

Experimental Identification and Control of a Twin-Rotor Helicopter Test Bench

Dev Kumar, Steevy Doumbe Moundy, Léa Couësme

▶ **To cite this version:**

Dev Kumar, Steevy Doumbe Moundy, Léa Couësme. Experimental Identification and Control of a Twin-Rotor Helicopter Test Bench. Independent research; original project at École Nationale Supérieure d'Arts et Métiers. 2026, pp.1-29. ⟨hal-05873038⟩

HAL Id: hal-05873038

https://hal.science/view/index/docid/5873038

Submitted on 19 Nov 2026

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Copyright - All rights reserved

Experimental Identification and Control of a Twin-Rotor Helicopter Test Bench

Sampled pole placement, integral linear-quadratic control
and reproducible model comparisons

Dev KUMAR, Steevy DOUMBE MOUNDY and Léa COUËSME

École Nationale Supérieure d'Arts et Métiers
Conservatoire National des Arts et Métiers

Paris and Saint-Denis, France

2025–2026 engineering project
Experimental records and numerical study

Abstract

A three-axis twin-rotor test bench is studied using the original identification and control results from two campaigns. The first design uses sampled elevation feedback with integral action, and the second uses continuous multivariable integral linear-quadratic control of elevation and yaw rate. The archived elevation transfer and state realization have different input gains. Exact zero-order-hold reconstruction shows that the numerical matrix reproduces the reported digital gains, poles and command scales. For the multivariable design, the printed Riccati matrix and poles belong to a synthesis matrix shifted by 0.5 s^{-1}. The physical poles and the reconstructed nominal responses account for the archived 5% times of about 3.1 s in elevation and 2.7 s in yaw rate. A Routh calculation proves stability for the fixed printed gain and every yaw lag from 10 to 20 s in the decoupled linear plant family. The same lag change increases simulated yaw-rate settling from 2.72 to 7.30 s. All technical figures from the 20-page source are reproduced. Additional Python tests distinguish the voltage required by reference tracking from recovery after a finite motor-input pulse.

Keywords twin-rotor bench, experimental identification, digital state feedback, pole placement, integral linear-quadratic control, Riccati equation, Routh criterion.

Study record

Source [1] is the 20-page engineering report from the 2025–2026 project. Page references to that report count the cover. The 22 retained technical panels include the cover schematic, the overview illustration, bench photographs, all original curves and implementation diagrams. Appendix A links each panel to its revised number. Appendix B preserves the source equations and the original comparison table. Appendix C describes the independently generated simulations.

Acknowledgements

The original project was supervised by Hervé Guillard, Nazih Mechbal and Marc Rébillat. The source credits École Nationale Supérieure d’Arts et Métiers and Conservatoire National des Arts et Métiers for the project.

Nomenclature

Symbol	Meaning
ε, θ, λ, v	Elevation, pitch, yaw angles and yaw rate $v = d\lambda/dt$
u_g, u_d, u_Σ, u_Δ	Motor voltages, their sum and their difference
H_e, H_t, H_l	Elevation, pitch and pitch-to-yaw transfer functions
K_e, K_t, K_l, b_e	Channel gains and elevation acceleration coefficient
ξ, ω_0, ω_d, τ_l	Damping ratio, natural and damped angular frequencies, yaw time constant
x, x_e, x_a, η	Plant states, sampled and continuous augmented states, integrated tracking error
A, B, C, D, A_a, B_a	Continuous plant and integral-augmented matrices
F, G, F_e, G_e	Sampled plant and integral-augmented matrices
E_r, r, d, W_2, W_3	Reference input map, reference, input disturbance, reachability matrices
K_p, K_i, Q, R, P	Feedback gains, quadratic weights and Riccati solution
T_s, T_c, α_c	Sampling period, placement time constant and LQI spectral shift
s, z, Δ_2, Δ_3	Continuous and sampled complex variables, yaw Routh determinants

Contents

Abstract 1
Study record 1
Acknowledgements 1
Nomenclature 1
1. Introduction 3
2. Experimental identification and plant models 3
2.1 Bench coordinates and motor mixing 3
2.2 Elevation step in the multivariable campaign 4
2.3 Differential drive and pitch 6
2.4 Pitch-to-yaw transfer 6
2.5 Five-state multivariable representation 7
3. Sampled elevation servo 8
3.1 Refined identification and coordinate order 8
3.2 Exact zero-order hold and reachability 8
3.3 Pole assignment and command generation 9
3.4 Nominal and faster responses in the original figures 11
3.5 Python comparison at fixed amplitude 12
3.6 Real-time implementation 13
4. Multivariable integral linear-quadratic control 13
4.1 Augmented reachability and integral tracking 13
4.2 Quadratic criterion and spectral shift 14
4.3 Numerical audit of the reported gain and poles 14
4.4 Comparison with the archived nominal response 15
4.5 Comparison with the archived disturbed response 16
5. Model sensitivity and controlled simulations 16
5.1 Decoupling and actuator demand 16
5.2 Continuous proof for the yaw-lag interval 17
5.3 Step tracking and the effort required by the shift 18
5.4 A finite input pulse with known removal time 19
6. Comparison and conclusions 20
6.1 What the two campaigns establish 20
6.2 Limits and measurements that would resolve them 21
6.3 Conclusion 21
References 22
Appendix A. Original figure concordance 23
Appendix B. Original equations and comparison table 24
Appendix C. Reproducibility 28
C.1 Embedded material 28
C.2 Calculation sequence 28

1. Introduction

A twin-rotor helicopter test bench combines actively driven elevation and pitch with yaw motion induced by the inclination of the rotor assembly. The same pair of motor voltages is used for both tasks. This makes the bench useful for studying how a design that starts from separate identification runs behaves when the commands share the available actuator voltage. The engineering project in [1] developed both a sampled elevation servo and a continuous multivariable servo for elevation and yaw rate. Here those two campaigns are brought together using the original figures and implementation diagrams.

The study tests three questions. Can the printed models reproduce the sampled controller gains and the reported continuous pole locations? What does the difference between the printed elevation models and the two yaw time constants change when the feedback gains are held fixed? Which aspects of the archived time histories agree with these reconstructions? The first two questions use algebra and simulations. The third uses the visible annotations and axes of the archived plots. The full bench recordings, calibration files and model workspaces were not available in the 20-page source [1]. Times quoted from its prose and values read from its images are identified as such.

The notation fixes the physical direction of the motor commands before defining the feedback signs. Elevation is the vertical rotation of the main arm. Pitch is the rotation of the rotor cradle. Yaw is the azimuth of the bench. The experiments described in [1] used the responses of individual channels to estimate low-order dynamics. Local motion about the test configuration is the working domain of these models. A reference in yaw rate changes the azimuth continuously as long as that rate is maintained.

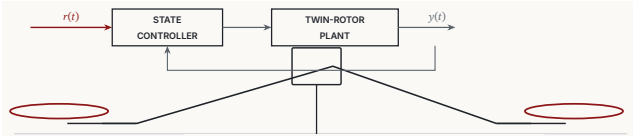

Figure 1.1 – Original feedback sketch from the cover of [1]. The two motors act on a common mechanical assembly.

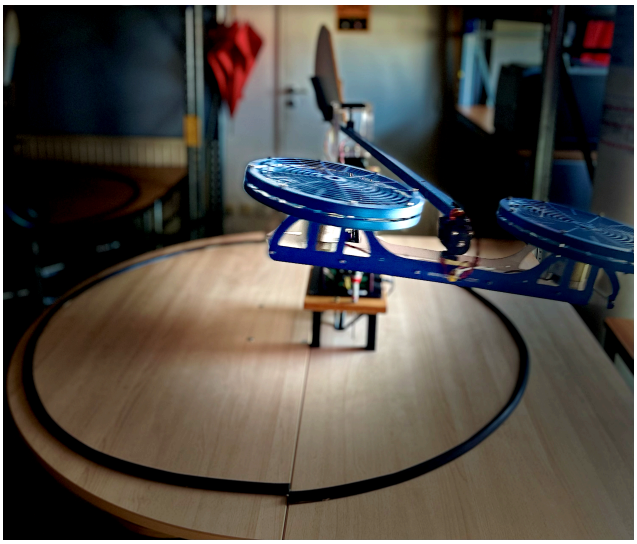

Figure 1.2 – Overview illustration accompanying the original description of the project [1, p. 3]. The bench photographs are in Figures 2.1–2.3.

The report first identifies the assumptions linking the three tests to the five-state model. It then reconstructs the sampled controller and the integral linear-quadratic controller. The original traces are reproduced alongside the relevant equations, including the simulations used for controller tuning and the dSPACE implementation figures. Sections 5 and 6 trace the effects of yaw time constant and finite input disturbances, and compare the two control designs within their respective operating scopes.

2. Experimental identification and plant models

2.1 Bench coordinates and motor mixing

Let u_g and u_d be the left and right voltage commands, with ε, θ and λ the elevation, pitch and yaw angles. The source uses a sum channel for elevation and a difference channel for pitch. It reports angle and rate signals available to the real-time controller. Sensor type, bias removal, rate estimation and voltage calibration are not

documented in the supplied figures. The lowercase commands here refer to model inputs about the chosen equilibrium.

$$\begin{aligned} u_\Sigma &= u_g + u_d, \qquad u_\Delta = u_g - u_d \\ u_g &= \tfrac{1}{2}(u_\Sigma + u_\Delta), \quad u_d = \tfrac{1}{2}(u_\Sigma - u_\Delta) \end{aligned} \tag{1}$$



Figure 2.1 – Side photograph of the experimental bench in [1]. Source Fig. 1.1(a), p. 8 of [1].



Figure 2.2 – Elevation configuration in the original photographic record. Source Fig. 1.1(b), p. 8 of [1].



Figure 2.3 – Horizontal configuration in the original photographic record. Source Fig. 1.1(c), p. 8 of [1].

The summation and difference are useful for identification because the same thrust change on both sides mainly excites elevation, while opposing changes mainly generate pitch. The yaw model is then excited by pitch. The resulting input decomposition is exact for the assembled linear matrices. On the bench the separation can be affected by rotor mismatch, amplitude-dependent thrust, friction and geometry. The original tests do not quantify those cross-channel effects. The decomposition and the shared voltage constraint are therefore treated separately in Sections 4 and 5.

2.2 Elevation step in the multivariable campaign

With pitch and yaw mechanically constrained, the first campaign in [1] applies equal voltage changes to the two motors. The reported response is fitted by the underdamped model

$$\begin{aligned} H_e(s) &= \frac{\varepsilon(s)}{u_g(s) + u_d(s)} \\ H_e(s) &= \frac{K_e \omega_0^2}{s^2 + 2\xi\omega_0 s + \omega_0^2}. \end{aligned} \tag{2}$$

Here K_e is the steady gain against the summed input, ω_0 is the undamped frequency and ξ is the damping ratio. If each motor is stepped by the same amount, the step in the model input is twice the per-motor step. The gain from a single plotted voltage therefore depends on whether the display

shows one command or the sum. Let a denote the step amplitude entering (2) and take the initial deviation and its derivative as zero. Inverting $aH_e(s)/s$ gives the step solution

$$\omega_d = \omega_0\sqrt{1-\xi^2}, \qquad y_\infty = aK_e$$
$$g(t) = \cos(\omega_d t) + \frac{\xi\omega_0}{\omega_d}\sin(\omega_d t) \tag{3}$$
$$y(t) = y_\infty\left[1 - e^{-\xi\omega_0 t}g(t)\right].$$

$$\phi = \arctan\left(\frac{\sqrt{1-\xi^2}}{\xi}\right)$$
$$y(t) = aK_e\left[1 - \frac{e^{-\xi\omega_0 t}}{\sqrt{1-\xi^2}}\sin(\omega_d t + \phi)\right]. \tag{4}$$

Differentiation removes the cosine term. Specifically, $\dot{y}$ is proportional to $e^{-\xi\omega_0 t}\sin\omega_d t$ for a positive step. The successive extrema alternate about the eventual plateau and are spaced by π/ω_d. If d_i is the magnitude of the residual from the final output at the i-th extremum, d_2/d_1 is the ratio of the exponential envelope evaluated half a period later. Inverting that ratio yields

$$\frac{d_2}{d_1} = \exp\left(-\frac{\pi\xi}{\sqrt{1-\xi^2}}\right)$$
$$\delta = \log(d_1/d_2), \quad \xi = \frac{\delta}{\sqrt{\pi^2+\delta^2}} \tag{5}$$
$$\omega_0 = \frac{\pi}{(t_2 - t_1)\sqrt{1-\xi^2}}.$$

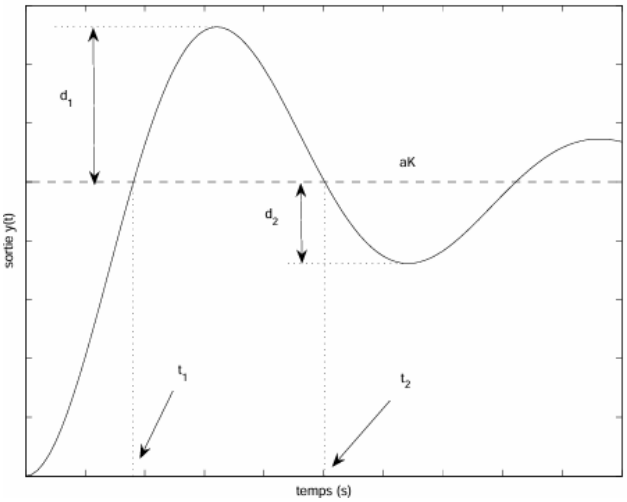

Figure 2.4 – Original diagram defining the step extrema and timing. The alternating-extremum interpretation supplies the factor π in (5). Source Fig. 1.2(a), p. 9 of [1].

