SSC26-P2-54

Generalized Attitude Control for Small Spacecraft

Patrick McKeen
Department of Aeronautics and Astronautics, Massachusetts Institute of Technology
77 Massachusetts Avenue, Cambridge, MA 02139; 336-671-3401
pmckeen@mit.edu

Niclas Scheuer
Department of Mechanical and Process Engineering, ETH Zürich
Rämistrasse 101, 8092 Zürich, Switzerland; +41 78 230 43 50
nscheuer@ethz.ch

Kerri Cahoy
Department of Aeronautics and Astronautics, Massachusetts Institute of Technology
77 Massachusetts Avenue, Cambridge, MA 02139
kcahoy@mit.edu

ABSTRACT

Many classical attitude-control designs assume unconstrained torque or are developed for a fixed actuator set, requiring redesign for each new spacecraft and its goals. This unfortunately discourages exploring alternatives even when better-suited options exist. We present a framework that treats the control law as a black box and places goal formulation, compensation, torque allocation, and momentum management in a common adapter around it, so a law from the literature can be adapted across heterogeneous and underactuated actuator sets without reimplementation.

Before any simulation, a controllability analysis screens whether a given actuator set can meet a goal. Controllability is set by the achievable torque envelope, and wheel desaturation by where that envelope meets the wheel-momentum directions. Over a full orbit the geomagnetic field sweeps the magnetorquer torque envelope through all three axes, so magnetorquers alone are controllable for full attitude in the orbit-averaged sense (though one-orbit closed-loop convergence for that case remains slow). Reduced-attitude objectives such as camera-boresight pointing relax the requirement further.

Torque allocation under actuator bounds uses linear (LP) or quadratic programming (QP). If a reaction wheel is present, the direction-preserving LP is the better default, as it preserves the wrapped law's stability assumptions. The least-squares QP performs better on a magnetorquer-only bus.

The framework is validated in 100-trial Monte Carlo campaigns (randomized attitudes, body rates, orbits, and pointing goals) across configurations from magnetorquer-only to three-wheel, including the 3MTQ+1RW (three magnetorquers and one reaction wheel) hybrid. Each campaign isolates one variable: one law across actuator sets, different laws through one interface, gain scaling across inertias, and a reaction-wheel failure absorbed by reallocation. Compensation and momentum management are built into the adapter rather than the control law. A single law in this framework delivers useful three-axis pointing across these architectures, where a fixed design applied outside its intended configuration fails to converge. Because goal, law, and allocation are separated, a shortfall can be charged to hardware geometry, control law, or allocator strategy rather than to the controller as a whole.

I. INTRODUCTION

Many satellites need to control their attitude to meet their mission goals, whether that means pointing an imager at a target, a radio at a ground station, or solar panels towards the sun; or holding a precise three-axis orientation for a maneuver or for inter-satellite alignment. While critical for most missions, ADCS is considered a difficult and arcane area of work. Development of a working ADCS can be painful, especially for resource-constrained teams. Even with the wealth of existing research and many control options available in literature [1, 2, 4, 8], catalogued in standard texts [15–19] and, for magnetic actuation, in dedicated surveys [13, 14], differences in experimental setup mean that evaluating an alternative architecture requires reimplementation, not just reconfiguration. Many published control laws and other ADCS components assume specific hardware, actuators, orbits, goals, and more. Translating that research to your spacecraft's particulars can take days or weeks of custom development.

This makes it hard to compare published control laws, or to benchmark a controller under development against them. A controller tuned for a specific actuator configuration often fails to converge when applied outside that

configuration. Even before tuning, the practitioner faces a wall of mismatched interfaces. The published law may, for instance:

- expect a quaternion error when the goal is a boresight direction;
- command magnetic dipoles when the spacecraft carries reaction wheels;
- assume disturbances the new environment lacks, or omit ones it has;
- carry gains tuned for a different inertia and actuator scale;
- leave wheel-momentum regulation unspecified; or
- rest on a stability guarantee that depends on a disturbance the new orbit lacks (e.g., gravity gradient).

Replicating the law on the new hardware, then tuning it to the new mission, can absorb weeks of effort per controller. A modular pipeline for adapting control laws can collapse most of this work into configuration changes. We present such an architecture. It wraps any control law as a black box and handles the translations between "desired torque" and "your actual actuators," between "your goal" and "the goal for which the law was designed," and between "the disturbances you need to counter" and "those included originally." For a given controller, this adapts it to the specific use case with minimal hand work. There are limits. It cannot deliver perfect control on an underactuated platform, and some gain tuning may still be needed (though we reduce that where we can). With this reduced friction, ADCS development shifts from a per-mission reimplementation project to a per-mission configuration exercise, and small teams can evaluate more options in the early-design phase.

The paper's primary contribution is this adapter architecture: a common wrapper that pulls goal formulation, compensation, allocation, and momentum management out of the control law, so the law need not know the spacecraft it runs on. Because each concern lives in its own stage, the pipeline also separates the reasons a controller can underperform so a shortfall can be charged to the right layer rather than blamed wholesale on "the hardware." Two pieces of theory support it, a task-dependent controllability test that screens a configuration–goal pairing before any simulation (Section III), and the conditions under which the wrapper preserves a base law's stability certificate (Section V-A). The evidence is a sequence of paired experiments that each hold one element fixed and vary another: the same law across actuator sets, different laws through one interface, the same law rescaled across inertias, and the same law after an actuator failure (Section IV).

Table 1: Notation used throughout.

