classic ODE $\dot{\alpha} = \lambda\alpha$ and the solution $x(t) = e^{\lambda t}x_0$.

- Let A have an eigenbasis defined by $AV = V\Lambda$, then express our initial state $x(0) = \sum_i \alpha_i v_i$. As each state evolves independently, we have $x(t) = \sum_i \alpha_i e^{\lambda_i t} v_i$ (if we represent $\alpha_i(t) = \alpha_i(0)e^{\lambda_i t}$, this is $e^{\Lambda t}\alpha(0)$, where $\alpha(0) = V^{-1}x(0)$). Hence, $e^{At} = Ve^{\Lambda t}V^{-1}$.

(b)
- Get good at doing partial fractions. Remember that repeated roots expand to terms which transform back to polynomial terms in t up to one degree less than their multiplicity.

- When complex eigenvalues show up, leave them in their quadratic form and look to pattern match the numerator $(bs + c)$ to the Laplace transform pairs:

$$\begin{aligned} e^{-\alpha t}\sin\omega t &\quad\leftrightarrows\quad \frac{\omega}{(s+\alpha)^2+\omega^2} \\ e^{-\alpha t}\cos\omega t &\quad\leftrightarrows\quad \frac{s+\alpha}{(s+\alpha)^2+\omega^2}. \end{aligned}$$

  Take however much cos you need to match the bs, then make up the difference in the c term with sin.

- Non-zero equilibrium requires finding $x \neq 0$ such that $\dot{x} = Ax = 0$, so we're looking for the null space of A. A condition for this is that $\det A = 0$ and in this case the null space is obvious, but the exam may not be so nice!

Question 6

(a) More diagonalisation. Drawing our little first order expansion diagram of messy, skewed real space and perfect, neatly orthogonal eigenspace:

$$\begin{aligned} \dot{x} &= f(x_0 + \delta x, u_0 + \delta u) \\ &= f(x_0, u_0) + \frac{\partial f}{\partial x}(x_0, u_0)\delta x + \frac{\partial f}{\partial u}(x_0, u_0)\delta u + O(\delta^2) \\ &\approx A\delta x + B\delta u \end{aligned}$$

Note that x_0 is preserved! We're only distorting the local volume around x_0. Now rename $x := \delta x, u := \delta u$ to get rid of the annoying delta terms. We'll use the convention $x = Tz$ to represent the fact that z is the preferred form of the system and x is the transformation of that. Giving the game away slightly by naming $\overline{A} = \Lambda$, but the transform that diagonalises A satisfies $AT = T\Lambda$, so a little algebraic manipulation to express x in terms of z yields $\dot{z} = T^{-1}\dot{x} = T^{-1}Ax + T^{-1}Bu = T^{-1}ATz + \overline{B}u = \Lambda z + \overline{B}u$. Remember that if we start on an axis in eigenspace we stay on that axis and all the eigenstates evolve independently of each other, whereas in real space every state is continuously influencing the evolution of every other state. When it comes back to the real world and we need the outputs, we can go fishing from real space into eigenspace and get $y = Cx = CTz = \overline{C}z$.

(b) We can round trip the system by going $x(0) \xrightarrow{T^{-1}} z(0) \xrightarrow{e^{\Lambda t}} z(t) \xrightarrow{T} x(t)$, giving our relationship for $e^{At} = Te^{\Lambda t}T^{-1}$.

Question 7

(b) Some guidelines for sketching:

- To get the rotation direction of an equilibrium point with complex eigenvalues, just pick a point with a small offset and plot $f(x)$ or $A\delta x$. Make life easy for yourselves by aligning your offsets with the axes and just plotting the matrix columns. Don't put massive amounts of effort in, just draw a bunch of swirls going into or out of the point depending on the sign of $\Re(\lambda)$.
- Otherwise, sketch the eigenbasis axes at the point with arrows showing the direction of motion on them, then draw some sweeping trajectories into, out of or saddling this point depending on the signs of λ_i.
- If the point has some special properties such as uniform expansion or circular motion, make sure to highlight this to the examiner with a little note.