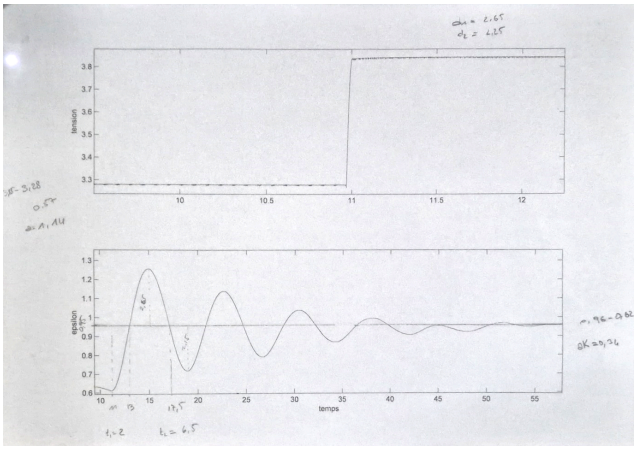

Figure 2.5 – Archived first-campaign motor and elevation traces. The lower panel shows an oscillatory response about a displaced level. The recorded panel and its original axes are retained. Source Fig. 1.2(b), p. 9 of [1].

The quoted graphical values are d_1=2.65 cm and d_2=2.15 cm, at 2.0 and 6.5 s. Distances in centimetres are readings on the printed graph in [1]. Only their dimensionless ratio enters (5). Insertion gives ξ=0.06641 and ω_0=0.69968 rad/s. The source prose gives ξ=0.05 and ω_0=0.7 rad/s. The five-state matrix in [1] uses a first-derivative coefficient of 0.112 s^{-1}, equivalent to ξ=0.080 at 0.7 rad/s. All three values describe low damping. They imply different decay envelopes even though the calculated oscillation periods are close.

The source gives aK_e=0.34 and K_e=0.29. These imply a simulated summed step of approximately 1.17 in the source input scale. Figure 2.6 holds the final elevation at 0.34 and compares the three damping choices. The first extrema differ modestly. The effect accumulates over several cycles, which is where a new recorded trace would best discriminate the models. The figure is a sensitivity check on the equations. It was generated without a digitised bench trace.

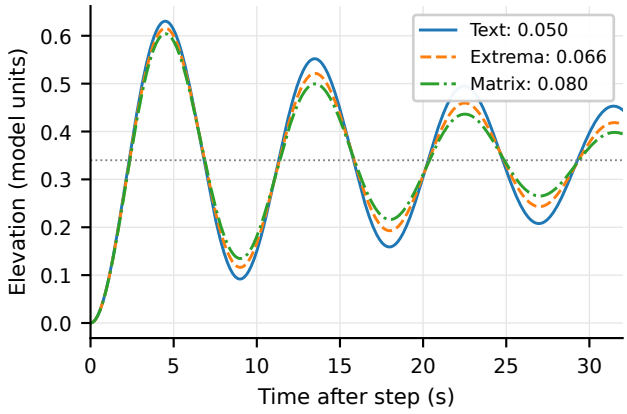

Figure 2.6 – Python. Comparison of the elevation step envelopes using the damping reported in the prose, inferred from the printed extrema and implied by the printed state matrix.

2.3 Differential drive and pitch

The source describes a pitch step experiment driven through the differential input. A simple local model of pitch is a double integrator, where a constant differential input generates constant pitch acceleration.

$$H_t(s) = \frac{\theta(s)}{u_g(s) - u_d(s)} = \frac{K_t}{s^2}$$
$$y(t) = \frac{K_t a}{2}(t - t_0)^2, \qquad t \geq t_0. \tag{6}$$

Subtracting the initial pitch and using the same-sign, post-step branch gives the square-root form. The chosen sign of the step must make the displacement under each square root nonnegative, or the trace must first be multiplied by the corresponding sign. If the initial angular velocity is nonzero, the square-root relation changes because an additional term linear in time enters the displacement.

$$\sqrt{y(t)} = \sqrt{K_t a/2}\,(t - t_0)$$
$$K_t = \frac{2}{a}\left(\frac{\sqrt{y_2} - \sqrt{y_1}}{t_2 - t_1}\right)^2. \tag{7}$$

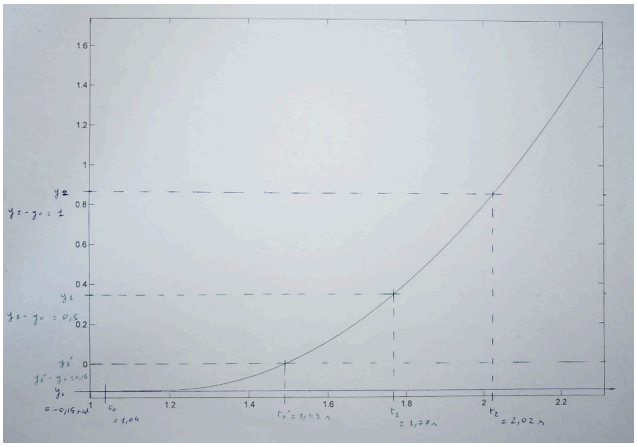

Figure 2.7 – Original pitch step record and graphical point selection for the double-integrator gain. The source uses the increasing portion after the onset of the command. Source Fig. 1.3, p. 10 of [1].

The printed estimate is K_t=1.336 in the source voltage and angle scale. Taking two sufficiently separated points reduces sensitivity to readout error in the denominator. The shape of the image supports a local acceleration model over the selected interval. It does not distinguish low viscous damping from a perfect integrator over longer times, and the double integrator should not be extrapolated to arbitrary arm angles or commands.

2.4 Pitch-to-yaw transfer

The yaw response in [1] uses the pitch angle as an input to yaw dynamics. The model contains an integrator in yaw angle and a first-order lag in yaw rate. The negative sign of the identified gain fixes the relationship between positive pitch and negative yaw rate under the source convention.

$$H_l(s) = \frac{\lambda(s)}{\theta(s)} = \frac{K_l}{s(1 + \tau_l s)}$$
$$\frac{v(s)}{\theta(s)} = \frac{K_l}{1 + \tau_l s}, \qquad v = \dot{\lambda}. \tag{8}$$

For a pitch step of amplitude a, with initial yaw rate zero, (8) gives $v(t) = aK_l(1 - e^{-t/\tau_l})$. Integrating from the onset yields the yaw displacement.

$$\hat{y}(s) = \frac{aK_l}{s^2(1 + \tau_l s)}$$
$$y(t) = aK_l(t - \tau_l + \tau_l e^{-t/\tau_l}). \tag{9}$$

At late times, the yaw-angle asymptote has slope aK_l and crosses the time axis at τ_l relative to the step

origin. The source uses graphical points A=(20,0), B=(60,102.5), a pitch amplitude of −0.491 and a lag of 20 s. Preserving the plot's original scale gives

$$aK_l = \frac{102.5 - 0}{60 - 20} = 2.5625$$
$$K_l = \frac{2.5625}{-0.491} \simeq -5.219. \tag{10}$$

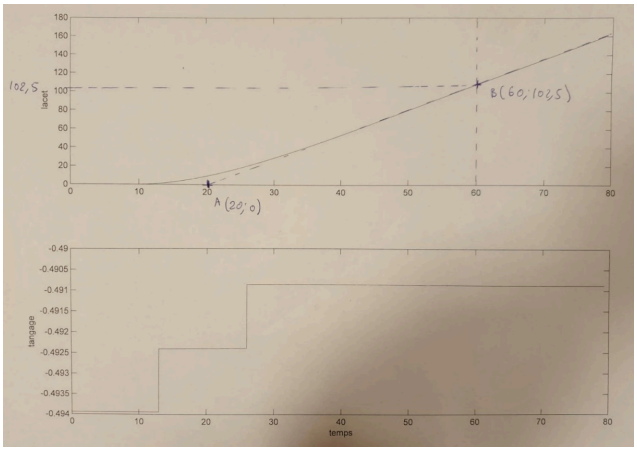

Figure 2.8 – Original measured yaw and input sequence. The upper signal tends towards a ramp. The lower signal contains more than one command level, and the numerical asymptote is interpreted in the source's chosen interval. Source Fig. 1.4, p. 11 of [1].

The yaw-rate row of the assembled state matrix has a diagonal entry of −0.1000 s^{-1} and a pitch coefficient of −0.5219 s^{-1} times the gain units. The two entries imply τ_l=10 s and the same K_l=−5.219. Changing the lag from 10 to 20 s at fixed static rate gain divides both entries by two. Figure 2.9 shows the difference for a single model step on the source numeric scale. The two model curves share the eventual slope and differ in the intercept and transient curvature. The effort to distinguish these two values from the yaw image is limited by the sequence of pitch commands and the lack of recorded onset and calibration data.

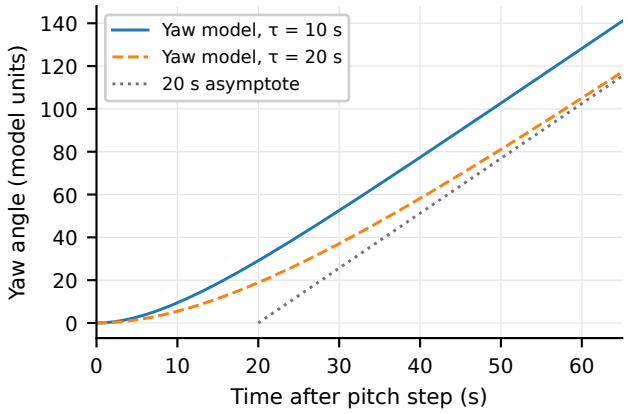

Figure 2.9 – Python. Yaw-angle evolution for the two archived time constants at the same pitch-step amplitude in the source model scale. The dotted line is the 20 s model's late-time asymptote.

2.5 Five-state multivariable representation

Define $x=(\varepsilon,\dot{\varepsilon},\theta,\dot{\theta},v)^T$, $u=(u_g,u_d)^T$ and $y=(\varepsilon,v)^T$. At zero initial state the three channel equations give

$$\ddot{\varepsilon} = -2\xi\omega_0\dot{\varepsilon} - \omega_0^2\varepsilon + K_e\omega_0^2 u_\Sigma$$
$$\ddot{\theta} = K_t u_\Delta \tag{11}$$
$$\dot{v} = -v/\tau_l + (K_l/\tau_l)\theta.$$

$$\dot{x} = Ax + Bu, \qquad y = Cx, \qquad D = 0$$
$$C_{1,:} = (0,1,0,0,0), \quad C_{2,:} = (0,0,0,0,1). \tag{12}$$

The first row of A in the printed realization is (−0.112,−0.49,0,0,0), and the fifth is (0,0,0,−0.5219,−0.1000). The second and fourth rows implement the angle derivatives and the third row is zero. The first and third rows of B are (0.1421,0.1421) and (1.336,−1.336), respectively. The full symbolic and numeric matrices appear in Appendix B. In Section 4, “nominal” refers to these numbers, as required to check the printed Riccati solution. Whenever the nominal yaw coefficient is changed, the lag and pitch entries are changed together to preserve the static rate gain.

