IAES International Journal of Robotics and Automation (IJRA)
Vol. 15, No. 4, December 2026, pp. 1109∼1124
ISSN: 2722-2586, DOI: 10.11591/ijra.v15i4.pp1109-1124

Full-state LQI hover stabilization of a quad tail-sitter UAV using Newton-Euler modeling and PyBullet simulation

Ariesta Martiningtyas Handayani[1], Isnan Nur Rifai[1], Andi Dharmawan[2]
[1]Department of Electrical Engineering and Informatics, Vocational School, Universitas Gadjah Mada, Yogyakarta, Indonesia
[2]Department of Computer Science and Electronics, Faculty of Mathematics and Natural Sciences, Universitas Gadjah Mada, Yogyakarta, Indonesia

Article Info

Article history:

Received Jun 13, 2026
Revised Oct 2, 2026
Accepted Nov 10, 2026

Keywords:

Full-state LQI
Hover stabilization
Newton-Euler model
PyBullet simulation
Quad tail-sitter UAV

ABSTRACT

This paper couples a Newton-Euler hover model with a full-state linear quadratic integral controller for a 1.6 kg quad tail-sitter. Altitude, attitude, and horizontal position are all regulated by Riccati-derived integral loops, so no gain is hand-tuned, and every plant, controller, allocation, and disturbance parameter is reported for reproduction. Evaluated in PyBullet with actuator saturation, anti-windup, and ramping, hover holds an altitude error of 0.96 mm and settles in 2.34 s, the anti-windup clamp sized from an explicit thrust-deficit budget rather than by trial. A 2 N longitudinal impulse gives 36.1 mm peak displacement, recovering to a 10% band in 8.0 s, and a zero-mean Dryden gust a longitudinal root-mean-square error of 0.55 mm. The control law costs 19.1 μs per step, and discrete-time analysis gives at least 7.8 dB gain and 34.4° phase margin. The principal finding is where the design fails: sweeping the gust amplitude places the operational boundary at saturation of the outer-loop attitude reference, a bound derived analytically as $mg\tan\theta_{\max} = 5.10$ N and reached within 76 to 95%. Altitude regulation survives beyond it, so position hold is lost while hover is not. Results are simulation-only and not hardware-validated.

This is an open access article under the CC BY-SA license.

Corresponding Author:

Ariesta Martiningtyas Handayani
Department of Electrical Engineering and Informatics, Vocational School, Universitas Gadjah Mada
Yogyakarta, Indonesia
Email: ariesta_mh@mail.ugm.ac.id

1. INTRODUCTION

Tail-sitter unmanned aerial vehicles combine vertical take-off and landing capability with the aerodynamic efficiency of fixed-wing flight, making them attractive for mapping, inspection, and delivery missions that require both hover precision and forward-flight endurance [1]-[4]. Airframe studies of hybrid layouts, unmanned prototypes, and reconfigurable wings have established the configuration as a practical platform [5]-[8]. Simulation-first validation matters because a hover failure is rarely recoverable: loss of attitude at low altitude ends in ground impact that destroys the airframe and the sensor payload it carries, and over an inspection or delivery site it endangers the structures and people beneath the vehicle. Airframes are costly to rebuild, each rebuild stops the survey campaign it serves, and controller iteration is far faster and safer in software. Hover regulation nevertheless suffers steady-state error under model mismatch and disturbance, and the problem is sharper for quad tail-sitters without control surfaces, where every corrective force and moment comes from rotor-thrust allocation alone, so all translational authority requires tilting the whole airframe.

Prior work falls into five strands, each limited in a way that matters for hover. Transition-oriented studies [9]-[13] prioritise envelope transfer, so hover is a boundary condition whose accuracy is rarely reported. Nonlinear and robust approaches [14]-[22] are robust but need large, strongly interacting tuning sets, making independent auditing difficult and leaving the control law inseparable from the tuning. Hover-specific and disturbance-focused studies [23]-[26] come closest and do report behaviour under perturbation, but only at a nominal level, without identifying the amplitude at which the design stops working. Linear regulators without integral action [27] give clean nominal transients but leave persistent offset under bias. Learning-based and adaptive-predictive designs form a fifth and fast-growing strand: reinforcement learning has been used to predict attitude-controller gains online [28], model-reference adaptive control has been applied to tail-sitter roll regulation under gust [29], and adaptive model predictive control driven by online system identification has been demonstrated on dual-system VTOL platforms [30]. These remove the hand-tuned gain set, but the resulting behaviour is fixed by a learned or online-identified model that cannot be audited from the published text alone, which is the reproducibility property the present work targets. Across all five strands the recurring omission is the same: the parameter set is incomplete and the amplitude at which hold is actually lost is never identified.

Two gaps follow. First, on reporting: no prior tail-sitter hover study combines a Newton-Euler hover model, integral augmentation on all six regulated outputs, practical anti-windup and reference ramping, and a parameter set complete enough for independent reproduction. Second, and more substantively, robustness is almost always reported as performance under a nominal disturbance, so the reader learns that a controller works but not where it stops working or why.

The choice of a linear design is deliberate and quantifiable. Exploiting roll-pitch and x-y symmetry, the controller is defined by twenty independent scalars, all in Table 1: five each for altitude, shared roll-pitch, and yaw, four for the shared horizontal outer loop, and one attitude-reference limit. This is far smaller and less strongly coupled than backstepping or sliding-mode synthesis requires, and every gain follows from a small Riccati solution verifiable by inspection rather than by trust in a tuning procedure. For hover-only regulation about a fixed operating point that transparency outweighs the extra margin a nonlinear design would give, particularly since the controller tolerates an inertia mismatch of 18.9 in roll.

The contributions are: i) an explicit 12-state Newton-Euler hover linearization with its assumptions stated; ii) a full-state LQI in which altitude, attitude, and horizontal position are all regulated by Riccati-derived integral loops, so no gain is hand-tuned; iii) a parameter set, including allocation coefficients and closed-form disturbance profiles, sufficient for independent reproduction; iv) evaluation under impulsive, zero-mean Dryden gust, and sensor-noise conditions; v) an operational boundary derived analytically and confirmed against a swept gust amplitude; and vi) a component-wise ablation quantifying what each safeguard contributes and costs.

This matters because it shows that a linear controller with modest augmentation holds hover to within one millimetre and recovers from disturbance, giving a low-complexity reproducible baseline for future transition controllers, and because it shows where such a controller stops working and why: the envelope is fixed by a designer-chosen saturation limit in the outer loop rather than by the optimal control law. The result belongs to robotics and automation rather than to aerodynamics: hover is the innermost regulation layer of an aerial robot, the quantity it stabilises is the pose from which every robot-centric sensor measures, and the millimetre-level hold reported here is what makes automated inspection of infrastructure and of manufacturing plant feasible from a free-flying platform.

The objectives follow directly: i) to derive the hover model and state its assumptions; ii) to design a full-state LQI covering every regulated output; iii) to report all parameters needed for reproduction; iv) to quantify performance under the three disturbance classes; and v) to locate the operational limit and identify the mechanism that sets it.

Table 1. Controller weights, integral gains, and limits. The outer loop carries its integral weight inside $\mathbf{Q}_o$, so it has no separate k_i

Channel	$\mathbf{Q}$ or gains	$\mathbf{R}$	k_i	Integral clamp
Altitude z	$\mathrm{diag}(220, 70)$	0.18	1.20	0.60
Roll ϕ / Pitch θ	$\mathrm{diag}(90, 12)$	0.35	0.25	0.18
Yaw ψ	$\mathrm{diag}(24, 6)$	0.55	0.45	0.25
Position x / y (outer)	$\mathbf{Q}_o = \mathrm{diag}(6, 2, 1.2)$	1.00	—	0.25

2. THE PROPOSED FULL-STATE LQI HOVER CONTROLLER

2.1. Hover dynamics model

Let $\mathbf{p} = [x, y, z]^T$ be inertial position, $\mathbf{v} = \dot{\mathbf{p}}$ inertial velocity, $\boldsymbol{\eta} = [\phi, \theta, \psi]^T$ Euler angles, and $\boldsymbol{\omega} = [p, q, r]^T$ body angular rates. With $\mathbf{R}(\boldsymbol{\eta})$ the body-to-inertial rotation, $\mathbf{I}_b = \mathrm{diag}(I_x, I_y, I_z)$, T the total rotor thrust along the body vertical $\mathbf{e}_3$, and $\boldsymbol{\tau} = [\tau_\phi, \tau_\theta, \tau_\psi]^T$ the control moments, the Newton-Euler equations are

$$m\dot{\mathbf{v}} = \mathbf{R}(\boldsymbol{\eta})\, T\mathbf{e}_3 - mg\mathbf{e}_3 - \mathbf{C}_{\mathrm{lin}}\mathbf{v}, \qquad \mathbf{I}_b\dot{\boldsymbol{\omega}} + \boldsymbol{\omega} \times (\mathbf{I}_b\boldsymbol{\omega}) = \boldsymbol{\tau} - \mathbf{D}_{\mathrm{ang}}\boldsymbol{\omega}, \tag{1}$$