- Don't put massive amounts of effort in unless instructed to, just tell the examiner a convincing story about the trajectories by linking together the regions of influence of the points you have computed.

(c) If the equilibrium point is stable for a parameter range $\gamma \in [a, b]$, but the point $x_0 = x_0(\gamma)$, then as long as you start the system at $x(0)$ *near enough* to $x_0(\gamma(0))$ and change $\gamma(t)$ *slowly enough*, the trajectory will remain *near* $x_0(t)$. How near? How slow? I don't know, but there always exists a distance small enough that the higher order effects are overwhelmed. Take 4F2 and find out just how little we can say in general.

- The concept of Lyapunov exponent: we might want to place a bound on the distance between two system trajectories at time t given that they started $\delta x(0) = x_1(0) - x_0(0)$ apart. Locally, we will have the bound $\|\delta x(t)\| \leq e^{\lambda_{\max} t}\|\delta x(0)\|$. In this course, we're mostly interested in $x_0(t) = x_0$ for an equilibrium point.

Examples Paper 2

Question 1

(b) Remember that in the general case we may have multiple eigenvalues $\{\lambda_i \colon \Re(\lambda_i) < 0\}$ with associated eigenvectors $Av_i = \lambda_i v_i$, then our general form for a stable x_0 is $\sum_i \alpha_i v_i$. If you have calculated the eigenvectors, write it out as $\alpha v_1 + \beta v_2 + \ldots$, maybe even explicitly in column notation.

Question 2

This question is an exercise in boxing/unboxing systems, manipulating the time derivative and converting to Laplace transforms. It's good practice to get comfortable with both the boxed $G(s)$ form and explicitly laying out the individual components of the system, as we will later be tapping off from $\hat{x}$ and feeding in to $\hat{\dot{x}}$ when we do observers.

- The identity $A(I + BA) = (I + AB)A$ may be useful to remember here: if in doubt, pull out the factor of A on either side of $A + ABA$.
- A related and equally useful identity that shows up when you have a system G in negative feedback with a controller K: $G(I + KG)^{-1} = (I + GK)^{-1}G$.

Question 3

No comments, which means it's time for a rambling aside about root loci. I can't find a source for any of this, so you'll just have to trust me.

Let's start with

$$d(s) + kn(s) = p(s, k) = \beta \prod_i (s - \alpha_i(k))$$

for our roots α_i as a function of k. I think the lecture notes do a good job of covering the behaviour of $k \to 0$ and $k \to \infty$, so I won't go over that again. Note that if $\deg d(s) > \deg n(s)$, β is determined entirely by the leading order term of $d(s)$.

Right, let's suppose we have $\alpha_i(k)$ for a particular value of k and that the particular α_i in question is isolated. What happens when we change k? From the same principle as used to derive the linearised equations for an implicit system, we have $p(s, k) = 0$ where we view $p(s, k) : \mathbb{C} \times \mathbb{R} \to \mathbb{C}$ as a restricted set of functions with special rules from something that looks like $\mathbb{R}^2 \times \mathbb{R} \to \mathbb{R}^2$: $x + iy, k \to u(x, y, k) + iv(x, y, k)$. So our complex derivative $\frac{\partial p}{\partial s}$ is actually a Jacobian J, and by the implicit function theorem we have:

$$\frac{\partial \alpha}{\partial k} = -J^{-1}\frac{\partial p}{\partial k}.$$

Well, $\frac{\partial p}{\partial k} = n(s)$ and $\frac{\partial p}{\partial s} = d'(s) + kn'(s)$. Let's start with some properties of J.

Holomporphicity. A holomorphic mapping is a function which is complex differentiable everywhere on its domain. It satisfies the Cauchy-Riemann equations $u_x = v_y$; $u_y = -v_x$ (where the subscript denotes partial derivatives), meaning that its Jacobian is a rotation and scaling. *Conformality.* A mapping is conformal if it preserves angles at a point. They are holomorphic with the constraint that the derivative is non-zero.