For summed and differential actuation the plant separates into a second-order elevation subsystem and a third-order pitch-to-yaw-rate subsystem. In voltage-to-angle coordinates the latter would acquire one more integrator from yaw kinematics. The rate state is retained because it is the regulated output. At a constant command $v_ref \neq 0$, the model has a stationary rate and a drifting yaw angle. Its full

arm configuration consequently need not remain near the original identification point over an arbitrary duration.

3. Sampled elevation servo

3.1 Refined identification and coordinate order

The second campaign concerns the elevation loop in isolation. Its input and response are shown in Figure 3.1. The command rises to a roughly constant level and the output shows a decaying oscillatory component. The source reports $\xi=0.0912$, $\omega_0=0.8038$ rad/s and $K=0.3299$ rad/V. It prints the following transfer and differential equation.

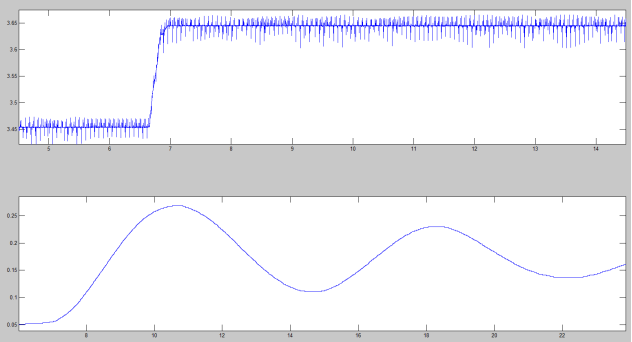

Figure 3.1 – Original multitrace identification record from the sampled elevation campaign. The top panel is the command and the bottom panel is the measured elevation response. Source Fig. 2.1, p. 13 of [1].

$$H_e(s)=\frac{K}{s^2/\omega_0^2+(2\xi/\omega_0)s+1}$$
$$H_e(s)=\frac{0.2131}{s^2+0.1465s+0.6461}. \tag{13}$$

$$\ddot{\varepsilon}+2\xi\omega_0\dot{\varepsilon}+\omega_0^2\varepsilon=K\omega_0^2u$$
$$\dot{x}=Ax+Bu,\quad y=Cx+Du. \tag{14}$$

The text before the symbolic realization in [1] puts the rate first. The matrix with first row (0,1) and the output row (1,0) instead implement elevation first. The scalar command in this section is the command used for that isolated realization. Its relationship to a single physical motor or to the summed input of Section 2 requires the bench calibration. Here $x=(\varepsilon,\dot{\varepsilon})^{\mathrm{T}}$ is used consistently, including for the two entries of each state feedback gain. The use of elevation first in the sampled loop and rate first in the five-state loop is deliberate.

For this state order the symbolic form has $A_{21}=-\omega_0^2$, $A_{22}=-2\xi\omega_0$, $B_2=K\omega_0^2$, and $D=0$. The printed numerical state realization has $A_{21}=-0.6461$, $A_{22}=-0.146$, and $B_2=0.2652$. By direct multiplication the quoted K and ω_0 give $K\omega_0^2=0.21315$. Their damping calculation gives $2\xi\omega_0=0.14661$. Differences at the level of the last reported digits can occur when intermediate parameters are rounded. The numerator difference of 0.0521 is about 24.4% relative to the numerator of the transfer and is too large to attribute to rounding of the given digits.

For clarity, the matrix realization and the transfer realization are kept as two different numerical plants, with the printed damping values 0.146 and 0.1465 respectively. The corresponding DC gains are 0.4105 and 0.3298 rad/V. The recorded elevation trace alone gives no independent choice between these realizations because command scaling is unavailable. The printed digital gains, discrete model and pole map give separate algebraic tests.

3.2 Exact zero-order hold and reachability

At $T_s=0.05$ s, the state advances under the held input according to (15). In the last equality, the upper-right block of the matrix exponential produces the input integral even for singular A. This is the expression evaluated numerically.

$$x[k+1]=Fx[k]+Gu[k],\quad y[k]=Cx[k]$$
$$F=e^{AT_s},\qquad G=\int_0^{T_s}e^{A\tau}B\,d\tau. \tag{15}$$

$$\mathcal{F}=\exp(T_s\mathcal{A})$$
$$\mathcal{A}=\begin{bmatrix}A & B\\ 0 & 0\end{bmatrix},\quad \mathcal{F}=\begin{bmatrix}F & G\\ 0 & I\end{bmatrix}. \tag{16}$$

For the printed numeric realization, the first row of F is (0.99919445,0.04980453) and the second is (−0.03217871,0.99192298). The entries of G are 0.0003306503 and 0.0132081621. For the transfer realization, the corresponding G entries are 0.0002656900 and 0.0106132135. Only the first pair rounds to the G printed in [1]. The entries of F remain close for the two plants because the denominator coefficients are close.

The integral state in the sampled calculation accumulates one error per sample according to $\eta[k+1]=\eta[k]+Cx[k]-r[k]$. This sign, used with negative state feedback, causes a positive step demand to produce positive control as the error accumulates. With $x_e=(x^T,\eta)^T$, its dynamics and reference column are

$$x_e[k+1]=F_e x_e[k]+G_e u[k]+E_r r[k]$$
$$F_e=\begin{bmatrix}F & 0\\ C & 1\end{bmatrix},\quad G_e=\begin{bmatrix}G\\ 0\end{bmatrix},\quad E_r=\begin{bmatrix}0\\ -1\end{bmatrix} \tag{17}$$
$$u[k]=-K_p x[k]-K_i\eta[k].$$

For a scaled integral $\eta_t=T_s\eta$, the same feedback requires the multiplier K_i/T_s on η_t. In the unscaled coordinates the integral mode is at z=1 before feedback. The source tests det[G,FG] using rounded numbers. Nonsingularity of the state-transition matrix does not establish reachability. A direct determinant evaluation and a structural check both establish the property here.

Write $W_2=[G,FG]$ and $W_3=[G_e,F_eG_e,F_e^2G_e]$. Cayley–Hamilton gives $F^2-(\mathrm{tr}\,F)F+(\det F)I=0$. In W_3, the bottom entries of the three columns are 0, CG and C(F+I)G. Subtracting (trF) times the second column from the third and adding (det F) times the first leaves the top part of column three equal to zero. Its bottom entry becomes Cadj(I−F)G. Thus

$$\begin{aligned}\det W_3 &= \det W_2\; C\mathrm{adj}(I-F)G\\ &= \det W_2\;\det(I-F)\;C(I-F)^{-1}G.\end{aligned} \tag{18}$$

This is nonzero when the sampled plant is reachable and has nonzero steady gain, provided I−F is nonsingular. The last factor is the discrete DC gain and matches the continuous one under a constant held input. Table 3.1 gives the exact sampled determinants for both plant definitions. The magnitude of an individual determinant depends on the units and state scaling. Its value is used here to reproduce and check the original rank calculation.

Table 3.1 – Sampled reachability at 0.05 s using unrounded zero-order-hold matrices.

Plant	DC gain (rad/V)	det W_2	det W_3
Matrix	0.4105	-8.724×10^{-6}	-5.762×10^{-9}
Transfer	0.3298	-5.633×10^{-6}	-2.989×10^{-9}

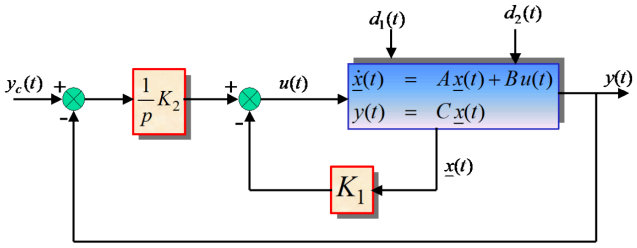

Figure 3.2 – Original sampled integral servo architecture. It identifies reference, state regulator, and possible plant-input and measured-output disturbance locations. Source Fig. 2.2, p. 14 of [1].

3.3 Pole assignment and command generation

The source selects pole-placement horizons at −1 and $-2\ \mathrm{s}^{-1}$, denoted T_c=1 and 0.5 s. The sampling guideline printed in [1] is retained in (19). With the faster real-part time constant of 0.5 s, T_s=0.05 s lies at its lower end. The same period is kept for both designs to isolate the effect of retuning.

$$\frac{\tau_{\min}}{10}\le T_s\le\frac{\tau_{\min}}{5},\qquad z_i=\exp(s_iT_s). \tag{19}$$

Let $\Phi(z)=\prod_i(z-z_i)$ be the monic third-order design polynomial. The single-input feedback row can be obtained by the Ackermann formula [2]. Reaching the poles presupposes the augmented reachability established by (18). The placement controls the polynomial of the augmented plant under the same implementation of the integral update that was used for synthesis.

$$K_e^{\mathrm{fb}} = e_3^T W_3^{-1} \Phi(F_e)$$
$$\det(zI - F_e + G_e K_e^{\mathrm{fb}}) = \Phi(z). \tag{20}$$

The source reports the nominal pole set $\{0.9512, 0.9505 \pm 0.0381i\}$ and the gains recorded in Table 3.2. The real parts of the target continuous poles are -1 or $-2\ \mathrm{s}^{-1}$, and a complex pair near ± 0.8 rad/s is consistent with the printed nominal set. Designing against this explicit target with the matrix realization gives gains close to the rounded rows from [1]. The comparison in Table 3.2 evaluates the printed gains unchanged on both plants, using $s_{eq}=\log(z)/T_s$ on the principal branch. This is a representation of the computed discrete poles, with oscillation frequencies in the unaliased baseband.

Table 3.2 – Printed feedback gains and principal continuous equivalents of the poles they place on each plant. Gains act on elevation, elevation rate and the unscaled error sum, in that order.

Plant and tuning	Printed gain row	Equivalent poles (s^{-1})
Matrix, nominal	(10.6434, 10.2527, 0.2882)	$-1.0000 \pm 0.8011i$ -0.9996
Matrix, faster	(40.3459, 20.0418, 1.5137)	$-2.0015 \pm 0.8014i$ -1.9969
Transfer, nominal	(10.6434, 10.2527, 0.2882)	$-0.8427 \pm 1.0451i$ -0.7210
Transfer, faster	(40.3459, 20.0418, 1.5137)	$-1.7033 \pm 1.6298i$ -1.3005

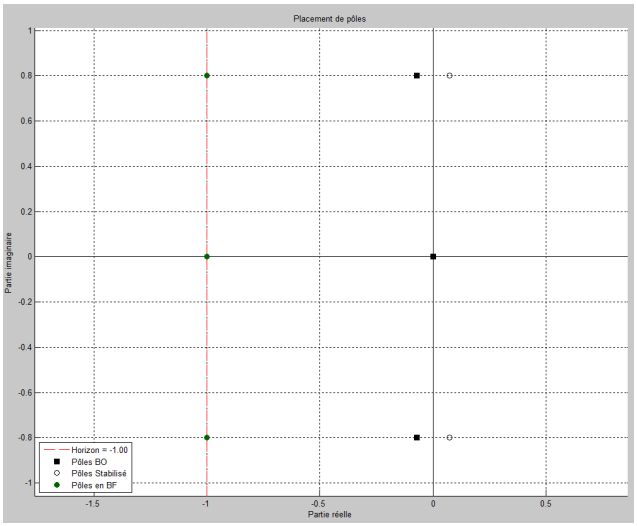

Figure 3.3 – Original open-loop and nominal assigned pole plot. The integral eigenvalue before feedback is at unity. Source Fig. 2.3(a), p. 15 of [1].

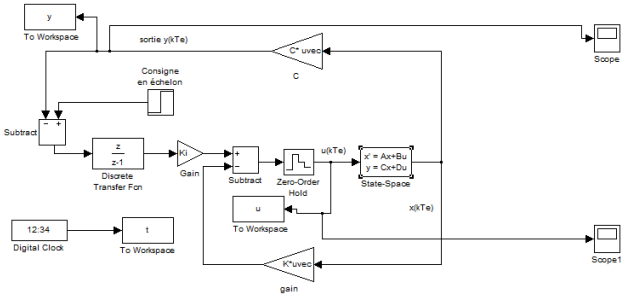

Figure 3.4 – Original Simulink model of the nominal discrete integral loop. Source Fig. 2.3(b), p. 15 of [1].