Symbol	Meaning
q_g, q_e	goal / error quaternion
σ, $\omega \in \mathbb{R}^3$	attitude-error (MRP) and body rate
$\hat{b}$, $\hat{u}$	body vector / world target (vector goal)
P, $P_\perp$	error / transverse-plane projection
τ_{ref}, τ_{comp}	control-law / compensated torque
$B_\tau(t)$	torque effectiveness matrix
$\mathbf{u} \in \mathscr{U}$	actuator command / feasible box
α, $\bar{\alpha}$	LP scale factor / trial mean
z	required controllable subspace
$\mathscr{T}(\mathbf{B}_b)$	achievable torque set
h	wheel momentum
$\mathbf{B}_b(t)$	body-frame magnetic field

This work utilizes the authors' prior development of generalized small-satellite ADCS package which was used for testing, and the methods described are included in the package from those works [10, 11]. Section II lays out the pipeline stage by stage: the goal reframing that lets a team demand only the pointing it actually needs (Subsection II-B), the black-box control-law interface (Subsection II-C), the compensation block for gyroscopic, disturbance, and frame-rotation feedforward (Subsection II-D), the allocation methods (primarily LP and QP) that map a desired torque to arbitrary actuators with a stability argument (Subsection II-E), and momentum management folded into allocation rather than run as a cascaded loop (Subsection II-F). Section III gives the controllability analysis across goal and actuator configurations; Section IV validates the framework across control laws, actuator configurations, goals, and failure scenarios; and Section V covers practical considerations, limitations, and future work.

II. FRAMEWORK ARCHITECTURE

II-A. The Pipeline

Figure 1 shows the pipeline. The stages may read as involved, but each reduces to a small matrix operation selected by a short decision tree, so the whole runs comfortably in real time.

The pipeline has five stages. **(1) Goal formulation** converts between full-attitude and reduced (vector-alignment) goals and produces the error signal for the law; as implemented in the authors' open-source framework [11][1] it also accepts named goals ("nadir," "sun-

[1] https://github.com/nscheuer/Generalized_ADCS

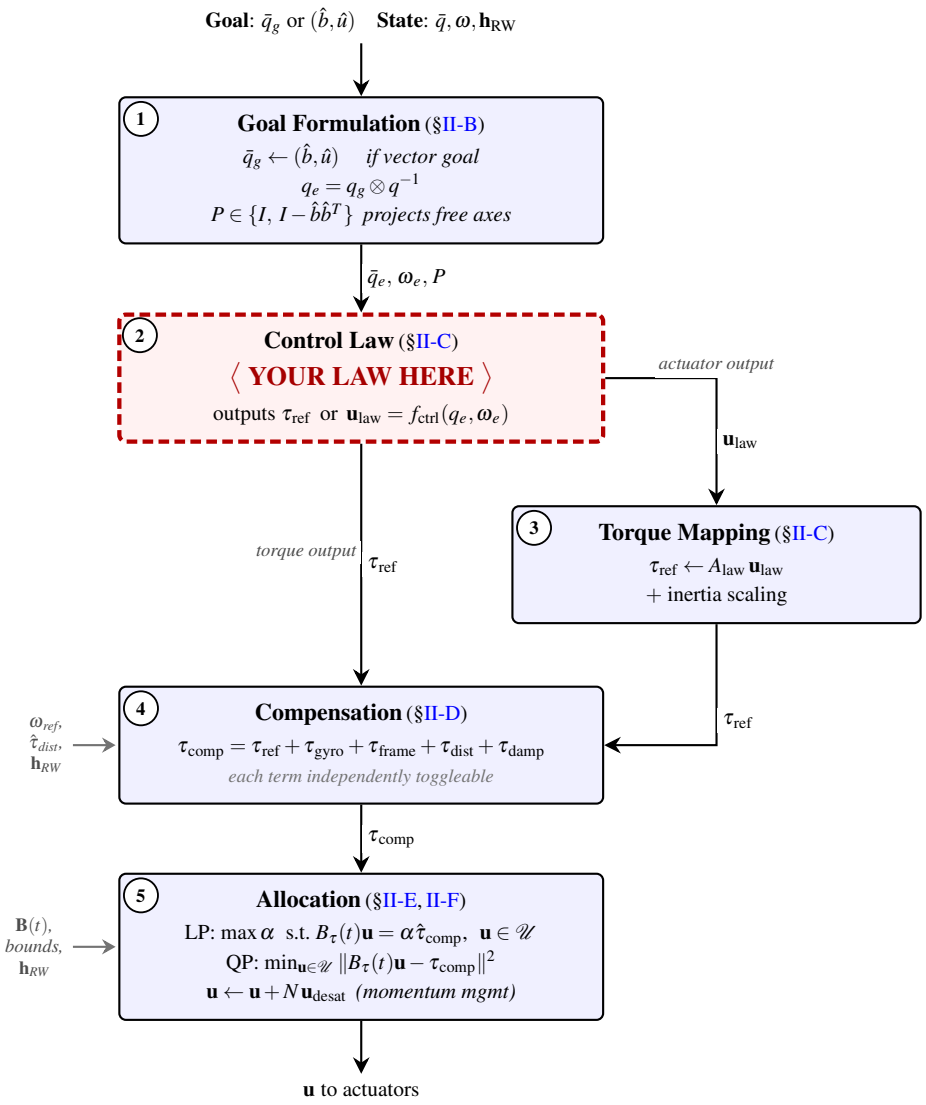


Figure 1: The generalized control pipeline. A user-supplied control law (Stage 2, dashed) sits inside a modular adapter: goal formulation (Stage 1), optional actuator-command-to-torque conversion (Stage 3), feedforward compensation (Stage 4), and bound-respecting allocation with implicit momentum management (Stage 5). Side inputs supply state estimates, field measurements, and actuator limits to the stages that need them. Only Stage 2 changes when the control law is swapped, and only Stage 5 when the allocator or actuator set changes.

ward," a ground location) and any required tracking rate. **(2) The control law** is a black box, inserted with at most minor reformulation. **(3) Torque mapping** converts an actuator-command output back to a desired torque (with inertia scaling) when the law does not already emit one. **(4) Compensation** adds independently toggleable feedforward terms (disturbance, gyroscopic, and frame-rotation) around the law, dropping any the law already handles to avoid double-correction. **(5) Allocation** maps the desired torque to bounded actuator commands (LP, QP, clipping, or other options such as QP with an energy constraint) and folds momentum desaturation into the same step rather than a cascaded loop. These steps are detailed in the subsections below.