- If $f(s), g(s)$ are holomorphic, then $h(s) = f(s) + g(s)$ is holomorphic (trivial).
- If $f(s), g(s)$ is holomorphic, then $h(s) = f(s)g(s)$ is holomorphic. For our transform now becomes:

$$u^*(x,y) = U(x,y)u(x,y) - V(x,y)v(x,y)$$
$$v^*(x,y) = V(x,y)u(x,y) + U(x,y)v(x,y),$$

with the associated Jacobian

$$J^* = \begin{bmatrix} U_x u + U u_x - V_x v - V v_x & U_y u + U u_y - V_y v - V v_y \\ V_x u + V u_x + U_x v + U v_x & V_y u + V u_y + U_y v + U v_y \end{bmatrix}$$
$$= \begin{bmatrix} U_x u + U u_x + U_y v + V u_y & U_y u + U u_y - U_x v - V u_x \\ -U_y u + V u_x + U_x v - U u_y & U_x u + V u_y + U_y v + U u_x \end{bmatrix} = \begin{bmatrix} a & b \\ -b & a \end{bmatrix}$$

- $f(s) = s$ is holomorphic ($J = I$).
- Conclusion: polynomials in s, such as $p(s,k)$ at fixed k, are holomorphic.

What does this mean for root loci? Well, if J is invertible then we can take $\frac{\partial \alpha}{\partial k}$ meaningfully, and we get a differentiable trajectory as a function of k: we come into a point in the opposite direction to how we leave it. Remember also that a rotation and a scaling is exactly what a complex number in $re^{j\theta}$ form does, so our derivative itself is just a complex number D and the inverse is just $1/D$. What about when J cannot be inverted? This requires that $\frac{\partial p}{\partial s} = 0$, and conformality is lost. This is the only way that certain structures could occur, such as kinks in the root locus diagram like at breakaway points. So let's close the loop now: let $k_0 = -\frac{d(\alpha_0)}{n(\alpha_0)}$ where conformality is lost, and plug into $p'(\alpha_0, k_0) = d'(\alpha_0) - \frac{d(\alpha_0)}{n(\alpha_0)} n'(\alpha_0) = 0$. Rearranging, we find that α_0 solves $d(s)n'(s) - n(s)d'(s) = 0$, which is just the numerator part of the condition $\frac{dL}{ds} = 0$ that you've already seen.

More generally, anywhere that $p(s,k)$ has repeated roots is degenerate at the repeated roots:

$$p(s) = (s - \alpha_0)^n \prod_i (s - \alpha_i)$$
$$p'(s) = (s - \alpha_0)^{n-1} \left(n + \sum_i \frac{s - \alpha_0}{s - \alpha_i} \right) \prod_i (s - \alpha_i)$$

including any repeated poles and zeros (think about how the behaviour changes near repeated poles around $k = 0$ with both positive and negative k). What happens in the vicinity of multiple roots? Can multiple roots move together, i.e. $\alpha_i(k) = \alpha_j(k)$ for $k_0 < k < k_1$? Can the same root appear for separate values of k? That last one is easy to answer: for if α_0 solves $d(s) + k_1 n(s) = d(s) + k_2 n(s) = 0$, then we have $d(\alpha_0)$ and $n(\alpha_0)$ on both sides and thus $k_1 = k_2$. So contours on the root locus diagram can never self-intersect.

Roots moving together would require a continuum of points at which $\alpha(k)$ fails to be differentiable with respect to k, yet we have only a finite number of these points as $\frac{dL}{ds}$ has a polynomial numerator. So we can conclude that any point at which roots come together also requires them to leave, i.e., that the point is isolated.

Finally: when can a trajectory stop? This would require $\frac{\partial \alpha}{\partial k} = 0$, which can only be when $\frac{\partial p}{\partial k} = n(s) = 0$. So the only locations where the trajectories stop moving are the open loop zeroes.

Question 4

- Let's tie Part I vibrations into root locus theory. When $\zeta = 1$ we have critical damping, which you will recall means that we have repeated roots and any less damping will lead to oscillations. On the root locus diagram, we have repeated roots at a breakaway point and any less damping will lead to the roots moving away from the real axis, gaining an imaginary component and leading to oscillations. Any more damping and the roots move away from each other, giving us the standard two mode decay of an overdamped system, and extreme values lead to one very slow and one very fast mode as one root goes to zero and the other to infinity.