where $\mathbf{C}_{\mathrm{lin}} = \mathrm{diag}(c_x, c_y, c_z)$ and $\mathbf{D}_{\mathrm{ang}} = \mathrm{diag}(d_\phi, d_\theta, d_\psi)$ collect velocity-proportional dissipative effects, taken isotropically as $c_x = c_y = c_z = C_{d,\mathrm{lin}}$ and $d_\phi = d_\theta = d_\psi = C_{d,\mathrm{ang}}$. The state is $\mathbf{x}_{12} = [x, y, z, \dot{x}, \dot{y}, \dot{z}, \phi, \theta, \psi, p, q, r]^T$ and the input $\mathbf{u} = [\delta T, \tau_\phi, \tau_\theta, \tau_\psi]^T$, with $\delta T = T - mg$ the thrust deviation from hover trim. Linearization about $\mathbf{v}_0 = \mathbf{0}$, $\boldsymbol{\omega}_0 = \mathbf{0}$, $\boldsymbol{\eta}_0 = \mathbf{0}$, $T_0 = mg$ uses six assumptions, stated so that the validity envelope is unambiguous. (A1) The airframe is rigid with constant mass and inertia. (A2) Attitude deviations are small, so $\sin\alpha \approx \alpha$, $\cos\alpha \approx 1$, products of angles vanish, and $\dot{\boldsymbol{\eta}} \approx \boldsymbol{\omega}$. (A3) The gyroscopic term $\boldsymbol{\omega} \times (\mathbf{I}_b\boldsymbol{\omega})$ is second order and discarded, as are rotor gyroscopic torque and spin inertia, the rotors being 1.25% of airframe mass. (A4) Aerodynamic loading is linear damping only; wing lift and drag, rotor inflow, blade flapping, propeller interaction, and ground effect are not modelled. (A5) Rotor thrust responds instantaneously along the body vertical, so motor and speed-controller lag are absent. (A6) The full state is available, so no observer appears. The result is $\dot{\mathbf{x}}_{12} = \mathbf{A}\mathbf{x}_{12} + \mathbf{B}\mathbf{u}$, $\mathbf{y} = \mathbf{C}\mathbf{x}_{12}$ with $\mathbf{y} = [x, y, z, \phi, \theta, \psi]^T$, and in 3×3 blocks

$$\mathbf{A} = \begin{bmatrix} \mathbf{0} & \mathbf{I} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & -m^{-1}\mathbf{C}_{\mathrm{lin}} & \mathbf{G} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{I} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & -\mathbf{I}_b^{-1}\mathbf{D}_{\mathrm{ang}} \end{bmatrix}, \ \mathbf{B} = \begin{bmatrix} \mathbf{0}_{3\times1} & \mathbf{0} \\ m^{-1}\mathbf{e}_3 & \mathbf{0} \\ \mathbf{0}_{3\times1} & \mathbf{0} \\ \mathbf{0}_{3\times1} & \mathbf{I}_b^{-1} \end{bmatrix}, \ \mathbf{C} = \begin{bmatrix} \mathbf{I} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{I} & \mathbf{0} \end{bmatrix}, \tag{2}$$

where $\mathbf{G}$ has entries $G_{12} = g$ and $G_{21} = -g$ and is zero elsewhere. The block $\mathbf{G}$ carries the only surviving cross-channel coupling: $\ddot{x} = g\theta - m^{-1}c_x\dot{x}$ and $\ddot{y} = -g\phi - m^{-1}c_y\dot{y}$, while $\ddot{z} = m^{-1}(\delta T - c_z\dot{z})$ and each rotational axis is an independent damped double integrator. With Table 2 values, $m^{-1}C_{d,\mathrm{lin}} = 0.156\ \mathrm{s}^{-1}$ and $\mathbf{I}_b^{-1}\mathbf{D}_{\mathrm{ang}} = \mathrm{diag}(3.333, 0.667, 0.667)\ \mathrm{s}^{-1}$. Both $(\mathbf{A}, \mathbf{B})$ and $(\mathbf{A}, \mathbf{C})$ attain full rank 12 for the design and the plant inertias alike, so the mismatch noted in Table 2 does not compromise controllability or observability. Horizontal motion is reachable only through the tilt path in $\mathbf{G}$, which is the structural reason for the inner-outer arrangement below and, as section 4.3 shows, for the operational limit.

Table 2. Physical plant and simulation parameters (SDF: Simulation description format)

Symbol	Parameter	Value	Symbol	Parameter	Value
m	Airframe mass	1.600 kg	$C_{d,\mathrm{lin}}$	Translational damping	0.250 N·s/m
m_r	Rotor mass (each, 4×)	0.005 kg	$C_{d,\mathrm{ang}}$	Rotational damping	0.020 N·m·s/rad
I_{xx}	Roll inertia, SDF plant	0.11333	$f_{\min}, f_{\max}$	Per-rotor thrust bounds	0, 18.0 N
I_{yy}	Pitch inertia, SDF plant	0.03021	g	Gravity	9.810 m/s^2
I_{zz}	Yaw inertia, SDF plant	0.08354	Δt	Integration time step	1/240 s
$\hat{I}_x$	Roll inertia, design	0.006	z_{ref}	Altitude command	3.0 m
$\hat{I}_y$	Pitch inertia, design	0.030	z_0	Initial altitude	0.05 m
$\hat{I}_z$	Yaw inertia, design	0.030	T_{ramp}	Reference ramp constant	0.70 s
ℓ_x	Rotor arm, longitudinal	0.145 m	γ	Rotor tilt angle	0.170 rad
ℓ_y	Rotor arm, lateral	0.230 m	k_ψ	Yaw allocation coefficient	0.046 m

Inertias in kg·m^2. The design inertias $\hat{I}$ differ from the plant values because gains were tuned against a simplified estimate before the SDF mass distribution was final. The mismatch is retained deliberately: roll is designed against an inertia 18.9 times smaller than the plant, so section 4 reports performance under substantial parametric uncertainty

2.2. Full-state LQI formulation

The controller is an inner-outer structure. The inner layer applies four independent linear-quadratic regulators to the altitude and attitude channels, each derived from the double-integrator model [31] $\dot{\mathbf{x}}_c = \mathbf{A}_c\mathbf{x}_c + \mathbf{B}_c u_c$ with $\mathbf{A}_c = [[0, 1], [0, 0]]$ and $\mathbf{B}_c = [0,\ 1/m]^T$ for altitude or $[0,\ 1/\hat{I}]^T$ for roll, pitch, and yaw. Each channel carries a clamped integral term supplying the offset rejection a pure state-feedback law cannot provide:

$$u_c = -\mathbf{K}_c \mathbf{x}_c - k_i \xi_c, \qquad \xi_c = \mathrm{sat}\left(\int e_c \, dt, \, [-\xi_{\max}, \xi_{\max}] \right). \tag{3}$$

The outer layer regulates horizontal position and is likewise an LQI. Writing $e_x = x - x_{\mathrm{ref}}$ and $\xi_x = \int e_x \, dt$, the $\mathbf{G}$ block of section 2.1 gives $\ddot{e}_x = g\theta - m^{-1} C_{d,\mathrm{lin}} \dot{e}_x$, so the augmented horizontal channel is

$$\mathbf{A}_o = \begin{bmatrix} 0 & 1 & 0 \\ 0 & -m^{-1} C_{d,\mathrm{lin}} & 0 \\ 1 & 0 & 0 \end{bmatrix}, \qquad \mathbf{B}_o = \begin{bmatrix} 0 \\ g \\ 0 \end{bmatrix}, \tag{4}$$

with tilt as the control. The pair $(\mathbf{A}_o, \mathbf{B}_o)$ has full rank 3, and solving the algebraic Riccati equation for the weights of Table 1 yields $\mathbf{K}_o = [k_x, \, k_v, \, k_i] = [3.091, \, 1.606, \, 1.095]$. The attitude references follow as

$$\theta_{\mathrm{ref}} = -\mathrm{sat}\big(\mathbf{K}_o [e_x, \dot{x}, \xi_x]^T, [-\theta_{\max}, \theta_{\max}]\big), \quad \phi_{\mathrm{ref}} = \mathrm{sat}\big(\mathbf{K}_o [e_y, \dot{y}, \xi_y]^T, [-\phi_{\max}, \phi_{\max}]\big), \tag{5}$$

with $\theta_{\max} = \phi_{\max} = 18^\circ$. The resulting closed-loop poles of the outer channel are -13.76, -1.68, and -0.46, all real, so the horizontal loop is overdamped by construction rather than by tuning. Every regulated output in the vehicle is therefore governed by an LQI: four two-dimensional Riccati solutions for altitude and attitude, and one three-dimensional solution shared by the two horizontal channels. This decentralised arrangement keeps every Riccati solution small and analytically transparent, while position-attitude coupling is handled explicitly through the outer loop and its reference limits.

The weights in Table 1 follow a consistent priority ordering. Altitude carries the largest position weight, sub-centimetre height holding being the primary objective, and its small control penalty lets thrust respond quickly during the climb. Roll and pitch receive equal weights well above yaw, since heading error matters far less for hover accuracy than tilt error, which couples into horizontal drift through $\mathbf{G}$. Velocity and rate weights sit near one third of the corresponding position or attitude weight, damping the transient without extending rise time. Integral gains stay an order of magnitude below the proportional terms, so offset removal is slower than the primary loop, and every integral state is bounded by an explicit clamp providing the anti-windup examined in section 4.2.

2.3. Thrust and moment mixing

The airframe layout and body-fixed axes are shown in Figure 1. The four rotors sit at $(\pm\ell_x, \pm\ell_y)$ and, unlike a planar quadrotor, each is canted about *two* axes: rotor i is mounted with roll $s_i\gamma$ and pitch $t_i\gamma$, where $s_i, t_i \in \{+1, -1\}$ follow the pattern $(+,-)$, $(-,+)$, $(-,-)$, $(+,+)$. Its thrust therefore acts along

$$\mathbf{n}_i = \mathbf{R}_x(s_i\gamma) \, \mathbf{R}_y(t_i\gamma) \, \mathbf{e}_3 = \begin{bmatrix} t_i \cos\gamma \sin\gamma \\ -s_i \sin\gamma \\ \cos^2\gamma \end{bmatrix}, \tag{6}$$

which at $\gamma = 0.17$ rad evaluates to $[\pm 0.1667, \, \mp 0.1692, \, 0.9714]^T$ and matches the directions measured in the simulator to within 2×10^{-8}. Two consequences follow and both are used later. Each rotor contributes only $\cos^2\gamma = 97.14\%$ of its force to the body vertical, which is the geometric half of the thrust deficit quantified in section 4.1; and because $\sum_i s_i = \sum_i t_i = 0$, the in-plane components cancel exactly in symmetric hover, so the cant produces yaw authority without net side force.