The numerical effect of the transfer realization in this test is a slower dominant modal rate with the same feedback. There are three discrete modes. Their damping, residues and interaction determine the time at which the response enters a specified band. A near equality of two measured settling times therefore does not imply an equality of the pole locations. A steady voltage test would be especially useful because its response is directly sensitive to the numerator difference.

3.4 Nominal and faster responses in the original figures

The nominal archived response is smooth and approaches roughly 0.79 rad. Its origin is near a delayed input step. The voltage trace rises to an early lobe near 1.2 V, falls, and then approaches approximately 1.9 V. The comparison trace in Figure 3.9 displays the same structure for the nominal design. The two voltage numbers in the original discussion refer to different portions of the transient. The 1.2 V reading describes the first local maximum. The final level near 1.9 V sets the maximum of the displayed nominal trace. The faster voltage reaches approximately 3.3 V early and approaches the same steady level. The output figure for the faster case shows an annotated point near 0.76 rad at 3 s.

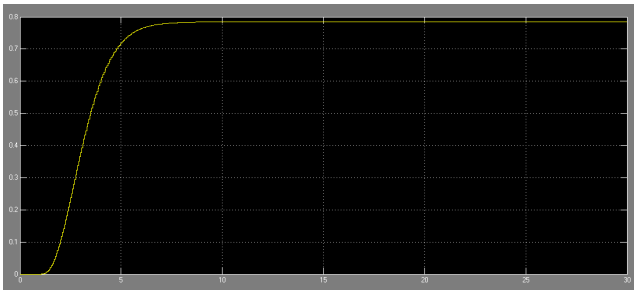

Figure 3.5 – Original nominal elevation simulation. The curve approaches a constant level near 0.79 rad. Source Fig. 2.4(a), p. 15 of [1].

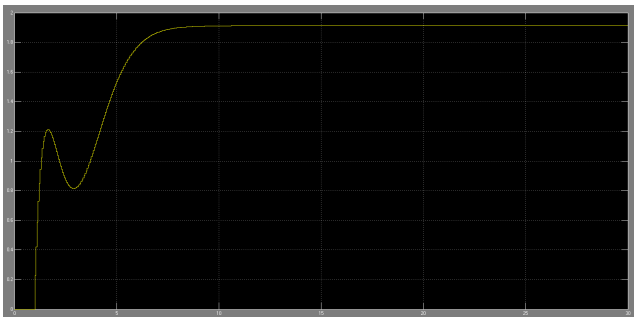

Figure 3.6 – Original nominal command. The first lobe and the later approximately 1.9 V plateau are distinct features of the same trace. Source Fig. 2.4(b), p. 15 of [1].

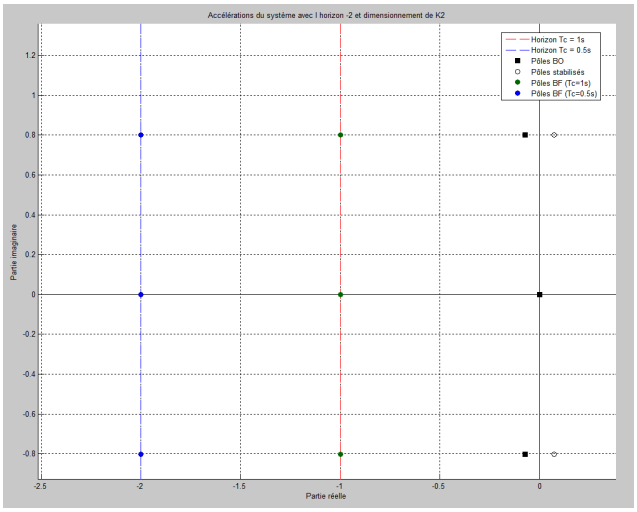

Figure 3.7 – Original pole migration for the nominal and accelerated designs. The faster placement has smaller discrete pole moduli. Source Fig. 2.5(a), p. 16 of [1].

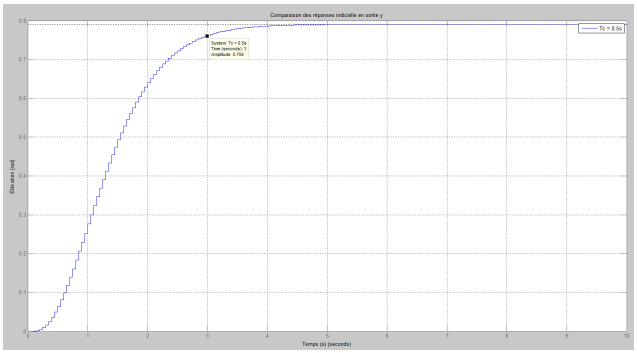

Figure 3.8 – Original faster elevation response and time annotation. The source text reports a response time of about 3 s for this design. Source Fig. 2.5(b), p. 16 of [1].

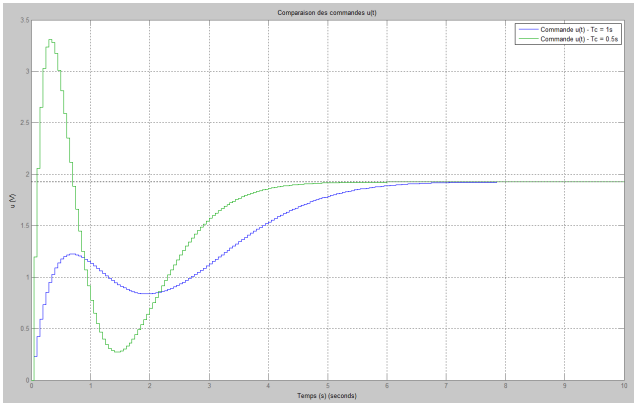

Figure 3.9 – Original comparison of the two held commands. The faster design has a pronounced early voltage demand. Both commands converge to approximately 1.9 V. Source Fig. 2.6, p. 16 of [1].

The reference scale is relevant for interpreting these absolute voltages. Taking the plateau as approximately π/4 rad is a reading proxy from the

original axes. At that proxy, the printed matrix predicts a 1.91 V steady command, from $u_\infty = r/H_e(0)$. Rescaling the computed traces gives a transient lobe of about 1.2 V for the nominal gain and an early peak of about 3.29 V for the faster gain. The transfer realization would instead need about 2.38 V at the same plateau. This check supports the use of the printed matrix for the archived simulation and is independent of the sampled pole comparison. The absolute voltage mapping on the physical bench still depends on the actuation and calibration records.

3.5 Python comparison at fixed amplitude

The additional test applies a 0.10 rad reference at 1 s with zero initial deviations and no actuator clamp. It uses the printed gains with the same sample period and update order for the two plant definitions. Settling is defined as the elapsed time after the step until every subsequent sampled error in the study window is within ±5%, ±2% or ±1% of the reference amplitude. Voltages in Table 3.3 are values of the scalar command in the model. The final command follows directly from $y(\infty) = H_e(0)u(\infty)$.

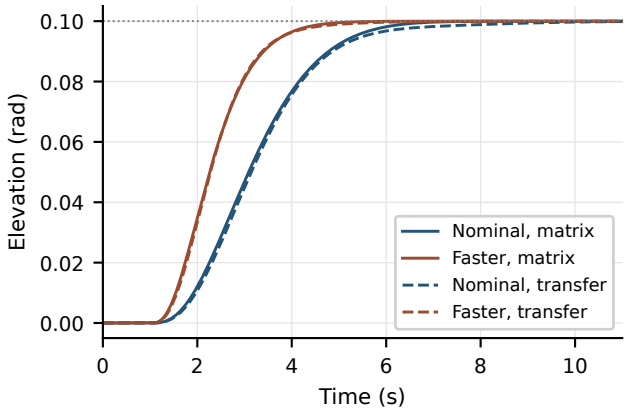

Figure 3.10 – Python. Elevation rise and convergence with fixed gains and the two plant realizations. The target is 0.10 rad.

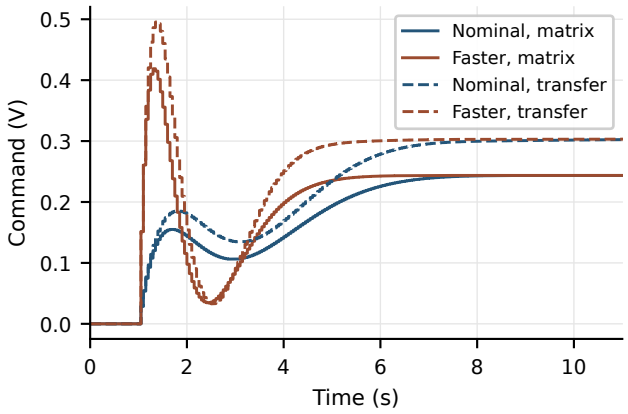

Figure 3.11 – Python. Held command for the four cases of Figure 3.10. The higher final command on the transfer realization comes from its smaller DC gain.

Table 3.3 – Python sampled step test. Time is measured from the 0.10 rad reference edge. The command values are the largest absolute sample and the final sample.

Realization	Tuning	5% (s)	2% (s)	1% (s)	Peak (V)	Final (V)
Matrix	Nominal	4.35	4.95	5.40	0.244	0.244
Matrix	Faster	2.85	3.30	3.65	0.419	0.244
Transfer	Nominal	4.55	5.75	7.25	0.303	0.303
Transfer	Faster	2.85	3.45	4.05	0.495	0.303

For the matrix realization the faster setting reduces the 5% time from 4.35 s to 2.85 s. The peak increases from 0.244 V to 0.419 V for the chosen small step. The squared-command integral from 1 to 11 s increases from 0.409 to 0.517 $V^2 \cdot s$. These values refer to a command-energy proxy. They do not require assumptions about motor current or winding resistance. The transfer realization needs 0.303 V to hold the same reference, compared with 0.244 V for the printed matrix. Under fixed gains, the response can be similar near the 5% threshold despite a larger difference by the 1% threshold. The four responses are stable on the sampled models.

The reconstructed 5% value is defined here from the applied step at 1 s. Source [1] quotes a nominal simulation time of about 6–6.2 s and a bench time of 5.78 s. The exact band, onset and averaging used for those times are not in the report. A different time

origin or a tighter band changes those values. Table 3.3 supplies the explicit timing definition for the reconstruction and retains the archived times as reported observations in Section 6.

3.6 Real-time implementation

Source [1] credits MATLAB/Simulink and a dSPACE DS1103/1104 rapid-prototyping interface for deployment of the elevation controller. The supplied diagrams show the separation of reference generation, feedback, conditioning and the external plant. The internal command block implements one integrated elevation error. The figures establish the described structure. They do not establish the actual interface model, derivative filtering, real-time sample jitter, absolute equilibrium bias or motor voltage limits.

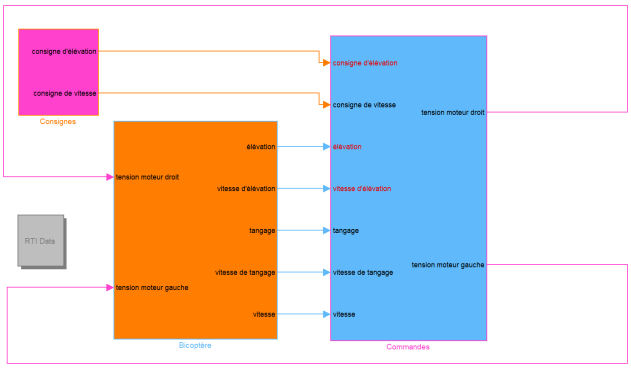

Figure 3.12 – Original top-level dSPACE implementation diagram for the sampled controller. Source Fig. 2.7(a), p. 17 of [1].

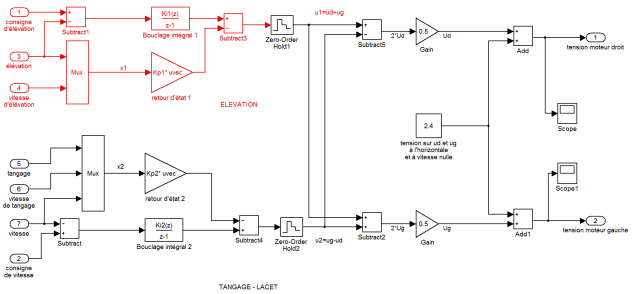

Figure 3.13 – Original internal command model with integral state feedback. Source Fig. 2.7(b), p. 17 of [1].