Question 5

- Final value theorem and steady state error: in Part I you were taught that zero steady state error requires the system to have an integral response. But this system has an integrator, and yet there's still a steady state error! Why?

  Suppose $L(s) = \frac{N(s)}{D(s)s^n}$, where we have pulled out all our excess integrators (n of them) explicitly from the remaining poles. The error

  $$\tilde{e} = \frac{1}{1 + kL(s)}\tilde{r} = \frac{D(s)s^n}{D(s)s^n + kN(s)}\tilde{r},$$

  and our reference signal $r(t) = t^a$ transforms to $\tilde{r}(s) = \frac{a!}{s^{a+1}}$. Applying the final value theorem:

  $$\lim_{t\to\infty} e(t) = \lim_{s\to 0} s\tilde{e}(s) = s\frac{D(s)s^n}{D(s)s^n + kN(s)}\frac{a!}{s^{a+1}} = \frac{D(s)a!}{D(s)s^n + kN(s)}\frac{s^{n+1}}{s^{a+1}}.$$

  The condition for this limit to go to 0 is that $n > a$, or, that we have one more integrator in our system than the degree of the reference signal.

Examples Paper 3

This paper tends to pose few practical problems, just that students don't know why they're doing it. The lecture notes are pretty chaotic, so here's how I view the content covered.

Theory, continuous time

Cayley-Hamilton and e^{At}

Cayley-Hamilton theorem allows us to write $A^k = \sum_{i=0}^{n-1} \alpha_{k,i} A^i$ for $k \geq n$, where the $\alpha_{k,i}$ is understood as being the contribution from the term k to the term i. We're not particularly interested in the values of $\alpha_{k,i}$, just that they exist. Write out the definition of e^{At}:

$$e^{At} = \sum_{k=0}^{\infty} \frac{t^k A^k}{k!},$$

and apply Cayley-Hamilton to the terms of $k \geq n$ to move their contribution down to the first n terms:

$$\begin{aligned} e^{At} &= \sum_{k=0}^{n-1} \left(\frac{t^k}{k!} A^k + \sum_{m=n}^{\infty} \frac{t^m}{m!} \alpha_{m,k} A^k \right) \\ &= \sum_{k=0}^{n-1} \left(\frac{t^k}{k!} + \sum_{m=n}^{\infty} \frac{t^m}{m!} \alpha_{m,k} \right) A^k \\ &= \sum_{k=0}^{n-1} \beta_k(t) A^k. \end{aligned}$$

Some facts about β_k that we will save for later:

- No β_k is identically zero over any finite interval $0 \leq t \leq T$, as each β_k is a power series and has at least one non-zero coefficient in $\frac{t^k}{k!}$.
- No β_k can be expressed as a linear combination of the other $\beta_{i\neq k}$, as each β_k is the only one that has a term of t^k.

Some facts that we will not do much with yet, but a future version of the companion notes might:

- $\beta_0(0) = 1$, $\beta_{k\neq 0}(0) = 0$.
- The relevant β_k are continuous. For e^{At} converges everywhere, so for β_k to not converge at a point we must be able to express A^k as a sum of other powers of A^m. In this case distribute A^k over the remaining terms until this can no longer be done, then the surviving β_m are power series which converge everywhere, hence are continuous.

Observability

Rather than take derivatives of y at a point in time, I prefer the alternative view of observability: if we watch the system evolve over a finite time interval $[0, t]$, can we deduce the initial condition x_0 (and hence all future output)? Note that this is independent of our notion of observing as we probably would in practice by finite time sampling (that comes later), and we jump straight to the idea of feeding the output into an observer system: these evolve continuously with time and so can "feel" everything from $y(\tau)$ over the interval.

With with our output at time τ: $y(\tau) = Ce^{A\tau}x_0$, we apply Cayley-Hamilton and find:

$$y(\tau) = \sum_{k=0}^{n-1} \beta_k(\tau) CA^k x_0.$$

There are two ways that a state x_0 could not show up in the output. Firstly, it could only be in the row space of CA^k for specific powers of k, but these powers are zeroed out by β_k: we have already established that this is not possible for all time over a finite interval. The only remaining way for information to disappear is that it lies in the null space for all CA^k, yielding our Q test and the corresponding decomposition that the observable subspace is the row space of Q and unobservable subspace is the null space of Q.