Writing $\mathbf{r}_i$ for the rotor position relative to the centre of mass, the exact map is $[F_z, \boldsymbol{\tau}]^T = \mathbf{A}_{\mathrm{ex}} \mathbf{f}$ with rows $\mathbf{n}_{i,z}$ and $\mathbf{r}_i \times \mathbf{n}_i$. The implemented allocation simplifies this by taking the rotor forces along the body vertical and retaining the cant only in the yaw row, so roll follows the lateral arm, pitch the longitudinal arm, and yaw the tilt-induced in-plane force: $[T, \tau_\phi, \tau_\theta, \tau_\psi]^T = \mathbf{M}[f_1, f_2, f_3, f_4]^T$ with

$$\mathbf{M} = \begin{bmatrix} 1 & 1 & 1 & 1 \\ -\ell_y & \ell_y & \ell_y & -\ell_y \\ -\ell_x & \ell_x & -\ell_x & \ell_x \\ -k_\psi & -k_\psi & k_\psi & k_\psi \end{bmatrix}, \qquad k_\psi = \sin\gamma \sqrt{\ell_x^2 + \ell_y^2}. \tag{7}$$

Substituting $\ell_x = 0.145$ m, $\ell_y = 0.230$ m, and $k_\psi = 0.046$ m gives a full-rank matrix of condition number 21.7. Rotor commands follow from $\mathbf{M}^{-1}$, clipped to $[0, 18]$ N.

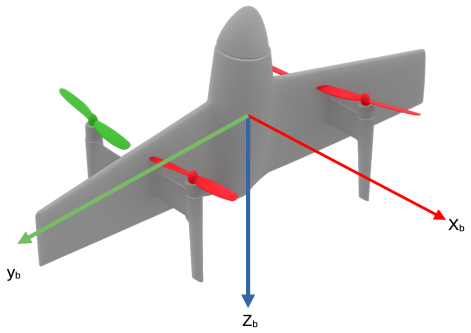

Figure 1. Quad tail-sitter UAV configuration and body-fixed reference axes used in the hover model

Row by row, $\mathbf{M}$ overstates vertical efficiency by 2.9% (1 against 0.9714), roll authority by 3.7% (0.230 against 0.2217 m) and pitch authority by 1.8% (0.145 against 0.1425 m), and understates yaw authority by 27% (0.046 against 0.0629 m). Only the first is a standing bias, and it appears directly in the altitude budget of section 4.1; the attitude errors act on channels whose references are zero in hover, so they perturb the effective loop gain rather than introduce offset, and the yaw discrepancy is invisible here because no scenario commands or disturbs heading.

$\mathbf{M}$ is retained deliberately; running the identical controller with $\mathbf{A}_{\mathrm{ex}}$ substituted shows what that costs. The exact allocation removes the geometric part of the thrust deficit, cutting the standing demand on the integrator from 0.665 to 0.196 N and the required authority from $\xi_{\max,z} = 0.554$ to 0.163; with integral action disabled the steady altitude error falls from 19.0 to 5.6 mm. With integral action enabled and the clamp sized to each case, the two reach the same place: 0.96 mm for $\mathbf{M}$ at $\xi_{\max,z} = 0.60$ against 0.35 mm for $\mathbf{A}_{\mathrm{ex}}$ at 0.18. Allocation fidelity and integral authority are therefore substitutes to first order, and the budget of section 4.1 is the exchange rate between them: a designer who cannot measure the cant angles precisely can buy the same accuracy by enlarging $\xi_{\max,z}$ instead, provided the actuators retain headroom to deliver it.

2.4. Implementation safeguards

Figure 2 shows the architecture as implemented, with numerical values attached to each block. Four implementation details complete it. Translational and rotational damping are retained in the plant and compensated in tuning to suppress oscillatory transients. Integrator anti-windup prevents the six integral states from growing without bound when rotor commands saturate. Reference ramping limits the initial altitude demand, avoiding an unrealistically abrupt jump from ground to the 3 m command. Command saturation is enforced after allocation so each rotor stays within feasible thrust bounds.

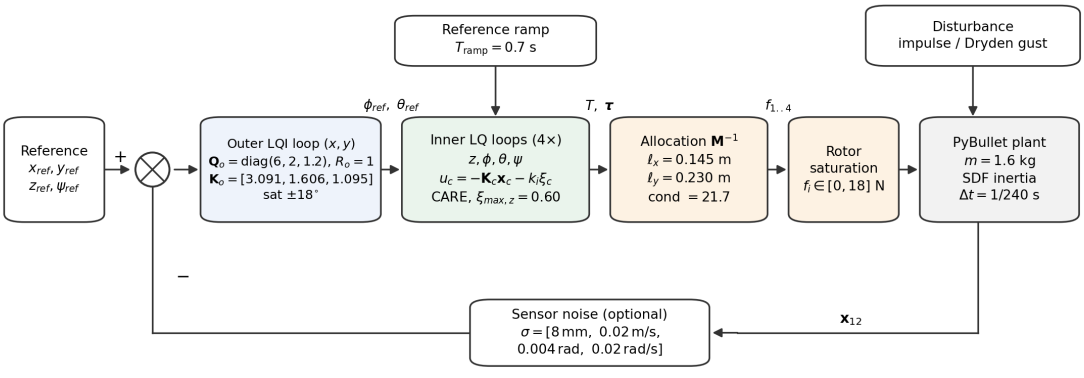

Figure 2. Implemented full-state LQI hover architecture: the decentralised inner-outer arrangement in which both layers are LQI, the attitude-reference saturation that sets the operational limit of section 4.3, the allocation and rotor-saturation stages, and the entry points of disturbance and sensor noise

3. METHOD

3.1. Simulation setup and reproducibility

All results come from PyBullet 3.2.7 under Python 3.11 with NumPy and SciPy, using a rigid-body tail-sitter defined in SDF, with PyBullet's own linearDamping and `angularDamping` both 0.01 in addition to the damping of Table 2. Every regulator gain comes from the continuous-time algebraic Riccati equation via `scipy.linalg.solve_continuous_are`, so no gain is hand-tuned after synthesis. The vehicle starts from rest at $z_0 = 0.05$ m with zero attitude, integral states held at zero until the ramp reaches 95% of final value; runs last 20 s; seeds are 11 for all disturbance cases and 100–111 for the noise Monte Carlo. Reproduction needs only the constants of Tables 1 and 2, the allocation of section 2.3, and the recursions below. In the terminology of experimental design this is a one-group posttest-only study: a single controller configuration is exercised against fixed, fully specified disturbance realisations, with the ablation variants of section 4.2 and the hand-tuned outer loop of section 4.1 serving as internal comparison groups rather than as an independent control group. The design supports attribution of an effect to a named controller component, which is what the ablation asks, but not generalisation to hardware, which would require hardware-in-the-loop or flight testing; section 5 states that limit. Scripts, SDF model, and meshes are available from the corresponding author on reasonable request, and are organised so that each figure and table in this paper is regenerated by one named script from the constants reported here, allowing a reader to reproduce a single result without rerunning the whole study, to substitute another controller behind the same allocation and disturbance code, or to reuse the disturbance recursions alone as a benchmark for a different airframe.

3.2. Disturbance profiles

Because reproducibility depends on the disturbance being specified rather than named, each profile is given exactly as used. Impulse: a constant $F_x = 2$ N at the centre of mass over $5.0 \le t \le 5.2$ s, an impulse of 0.4 N·s. Dryden gust: a Dryden-form shaping filter driven by white noise generates $\mathbf{v}_g = [u_g, v_g, w_g]^T$, the longitudinal component using the first-order spectrum and the lateral and vertical components the second-order form, discretised at Δt with $V = 12$ m/s, length scales $L_u = 35$, $L_v = 20$, $L_w = 8$ m, and intensities $\sigma_u = 0.6$, $\sigma_v = 0.8$, $\sigma_w = 0.5$ m/s. The longitudinal channel is a first-order filter and reproduces σ_u exactly; the lateral and vertical channels are second-order filters whose noise-injection scaling instead gives a long-run RMS of 1.20 and 0.83 m/s, above σ_v and σ_w (obtained by solving each filter's discrete Lyapunov equation), so any single realisation, including the 20 s run of section 4.3, can differ further still. A 200-step burn-in precedes mean subtraction, so the gust is strictly zero-mean. Velocity becomes force through a diagonal admittance $\mathbf{F}_g = k_g \,\mathrm{diag}(1, 0.909, 0.182)\,\mathbf{v}_g$, nominal $k_g = 0.044$ N per m/s: an equivalent-force approximation, not a lift-and-drag model, which is why section 4.3 sweeps k_g. Sensor noise: zero-mean Gaussian noise on the measured state each step, standard deviations 8 mm in position, 0.02 m/s in velocity, 0.004 rad in attitude, 0.02 rad/s in angular rate, over twelve independent runs.

3.3. Evaluation metrics

Reported metrics are altitude steady-state error, altitude settling time to a 2% band, peak commanded thrust, peak horizontal displacement, RMS longitudinal error, Euler-angle excursions, and peak integral-state norm. Loss of position hold is defined as a peak horizontal excursion exceeding 1 m, one third of the hover altitude. Two further metrics answer questions the transient traces cannot. Computational cost is the wall-clock time of one evaluation of the complete control law, measured over 2×10^5 repetitions and compared against the Δt budget. Robustness is reported per channel on the discretised loop, as the peak sensitivity $M_S = \max_\omega |S|$ with $S = (1 + L)^{-1}$, its reciprocal the modulus margin, and the classical phase margin; M_S is preferred to a bare gain margin because each loop contains a double integrator together with integral action, so its phase crosses -180° at low frequency and a first-crossing gain margin is not meaningful, whereas M_S bounds both classical margins through $\mathrm{GM} \ge M_S/(M_S - 1)$ and $\mathrm{PM} \ge 2\arcsin(1/2M_S)$.