The source reports 5.78 s for the bench response time and an average of 4.41 s for recovery after manual disturbances. These values are retained with their reported precision. The source gives neither the number of repetitions nor individual disturbance profiles, so a variance or confidence interval cannot be calculated. Time synchronisation, angle differentiators and the command actually issued after any limiting would be required to attribute the difference between the simulated and reported bench response times. No new deployment or experimental run is claimed here.

4. Multivariable integral linear-quadratic control

4.1 Augmented reachability and integral tracking

The full control objective is to regulate $y=(\varepsilon,v)^{T}$. Two integrated errors are appended to the five original states. Let $r=(r_{\varepsilon},r_{v})^{T}$, $\dot{\eta}=Cx-r$ and $x_{a}=(x^{T},\eta^{T})^{T}$. The input matrices have compatible zero blocks for n=5 states, m=2 motor commands and p=2 regulated outputs. The reference enters through the error states. An additive input disturbance d can enter the plant through the same motor channels.

$$\dot{x}_{a}=A_{a}x_{a}+B_{a}(u+d)+E_{r}r$$
$$A_{a}=\begin{bmatrix}A & 0_{n\times p}\\ C & 0_{p\times p}\end{bmatrix},$$
$$B_{a}=\begin{bmatrix}B\\ 0_{p\times m}\end{bmatrix},\qquad(21)$$
$$E_{r}=\begin{bmatrix}0_{n\times p}\\ -I_{p}\end{bmatrix}.$$

The augmented pair has reachability rank seven for the numerical plant, using the columns of $[B_{a},A_{a}B_{a},\dots,A_{a}^{6}B_{a}]$. The absence of an extra uncontrollable mode at the origin can also be checked from the plant equations. For Cx=0, both

elevation and yaw rate vanish. If Ax+Bu=0, the kinematic rows give $\dot{\varepsilon}=\dot{\theta}=0$. The yaw row forces $\theta=0$ because $K_l/\tau_l\neq 0$. The elevation and pitch rows then give $u_\Sigma=u_\Delta=0$. Thus the augmented controllability condition at zero, equivalently nonsingularity of the square system matrix with blocks A,B;C,0, holds. At nonzero eigenvalues reachability follows from the uncoupled pairs. The elevation subsystem is driven in its acceleration state. The differential channel drives yaw rate through pitch, forming a reachable chain when $K_t K_l\neq 0$.

At a stable equilibrium with constant r and feasible u, the stationary integral states imply Cx=r. This is the origin of the zero-offset property. The implication assumes that the required input can be applied and that the integrators remain active. A constant sensor offset would change the physical equilibrium even if the measured error vanishes.

4.2 Quadratic criterion and spectral shift

The source uses the cost and weights

$$J=\int_0^\infty (x_a^T Q x_a + u^T R u)\,dt \qquad Q=\mathrm{diag}(0,0,0,0,0,40,40),\quad R=I_2. \tag{22}$$

The positive weights penalise the two accumulated errors and the sum of squared motor voltages. The chosen units of the error states matter. An elevation error and a yaw-rate error contribute in different physical units. The numbers 40 and 40 are the reported design weights in their original coordinates. The cost includes neither motor-current measurement nor a hard bound on voltage. Multiplying Q and R by the same positive factor scales P and leaves the feedback row unchanged.

The design shifts A_a by $\alpha_c I_7$, with $\alpha_c=0.5\ \mathrm{s}^{-1}$. Let $A_\alpha=A_a+\alpha_c I_7$. For regulation about an equilibrium, the change of variables $z=e^{\alpha_c t}x_a$, $w=e^{\alpha_c t}u$ transforms the unforced state equation into $\dot{z}=A_\alpha z+B_a w$. One can therefore use the ordinary infinite-horizon regulator [2, 3] in (z,w). In the physical coordinates the integrand of (22) then has the weight $e^{2\alpha_c t}$.

$$A_\alpha = A_a + \alpha_c I_7 \qquad A_\alpha^T P + P A_\alpha - P B_a R^{-1} B_a^T P + Q = 0 \qquad K = R^{-1} B_a^T P = [K_p\ K_i],\quad u=-K x_a. \tag{23}$$

The archived continuous Riccati equation is written with A_a in its linear terms. The displayed numerical P satisfies the version using A_α. Equivalently, using A_a in (23) requires adding $2\alpha_c$ P to the left-hand side. Although Q is semidefinite, the two integral coordinates observe all eigenvectors of the augmented plant in the model. Indeed, an unobserved eigenvector has $\eta=0$, so the lower block of the eigenvalue equation gives Cx=0. The elevation kinematic row forces the elevation rate to vanish. The yaw row forces pitch to vanish, and the pitch kinematic row then forces its rate to vanish. The same eigenvectors apply after a scalar shift. This verifies the detectability hypothesis associated with the stabilising solution [3]. The source matrix P and the two printed gain matrices are reproduced in Appendix B.

If $A_c l=A_a-B_a K$, substitution in the Riccati equation gives the spectral relation and Lyapunov identity (24).

$$\mathrm{spec}(A_\alpha - B_a K) = \mathrm{spec}(A_{\mathrm{cl}}) + \alpha_c \qquad A_{\mathrm{cl}}^T P + P A_{\mathrm{cl}} = -Q - K^T R K - 2\alpha_c P. \tag{24}$$

For zero reference perturbation and no disturbance in the nominal model, the last equation gives $\dot{V}\le -2\alpha_c V$ with $V=x_a^{\mathrm{T}} P x_a$. This establishes a nominal state-decay bound in the quadratic norm. The equality is written for the physical closed-loop matrix. The pole calculation for the synthesis frame concerns the shifted matrix, which explains why the printed real parts are greater by $0.5\ \mathrm{s}^{-1}$.

4.3 Numerical audit of the reported gain and poles

Substitution of the rounded P printed in [1] into the shifted and unshifted equations gives residual Frobenius norms of 5.25×10^{-3} and 6.21×10^{2}, respectively. Rounded entries do not solve the shifted equation to machine accuracy. A new solution of that equation has a residual below 2×10^{-11}. Its gain differs from the printed gain by at

most 1.16×10^{-4} per entry. The same numbers also distinguish the frames of the reported poles. The distance from each printed pole to its nearest shifted calculated pole is at most 4.1×10^{-5} s^{-1}. The printed poles are approximately 0.5 s^{-1} to the right of the corresponding physical ones.

Table 4.1 – Paired poles of the two continuous matrices at the shifted solution on the 10 s plant.

Physical (s^{-1})	Synthesis frame (s^{-1})
−2.0259 ± 0.5786i	−1.5259 ± 0.5786i
−1.5098	−1.0098
−1.1338 ± 1.3929i	−0.6338 ± 1.3929i
−1.0607 ± 0.9945i	−0.5607 ± 0.9945i

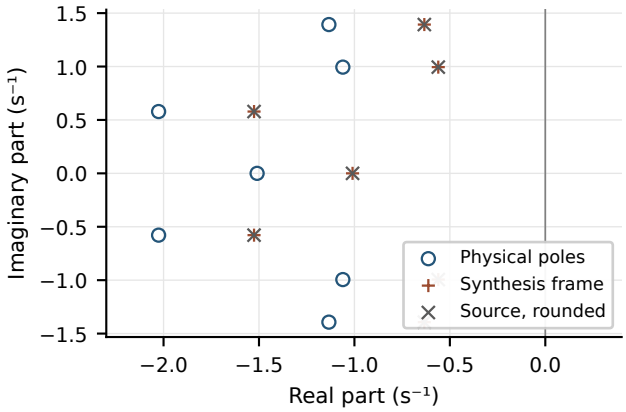

Figure 4.1 – Python. Physical poles of the closed loop, synthesis-frame poles of the shifted regulator and rounded eigenvalues printed in [1]. Coincident crosses identify the frame of the printed values.

At nominal yaw lag, the physical spectral abscissa is about −1.061 s^{-1} for the shifted design. Solving the CARE without the shift, with the same entries in Q and R, gives a spectral abscissa of −0.469 s^{-1} on the physical plant. Positive shift in the synthesis increases the penalty on slow error decay. The resulting tracking and voltage changes are measured on the same reference in Section 5.

4.4 Comparison with the archived nominal response

Figure 4.2 contains the original four simulated histories. Its marker near 1.002 s indicates the start of the changes in both outputs. The elevation trace reaches 0.497253 rad at 4.091 s and reads 0.52354 rad at 10 s. The yaw-rate trace reaches 0.950518 rad/s at 3.713 s and reads 1.00007 rad/s at 10 s. The two intermediate values are approximately 95% of the corresponding final displayed values. This gives elapsed times of 3.089 and 2.711 s read between markers. The source reports rounded 5% response times of 3.1 and 2.7 s.

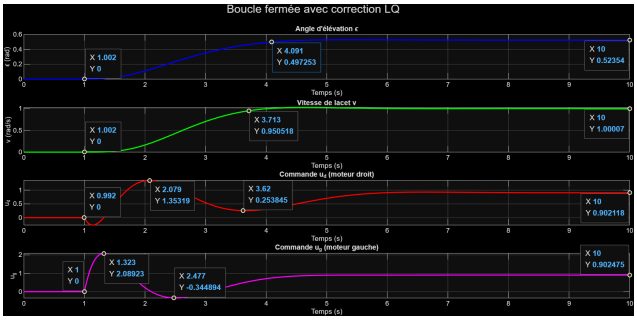

Figure 4.2 – Original nominal LQI simulation. Elevation, yaw rate, right-motor command and left-motor command are displayed in that order. The original markers remain in the screenshot. Source Fig. 2.8, p. 18 of [1].

The independent reconstruction uses the displayed endpoint amplitudes 0.52354 rad and 1.00007 rad/s as reference proxies and a 1.0 s step time. It evaluates the shifted gain on the physical 10 s plant. The calculated 5% settling times are 3.100 and 2.715 s. The largest positive commands over the first 10 s are 1.351 and 2.085 V, close to the 1.353 and 2.089 labels in Figure 4.2. At 10 s, the calculation gives 0.9020 and 0.9024 V, while the figure shows 0.902118 and 0.902475. In the reconstruction the larger transient belongs to the second command, labelled right under the algebraic convention of (1). The figure assigns it to the left. The output timings and the unordered command magnitudes support use of the 10 s matrix and shifted CARE in the archived simulation. The signal-routing and sign conventions would be needed to resolve the exchange of command labels.

The source describes the visible response as non-overshooting. In the reconstructed model the continuous sample peaks exceed the final levels by approximately 1.2% in elevation and 2.7% in yaw rate. They are smaller than the 5% band. The discrepancy matters when reporting peak excursions and has little effect on the printed settling times. The figure does not document the configured voltage limits or whether the display is before or

after an actuator limiter. Agreement with these plot markers is evidence about the archived simulation. The dSPACE observations in [1] support bench stability qualitatively.

4.5 Comparison with the archived disturbed response

The source describes simultaneous two-second rectangular disturbances in the symmetric and differential channels. Figure 4.3 preserves the plotted outputs and commands. Around the 10 s marker, the elevation baseline is 0.523538 rad and a later marker is 0.893975 rad at 11.984 s. The elevation reading at 14.952 s is 0.530071 rad. The yaw-rate trace reads 1.00006 rad/s near 10.019 s, dips to 0.747798 rad/s at 11.286 s and reads 0.995749 rad/s at 15.01 s. The plots show recovery of both outputs within roughly 5 s from the marked disturbance onset.

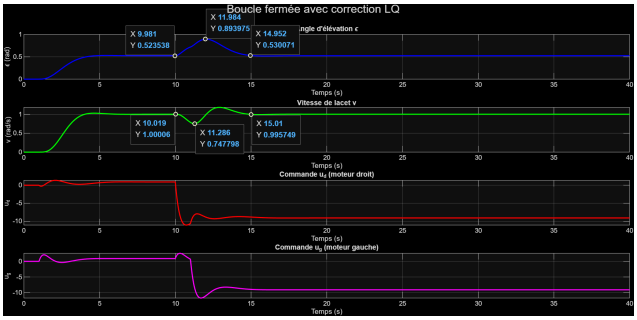

Figure 4.3 – Original disturbed LQI simulation. The outputs return close to the initial targets, while the plotted motor commands settle at levels different from their pre-event values. Source Fig. 2.9, p. 19 of [1].