Controllability

Suppose we have a target state that we wish to achieve at some time t after we set it going $x^*(t)$ and an initial state x_0. We need to find an input $u(\tau)$ such that we make up the difference between how the initial state evolves and the target state:

$$x^*(t) - e^{At}x_0 = \Delta = \int_0^t e^{A(t-\tau)} B u(\tau)\, d\tau.$$

Apply Cayley-Hamilton, and exploit linearity to swap the order of summation and integration:

$$\Delta = \sum_{k=0}^{n-1} A^k B \int_0^t \beta_k(t-\tau) u(\tau)\, d\tau,$$

i.e., the attainable perturbation from the natural evolution of the initial state is a sum of the columns of P, and the remaining question is whether we can find u such that we can exactly control Δ. The first way that a column could not contribute to the final state regardless of the input is if the associated β_k was equal to zero over the whole interval, but this has already been excluded.

Whilst I would like to denote the time reversed version of β_k as $\overleftarrow{\beta}$ or $\beta^{\leftarrow}$, this makes typesetting difficult so I will instead use $b_k(\tau) = \beta_k(t-\tau)$. Then our integral of the i^{th} input can be seen to be an inner product

$$(b_k, u_i) = \int_0^t b_k(\tau) u_i(\tau)\, d\tau.$$

For $r = \text{rank}(P)$ arbitrary columns in P that span the reachable space (column space of P), we want to fashion an input that excites that column alone and nothing else. There's certainly something to excite: b_k is continuous and not identically zero, so it is non-zero with consistent sign and magnitude bounded below over a subinterval of $[0, t]$, giving the existence of a target location for us to apply an input function and get something out of the integral. Suppose that one such column is the i^{th} column of BA^j and we need an amount α of this column to build Δ. Then we wish for a function $\tilde{u}_j$ such that $(\tilde{u}_j, b_j) = 1$ and $(\tilde{u}_j, b_k) = 0$ for $j \neq k$, or that $\tilde{u}_j$ is orthogonal to the other b_k, and we will apply $\alpha\tilde{u}_j$ to u_i.

To get our functions $\tilde{u}_j$, we will use Gram-Schmidt orthogonalisation for each function in turn: reorder the functions to put b_j as the n^{th} function in the set $\{b_k\}$ and progressively remove components of the remaining functions:

$$\begin{aligned} a_1 &= b_1 \\ a_i &= b_i - \sum_{k<i} \frac{(a_k, b_i)}{(a_k, a_k)} a_k, \end{aligned}$$

then set $\tilde{u}_j = \frac{a_n}{(b_j, a_n)}$ to achieve our unity scaling. Remember that the intermediate functions a_i are there to account for the fact that the intermediate (b_i, b_j) terms may not be zero, so it is not enough

to just progressively remove the component of b_i in isolation as we may reintroduce previously removed components in the process. Here we reach our final possibility for a column of P to go missing from the final state: that $a_n = 0$, i.e., there is nothing left over after removing the components of the other b_k. This would require b_j to be a linear combination of the other b_k, but fortunately we have had the foresight to establish that this is not the case already.

Thus we can conclude that for any time interval $[0, t]$, we can perturb our state at t by any amount Δ that can be expressed as a linear combination of the columns of P.

Gramians

Note that the tests on P, Q are an existence test: they tell us that for any finite time interval $[0, t]$ we *can* identify the observable states and we *can* reach the controllable (hence the alternative name reachable) states. What they don't tell us is whether our proposed strategy is a good idea:

- Are we looking for a signal that is drowned out by other signals?
- Are our chosen P columns particularly stubborn, and would we do better to mix and match different amounts of all the columns at varying times?
- If we accepted a longer observation period or slower movement to our final state, could we make life easier?