4. RESULTS AND DISCUSSION

4.1. Nominal hover and impulse response

Table 3 summarises both cases. In nominal hover the altitude error settles at 0.96 mm, reaching the 2% band in 2.34 s without overshoot. A standing thrust deficit must be absorbed for this to happen, and quantifying it fixes the required integral authority. The plant carries the four rotor masses, a weight of $1.620 \times 9.81 = 15.892$ N against the 15.696 N the controller assumes, and each rotor is canted by γ about two axes, so only

$\cos^2\gamma = 97.14\%$ of rotor force acts vertically. Holding altitude therefore demands $15.892/\cos^2\gamma = 16.360$ N, the steady value in Table 3, while the feedforward mg supplies 15.696 N. Only the integral channel can close the outstanding 0.665 N without a standing error, which requires

$$\xi_{\max,z} \geq \frac{T_{\text{req}} - mg}{k_{i,z}} = \frac{0.665}{1.20} = 0.554. \tag{8}$$

The clamp is set to 0.60; any smaller value caps the integral channel below the deficit and forces the remainder onto proportional action, costing $(0.665 - k_{i,z}\xi_{\max,z})/K_{z,1}$ in standing error. At $\xi_{\max,z} = 0.30$ formula and simulation agree at 8.71 mm, and with integral action removed at $0.665/34.96 = 19.01$ mm, matching Table 4, so the rule reproduces the ablation to three digits. Commanded thrust during the climb reaches 106.27 N against the 72 N the four rotors can deliver, so actuators saturate through the early transient and the anti-windup clamps are active in the nominal case too. The attitude and horizontal channels sit at the numerical floor because the nominal manoeuvre is symmetric and noise-free: they show that the altitude loop does not excite the lateral channels, but they do not measure lateral accuracy, which section 4.4 supplies once noise breaks the symmetry. Figure 3 shows the altitude and thrust traces, with the altitude response in Figure 3(a) and the total commanded thrust in Figure 3(b).

Table 3. Nominal hover and 2 N longitudinal impulse performance

Nominal hover metric	Value	Impulse metric ($t = 5$ s)	Value
Altitude steady-state error	0.00096 m	Peak longitudinal displacement	0.0361 m
Altitude settling time (2% band)	2.34 s	Peak lateral displacement	6.8×10^{-10} m
Altitude overshoot	0.000 m	Peak pitch excursion	0.1398 rad
Steady total thrust	16.36 N	Altitude deviation	0.0013 m
Peak commanded thrust (climb)	106.27 N	Residual displacement at 20 s	0.0005 m
Peak per-rotor command	18.0 N (saturated)	Recovery time (10% band)	8.04 s
Peak roll/pitch/yaw deviation	$< 2 \times 10^{-18}$ rad	Oscillation period	≈ 1.0 s

Table 4. Component-wise ablation across three scenarios. The Dryden columns use an admittance of $k_g = 0.880$ N per m/s, twenty times nominal, because the nominal gust is too weak to separate the variants; this level is still well inside the stable envelope, with peak roll at 13° against the 18° limit

	Nominal hover			2 N impulse		Dryden gust	
Variant	SSE_z (m)	Settling (s)	Peak $\|\|\xi\|\|$	Peak x (m)	Peak pitch (rad)	RMS x (m)	Peak roll (deg)
Full controller	0.00096	2.34	0.600	0.0361	0.1398	0.01093	13.00
No integral action	0.01901	2.53	0.600	0.0362	0.1385	0.01296	13.00
No anti-windup	0.00173	2.33	0.638	0.0361	0.1398	0.01093	13.00
No damping	0.00096	2.31	0.600	0.0371	0.1430	0.01233	13.67
No reference ramping	0.00096	2.22	0.600	0.0361	0.1398	0.01101	14.53

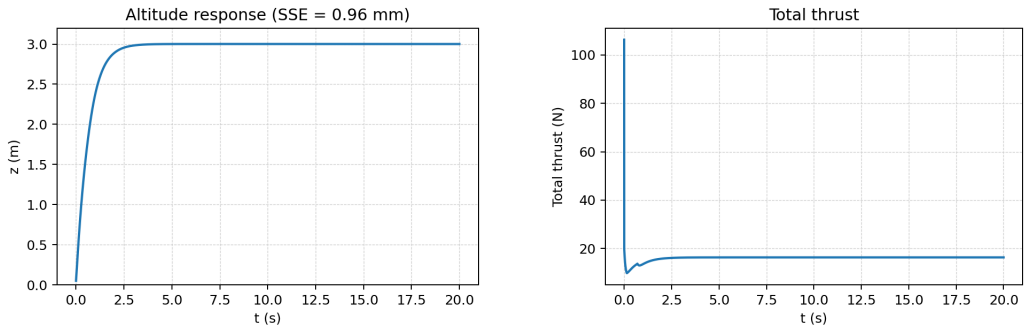

(a) (b)

Figure 3. Principal nominal-hover responses: (a) the altitude response reaching the 3 m command, and (b) the total commanded thrust over the same run. The longitudinal and lateral traces are omitted because they remain at the numerical floor of the simulation throughout, as discussed in section 4.1

Under the 2 N impulse the vehicle is displaced 36.1 mm at peak with a pitch excursion of 0.1398 rad. Altitude is essentially unaffected, the vertical thrust component falling by under 1% at that tilt, leaving a 1.3 mm altitude excursion, and lateral displacement stays at the numerical floor, confirming that the allocation keeps roll and pitch decoupled. Recovery decays with a period near 1.0 s, re-enters a 10% band 8.04 s after the impulse, and leaves 0.5 mm of residual motion at 20 s.

The outer-loop weighting produces this. Because $\mathbf{K}_o$ comes from a Riccati solution with all-real closed-loop poles, the horizontal loop commands a larger initial tilt and spends that authority arresting the motion rather than overshooting it. Replacing only the outer layer by a hand-tuned proportional-derivative law of comparable bandwidth ($k_p = 1.2$, $k_d = 0.5$, clamped integral gain 0.35) isolates what the Riccati solution buys: peak displacement falls from 68.5 to 36.1 mm, the residual at 20 s from 16.9 to 0.5 mm, and the longitudinal RMS error from 25.2 to 7.9 mm, while the hand-tuned loop never re-enters a 10% band within the window. The cost is more tilt during recovery, $8.0°$ against $4.9°$, still under half the $18°$ limit. Figure 4 shows the flight path and per-axis traces, with the three-dimensional path in Figure 4(a), the longitudinal position in Figure 4(b), and the pitch history in Figure 4(c).

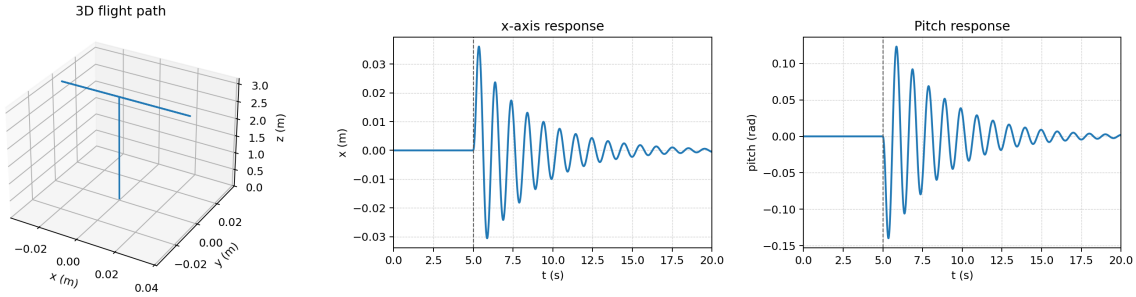

(a) (b) (c)

Figure 4. Response to the 2 N longitudinal impulse applied at $t = 5$ s: (a) the flight path, showing the vertical climb followed by the longitudinal excursion at hover altitude, with horizontal axes scaled to the data; (b) longitudinal position; and (c) pitch. The altitude trace is omitted, being indistinguishable from Figure 3(a) at this disturbance level

4.2. Component-wise ablation

Table 4 separates each element across three scenarios and Figure 5 overlays all variants, with the altitude near steady state in Figure 5(a), the longitudinal recovery after the impulse in Figure 5(b), and the lateral response under the Dryden gust in Figure 5(c). Removing integral action more than doubles the nominal altitude error, 8.7 to 19.0 mm, and leaves it at 19.0 mm under the Dryden gust as well: the clearest single effect in the study, confirming that the integral states rather than the Riccati solution alone produce the sub-centimetre result. Removing damping leaves the nominal error unchanged but degrades the disturbed response, raising peak displacement from 36.1 to 37.1 mm and giving the least dissipated trace in Figure 5. Removing reference ramping changes neither, since the ramp shapes only the early climb and its role is confined to actuator demand at take-off.

The evidence needs two qualifications. Scenarios differ sharply in discriminating power: under the impulse the horizontal differences are small, 36.2 against 36.1 mm, because the outer loop is itself an LQI whose integral weight sits inside $\mathbf{Q}_o$ rather than in a separate clamped gain, whereas under a continuous gust at $k_g = 0.880$ N per m/s the no-integral variant raises the longitudinal RMS error from 10.9 to 13.0 mm and the no-damping variant to 12.3 mm; persistent disturbance is the condition in which the horizontal safeguards can be measured. The gust level must also be chosen deliberately: at the nominal admittance the five variants separate by less than 0.1 mm, so an ablation there would show nothing, while near the operational boundary single-run peak metrics turn erratic as the trajectory grows sensitive to each variant's phase in the gust sequence. The value used here sits between those regimes.