Both command traces eventually remain at markedly changed levels after the displayed output recovery. In the nominal linear model with constant references, a finite motor-input pulse that returns to zero must lead back to the original equilibrium command. The numerical relation follows from stability and uniqueness of the closed-loop equilibrium for fixed r and d. The screenshot and the prose leave the disturbance definition, signal logging points and command offsets ambiguous. The figure supports a comparison of output marker times. It cannot fix the time from pulse removal or verify that the plotted disturbance had returned to zero. Section 5.4 defines and evaluates a separate two-second pulse with an explicit return to zero.

5. Model sensitivity and controlled simulations

5.1 Decoupling and actuator demand

The sign symmetry of the printed gain is visible in the two rows of K_p and K_i in Appendix B. The coefficients on the elevation states and integral agree between motors. The coefficients on pitch, yaw rate and the yaw integral reverse sign. Adding the two feedback commands eliminates the yaw terms. Subtracting them eliminates the elevation terms. The elevation coefficients in (25) multiply the angular rate, the angle and its integrated error. The differential coefficients multiply the pitch rate, the pitch angle, the yaw rate and its integrated error.

$$\begin{aligned} u_\Sigma &= -2(k_{\dot\varepsilon}\dot\varepsilon + k_\varepsilon\varepsilon + k_{\eta_e}\eta_e) \\ u_\Delta &= -2(k_{\dot\theta}\dot\theta + k_\theta\theta + k_v v + k_{\eta_v}\eta_v). \end{aligned} \tag{25}$$

The equal motor weights in (22) do not create a cross-term in these inputs. Expansion of the inverse transform gives the exact identities

$$\begin{aligned} u_g^2 + u_d^2 &= \frac{u_\Sigma^2 + u_\Delta^2}{2} \\ \max(|u_g|, |u_d|) &= \frac{|u_\Sigma| + |u_\Delta|}{2}. \end{aligned} \tag{26}$$

For the second identity, one of $u_\Sigma+u_\Delta$ and $u_\Sigma-u_\Delta$ has magnitude $|u_\Sigma|+|u_\Delta|$, depending on the signs. Thus two equal symmetric absolute limits $|u_g|,|u_d|\le U_{max}$ would impose the shared constraint $|u_\Sigma|+|u_\Delta|\le 2U_{max}$. This property persists even though the unconstrained linear model decouples. A fixed equilibrium bias or unequal motor limits changes the feasible set in the deviation coordinates. The archival plots do not provide those bounds and the

additional simulations use an unconstrained command.

At a steady reference with no additive disturbance, the pitch dynamics imply $u_{\Delta}=0$. The elevation equation gives $u_{\Sigma}=\omega_0^2\, r_{\varepsilon}/b_e$, where $b_e=0.1421$ is the coefficient of the sum in elevation acceleration. The yaw equation requires $\theta=r_v/K_l$. For the small step used below, the resulting steady motor commands are 0.1724 V each and the required pitch is about −0.0192 rad. These are model deviations about the equilibrium. The omitted yaw position continues to change.

5.2 Continuous proof for the yaw-lag interval

The prose-based and matrix-based yaw lags in [1] define a natural sensitivity interval. Here the exactly rounded printed gain is fixed while $\tau=\tau_l$ varies continuously from 10 to 20 s. The signed static yaw gain K_l, other plant parameters and the voltage convention are held fixed. By (25), elevation and yaw dynamics occupy diagonal blocks after permuting the augmented state. The overall characteristic polynomial is the product of the elevation cubic and the yaw quartic, with their respective leading constants.

Let $a=2K_tk_d$ and $b=2K_tk_\theta$. For differential motion, $v=K_l\theta/(1+\tau\, s)$ and $s\eta_v=v$ at zero reference. The pitch equation contributes $s^2\theta=K_t\, u_{\Delta}$. Eliminating η_v and v gives the yaw polynomial

$$p_y(s)=(1+\tau s)s(s^2+as+b) \\ +2K_tK_l(k_vs+k_{\eta_v}). \tag{27}$$

For $p_y=a_4s^4+a_3s^3+a_2s^2+a_1s+a_0$ the ordered coefficients are

$$a_4=\tau,\quad a_3=1+a\tau,\quad a_2=a+b\tau \\ a_1=b+2K_tK_lk_v,\quad a_0=2K_tK_lk_{\eta_v}. \tag{28}$$

They are all positive throughout the interval. Positivity of the first column of the Routh array [2] reduces to the two additional inequalities

$$\Delta_2=a_3a_2-a_4a_1>0 \\ \Delta_3=a_1\Delta_2-a_3^2a_0>0. \tag{29}$$

Both $\Delta_2(\tau)$ and $\Delta_3(\tau)$ are quadratic in τ. Evaluating them at the interval endpoints and at any internal stationary point gives a minimum of about 8.336×10^3 and 1.357×10^6, respectively, at 10 s. Table 5.1 includes the coefficients and endpoint values. The elevation block has polynomial $s^3+3.63125s^2+5.31705s+3.19191$, whose cubic Routh difference is approximately 16.116. These facts certify asymptotic stability of the fixed printed gain for every yaw lag in the interval in this linear plant family.

$$p_e(s)=s^3+(2\xi\omega_0+2b_ek_{\dot\varepsilon})s^2 \\ +(\omega_0^2+2b_ek_{\varepsilon})s+2b_ek_{\eta_e}. \tag{30}$$

Table 5.1 – Routh determinants for the yaw polynomial with the printed gain. Coefficients are shown in ascending powers of τ. Minima are evaluated on 10–20 s.

Determinant	Constant	τ coefficient	τ² coefficient	At 10 s	At 20 s
Δ_2	6.219	−176.445	100.942	8.336×10^3	3.685×10^4
Δ_3	1295.692	−42602.536	17814.856	1.357×10^6	6.275×10^6

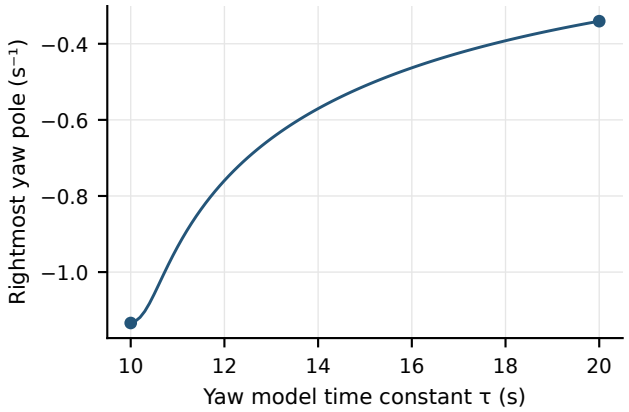

Figure 5.1 – Python. Rightmost yaw-block pole for 101 time constants using the rounded source gain. The Routh inequalities establish stability on the continuous interval.

The rightmost yaw pole moves from approximately $-1.134\ s^{-1}$ at 10 s to $-0.341\ s^{-1}$ at 20 s. Figure 5.1 shows the sensitivity of the decay rate. The positive Routh determinants guarantee stability over the interval. They give no fixed settling time and do not cover unknown dynamics such as cross-axis terms, sample delay, estimator dynamics or saturation.

5.3 Step tracking and the effort required by the shift

Three cases use the same initial zero deviations and references of 0.10 rad in elevation and 0.10 rad/s in yaw rate, applied at 1 s. The first uses a new unshifted CARE with τ=10 s. The second uses the reconstructed shifted CARE at the same lag. The third applies the same shifted gain to the 20 s plant. It tests model mismatch without gain retuning. The precise integration and time definitions are in Appendix C.

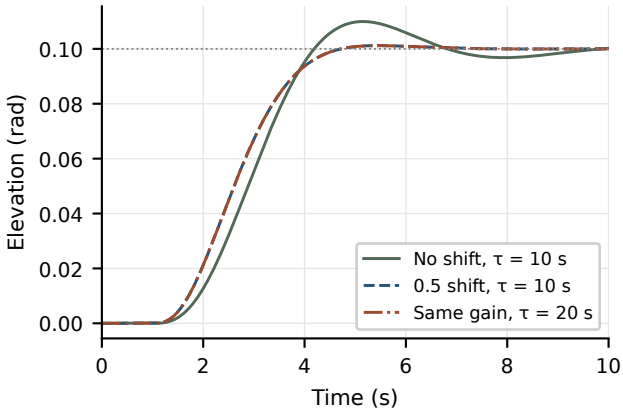

Figure 5.2 – Python. Elevation for the two Riccati designs and a fixed-gain yaw-lag change. The two shifted traces coincide because the model blocks decouple.

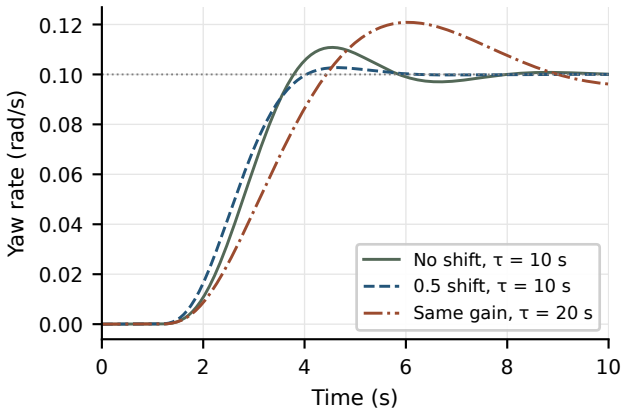

Figure 5.3 – Python. Yaw-rate tracking with the same reference and three plant–controller pairs. The 20 s case has a longer oscillatory tail with the fixed shifted gain.

Table 5.2 – Continuous reference test. The 5% times and overshoots are for elevation and yaw rate. The last column is peak absolute command on either motor.

Design and lag	t_5%, ε (s)	t_5%, v (s)	ε peak (%)	v peak (%)	u peak (V)
α = 0, τ = 10 s	5.12	4.34	10.0	10.8	0.205
α = 0.5, τ = 10 s	3.10	2.72	1.2	2.7	0.306
Same gain, τ = 20 s	3.10	7.30	1.2	20.9	0.309

At 10 s lag, using the shifted design reduces the elevation 5% settling time from 5.12 to 3.10 s and the yaw-rate time from 4.34 to 2.72 s. The peak motor demand in this reference test increases from 0.205 to 0.306 V. The squared-voltage integral from the step to 11 s increases from about 0.471 to 0.494 $V^2{\cdot}s$ for both motors together. The first feedback adjustment is consequently sharper, while the eventual equilibrium command is unchanged. In the same comparison, elevation overshoot falls from 10.0 to 1.2% and yaw-rate overshoot from 10.8 to 2.7%.

Replacing the lag in the plant by 20 s at the same shifted gain leaves elevation unchanged in the decoupled linear model. Yaw-rate 5% settling rises to 7.30 s and yaw-rate overshoot to about 20.9%. This test separates stability and performance within the certified time-constant interval. Because the yaw angle is the integral of the rate, even 0.10 rad/s leads to a simulated total yaw change of about 2.54 rad by 28 s. The rate-regulation result concerns the model and the limited trajectory description in its measured states.

5.4 A finite input pulse with known removal time

A second test uses the same references and a simultaneous rectangular voltage injection in both modal input channels from 12 to 14 s. Its specified symmetric and differential amplitudes are +0.40 and +0.04 V, corresponding to +0.22 and +0.18 V at the two motor inputs by (1). This is an additive plant-input test. Model output error and measured sensor error would excite the dynamics differently. After 14 s the added voltages are set to zero exactly. The disturbance amplitudes are analysis choices and are independent of the archived perturbation.

For each design the plotted pulse response is the perturbed trace minus the trace computed with identical references and no pulse. In the linear model this isolates the effect of the input perturbation. With δx_a the difference and $\delta u = -K\delta x_a$ the resulting command correction,

$$\begin{aligned} \delta\dot{x}_a &= (A_a - B_a K)\delta x_a + B_a d \\ \delta y &= C\,\delta x, \qquad \delta u = -K\delta x_a. \end{aligned} \tag{31}$$

Following removal of the pulse, a stable model returns to the same outputs, states and commands as the no-pulse run. A permanent additive offset would still allow y=r in feasible steady state, but would change the controller's equilibrium command in compensation. This distinction explains the importance of the persistent command change in Figure 4.3.