II-B. Goal Formulation

The first stage of the pipeline translates the user's pointing objective into a standardized error signal for the control law. This is perhaps the most involved step, because while nearly all control laws produce similar outputs (a desired torque or actuator command), their *inputs* vary widely. Some laws expect a full three-axis quaternion error; others work natively with vector-alignment errors. Some accept angular velocity feedback; others do not. The goal formulation block resolves these mismatches so that goals can be paired with nearly any law.[2]

II-B.1 Goal Types and Normalization

User-specified goals are normalized into one of three canonical types (shown here as constant, though each may be time-varying—e.g., a vector tracking a ground station during a pass, or a quaternion holding the $+Z$ axis at zenith and $+X$ along ram):

- **Full attitude**: A desired orientation $\bar{q}_g \in \mathbb{S}^3$ (body-to-inertial by default, or to any specified reference frame), constraining all three rotational degrees of freedom. The user may specify this as a quaternion, direction cosine matrix, Euler angles, or modified Rodrigues parameters; the pipeline converts to a canonical unit quaternion.
- **Reduced attitude**: A vector-alignment goal requiring a body-frame direction $\hat{\mathbf{b}}$ to align with a world-frame direction $\hat{\mathbf{u}}$. The solution set is a one-parameter family of orientations (a circle S^1 in $SO(3)$), leaving rotation about the boresight unconstrained. The world-frame direction can be specified as a named direction (nadir, sunward, ram, etc.), an explicit vector, or a target location (latitude/longitude/altitude or ECI coordinates), all of which are resolved to a unit vector $\hat{\mathbf{u}}$ using the current orbital state.
- **No goal**: No pointing constraint during this period, useful for trajectory-planning "free time" or momentum management. An underactuated platform can use no-goal intervals to desaturate its wheels when the field geometry is favorable, rather than fighting a pointing constraint at the same time (Section II-F).

Table 2: Goal-to-Law Conversion Logic

	Full-Attitude Law	**Reduced-Attitude Law**
Full Goal	Direct	Alternating sub-goals
Reduced Goal	Quaternion set selection	Direct
No Goal	Identity error	Aligned vectors

Goal type directly affects controllability requirements: full-attitude goals require rank 6 controllability, while reduced-attitude goals require only rank 4 (see Section III). Missions that over-specify goals, demanding full attitude when only vector alignment is needed, artificially constrain the solution space.

II-B.2 Control Law Interface Types

Control laws declare one of two attitude interface types:

- **Full interface**: The law receives a quaternion error $\bar{q}_e$ (or equivalent) and angular velocity error ω_e, and outputs a desired torque τ_{ref}. Most PD, LQR, and sliding-mode controllers use this interface.
- **Reduced interface**: The law receives a body-frame vector $\hat{\mathbf{b}}$, a target vector $\hat{\mathbf{r}}$ (in either body or world frame, depending on the law), and ω_e, computing its own internal error. Laws such as those in [7, 9] use this interface.

The combination of goal type and law interface type determines the conversion logic, summarized in Table 2.

II-B.3 Direct Conversions

When goal type matches law interface type, conversion is straightforward. For a full goal with a full-attitude law, the quaternion error is

$$q_e = q_g \otimes q^{-1}, \tag{1}$$

[2] We cover the common cases; some variations surely remain.

with the scalar component enforced positive to ensure short-path rotation. For a reduced goal with a reduced-attitude law, the pipeline passes $\hat{\mathbf{b}}$ and $\hat{\mathbf{u}}$ (or $\hat{\mathbf{u}}$ rotated to body frame, depending on the law's convention) directly.

II-B.4 Reduced Goal to Full-Attitude Law

When a reduced-attitude goal is paired with a full-attitude law, the pipeline must select a specific quaternion from the set of orientations satisfying $A(q)\hat{\mathbf{b}} = \hat{\mathbf{u}}$. As derived in [9], this set can be parameterized as

$$f(\gamma) = x\cos\gamma + y\sin\gamma, \tag{2}$$

where x and y are basis quaternions defined by $\hat{\mathbf{b}}$ and $\hat{\mathbf{u}}$:

$$x = \begin{bmatrix} \cos\frac{\theta}{2} \\ \hat{\mathbf{x}}\sin\frac{\theta}{2} \end{bmatrix}, \quad y = \begin{bmatrix} 0 \\ \hat{\mathbf{y}} \end{bmatrix}, \tag{3}$$

with $\theta = \arccos(\hat{\mathbf{b}}\cdot\hat{\mathbf{u}})$, $\hat{\mathbf{x}} = (\hat{\mathbf{b}}\times\hat{\mathbf{u}})/\|\hat{\mathbf{b}}\times\hat{\mathbf{u}}\|$ (the minimum-rotation axis), and $\hat{\mathbf{y}} = (\hat{\mathbf{b}}+\hat{\mathbf{u}})/\|\hat{\mathbf{b}}+\hat{\mathbf{u}}\|$ (the halfway vector). The quaternion x corresponds to the minimum-angle rotation, and y to a 180° rotation about the halfway vector.

The default selection method minimizes geodesic distance from the current orientation. Since $q\cdot f(\gamma) = (q\cdot x)\cos\gamma + (q\cdot y)\sin\gamma$, the optimal parameter is

$$\gamma^* = \operatorname{atan2}(q\cdot y,\, q\cdot x). \tag{4}$$

This closed-form solution requires only two dot products per timestep. The framework supports alternative selection methods (e.g., Lyapunov-optimal or constraint-aware) through a configurable interface.