The anti-windup result turns on exactly where the clamp sits relative to two thresholds. With $\xi_{\max,z} = 0.60$ the clamp sits above the 0.554 the standing deficit requires but below the 0.638 the integrator reaches unaided, so it still binds during the saturated climb while leaving steady-state accuracy intact. Removing

it makes the nominal error slightly *worse*, 0.96 to 1.73 mm, because the integrator overshoots and unwinds slowly. Setting it too tight is the more damaging error: at 0.30 the nominal error rises to 8.71 mm, as section 4.1 predicts. This is a trade-off rather than a defect: final steady-state error alone is a misleading figure of merit when saturation and integral action interact, and the safeguard should be reported as a design choice with a measurable price, not a uniform improvement. Figure 6 makes the mechanism visible, with rotor commands pinned at the 18 N bound through the early climb while the integral states accumulate against their clamps.

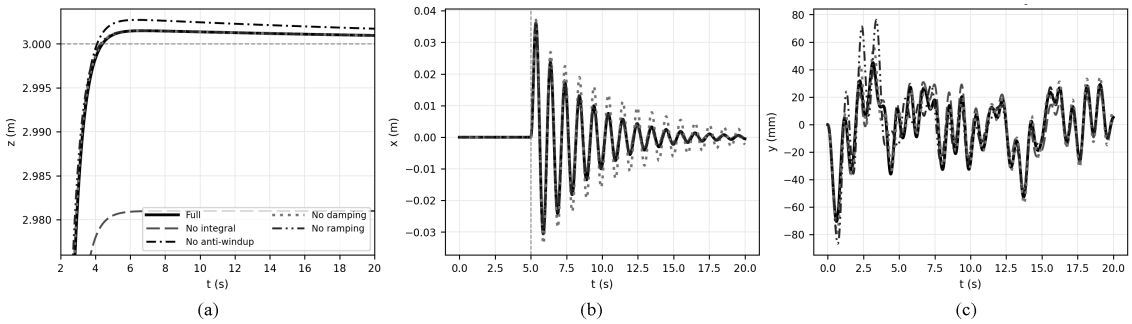

Figure 5. Ablation overlay for the full controller and the four ablated variants: (a) altitude near steady state, where the three distinct offset levels separate the variants directly; (b) longitudinal recovery after the 2 N impulse, in which the no-damping variant is visibly the least dissipated; and (c) lateral response under a Dryden gust at $k_g = 0.880$ N per m/s. Panel (a) is plotted on a millimetre scale, since on a full 0–3 m axis the five traces are indistinguishable

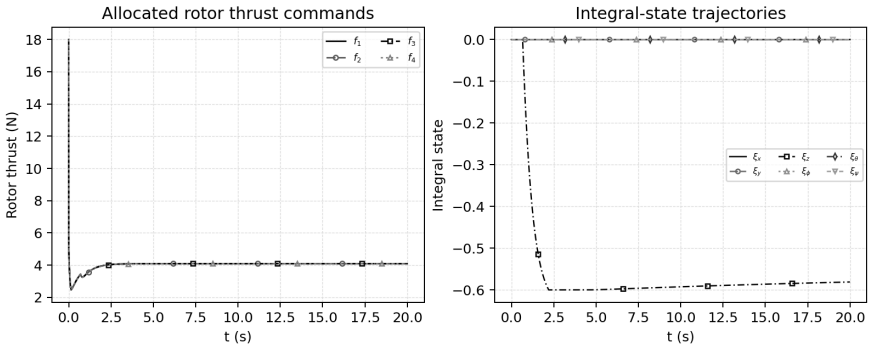

Figure 6. Allocated rotor thrust commands and the six integral-state trajectories during nominal hover, showing actuator saturation and anti-windup action directly

4.3. Dryden gust response and the operational limit

Figure 7 reports the response to the zero-mean Dryden gust at the nominal admittance of section 3.2, which is the level at which the disturbance is benign in absolute terms: measured gust RMS components of 0.60, 0.95, and 0.95 m/s for this run (the lateral and vertical values already sit above their design intensities of $\sigma_v = 0.8$ and $\sigma_w = 0.5$ m/s for the reason given in section 3.2, so this is not an especially quiet draw of the gust) give only 47 mN of equivalent force, and the ablation variants are barely distinguishable, the full controller and the no-integral variant differing by 0.10 mm in longitudinal RMS error, 0.55 against 0.65 mm, even though the no-integral altitude offset stays at 19.0 mm against 1.0 mm for the full controller. The Dryden columns of Table 4 are measured at the twenty-fold admittance of section 4.2 instead, precisely because this nominal level cannot separate the variants. Because the gust is strictly zero-mean it contains no sustained component for an integrator to remove, so integral augmentation shows its value in the altitude channel, where the residual thrust bias persists, and not in the horizontal channels. Within that figure, Figure 7(a) gives the gust velocity components and Figure 7(b) the resulting flight path, while the response of the five ablation variants appears in the longitudinal, lateral and altitude channels in Figure 7(c), Figure 7(d) and Figure 7(e), and in the roll, pitch and yaw channels in Figure 7(f), Figure 7(g) and Figure 7(h). The nominal admittance is therefore too small to constitute a stress test, and the admittance k_g is swept below until the controller actually fails.

A concrete prediction is available before any sweep is run. To hold altitude the vehicle must satisfy $T\cos\theta = mg$, so the largest horizontal force it can generate while remaining at altitude is

$$F_{h,\max} = mg\tan\theta_{\max} = 1.6 \times 9.81 \times \tan 18^\circ = 5.10\ \text{N}. \tag{9}$$

This is a static bound: it is the most the vehicle can push horizontally while still holding altitude, whatever the disturbance.

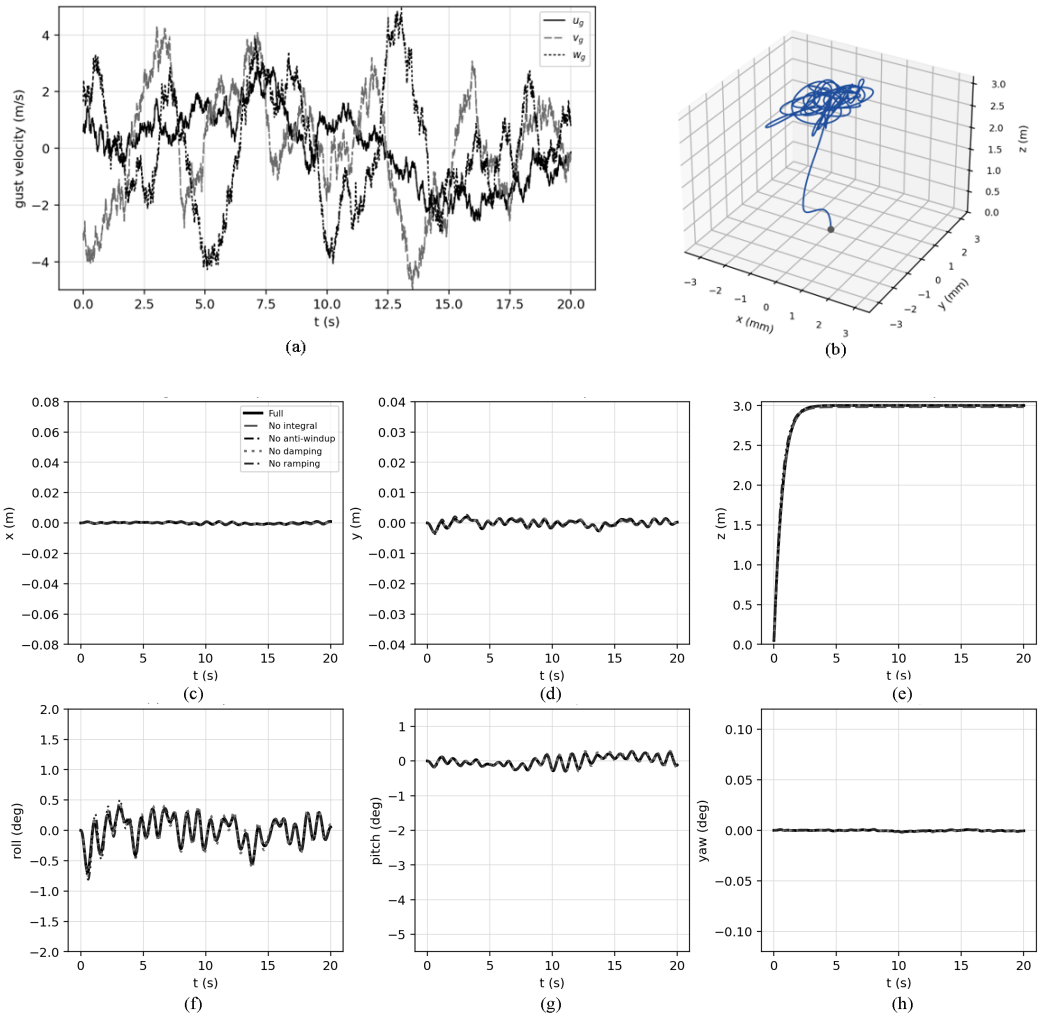

Figure 7. Zero-mean Dryden gust at nominal admittance: (a) gust velocity components, normalised to a -5 to 5 m/s scale; (b) the resulting flight path, plotted on a millimetre horizontal scale since the excursion spans only a few millimetres at this admittance; and the six-channel response of the five ablation variants, namely (c) longitudinal, (d) lateral, and (e) altitude position, and (f) roll, (g) pitch, and (h) yaw attitude

Figure 8 tests it, with the operational boundary in Figure 8(a) and the attitude-reference saturation behind it in Figure 8(b). Scaling the admittance k_g, the vehicle drifts laterally rather than longitudinally, because the Dryden lateral intensity σ_v exceeds the longitudinal σ_u, so the limiting channel is roll. Position hold is lost between $k_g = 1.76$ and 2.20 N per m/s, a peak gust force of 3.89 to 4.86 N, that is 76 to 95% of $F_{h,\max}$. Panel (b) shows the mechanism: peak roll rises monotonically until it pins at the 18° reference limit exactly where hold is lost, after which the outer loop cannot command further tilt and the residual force accelerates the vehicle without bound.