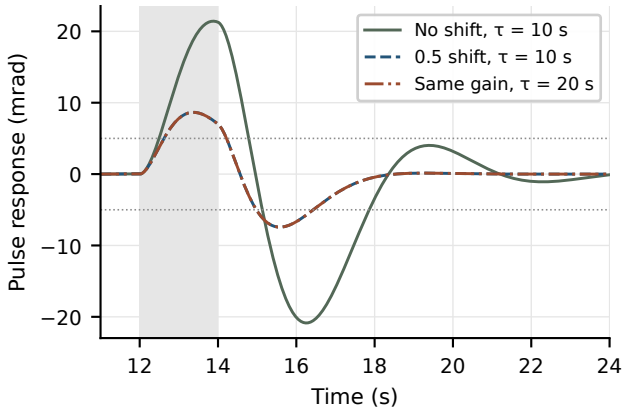

Figure 5.4 – Python. Isolated elevation error due to the input pulse. Shading gives the 12–14 s injection window, and the dotted lines show ±5 mrad.

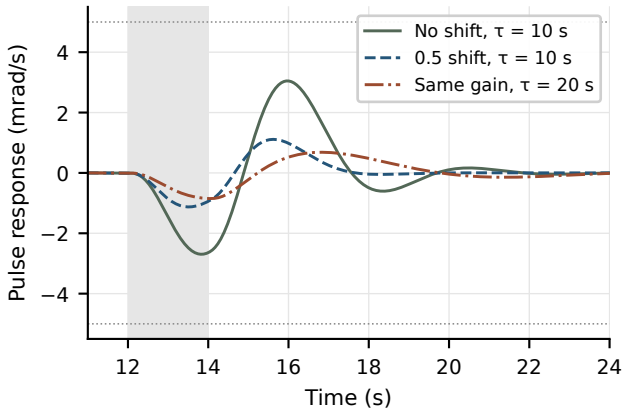

Figure 5.5 – Python. Isolated yaw-rate error due to the same injection. The yaw-rate errors remain inside the specified ±5 mrad/s band in these cases.

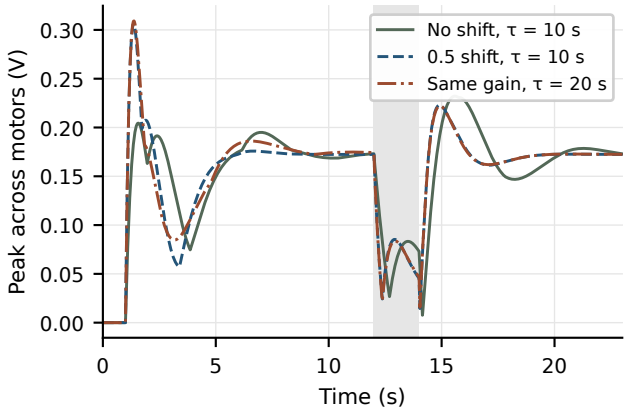

Figure 5.6 – Python. Larger absolute controller command across the motors for each case, including its response to the pulse. The additive disturbance itself is excluded from this command measure.

Table 5.3 – Finite-input test. Peak deviations are absolute. Recovery is measured from removal using ±5 mrad and ±5 mrad/s bands. The command peak is taken at or after pulse onset.

Design and lag	Peak ε (mrad)	Peak v (mrad/s)	Return ε (s)	Return v (s)	u peak (V)
α = 0, τ = 10 s	21.43	3.05	3.87	0.00	0.232
α = 0.5, τ = 10 s	8.66	1.13	2.44	0.00	0.223
Same gain, τ = 20 s	8.66	0.85	2.44	0.00	0.222

The unshifted controller has a peak elevation deviation of 21.4 mrad in this test. The shifted case gives 8.66 mrad. After pulse removal, return to a fixed ±5 mrad band takes 3.87 s and 2.44 s respectively. The yaw-rate error is always below 5 mrad/s for the selected amplitude, so the band metric is zero even though the yaw response is nonzero. Its peak is reported separately. The difference between the 10 and 20 s simulated yaw peaks depends on the input location, the pulse duration and the gains. It does not imply that the slower plant has better yaw-rate tracking, as seen in Figure 5.3.

The controller peak around the injection and the reference-induced peak describe different portions of the test. Under a common motor bound the relevant quantity is the required combined actuator voltage in the actual biased plant coordinates. The linear exercise instead records the model's controller command. It gives a consistent comparison of retuning under an identical input disturbance.

6. Comparison and conclusions

6.1 What the two campaigns establish

The sampled elevation design gives an explicit target for each closed-loop pole. The altered gains illustrate the speed–voltage trade-off on the same local mode. Modeling the plant with the printed numerical matrix reproduces the reported sampling matrix, pole locations and command scales. Applying the same gains to the transfer realization preserves stability in the two tested tunings and shifts the transient modes. The voltage required at steady state follows the transfer gain and is a direct way to resolve the two realizations with future calibrated records. The nominal bench time and manual recovery time are existing reported observations from [1].

The continuous design coordinates elevation and yaw-rate regulation from a single two-motor state law. Its performance depends on the integral coordinates and the weighting of demand and slow error. The MIMO formulation handles both outputs in the synthesis, and the ideal model admits exact separation under the sum and difference transformation. Shared command limits and physical interaction can remove this separation. The positive shift in the CARE accelerates the two reconstructed responses at the expense of higher transient actuator demand. The reported simulation pole set belongs to the shifted synthesis matrix. The physical poles lie 0.5 s^{-1} farther left and yield the nominal response times seen in the original LQI plot.

Table 6.1 – Comparison of the two archived controller campaigns. Timings and conditions are those reported in [1] and the original figures.

Criterion	Sampled integral pole placement	Continuous integral LQR
Scope	Elevation, one command channel	Elevation and yaw rate, two motor commands
Design inputs	Target discrete poles, T_s and T_c	Q, R and spectral shift α_c
Steady offsets	One accumulated tracking error	Two accumulated tracking errors
Nominal simulation	About 6–6.2 s, about 3 s with faster tuning	About 3.1 s for elevation and 2.7 s for yaw rate
Command	Nominal plateau about 1.9 V. Faster local peak about 3.3 V.	Quadratic command penalty. Limits not specified in the archive.
Disturbance	Mean bench recovery 4.41 s, as reported	Roughly 5 s from the simulation plot’s marked onset
Design interpretation	Explicit modal assignment and a small sampled loop	Quadratic synthesis across both inputs and outputs
Main dependency	Actuator demand grows with more aggressive gains	Timing depends on identified yaw lag and the weights

The original criteria table is also reproduced in Appendix B. A direct controller ranking would require identical reference signals, available motor voltage and performance definitions. The SISO and MIMO bench observations have different output scopes and their disturbances were not calibrated to the same profile. The same limits apply to comparing the 4.41 s manual-disturbance average and the 5 s recovery shown in a continuous simulation. The additional simulations give internally consistent comparisons within each controller family.

6.2 Limits and measurements that would resolve them

The archived nominal time histories support detailed comparisons against multiple markers. The reported motor labels are exchanged relative to the reconstruction under (1). Logging the input order at the controller, the motor driver outputs and the yaw sign at the sensor would determine which convention is in use. The disturbed LQI history requires the actual injected signal and measurement point to assess the return after removal. Its persistent command plateaux show why outputs alone are insufficient to identify the disturbance experiment.

A calibrated repeat of the yaw identification should retain the positive and negative pitch histories, onset times and yaw-rate output. Fitting the first-order rate equation over an interval with fixed pitch would avoid uncertainty in a graphical ramp intercept. Repeating the refined elevation experiment with a logged sum of motor voltages and a calibrated angular output would give an independent check of the DC gain. Exciting one modal command and recording both outputs would provide a direct estimate of residual coupling. These are specific tests of the model discrepancies observed here.

The original source suggests state estimation, a pitch output, and robust or predictive extensions. An observer or Kalman filter would first need the raw angle signals and information about noise and timing. A pitch-angle target and a yaw-rate target compete for the same differential input in steady state because $v=K_l\theta$. Independent references for these two outputs would therefore require compatible trajectories, revised objectives or an additional actuator. A constrained controller could explicitly use the motor limits once their calibration and equilibrium biases were established. The yaw-rate study also requires an admissible yaw-angle range for tests over extended time.

6.3 Conclusion

The two identification campaigns support a local second-order model for elevation, a pitch double integrator and a first-order pitch-to-yaw-rate model.

The printed values lead to two substantive ambiguities, in the refined elevation numerator and the yaw lag. The digital gains and archived command traces agree with the printed elevation matrix. The nominal multivariable plot and its published gains agree with the 10 s yaw matrix and a shifted CARE. With the rounded continuous gain fixed, the linear model is stable for every lag from 10 to 20 s. At 20 s the same gain gives a conspicuously slower and more oscillatory yaw response. The original plots, equations and implementation diagrams remain accessible in the manuscript and the independent calculation can be repeated from its embedded Python source.

References

[1] D. Kumar, S. Doumbe Moundy and L. Couësme. *Experimental Identification and Modern Control of a Twin-Rotor Helicopter*. École Nationale Supérieure d’Arts et Métiers and Conservatoire National des Arts et Métiers, 2025–2026, 20 pp. Unpublished engineering project report. Complete source embedded in this manuscript.

[2] K. J. Åström and R. M. Murray. *Feedback Systems. An Introduction for Scientists and Engineers*. 2nd ed., Princeton University Press, 2021. ISBN 978-0-691-19398-4.

[3] B. D. O. Anderson and J. B. Moore. *Optimal Control. Linear Quadratic Methods*. Prentice Hall, 1990. ISBN 978-0-13-638560-8.

Appendix A. Original figure concordance

Each entry links to a reproduction of the named source panel in the manuscript. Original axis labels, graphical readouts and interface screenshots are retained. The title logo is included with the intact source PDF. The remaining graphical content is included visibly in the text.

Table A.1 – Source-panel concordance. Page numbers count the cover.

Source identifier	Source p.	Subject	Open panel
Cover	1	Original cover feedback schematic	Figure 1.1
Project at a glance	3	Overview illustration of the platform	Figure 1.2
Fig. 1.1(a)	8	Platform side view	Figure 2.1
Fig. 1.1(b)	8	Elevation motion	Figure 2.2
Fig. 1.1(c)	8	Horizontal equilibrium	Figure 2.3
Fig. 1.2(a)	9	Underdamped identification diagram	Figure 2.4
Fig. 1.2(b)	9	Measured elevation step	Figure 2.5
Fig. 1.3	10	Measured pitch step	Figure 2.7
Fig. 1.4	11	Measured yaw trace and pitch sequence	Figure 2.8
Fig. 2.1	13	Refined elevation step	Figure 3.1
Fig. 2.2	14	Discrete servo architecture	Figure 3.2

Table A.2 – Source-panel concordance. Page numbers count the cover.

Source identifier	Source p.	Subject	Open panel
Fig. 2.3(a)	15	Nominal pole map	Figure 3.3
Fig. 2.3(b)	15	Nominal Simulink model	Figure 3.4
Fig. 2.4(a)	15	Nominal elevation response	Figure 3.5
Fig. 2.4(b)	15	Nominal command	Figure 3.6
Fig. 2.5(a)	16	Nominal and faster pole locations	Figure 3.7
Fig. 2.5(b)	16	Faster elevation response	Figure 3.8
Fig. 2.6	16	Comparative command histories	Figure 3.9
Fig. 2.7(a)	17	Top-level dSPACE interface	Figure 3.12
Fig. 2.7(b)	17	Internal servo implementation	Figure 3.13
Fig. 2.8	18	Nominal multivariable LQI histories	Figure 4.2
Fig. 2.9	19	Perturbed multivariable LQI histories	Figure 4.3

Appendix B. Original equations and comparison table

These excerpts preserve the printed equations, matrices, gains and original table. The source's notations and numerical discrepancies remain visible here. The main text gives the coordinates and numerical values used in the reconstructions. The complete source PDF is also embedded, including inline notation and the original section layout.