When $\hat{\mathbf{b}}$ and $\hat{\mathbf{u}}$ are nearly anti-parallel ($\hat{\mathbf{b}}\cdot\hat{\mathbf{u}} \approx -1$), $\hat{\mathbf{x}}$ and $\hat{\mathbf{y}}$ become ill-defined. A smooth regularization is applied:

$$\hat{\mathbf{x}} = \text{normalize}(\hat{\mathbf{b}}\times\hat{\mathbf{u}} + \varepsilon\,\hat{\mathbf{e}}_\perp), \quad \hat{\mathbf{y}} = \hat{\mathbf{b}}\times\hat{\mathbf{x}}, \tag{5}$$

where $\hat{\mathbf{e}}_\perp$ is any unit vector perpendicular to $\hat{\mathbf{b}}$ and ε is a small regularization parameter. Constructing $\hat{\mathbf{y}}$ from the cross product guarantees orthogonality of the basis regardless of the regularization perturbation.

Other criteria for γ are possible: maximizing controllability (control torque perpendicular to the field in magnetic-only operation), easing momentum management, or accounting for the current angular velocity.

II-B.5 Full Goal to Reduced-Attitude Law

When a full-attitude goal must be achieved by a reduced-attitude law, the pipeline decomposes the goal into two reduced sub-goals. Two non-parallel body vectors $\hat{\mathbf{b}}_1$, $\hat{\mathbf{b}}_2$ are selected (defaulting to orthogonal principal axes), and their required world-frame directions are computed as $\hat{\mathbf{u}}_i = A(\bar{q}_g)\hat{\mathbf{b}}_i$. Each sub-goal's solution set is a circle in $SO(3)$; the two circles intersect only at $\pm\bar{q}_g$, so alternating between them drives the system toward the unique full-attitude solution. The switching strategy is configurable (it can be threshold-based, time-based, or convergence-based), with the default alternating every control cycle. This conversion is primarily included for generality. In practice, full-attitude goals are typically paired with full-attitude laws.

II-B.6 Angular Velocity Handling

For tracking goals where the target direction changes over time, the pipeline computes a reference angular velocity ω_{ref} that the spacecraft must maintain to track the goal. For named goals with known analytical rates (e.g., nadir pointing at orbital rate, $\omega_{\text{ref}} = (\mathbf{r}\times\mathbf{v})/\|\mathbf{r}\|^2$), the analytical formula is used directly. Otherwise, the reference rate is obtained by finite-differencing the goal at the next scheduled timestep.

For reduced-attitude goals, a unified finite-difference formula handles both static and time-varying body vectors. Define $\hat{\mathbf{c}} = A(q)\hat{\mathbf{b}}(t+\Delta t)$ as the direction the next-step body vector currently points in the world frame. The reference angular velocity is

$$\omega_{\text{ref}} = \hat{\mathbf{n}}\frac{\delta\theta}{\Delta t}, \tag{6}$$

where

$$\hat{\mathbf{n}} = \frac{\hat{\mathbf{c}}\times\hat{\mathbf{u}}(t+\Delta t)}{\|\hat{\mathbf{c}}\times\hat{\mathbf{u}}(t+\Delta t)\|}, \quad \delta\theta = \arccos(\hat{\mathbf{c}}\cdot\hat{\mathbf{u}}(t+\Delta t)).$$

When only $\hat{\mathbf{u}}$ varies (the common case), $\hat{\mathbf{c}} \approx \hat{\mathbf{u}}(t)$ near the goal, recovering the standard tracking rate. When $\hat{\mathbf{b}}$ also varies, the formula naturally includes the correction without any special handling. For full-attitude tracking, the reference angular velocity is extracted from the axis-angle decomposition of the incremental quaternion $\delta q = q_g(t+\Delta t)\otimes q_g(t)^{-1}$, divided by Δt.

II-B.7 Projection of Unconstrained Axes

Reduced-attitude goals leave rotation about $\hat{\mathbf{b}}$ unconstrained. To prevent the control law from fighting this free degree of freedom, the angular velocity error is pro-

jected before reaching the law:

$$\omega_e = P\left(\omega - \omega_{\text{ref,body}}\right), \quad P = \begin{cases} I & \text{full attitude} \\ I - \hat{\mathbf{b}}\hat{\mathbf{b}}^T & \text{reduced attitude} \\ 0 & \text{no goal} \end{cases} \tag{7}$$

where $\omega_{\text{ref,body}} = A(\bar{q})\,\omega_{\text{ref}}$ converts the world-frame reference to body frame. The projection is applied before the law, not inside it, so the law is unmodified. The matrix P is also passed to the compensation block (Section II-D), enabling axis-aware damping injection for laws that do not accept angular velocity inputs.

Control laws additionally declare how they consume angular velocity: as a projected error $P(\omega - \omega_{\text{ref}})$ for laws with linear damping, as the unprojected difference $\omega - \omega_{\text{ref}}$ for laws that use angular velocity nonlinearly, or not at all (in which case the compensation block injects damping on the constrained axes). The internal quaternion convention is Hamilton scalar-first with $\bar{q}_e = \bar{q}_g \otimes \bar{q}^{-1}$ (conventions and error-state kinematics follow [20]); a thin conversion layer adapts to alternative conventions as needed by specific control law implementations.

II-C. Control Law Interface

The control law is treated as a black box within the pipeline. It receives the error signals produced by the goal formulation block and outputs a desired torque $\tau_{\text{ref}} \in \mathbb{R}^3$. Almost any law from the literature—PD, sliding mode, LQR, Lyapunov-based, adaptive, or the many magnetic-specific designs catalogued in recent surveys [13, 14]—can be inserted without modification to the surrounding pipeline, provided it declares its interface type (full or reduced attitude, per Section II-B.2) and its angular velocity requirements. Published laws often carry environment-specific assumptions (e.g., magnetic controllers relying on gravity-gradient stabilization [4]) which the framework does not remove, but makes cheap to test by swapping the law and re-running.

Some published control laws output actuator commands rather than torque (e.g., magnetic dipole moments for a magnetorquer-only system [3] or wheel torques for a specific reaction wheel configuration [4]). For these, the framework maps the law's actuator commands back to an equivalent torque using the law's assumed actuator model [8], then re-allocates to the actual hardware through the allocation stage (Section II-E). In many cases, such laws contain an internal step where a desired torque is computed before allocation. In such cases, this intermediate torque can be extracted directly, bypassing the back-mapping.