The measured boundary sits below the static bound because a zero-mean gust reverses: the outer loop must arrest accumulated momentum as well as oppose the instantaneous force, and tilt spent on overshoot is tilt unavailable for that. Reaching 76 to 95% of $F_{h,\max}$ shows the horizontal LQI spends nearly all its available

tilt opposing the disturbance. The failure is not an instability of the linear regulators: the inner loops stay well behaved and altitude regulates to within 23 mm even at the harshest setting, so what fails is position hold, not hover. The boundary is set by a saturation limit the designer chose, not by any property of the LQI synthesis, so it moves with $\theta_{\max}$ at the cost of the small-angle validity assumption (A2) depends on.

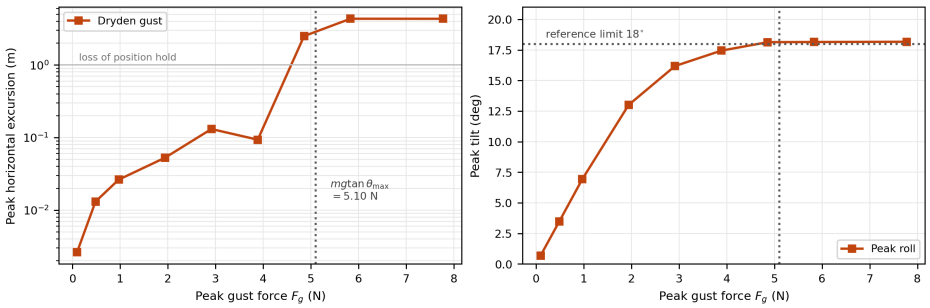

Figure 8. Operational boundary under a swept Dryden gust amplitude, plotted against the peak gust force the controller must counteract. (a) Peak horizontal excursion, with the 1 m loss-of-hold criterion and the analytical bound $F_{h,\max} = 5.10$ N. (b) Peak roll saturating at the 18° attitude-reference limit exactly where hold is lost

4.4. Sensor noise and measured lateral accuracy

Across twelve runs with the noise of section 3.2 the altitude error is 8.76 ± 8.19 mm. Throughout this section the quoted interval is the sample mean $\pm$ one sample standard deviation over the twelve runs, and the accompanying 95% confidence interval on the mean follows from Student's t distribution with eleven degrees of freedom; for the altitude error it is $[3.56,\ 13.97]$ mm. That mean is an order of magnitude larger than the 0.96 mm of the noise-free run, and the spread is comparable to the mean itself. Noise therefore does not bias regulated altitude, but it makes any single run an unreliable estimate of accuracy, which is why the result is reported as a distribution. This scenario also supplies the lateral figures the nominal case cannot: because noise breaks the symmetry the horizontal channels are genuinely excited, giving a longitudinal RMS error of 8.17 ± 0.13 mm (95% CI $[8.09,\ 8.26]$), peak longitudinal and lateral excursions of 31.4 ± 2.1 mm (95% CI $[30.1,\ 32.7]$) and 31.2 ± 3.0 mm (95% CI $[29.4,\ 33.1]$), and peak roll and pitch of 1.76° and 1.44°. The confidence interval on the altitude error is wide because that channel is dominated by where in its cycle the integrator happens to sit when the run ends, whereas the RMS and peak horizontal metrics average over the whole window and are correspondingly tight. These are the values to quote for lateral accuracy under realistic feedback, in place of the 10^{-18} figures of Table 3, which reflect only the numerical floor of a symmetric noise-free run. Tightening the altitude integral clamp from $\xi_{\max,z} = 0.60$ to 0.30 under this same twelve-run noise condition reduces the altitude overshoot from 32.0 ± 1.5 to 23.1 mm but narrows the margin against the steady-state bias identified in section 4.1, illustrating that the clamp setting trades noise rejection against bias rejection rather than improving on both.

4.5. Computational cost, robustness margins, and parametric sensitivity

One evaluation of the complete control law—both outer loops, four inner loops, six clamped integrators, and the allocation—takes $19.1\,\mu$s on an Apple M4 under Python 3.13 with NumPy 2.1, averaged over 2×10^5 repetitions: 0.46% of the $4167\,\mu$s available at $\Delta t = 1/240$ s, a real-time factor of 218, reached in interpreted Python with no compilation or vectorisation effort. The cost is dominated by six scalar integrator updates and four 2×1 inner products, and is independent of state, since no iteration, optimisation, or model propagation occurs online. This is the practical argument for the linear design: an adaptive predictive scheme of the kind in [30] must solve an optimisation every step, and a learned policy [28] must evaluate a network, neither of which reduces to twenty constants a reviewer can check by hand.

Robustness is summarised in Table 5, computed on the zero-order-hold discretisation at Δt rather than on the continuous loop, because the roll channel of the *design* model crosses over near 978 rad/s while the Nyquist rate is 754 rad/s, so a continuous-time margin there would describe a loop the sampled implementation cannot realise. Evaluated against the actual SDF plant every channel is stable, the largest closed-loop pole magnitude being 0.99993, and the smallest modulus margin is 0.591 on pitch, which alone guarantees at least

7.8 dB of gain margin and 34.4° of phase margin; the smallest classical phase margin is 65.3° and every gain crossover lies below Nyquist.

Table 5. Discrete-time robustness per channel at $\Delta t = 1/240$ s, against the actual SDF plant and against the design model. ρ_{cl} is the largest closed-loop pole magnitude, M_S the peak sensitivity, $1/M_S$ the modulus margin, and ω_c the gain crossover

	Actual SDF plant			Design model		
Channel	ρ_{cl}	$1/M_S$	ω_c (rad/s)	ρ_{cl}	$1/M_S$	ω_c (rad/s)
Altitude z	0.99985	0.971	14.1	0.99985	0.971	14.1
Roll ϕ	0.99993	0.892	52.0	3.07504	1.039	above Nyquist
Pitch θ	0.99993	0.591	202.5	0.99993	0.588	204.0
Yaw ψ	0.99971	0.916	40.3	0.99971	0.767	113.1
Outer x, y	0.99807	0.967	15.8	0.99807	0.967	15.8

The same table answers the sensitivity question directly, because the inertia mismatch of Table 2 is itself a large parametric perturbation applied to the real loop: the controller is synthesised for $I_x = 0.006\,\mathrm{kg\,m^2}$ while the plant has 0.1133, a factor of 18.9, with yaw mismatched by 2.78 and pitch matched to 1%. Repeating the calculation on the design model shows the mismatch is not merely tolerated but stabilising. Against the design inertia the roll loop would be unstable once sampled, with a closed-loop pole magnitude of 3.08, because the synthesis places the crossover above Nyquist; the eighteen-fold heavier real airframe slows that loop to a 52.0 rad/s crossover and a modulus margin of 0.892. A designer transferring these gains to an airframe whose roll inertia is closer to the design value must therefore raise the control rate or re-solve the Riccati equation with the true inertia—the one place where the reported gain set is not safe to reuse unchanged, and a failure mode invisible in every time trace here.

4.6. Comparison with representative literature

Table 6 places the work against representative tail-sitter and related VTOL hover studies. The comparison is cautious, since these papers use different vehicles, models, and environments, but the studies closest in scope report hover behaviour under wind or perturbation only at a nominal level: goodness-of-fit under hover perturbation [24], wind estimation validated against flight data [25], and agile tracking under wind with no stated failure amplitude [26]. The present work adds a fully specified controller, a verified allocation matrix, and an operational boundary predicted from the design parameters and then confirmed empirically. Against the nonlinear and adaptive strands the comparison is a trade rather than a ranking, and not a controlled experiment: those studies use other airframes and disturbance definitions, so only the shape of the trade transfers. Integral backstepping sliding-mode control [14] and robust H_∞ synthesis [16] buy guaranteed performance across an envelope this linear design cannot claim, and adaptive schemes [29], [30] additionally track parameter drift that fixed gains cannot follow. What they give up is the property measured in this section: a control law that is twenty auditable constants, carries the per-step cost of section 4.5 with no online optimisation, and whose failure amplitude can be predicted in closed form before any simulation is run. For hover-only regulation about a fixed operating point that trade favours the linear design; for transition, where the operating point moves and the small-angle model fails, it does not.

What the measured numbers mean for deployment can now be stated concretely, which is the level at which a robotics or automation practitioner selects a platform. Infrastructure inspection is the most demanding of the three target missions, because a camera or lidar held at a standoff of a few metres converts vehicle motion directly into ground-sample error: the 0.96 mm noise-free altitude hold and the 31.4 ± 2.1 mm horizontal excursion measured under realistic feedback are both far below the centimetre-scale features such surveys look for, so the limiting factor is the sensor rather than the airframe. Delivery imposes the opposite requirement, a bounded terminal transient rather than sustained precision, and the 36.1 mm peak displacement with re-entry to a 10% band in 8.0 s after a 2 N impulse indicates that a package release or a docking contact of that magnitude is absorbed without loss of hold. Mapping sits between the two and is governed by the gust boundary rather than by either: position hold survives to 3.89–4.86 N of gust force, which at the nominal admittance of section 3.2 corresponds to the benign conditions in which photogrammetric flights are normally scheduled, and degrades into drift rather than into loss of attitude, so a mission abort is triggered by a position-error monitor rather than by a crash. Across all three the automation-relevant conclusion is the same: the quantity that fails first is

the pose the payload is referenced to, not the vehicle, and it fails at an amplitude the designer can compute in closed form before flying.