You'll need to remember the minimum energy control law, as you don't have the tools yet to even start deriving it in the exam if you forget it. You've seen the verification that it does indeed get us from A to B and is the minimum energy control law, so for now that should be enough. When you've taken 4M12 and know how to use variational methods, the extra lecture slides cover how to derive it.

Theory, discrete time

In discrete time, our transition matrices are fixed by Δt to be $\Phi = e^{A\Delta t}$ and our selected input shape functions $u = U(\tau)u_k$ over the k^{th} interval $\Gamma = \int_0^{\Delta t} e^{A(\Delta t - \tau)} BU(\tau)\, d\tau$, to yield the update rule $x_{k+1} = \Phi x_k + \Gamma u_k$.

Now it is possible for things to go missing: remember that before we relied on the fact that we were dealing with a function over an interval, so the fact that $\beta_k(\tau_0) = 0$ was countered by the fact that we knew there existed other points in the interval where it was non-zero. Our output at point k arising from x_0 is given by

$$y_k = C\Phi^k x_0,$$

so we again apply Cayley-Hamilton to terms of $k \geq n$ and shift it down into the lower range:

$$y_k = \sum_{i=0}^{n-1} \phi_{k,i} C\Phi^i x_0$$

for contributions $\phi_{k,i}$. We can therefore see that our observable states span the row space of a new Q matrix formed of $C\Phi^k$ powers, and that if a state hasn't shown up after n iterations it's not showing up at all!

Controllability: the input from u_0 gets filtered into x_k as $\Phi^{k-1}\Gamma u_0$, which again we can apply Cayley-Hamilton to and conclude that our reachable states are the columns of $P = [\Phi^k \Gamma]$. Fashioning the required input is much easier in this case, as we just need to pick the columns of P and apply a pulse to the i^{th} element of u at k timesteps before our target. Note that in both cases the relevant columns of P and rows of Q mean that we might not be able to achieve or see certain states for all $k < n$, although everything after that is doable.

Observable and controllable subspaces

There's an arrow missing in the diagram from the observable/uncontrollable subspace to the observable/controllable subspace. Arrows indicate influence (non-zero blocks of $\tilde{A}$): controllable states can't influence uncontrollable states, and unobservable states can't influence observable states.

Remember that when doing your decompositions, if you have $\mathbb{R}^n = U \oplus V$ for subspaces U, V of dimension a, b with a chosen basis $\{u_i\}_{i=1}^a$, $\{v_j\}_{j=1}^b$, then to represent a vector as a linear combination of these basis vectors $x = \sum_{i=1}^a \alpha_i u_i + \sum_{j=1}^b \beta_j v_j$:

$$x = \begin{bmatrix} \uparrow & & \uparrow & \uparrow & & \uparrow \\ u_1 & \dots & u_a & v_1 & \dots & v_b \\ \downarrow & & \downarrow & \downarrow & & \downarrow \end{bmatrix} \begin{bmatrix} \alpha_1 \\ \vdots \\ \alpha_a \\ \beta_1 \\ \vdots \\ \beta_b \end{bmatrix} = \begin{bmatrix} T_u & T_v \end{bmatrix} \begin{bmatrix} \tilde{x}_u \\ \tilde{x}_v \end{bmatrix} = T\tilde{x},$$

hence our transform $x \to \tilde{x}$ being T^{-1}.

Question 4

Just because you can get to a state x^* using a subset of the inputs doesn't mean you can keep it there continuously, as you'll see when using T_2 alone for control. But what can you do?

- Keep some states there continuously: we need to set $\dot{x} = Ax + Bu = 0$, so any states where Ax^* lies in the column space of B can be kept there.
- Pick a time t, and apply an input $u^*(\tau)$ such that $\int_0^t e^{A(t-\tau)} B u^*(\tau)\, d\tau = (I - e^{At})x^*$. You can therefore return to x^* with period t, for arbitrarily small t. For some systems, the *describing function* approach covered in 4F2 may help you see how the universe foils our plans for infinitely high frequency oscillations and can be used to work out the frequency of limit cycles in the presence of nonlinearities.
- I have a strong suspicion—but not enough time to prove it—that by picking t small enough we can stay arbitrarily close to x^*.