When applying a control law designed for a different-sized spacecraft, the proportional gains should scale with the inertia tensor (e.g., $k_p \propto \|J\|$) and derivative gains accordingly. The framework does not perform this scaling automatically but provides guidance, as appropriate scaling depends on the specific law's structure. We validate this scaling directly in Section IV-D: with inertia-scaled gains a single tuned law holds full convergence across a twentyfold inertia range, while the unscaled canonical gains collapse on the largest bus.

II-D. Compensation

The compensation block adds feedforward and correction terms *around* the control law, modifying the desired torque without the law's knowledge. Each correction is independently toggleable—if a law already accounts for a particular effect internally, the corresponding correction should be disabled to avoid double-counting. The compensated torque is

$$\tau_{\text{comp}} = \tau_{\text{ref}} + \tau_{\text{gyro}} + \tau_{\text{frame}} + \tau_{\text{dist}} + \tau_{\text{damp}}, \tag{8}$$

where each term can be individually enabled or disabled.

Gyroscopic compensation cancels the nonlinear coupling term $\tau_{\text{gyro}} = \omega \times (J\omega + \mathbf{h}_{\text{RW}})$, so that the control law operates on effectively decoupled dynamics. This is particularly valuable for laws derived under linearized assumptions.

Frame rotation feedforward compensates for the angular acceleration of the reference frame for tracking goals (e.g., LVLH frame rotation for nadir tracking), using the reference angular velocity ω_{ref} passed from the goal formulation block.

Disturbance feedforward subtracts estimated disturbance torques $\tau_{\text{dist}} = -\hat{\tau}_{\text{disturbance}}$ (gravity gradient, aerodynamic, magnetic residual dipole, solar radiation pressure) from the allocation input. Estimates may come from physics-based models, augmented state estimation, or disturbance observers.

Damping injection is engaged when a control law does not accept angular velocity feedback. The compensation block adds $\tau_{\text{damp}} = -k_d\,P\,(\omega - \omega_{\text{ref,body}})$, where P is the projection matrix from the goal formulation block. The projection ensures that damping acts only on the constrained rotational axes. This allows proportional-only control laws to be used with less oscillation, while preserving the unconstrained freedom of reduced-attitude goals. It is suggested (but not required) to add some additional damping along the boresight axis when possible, with $\tau_{\text{damp,bs}} = -k_{d,bs}\hat{\mathbf{b}}\hat{\mathbf{b}}^T(\omega - \omega_{\text{ref,body}})$, to reduce velocity around the boresight axis. This can also

be achieved with $\tau_{\text{damp all}} = -k_{d,all}(\omega - \omega_{\text{ref,body}})$.

II-E. Torque Allocation

The allocation block maps the compensated torque τ_{comp} to actuator commands $\mathbf{u}$ that respect hardware bounds. Define the torque effectiveness matrix

$$B_\tau(t) = \begin{bmatrix} A_{\text{RW}} & M_{\text{eff}}(t) & A_{\text{other}}(t) \end{bmatrix}, \tag{9}$$

where A_{RW} stacks the reaction wheel torque axes (static for fixed-axis wheels), $M_{\text{eff}}(t) = -[\mathbf{B}_b(t)]_\times A_{\text{MTQ}}$ is the time-varying magnetorquer torque mapping, and $A_{\text{other}}(t)$ accommodates additional actuators such as thrusters or CMGs; the time argument is carried for generality, since some actuators have body-frame effectiveness that varies in time (a magnetorquer with the field, a CMG with its gimbal state) while fixed thrusters are constant. The command vector is $\mathbf{u} = [\mathbf{u}_{\text{RW}}^T, \mathbf{u}_{\text{MTQ}}^T, \mathbf{u}_{\text{other}}^T]^T$ with box constraints $\mathscr{U} = \{\mathbf{u} : \mathbf{u}_{\min} \le \mathbf{u} \le \mathbf{u}_{\max}\}$.

The framework supports multiple allocation formulations, reflecting ongoing research into the trade-offs between them. The LP formulation maximizes torque along the desired direction:

$$\begin{aligned} \max_{\mathbf{u},\alpha} \quad & \alpha \\ \text{s.t.} \quad & B_\tau(t)\,\mathbf{u} = \alpha\,\hat{\tau}_{\text{comp}} \\ & \mathbf{u} \in \mathscr{U}, \quad \|\tau_{\text{comp}}\| \ge \alpha \ge 0 \end{aligned} \tag{10}$$

where $\hat{\tau}_{\text{comp}} = \tau_{\text{comp}}/\|\tau_{\text{comp}}\|$. This preserves torque direction exactly, which is important for stability of Lyapunov-based control laws. The achieved torque has it's magnitude scaled but is always parallel to the desired torque and thus never adds a perpendicular component that could violate $\dot{V} \le 0$. However, when the torque direction is not feasible, it scales the magnitude to zero, as discussed in Section IV-F.

The QP formulation instead minimizes overall torque error:

$$\min_{\mathbf{u}\in\mathscr{U}} \quad \|B_\tau(t)\,\mathbf{u} - \tau_{\text{comp}}\|^2. \tag{11}$$

This achieves a larger fraction of the desired torque magnitude but may introduce direction errors near the boundary of the achievable envelope. The stability implications of these direction errors depend on the specific control law and require case-by-case analysis, further discussed in Section IV-F.

A key feature of the allocation layer is that it handles actuator failures natively. If a reaction wheel fails mid-operation, its bounds can be set to zero; the allocation is then simply solved with the remaining actuators, and the control law continues unchanged. The pointing performance degrades gracefully to whatever the remaining hardware can achieve. A more detailed analysis of the LP and QP formulations, including a closed-loop campaign comparing direction-preserving (LP1) and priority-restoring (LP2) variants, is presented in Section IV-F.