Table 6. Comparison against representative tail-sitter and VTOL hover-control literature

Reference	Controller focus	Integral action	Disturb. test	Reported quantitative hover metric
[1]	Optimal transition control	no	no	Transition emphasis; hover metric not reported
[24]	Hover stability, flight validated	no	yes	fit >93% pitch/roll, 79.8% yaw under hover perturbation
[25]	Wind estimation in hover	no	yes	Wind estimate validated on flight data; estimation, not regulation
[26]	Agile tracking under wind	yes	yes	Tracking error under wind; no failure amplitude stated
[14]	Integral backstepping sliding mode	yes	yes	Quadrotor trajectory tracking; hover metric not reported
This work	Full-state LQI	yes*	yes*	0.96 mm altitude error, 36.1 mm impulse recovery settling in 8.0 s, boundary at 76–95% of $mg\tan\theta_{\max} = 5.10$ N

*Integral action is applied to all six regulated outputs (position, altitude, and attitude), and the disturbance test spans three classes (impulsive, zero-mean Dryden gust, and sensor noise)

5. CONCLUSION

This paper presented a full-state LQI hover controller for a quad tail-sitter UAV, derived from a Newton-Euler hover model and evaluated in PyBullet. Beyond the metrics, the practical contribution is an auditable baseline: every plant parameter, controller weight, integral limit, allocation coefficient, and disturbance recursion is reported, so the design can be reproduced, criticised, or extended without contacting the authors. This matters because hover controllers are rarely deployed alone, forming the innermost layer of transition and cruise controllers, of payload stabilisation for inspection and mapping, and of the sensor-pointing loops that set mission data quality.

The clamp setting adopted here favours steady-state accuracy over noise rejection, a trade section 4.4 quantifies; an application dominated by noise rather than by bias would justify the opposite choice, using the same budget with the inequality reversed.

Eight limitations bound these results. The study is simulation-only, and assumptions (A4) and (A5) are the most likely to fail on a real airframe. The plant is a rigid body with linear damping, so rotor inflow, wing-surface loading, ground effect, and propeller interaction are absent. The small-angle assumption is not expected to hold through transition, and the 18° limit setting the operational boundary is itself near that edge, so the envelope cannot simply be widened by raising it. The Dryden gust couples through the equivalent admittance defined in section 3.2 rather than a lift-and-drag model, which is why section 4.3 sweeps it. The outer-loop synthesis treats tilt as instantaneous, leaving the attitude dynamics that deliver it unmodelled, and the residual oscillation in the impulse recovery follows from that omission. Rotor actuation is likewise ideal: commanded thrust is applied within the same step, so motor and ESC lag, rate limits, and thrust-response dynamics are absent; since such a lag would consume the measured phase margin—at the pitch crossover of 202.5 rad/s the 34.4° guarantee corresponds to roughly 3 ms of tolerable delay, the order of a small brushless rotor's rise time—this is the omission most likely to change Table 5 on hardware. Those margins are computed per channel on the decentralised loops, so they do not capture simultaneous perturbations in several channels. Finally, apart from the sensor-noise case each scenario is a single deterministic run at a fixed seed.

Three practical risks follow from those measurements. Actuator saturation is not an edge case but the normal condition of the climb, as the thrust figures of section 4.1 show: the allocation saturates and the anti-windup clamps bind in every run, the nominal one included, so a designer who sizes the clamp without allowing for that inherits the 8.71 mm bias rather than the 0.96 mm reported here. Aerodynamic mismatch is represented only by the retained inertia error, and its sign matters: Section 4.5 shows the same gains would be unstable once sampled on an airframe matching the design inertia. Sensor noise does not bias regulated altitude but destroys the precision of any single run, widening it to the interval of section 4.4, so a single-flight acceptance test would be uninformative. The same transparency has a pedagogical consequence: because the

controller reduces to twenty constants, a Riccati solution verifiable by inspection, and closed-form disturbance recursions, it is usable as a teaching case in UAV control courses, where a student can re-derive every gain, reproduce every figure, and then break the controller deliberately through the ablation of section 4.2 to see what each safeguard is for.

The most transferable finding is structural rather than numerical: the envelope is fixed by attitude-reference saturation in the outer loop, at a horizontal force of $mg \tan \theta_{\max}$, and the vehicle fails there rather than through any deficiency of the linear-quadratic synthesis. A designer needing a larger envelope must therefore revisit that limit and the small-angle model behind it, not the weighting matrices, and the measured boundary reaching 76 to 95% of the static bound shows little is left to recover by horizontal tuning alone. Scaling the architecture to a different rotor count is a question of allocation rather than of synthesis. For N rotors the mixing matrix becomes $\mathbf{M} \in \mathbb{R}^{4\times N}$, and the four inner loops and two outer loops are unchanged provided $\mathbf{M}$ retains rank four, which is the formal statement that the airframe still has independent authority over thrust and the three moments. For $N > 4$ the inverse used here is replaced by a pseudoinverse or, once the per-rotor bounds imposed by the actuator limits are enforced, by a constrained allocation, and the conditioning figure quoted in section 2.3 should be recomputed because it governs how much of the commanded moment survives saturation. The gains themselves do not transfer: every entry of the controller gain matrices follows from the mass and inertia of this airframe through the Riccati solutions of section 2.2, so a new platform requires those solutions to be re-run rather than the weights to be re-tuned, and section 4.5 shows why this matters, since a lighter roll axis moves the loop crossover toward Nyquist. Future work should pursue hardware-in-the-loop validation and flight testing; extension to transition and cruise via gain scheduling or a linear-parameter-varying formulation; a coupled position-attitude synthesis; and robustness testing against validated turbulence spectra.

ACKNOWLEDGMENTS

The authors would like to express their gratitude to the Department of Electrical Engineering and Informatics at Universitas Gadjah Mada for providing the necessary resources and support to conduct this research, and acknowledge the support of the Vocational School of Universitas Gadjah Mada through the 2025 Dana Masyarakat research scheme. We are also grateful to our colleagues and reviewers for their valuable feedback and insights that significantly improved the quality of this work.

FUNDING INFORMATION

This work was supported by the Vocational School of Universitas Gadjah Mada under the 2025 Dana Masyarakat research scheme.

AUTHOR CONTRIBUTIONS STATEMENT

This journal uses the Contributor Roles Taxonomy (CRediT) to recognize individual author contributions, reduce authorship disputes, and facilitate collaboration.

Name of Author	C	M	So	Va	Fo	I	R	D	O	E	Vi	Su	P	Fu
Ariesta Martiningtyas Handayani	✓	✓	✓		✓	✓		✓	✓				✓	✓
Isnan Nur Rifai		✓	✓	✓		✓				✓	✓			
Andi Dharmawan	✓				✓					✓		✓		

C : **C**onceptualization
M : **M**ethodology
So : **So**ftware
Va : **Va**lidation
Fo : **Fo**rmal Analysis
I : **I**nvestigation
R : **R**esources
D : **D**ata Curation
O : Writing - **O**riginal Draft
E : Writing - Review & **E**diting
Vi : **Vi**sualization
Su : **Su**pervision
P : **P**roject Administration
Fu : **Fu**nding Acquisition

CONFLICT OF INTEREST STATEMENT

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

INFORMED CONSENT

Not applicable.

ETHICAL APPROVAL

Not applicable.

DATA AVAILABILITY

The authors confirm that the data supporting the findings of this study are available within the article. Additional derived simulation data are available from the corresponding author upon reasonable request.

REFERENCES

[1] A. Oosedo, S. Abiko, A. Konno, and M. Uchiyama, "Optimal transition from hovering to level-flight of a quadrotor tail-sitter UAV," *Autonomous Robots*, vol. 41, no. 5, pp. 1143–1159, Jun. 2017, doi: 10.1007/s10514-016-9599-4.

[2] A. Oosedo, S. Abiko, A. Konno, T. Koizumi, T. Furui, and M. Uchiyama, "Development of a quad rotor tail-sitter VTOL UAV without control surfaces and experimental verification," in *2013 IEEE international conference on robotics and automation*, 2013, pp. 317–322, doi: 10.1109/ICRA.2013.6630594.

[3] Y. Wang, X. Lyu, H. Gu, S. Shen, Z. Li, and F. Zhang, "Design, implementation and verification of a quadrotor tail-sitter VTOL UAV," in *2017 International conference on unmanned aircraft systems (ICUAS)*, 2017, pp. 462–471, doi: 10.1109/ICUAS.2017.7991419.

[4] M. Hochstenbach, C. Notteboom, B. Theys, and J. De Schutter, "Design and control of an unmanned aerial vehicle for autonomous parcel delivery with transition from vertical take-off to forward flight: VertiKUL, a quadcopter tailsitter," *International Journal of Micro Air Vehicles*, vol. 7, no. 4, pp. 395–405, Dec. 2015, doi: 10.1260/1756-8293.7.4.395.

[5] H. Gu, X. Lyu, Z. Li, S. Shen, and F. Zhang, "Development and experimental verification of a hybrid vertical take-off and landing (VTOL) unmanned aerial vehicle (UAV)," in *2017 International Conference on Unmanned Aircraft Systems (ICUAS)*, 2017, pp. 160–169, doi: 10.1109/ICUAS.2017.7991420.

[6] R. Bapst, R. Ritz, L. Meier, and M. Pollefeys, "Design and implementation of an unmanned tail-sitter," in *2015 IEEE/RSJ international conference on intelligent robots and systems (IROS)*, 2015, pp. 1885–1890, doi: 10.1109/IROS.2015.7353624.

[7] K. Z. Y. Ang *et al.*, "Development of an unmanned tail-sitter with reconfigurable wings: U-Lion," in *11th IEEE International Conference on Control & Automation (ICCA)*, 2014, pp. 750–755, doi: 10.1109/ICCA.2014.6871015.

[8] S. Yoshinaga, A. Sakaguchi, and K. Yamamoto, "A novel tailsitter UAV with configurable wings," *Journal of Robotics and Mechatronics*, vol. 36, no. 5, pp. 1037–1042, Oct. 2024, doi: 10.20965/jrm.2024.p1037.