Source p. 8 · Input combination and elevation channel

$$u_\Sigma = u_g + u_d, \qquad u_\Delta = u_g - u_d.$$

$$\boxed{H_e(s) = \frac{\epsilon(s)}{u_g(s) + u_d(s)}}.$$

Source p. 9 · Elevation transfer, step solution and extrema

Section 2.2 recomputes the damping from the printed extrema and compares it with the two archived choices.

$$\boxed{H_e(s) = \frac{K_e}{1 + \frac{2\xi}{\omega_0}s + \frac{1}{\omega_0^2}s^2}}.$$

$$aK_e = 0.34, \qquad K_e = 0.29,$$

$$d_1 = 2.65\ \text{cm}, \qquad d_2 = 2.15\ \text{cm}, \qquad \boxed{\xi = 0.05},$$

$$y(t) = aK_e\left[1 - \frac{e^{-\xi\omega_0 t}}{\sqrt{1-\xi^2}}\sin\left(\omega_0\sqrt{1-\xi^2}\,t + \phi\right)\right], \qquad \phi = \arctan\left(\frac{\sqrt{1-\xi^2}}{\xi}\right).$$

$$t_1 = 2\ \text{s}, \qquad t_2 = 6.5\ \text{s}, \qquad \boxed{\omega_0 = 0.7\ \text{rad s}^{-1}}.$$

$$\frac{d_2}{d_1} = \exp\left(-\frac{\pi\xi}{\sqrt{1-\xi^2}}\right), \qquad \omega_0 = \frac{\pi}{\sqrt{1-\xi^2}}\frac{1}{t_2 - t_1}.$$

$$\boxed{H_t(s) = \frac{\theta(s)}{u_g(s) - u_d(s)}}.$$

Source p. 10 · Pitch gain and yaw step solution

$$\boxed{H_t(s) = \frac{K_t}{s^2}},$$

$$\boxed{H_l(s) = \frac{\lambda(s)}{\theta(s)}}.$$

$$y(t) = K_t\frac{(t-t_0)^2}{2}a.$$

$$\boxed{H_l(s) = \frac{K_l}{s(1+\tau_l s)}}.$$

$$K_t = \frac{2}{a}\left(\frac{\sqrt{y_2} - \sqrt{y_1'}}{t_2 - t_1'}\right)^2, \qquad \boxed{K_t = 1.336}.$$

$$\hat{y}(s) = \frac{K_l}{s(1+\tau_l s)}\frac{a}{s} \quad \Longrightarrow \quad y(t) = aK_l\left(t - \tau_l + \tau_l e^{-t/\tau_l}\right).$$

Source p. 11 · Yaw gain and plant-state definitions

$$K_l a = \frac{102.5 - 0}{60 - 20}, \qquad \boxed{K_l = -5.219}.$$

$$U = \begin{bmatrix} U_g \\ U_d \end{bmatrix}, \qquad Y = \begin{bmatrix} \epsilon \\ \nu \end{bmatrix}.$$

$$\ddot{\epsilon} = -2\xi\omega_0\dot{\epsilon} - \omega_0^2\epsilon + K_e\omega_0^2(U_g + U_d),$$

$$X = \begin{bmatrix} \dot{\epsilon} & \epsilon & \dot{\theta} & \theta & \nu \end{bmatrix}^T.$$

$$\ddot{\theta} = K_t(U_g - U_d),$$

$$\dot{X} = AX + BU, \qquad Y = CX,$$

$$\dot{\nu} = -\frac{1}{\tau_l}\nu + \frac{K_l}{\tau_l}\theta, \qquad \nu = \dot{\lambda}.$$

Source p. 12 · Symbolic and numerical plant matrices

Section 2.5 retains this matrix as the nominal multivariable realization. Its yaw time constant is 10 s.

$$A = \begin{bmatrix} -2\xi\omega_0 & -\omega_0^2 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & K_l/\tau_l & -1/\tau_l \end{bmatrix}, \qquad A = \begin{bmatrix} -0.1120 & -0.4900 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & -0.5219 & -0.1000 \end{bmatrix},$$

$$B = \begin{bmatrix} K_e\omega_0^2 & K_e\omega_0^2 \\ 0 & 0 \\ K_t & -K_t \\ 0 & 0 \\ 0 & 0 \end{bmatrix}, \qquad C = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}.$$

$$B = \begin{bmatrix} 0.1421 & 0.1421 \\ 0 & 0 \\ 1.3360 & -1.3360 \\ 0 & 0 \\ 0 & 0 \end{bmatrix}, \qquad C = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}.$$

Source p. 13 · Refined elevation model and realization

Section 3.1 uses the order elevation, elevation rate for the displayed realization and audits its input gain.

$$H_e(s) = \frac{K}{s^2/\omega_0^2 + (2\xi/\omega_0)s + 1} = \boxed{\frac{0.2131}{s^2 + 0.1465s + 0.6461}},$$

$$\boxed{\ddot{\epsilon}(t) + 2\xi\omega_0\dot{\epsilon}(t) + \omega_0^2\epsilon(t) = \omega_0^2 K\, u(t)}.$$

$\boxed{\xi = 0.0912}$, $\boxed{\omega_0 = 0.8038\ \mathrm{rad\,s^{-1}}}$, $\boxed{K = 0.3299\ \mathrm{rad\,V^{-1}}}$.

With $x = [\epsilon\ \dot{\epsilon}]^T$,

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

$$A = \begin{bmatrix} 0 & 1 \\ -\omega_0^2 & -2\xi\omega_0 \end{bmatrix}, \quad B = \begin{bmatrix} 0 \\ \omega_0^2 K \end{bmatrix}, \quad C = \begin{bmatrix} 1 & 0 \end{bmatrix}, \quad D = 0,$$

Source p. 14 · Sampled model and servo augmentation

Section 3.2 evaluates both exact sampled reachability determinants.

$$\boxed{A = \begin{bmatrix} 0 & 1 \\ -0.6461 & -0.146 \end{bmatrix}, \quad B = \begin{bmatrix} 0 \\ 0.2652 \end{bmatrix}, \quad C = \begin{bmatrix} 1 & 0 \end{bmatrix}, \quad D = 0}.$$

$$x_{k+1} = Fx_k + Gu_k, \qquad y_k = Cx_k,$$

$$F_e = \begin{bmatrix} F & 0 \\ C & 1 \end{bmatrix}, \qquad G_e = \begin{bmatrix} G \\ 0 \end{bmatrix},$$

$$F = e^{AT_s}, \qquad G = \int_0^{T_s} e^{A\tau} B\, d\tau.$$

$$F_e = \begin{bmatrix} 0.9992 & 0.0498 & 0 \\ -0.0322 & 0.9919 & 0 \\ 1 & 0 & 1 \end{bmatrix}, \qquad G_e = \begin{bmatrix} 0.0003 \\ 0.0132 \\ 0 \end{bmatrix}.$$

$$\boxed{F = \begin{bmatrix} 0.9992 & 0.0498 \\ -0.0322 & 0.9919 \end{bmatrix}, \qquad G = \begin{bmatrix} 0.0003 \\ 0.0132 \end{bmatrix}}.$$

$$\det\begin{bmatrix} G & FG \end{bmatrix} = 0.0003(0.0131) - 0.0009(0.0132) = -7.9 \times 10^{-6} \neq 0.$$

Source p. 15 · Pole locations, design rule and feedback gains

$$\frac{\tau_{\min}}{10} \leq T_s \leq \frac{\tau_{\min}}{5}.$$

$$\boxed{P_{BF,z} = \{0.9512,\ 0.9505 \pm 0.0381i\}}.$$

$$\boxed{K_p = \begin{bmatrix} 10.6434 & 10.2527 \end{bmatrix}, \qquad K_i = 0.2882}.$$

$$\boxed{K_p = \begin{bmatrix} 40.3459 & 20.0418 \end{bmatrix}, \qquad K_i = 1.5137}.$$

Source p. 17 · Multivariable augmentation and criterion

$$A_a = \begin{bmatrix} A & 0_{n\times p} \\ C & 0_{p\times p} \end{bmatrix}, \qquad B_a = \begin{bmatrix} B \\ 0_{p\times m} \end{bmatrix}.$$

$$A_a = \begin{pmatrix} -0.1120 & -0.4900 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -0.5219 & -0.1000 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \end{pmatrix},$$

$$B_a = \begin{pmatrix} 0.1421 & 0.1421 \\ 0 & 0 \\ 1.3360 & -1.3360 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & 0 \end{pmatrix}.$$

$$\boxed{J = \int_0^\infty \left(x^T Q x + u^T R u\right) dt},$$

$$Q = \operatorname{diag}(0, 0, 0, 0, 0, 40, 40), \qquad R = I_2.$$

$$A_{a,\alpha} = A_a + \alpha_c I_7.$$

Source p. 18 · Riccati equation, matrix, gains and reported poles

Section 4 solves the Riccati equation with the shifted state matrix and identifies the frame of the printed eigenvalues.

$$PA_a + A_a^T P - PB_a R^{-1} B_a^T P + Q = 0.$$

$$P = \begin{pmatrix} 87.1427 & 119.5267 & 0 & 0 & 0 & 79.0374 & 0 \\ 119.5267 & 278.1660 & 0 & 0 & 0 & 207.9665 & 0 \\ 0 & 0 & 1.7422 & 4.5466 & -11.5468 & 0 & -7.6855 \\ 0 & 0 & 4.5466 & 17.7038 & -53.7353 & 0 & -40.1133 \\ 0 & 0 & -11.5468 & -53.7353 & 189.5634 & 0 & 162.1503 \\ 79.0374 & 207.9665 & 0 & 0 & 0 & 212.2804 & 0 \\ 0 & 0 & -7.6855 & -40.1133 & 162.1503 & 0 & 170.8591 \end{pmatrix}.$$

The optimal gain $K = R^{-1} B_a^T P = [K_p\ K_i]$ is

$$K_p = \begin{pmatrix} 12.3830 & 16.9847 & 2.3276 & 6.0742 & -15.4265 \\ 12.3830 & 16.9847 & -2.3276 & -6.0742 & 15.4265 \end{pmatrix},$$

$$K_i = \begin{pmatrix} 11.2312 & -10.2679 \\ 11.2312 & 10.2679 \end{pmatrix}$$

$$\boxed{\{-0.5607 \pm 0.9945i,\ -0.6338 \pm 1.3929i,\ -1.0098,\ -1.5259 \pm 0.5786i\}},$$

Source p. 19 · Comparative table

Table 2.1: Comparison of integral pole placement and integral LQR.

Criterion	Integral state feedback	Integral LQR
Controlled scope	SISO elevation channel	MIMO elevation and yaw-rate channels
Primary tuning variables	Desired pole locations / T_c	Weighting matrices Q, R, and spectral shift α_c
Steady-state error	Zero through one integral state	Zero through two integral states
Nominal reported response	About 6 s for $T_c = 1$ s; about 3 s for $T_c = 0.5$ s	About 3.1 s elevation and 2.7 s yaw rate
Control effort	Explicitly rises as poles move left; about 3.3 V for the faster design	Limited systematically through R; no simulated saturation
Disturbance rejection	About 4.41 s in real manual tests	About 5 s for simultaneous simulated channel disturbances
Design transparency	Direct modal interpretation and straightforward implementation	Systematic optimal compromise, naturally multivariable
Main limitation	Gains become aggressive as response speed is increased	Performance depends on physically meaningful weight selection and model fidelity

Appendix C. Reproducibility

C.1 Embedded material

The PDF attachments are **Helicopter_Control_Report.pdf**, the complete 20-page source [1], **helicopter_study.py**, the independent simulation script, and **helicopter_results.json**, the values behind the numerical tables and checks. Extract the attachments from the PDF attachment pane of a desktop PDF reader. The original Simulink models are represented by the supplied screenshots. The archived PDF contains no model file or source-code listing.

```
python helicopter_study.py --output results
```

Python, NumPy, SciPy and Matplotlib versions are recorded in the JSON file. The script writes the eleven vector figures and the full numerical results to the chosen directory. It does not need the archived images or any hardware connection to run.

C.2 Calculation sequence

The sampled plant uses the augmented block exponential in (16) with an exact 0.05 s hold. The plant and discrete integrator are updated from the same current sample and reference. Pole-placement gains are checked both by recomputing poles from the source gain and by solving for the explicit exponential target in Section 3.3.

The continuous integration advances the closed-loop state with an exact held-input transition over intervals of 0.005 s. The discrete grid includes the reference edge and both pulse edges. The quantities held constant over each interval are r and any d, not the feedback u. The transition already includes the continuously acting feedback in its matrix. A finite-grid settling time is the first sample after the event such that all later samples in the evaluation window remain within the band. The pulse tests compare trajectories with identical references and define recovery using the difference between the two runs.

The script checks reachability, the shifted Riccati residual, and the equivalence of the frequency shift in the two matrices. It tests the yaw characteristic coefficients against the state matrix at 101 lags. The continuous-interval proof uses the polynomial minima in Section 5.2. The script verifies convergence towards the independently solved steady state for the three simulated cases. Numerical and archived results have separate captions and origins throughout the manuscript.