II-F. Momentum Management

Reaction wheels accumulate momentum from persistent disturbance torques, eventually reaching saturation limits. Desaturation (momentum dumping) uses other actuators (e.g., MTQs) to reduce stored momentum. In the pipeline, desaturation is integrated into the allocation layer rather than implemented as a separate operational mode.

For **overactuated configurations** (e.g., 3RW+3MTQ, where B_τ has six columns for three torque axes), the nullspace of the allocation matrix provides degrees of freedom for momentum management with no impact on the primary pointing objective [8]. The allocation first solves for the primary torque, then projects a desaturation command into the nullspace:

$$\mathbf{u} = \mathbf{u}_{\text{primary}} + N\left(N^T \mathbf{u}_{\text{desat}}\right), \tag{12}$$

where N spans the nullspace of B_τ and $\mathbf{u}_{\text{desat}}$ drives wheel speeds toward a target (typically near zero, but not zero due to stiction concerns). This is possible because $B_\tau N = 0$: any command in the nullspace produces zero net torque on the spacecraft while redistributing actuator effort without changing the net spacecraft torque.

For **underactuated configurations** (e.g., 3MTQ+1RW), nullspace desaturation is not always available. The framework supports two approaches: weighted desaturation, which blends a desaturation objective into the QP cost function at the expense of some pointing accuracy and suboptimal slews; and scheduled desaturation, which exploits the time-varying MTQ controllability to desaturate during periods of favorable magnetic field geometry. The trade-off between pointing and desaturation is controlled by user-configurable weights, and the framework can assess whether cumulative MTQ authority over an orbit is sufficient to counteract expected momentum buildup.

III. CONTROLLABILITY ANALYSIS

Before adapting a control law to a given actuator set, a mission designer needs to answer a prior question: *can this configuration achieve this goal at all?* The framework answers it with a compact controllability analysis

Table 3: Required controllable subspaces. The dimension $\dim z$ depends on the pointing objective and whether wheel momentum must be regulated.

Task	$\dim z$	Components
Full attitude	6	σ, ω
Vector, transverse rate	4	$P_\perp\sigma$, $P_\perp\omega$
Vector + axial damping	5	$+\hat{b}^\top\omega$
Full attitude + momentum	$6{+}m$	$+h$ (m wheels)

that depends only on the actuator geometry, the magnetic field, and the pointing objective—not on the choice of control law. This analysis demonstrates that attitude controllability reduces to a dimensional condition on the achievable torque set, so the question "can my hardware do this?" becomes "does my torque envelope span the directions this goal requires?" Because the condition is purely geometric, it stands on its own as a design-time screen. The rank tables below let a mission designer accept or reject some configurations (and the approach can be general numbers of MTQs and RWs) against the desired goals and predicted desaturation capability, before running a closed-loop simulation.

III-A. Task-Relevant State and the Required Controllable Subspace

Controllability is assessed not against the full state but against a *required controllable subspace* z, a selection of the state that captures only what the mission goal actually constrains. For full-attitude pointing, z includes all three attitude and three rate components ($z=[\sigma^\top,\,\omega^\top]^\top\in\mathbb{R}^6$, where $\sigma\in\mathbb{R}^3$ is the attitude-error vector in Modified Rodrigues Parameters and $\omega\in\mathbb{R}^3$ the body rate). For vector (reduced-attitude) pointing, only the two degrees of freedom orthogonal to the pointing axis $\hat{b}$ (a body-frame unit vector) matter, captured by the projection $P_\perp=I_3-\hat{b}\hat{b}^\top$, giving $z\in\mathbb{R}^4$; an optional fifth component regulates the axial rate. Including reaction wheel momentum h in z distinguishes formulations that require active momentum management from those that treat the wheels purely as torque-limited actuators. Table 3 enumerates these cases.

III-B. Linearized Dynamics and the Torque Effectiveness Matrix

Linearizing the spacecraft dynamics about an equilibrium (Modified Rodrigues Parameter $\sigma_0=0$, $\omega_0=0$, wheel momentum $h_0=\bar{h}$) yields a state-space system whose input matrix collapses to a single object: the *torque effectiveness matrix* $B_\tau(t)$ of (9). The achievable torque set is the Minkowski sum

$$\mathscr{T}(\mathbf{B}_b(t))=\mathscr{T}_{\text{rw}}\oplus\mathscr{T}_{\text{mtq}}(\mathbf{B}_b(t))\oplus\mathscr{T}_{\text{other}}(t), \tag{13}$$

with $\mathscr{T}_{\text{mtq}}(\mathbf{B}_b(t))$ confined at each instant to the plane perpendicular to $\mathbf{B}_b(t)$.

The linearized attitude subsystem has the block structure

$$\dot{x}=\underbrace{\begin{bmatrix}0 & \frac{1}{2}I_3\\ 0 & 0\end{bmatrix}}_{A}x+\underbrace{\begin{bmatrix}0\\ J^{-1}\end{bmatrix}B_\tau(t)}_{B(t)}u,\quad x=\begin{bmatrix}\sigma\\ \omega\end{bmatrix}, \tag{14}$$

where the $\frac{1}{2}I_3$ block is the linearization of the quaternion or MRP kinematics about the origin. The structure is the same for either parameterization because both linearize identically to a Cartesian description of small rotations.

III-C. Rank Reduces to a Torque-Space Condition

For the time-invariant case (frozen field, fixed actuator geometry), the standard controllability matrix $\mathscr{C}=[B\;\;AB\;\;A^2B\;\cdots\;A^{n-1}B]$ must have full rank $\operatorname{rank}\mathscr{C}=n$, where $n=\dim z$ is the dimension of the required controllable subspace (Table 3). The block structure of A in (14) makes $\mathscr{C}$ block-triangular in B_τ, and a direct calculation yields

$$\operatorname{rank}\mathscr{C}=2\operatorname{rank}B_\tau=2\dim\mathscr{T}(\mathbf{B}_b), \tag{15}$$

because each independent torque direction controls one attitude component (through double integration) and one rate component (at first order). Hence the factor of two: **full attitude controllability requires the torque set to span all three spatial dimensions**. The same calculation extends to the reduced-attitude case with the projection $P_\perp$ applied to both halves of z.