[9] J. Zhong, C. Wang, and H. Zhang, "Transition control of a tail-sitter unmanned aerial vehicle with L_1 neural network adaptive control," *Chinese Journal of Aeronautics*, vol. 36, no. 7, pp. 460–475, Jul. 2023, doi: 10.1016/j.cja.2023.04.002.

[10] Y. Ke, K. Wang, K. Gong, S. Lai, and B. M. Chen, "Model based robust forward transition control for tail-sitter hybrid unmanned aerial vehicles," in *2017 13th IEEE International Conference on Control & Automation (ICCA)*, 2017, pp. 828–833, doi: 10.1109/ICCA.2017.8003167.

[11] D. Kubo and S. Suzuki, "Tail-sitter vertical takeoff and landing unmanned aerial vehicle: transitional flight analysis," *Journal of Aircraft*, vol. 45, no. 1, pp. 292–297, Jan. 2008, doi: 10.2514/1.30122.

[12] S. Q. Sohail, F. Akram, N. Hussain, and A. Shahzad, "Design and transition of a quad rotor tail-sitter VTOL UAV with experimental verification," in *2021 International Bhurban Conference on Applied Sciences and Technologies (IBCAST)*, 2021, pp. 244–251, doi: 10.1109/IBCAST51254.2021.9393187.

[13] T. K. Priyambodo, A. Majid, and Z. S. S. Shouran, "Validation of quad tail-sitter VTOL UAV model in fixed wing mode," *Journal of Robotics and Control (JRC)*, vol. 4, no. 2, pp. 179–191, Mar. 2023, doi: 10.18196/jrc.v4i2.17253.

[14] S. H. Tsai, Y. P. Chang, H. Y. Lin, and L. M. Chang, "Design and implementation of integral backstepping sliding mode control for quadrotor trajectory tracking," *Processes*, vol. 9, no. 11, art. no. 1951, Nov. 2021, doi: 10.3390/pr9111951.

[15] N. Cheng and C. Wang, "A switching mode control scheme for the hovering control of quadrotor unmanned aerial vehicles," *Mathematics*, vol. 11, no. 4, art. no. 994, Feb. 2023, doi: 10.3390/math11040994.

[16] N. T. Hegde, V. I. George, C. G. Nayak, and A. C. Vaz, "Application of robust H_∞ controller in transition flight modeling of autonomous VTOL convertible quad tiltrotor UAV," *International Journal of Intelligent Unmanned Systems*, vol. 9, no. 3, pp. 204–235, Jun. 2021, doi: 10.1108/IJIUS-09-2020-0041.

[17] Z. Jinlong *et al.*, "Control design of the quadrotor aircraft based on the integral adaptive improved integral backstepping sliding mode scheme," *Engineering, Technology & Applied Science Research*, vol. 14, no. 5, pp. 17106–17117, Oct. 2024, doi: 10.48084/etasr.8361.

[18] F. Thomas and M. S. Johnson, "Composite nonlinear H_∞ based output feedback controller for an un-crewed helicopter in its hover flight," *Aircraft Engineering and Aerospace Technology*, vol. 93, no. 1, pp. 159–170, Feb. 2021, doi: 10.1108/AEAT-01-2020-0009.

[19] O. Saleem, A. Tanveer, and J. Iqbal, "Adaptive reconfigurable learning algorithm for robust optimal longitudinal motion control of unmanned aerial vehicles," *Algorithms*, vol. 18, no. 4, art. no. 180, Apr. 2025, doi: 10.3390/a18040180.

[20] N. Raj, R. Banavar, Abhishek, and M. Kothari, "Attitude control of a novel tailsitter: swiveling biplane-quadrotor," *arXiv:1907.08587*, Jul. 2019.

[21] L. Bauersfeld, L. Spannagl, G. J. J. Ducard, and C. H. Onder, "MPC flight control for a tilt-rotor VTOL aircraft," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 57, no. 4, pp. 2395–2409, Aug. 2021, doi: 10.1109/TAES.2021.3061819.

[22] Z. Jinlong, W. Jianhong, W. Ruchun, D. Yongjun, and L. Xi, "Control design of quadrotor aircraft based on improved integral backstepping sliding mode," *Asian Journal of Control*, vol. 27, no. 1, pp. 427–446, Jan. 2025, doi: 10.1002/asjc.3444.

[23] T. Matsumoto *et al.*, "A hovering control strategy for a tail-sitter VTOL UAV that increases stability against large disturbance," in *2010 IEEE international conference on robotics and automation*, 2010, pp. 54–59, doi: 10.1109/ROBOT.2010.5509183.
[24] H. Wang, S. Su, X. Qiu, Y. Liang, P. Yu, and X. Shan, "The hovering stability of the Egretta tail-sitter VTOL UAV," *International Journal of Aerospace Engineering*, vol. 2022, art. no. 9534180, Aug. 2022, doi: 10.1155/2022/9534180.
[25] Y. Demitri, S. Verling, T. Stastny, A. Melzer, and R. Siegwart, "Model-based wind estimation for a hovering VTOL tailsitter UAV," in *IEEE International Conference on Robotics and Automation*, 2017, pp. 3945-3952, doi: 10.1109/ICRA.2017.7989455.
[26] X. Zou, Z. Liu, Z. Jia, and B. Wang, "Accurate tracking of agile trajectories for a tail-sitter UAV under wind disturbances environments," *Drones*, vol. 9, no. 2, art. no. 83, Feb. 2025, doi: 10.3390/drones9020083.
[27] A. H. Mhmood and H. I. Ali, "Optimal H_∞ PID model reference controller design for roll control of a tail-sitter VTOL UAV," *Engineering and Technology Journal*, vol. 39, no. 4A, pp. 552–564, Apr. 2021, doi: 10.30684/etj.v39i4a.1861.
[28] S. Sönmez, *et al.*, "Reinforcement learning-based PD controller gains prediction for quadrotor UAVs," *Drones*, vol. 9, no. 8, art. no. 581, Aug. 2025, doi: 10.3390/drones9080581.
[29] A. R. Ajel, A. J. Humaidi, I. K. Ibraheem, and A. T. Azar, "Robust model reference adaptive control for tail-sitter VTOL aircraft," *Actuators*, vol. 10, no. 7, art. no. 162, Jul. 2021, doi: 10.3390/act10070162.
[30] M. Osman, *et al.*, "An adaptive SINDy-Lyapunov model predictive control framework for dual-system VTOL UAVs," *International Journal of Robust and Nonlinear Control*, vol. 36, no. 5, pp. 2388–2417, Nov. 2025, doi: 10.1002/rnc.70272.
[31] K. Ogata, *Modern Control Engineering*, 5th ed. Upper Saddle River, NJ, USA: Prentice-Hall, 2010.

BIOGRAPHIES OF AUTHORS



Ariesta Martiningtyas Handayani received the B.Sc. degree in Electronics and Instrumentation in 2010, the M.Sc. degree in Computer Science in 2015, and the Ph.D. degree in Computer Science in 2026, all from Universitas Gadjah Mada, Yogyakarta, Indonesia. Since 2011, she has been with the Department of Electrical Engineering and Informatics, Vocational School, Universitas Gadjah Mada, where she is currently a faculty member. Her research interests include robotic dynamics modeling, autonomous unmanned systems, and control systems. Her published work centres on linear-quadratic control for unmanned aerial vehicles across several airframe classes: flight-trajectory control for fixed-wing UAVs, an automatic vertical take-off and landing system for a quadrotor combining a Madgwick quaternion Kalman filter with a linear quadratic regulator, hover attitude control for a quad tilt-wing vehicle, and an autonomous take-off system with stall prevention for fixed-wing aircraft. She has also published on image-processing applications in electronics and instrumentation. The present paper continues that line of work, extending the linear-quadratic approach to a quad tailsitter and adding Riccati-derived integral action on every regulated output. She can be contacted at email: ariesta_mh@mail.ugm.ac.id.



Isnan Nur Rifai received the B.Sc. degree in Electronics and Instrumentation and the M.Eng. degree in Electrical Engineering from Universitas Gadjah Mada in 2007 and 2012, respectively, and the Ph.D. degree in Mechanical Engineering from Chiba University, Japan, in 2023. Since 2011, he has been a lecturer and researcher with the Department of Electrical Engineering and Informatics, Vocational School, Universitas Gadjah Mada. His research interests include control systems, electronic systems, and instrumentation, particularly the measurement chains that determine how well a control law designed in simulation transfers to physical hardware. His published work is concentrated in electrical impedance measurement and front-end instrumentation, including a field-programmable gate array based planar-sensor electrical impedance tomography system, square-wave open electrical impedance tomography for imaging sodium concentration in the dermis layer, and adaptive signal conditioning based on a programmable gain amplifier for multi-frequency bioimpedance measurement. That instrumentation background informs the treatment of sensor noise and measurement realism in the present study. He can be contacted at email: isnan.nur@ugm.ac.id.



Andi Dharmawan received the B.Sc. degree in Electronics and Instrumentation in 2006, the M.Sc. degree in Computer Science in 2009, and the Ph.D. degree in Computer Science in 2017, all from Universitas Gadjah Mada, Yogyakarta, Indonesia. He is currently an Associate Professor with the Department of Computer Science and Electronics, Universitas Gadjah Mada. His research interests include robotics control systems, autonomous unmanned systems, and automation. His published work spans several classes of autonomous platform and their supporting computation, including balance and walking control for humanoid robots, an autonomous underwater vehicle for subaquatic operation, camera-based simultaneous localization and mapping, LiDAR-based sensor fusion for indoor autonomous navigation, and field-programmable gate array architectures for efficient onboard numerical computation. Work on estimation and navigation for unmanned platforms connects directly to the hover regulation problem treated here, since the pose held by the controller is the reference from which such estimators measure. He can be contacted at email: andi_dharmawan@ugm.ac.id.