The time-varying (LTV) case is handled via the finite-horizon controllability Gramian $W_c(t_0,t_f)=\int_{t_0}^{t_f}\Phi(t_f,\tau)B(\tau)B(\tau)^\top\Phi(t_f,\tau)^\top\,\mathrm{d}\tau$ (Φ the state-transition matrix), which is non-singular iff the system is controllable on $[t_0,t_f]$. For our problem the time-varying part of $B(t)$ is confined to $B_\tau(t)$, and the geometric statement equivalent to a non-singular Gramian is that the union of instantaneous torque sets spans the required space:

$$\dim\left(\bigcup_{t\in[t_0,t_f]}\mathscr{T}(\mathbf{B}_b(t))\right)=\dim z/2. \tag{16}$$

Equation (16) is the practical content of the Gramian condition: instead of evaluating an integral, count the dimensions the time-varying torque polytope covers over the orbit—no required direction may be absent for the entire horizon.

The same geometry yields the condition for reaction wheel desaturation without disturbing attitude. With the

attitude converged, canceling the wheel-induced torque on the body with a magnetorquer torque requires

$$\mathscr{T}_{\text{mtq}}(t) \cap \left(-\mathscr{T}_{\text{rw}}\right) \neq \{0\} \tag{17}$$

at some point in the orbit. When the intersection is nontrivial, wheel momentum can be reduced at zero net body torque; when it is empty for all t, any change in h necessarily perturbs the attitude, and momentum management must be scheduled to favorable times or weighted into the cost function.

III-D. Numerical Results Across the Configuration Grid

Table 4 reports rank-of-the-controllability-matrix results for representative configurations. Each cell shows $\text{rank}\,\mathscr{C}\,/\,n$, where n is the dimension of the required controllable subspace (Table 3). Three task formulations appear: full attitude ($n=6$, columns labeled **Full**), reduced-attitude vector pointing with transverse rate only ($n=4$, **Vec**), and vector pointing with axial-rate damping ($n=5$, **VecD**). Rows labeled $h-$ exclude wheel momentum from z; $h+$ rows include it (one additional dimension per wheel, as wheel momentum is also a controlled state). Magnetorquers are added along principal body axes; reaction wheels along $\hat{z}$, $\hat{y}$, $\hat{x}$ in order, with an optional fourth wheel along $[1,1,1]/\sqrt{3}$.

We highlight four findings.

Time variation enhances controllability. Orbital field variation consistently expands the controllable set (LTI vs. LTV columns). The headline case is 3MTQ+0RW: under a frozen field it is uncontrollable for full attitude ($4/6$), the magnetorquer torque trapped in the plane perpendicular to $\mathbf{B}_b$, but under a time-varying field it becomes fully controllable ($6/6$) as that plane sweeps through all orientations over the orbit (Figure 2, left). This is the formal basis for the familiar claim that a single orbit supplies enough field variation for full three-axis magnetic control. The amount of variation depends on orbit geometry: near-equatorial orbits see the field sweep less over an orbit, making magnetorquer-only three-axis control slower and harder while higher-inclination orbits sweep it more and ease the problem.

The enhancement is bounded by actuator geometry. Time variation is not unconditional. A single magnetorquer with dipole axis $\hat{m}$ produces torque $\hat{m}\times\mathbf{B}_b$, always orthogonal to $\hat{m}$: as the field sweeps, the torque traces the plane perpendicular to $\hat{m}$, raising the rank from 2 to 4 but never to 6, since no field orientation produces torque about $\hat{m}$ itself (Figure 2, right). Two or more nonparallel magnetorquers escape this because their swept planes union to all of $\mathbb{R}^3$, reaching $6/6$ under LTV. Time variation cannot manufacture a torque direction the actuator is structurally incapable of producing.

Reduced attitude relaxes the requirement. The **Vec** and **VecD** thresholds sit below **Full** throughout. A magnetorquer-only spacecraft, uncontrollable for full attitude under a frozen field, is fully controllable for vector pointing ($4/4$) once time variation is included, and two magnetorquers suffice. Since many missions need only single-axis pointing (camera boresight, antenna, solar array), the practical bar is often the reduced one: establishing the true pointing objective before sizing actuators avoids making an adequate configuration look uncontrollable. The portability study of Section IV-B shows the same ordering in convergence: on 3MTQ+1RW, relaxing the goal from full attitude to vector pointing raises convergence from 42% to 80% under an identical 1000 s horizon.

Momentum management costs a dimension instantaneously, recovered over the orbit. Including wheel momentum in z (the $h+$ columns) raises the required dimension by one per wheel via the desaturation condition (17). Under a frozen field, 3MTQ+3RW reaches only $8/9$. The magnetorquers, confined to a single plane, cannot independently bleed all three wheels, but under a time-varying field it recovers to $9/9$ as the orbital field supplies the missing directions. The desaturation story is thus the momentum-space counterpart of the attitude-controllability story, turning on the same intersection condition (17).

Taken together, these single out **3MTQ+1RW** as a small-satellite sweet spot: fully controllable for full attitude under both frozen and time-varying fields ($6/6$), gaining independent momentum management over the orbit ($6/7$ under LTI rising to $7/7$ under LTV), and degrading gracefully (if the wheel fails it reverts to magnetorquer-only operation, still fully controllable for full attitude under LTV) all at lower size, weight, and power than a three-wheel bus. The configuration sweep in Section IV-B exercises exactly this point, bracketing 3MTQ+1RW between the underactuated 3MTQ+0RW and the fully actuated 3MTQ+3RW.

IV. RESULTS

Five experiments exercise the framework, each holding one element fixed and varying another: whether one law

Table 4: Controllability ranks $\mathrm{rank}\,\mathscr{C}/n$ across configurations, under a frozen field (LTI) and an orbit-averaged time-varying field (LTV). Tasks: full attitude (Full, n=6), vector pointing with transverse rate (Vec, n=4), and with axial-rate damping (VecD, n=5). $h-$/$h+$ exclude/include wheel momentum in z (one dimension per wheel). Field variation lifts the rank across a row (e.g. 3MTQ+0RW full attitude: $4/6 \to 6/6$).

	LTI (frozen field)						LTV (time-varying field)					
	$h-$			$h+$			$h-$			$h+$		
Config	Full	Vec	VecD	Full	Vec	VecD	Full	Vec	VecD	Full	Vec	VecD
1MTQ+0RW	2/6	2/4	2/5	2/6	2/4	2/5	4/6	2/4	4/5	4/6	2/4	4/5
2MTQ+0RW	4/6	4/4	4/5	4/6	4/4	4/5	6/6	4/4	5/5	6/6	4/4	5/5
3MTQ+0RW	4/6	4/4	4/5	4/6	4/4	4/5	6/6	4/4	5/5	6/6	4/4	5/5
3MTQ+1RW	6/6	4/4	5/5	6/7	4/5	5/6	6/6	4/4	5/5	7/7	4/5	5/6
3MTQ+2RW	6/6	4/4	5/5	7/8	4/6	5/7	6/6	4/4	5/5	8/8	4/6	5/7
3MTQ+3RW	6/6	4/4	5/5	8/9	4/7	5/8	6/6	4/4	5/5	9/9	4/7	5/8

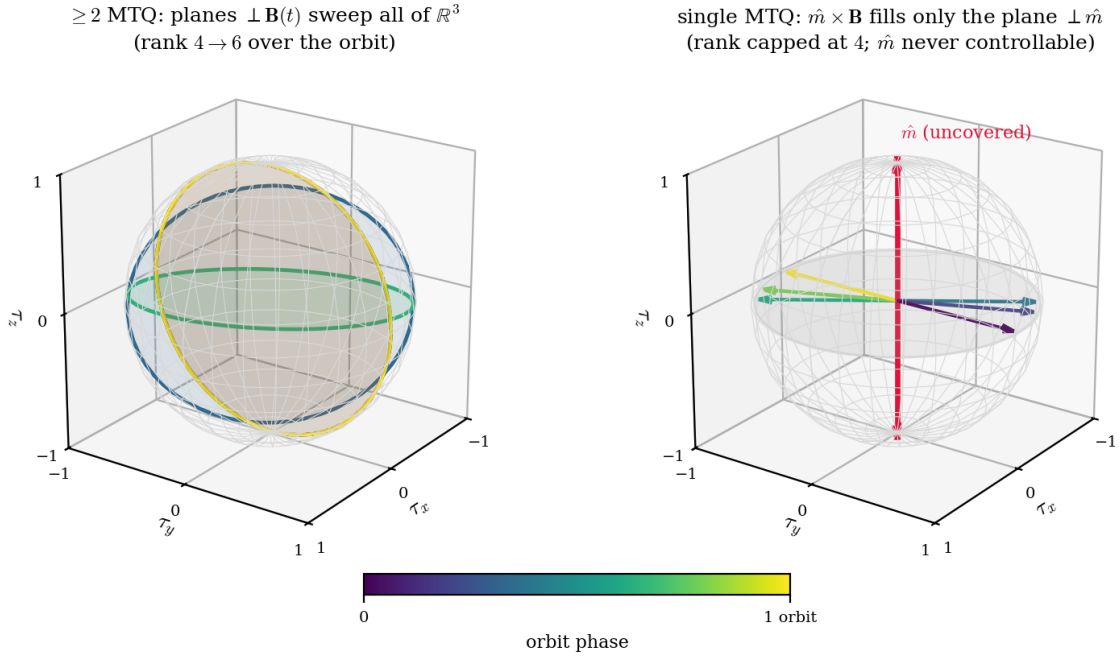


Figure 2: A geometric reading of the rank tables: each magnetorquer's torque lies in the plane perpendicular to the body-frame field $\mathbf{B}(t)$. *Left:* with ≥ 2 magnetorquers the plane reorients as the field sweeps the orbit and the union spans $\mathbb{R}^3$, lifting attitude rank from 4 (frozen) to 6 (time-varying). *Right:* one magnetorquer gives torque $\hat{m} \times \mathbf{B} \perp \hat{m}$, so the swept set fills only the plane $\perp \hat{m}$ and $\hat{m}$ is never controllable, leaving the rank capped. Colored by orbit phase.

ports across actuator sets (Section IV-B), whether different laws run through one interface (Section IV-C), whether one law rescales across inertias (Section IV-D), whether the framework absorbs an actuator failure (Section IV-E), and what the allocator choice itself trades away (Section IV-F).

IV-A. Simulation Setup

All campaigns in this section share the spacecraft and environment model of Table 5; per-campaign deviations are noted where they occur.

Initial conditions are drawn per campaign, with seeds paired across the conditions under comparison. The portability (§IV-B) and law-swap (§IV-C) studies each draw a uniformly random initial orientation, an ini-

Table 5: Simulation parameters for the framework campaigns. Per-campaign deviations (notably the two actuator-limit scales) are noted in the rows and in the text.

Quantity	Value
Inertia J	$\mathrm{diag}(0.022, 0.022, 0.004)$ kg m^2 (all Monte Carlo campaigns; a 3U-class CubeSat in the spirit of BeaverCube [12])
Magnetorquers	3, principal body axes; max dipole 0.2 A m^2 (allocator, failure, and direction-error studies) or 0.4 A m^2 (portability and law-swap campaigns)
Reaction wheels	along $+\hat{z}, +\hat{y}, +\hat{x}$ (order added); max momentum 16.2 mN m s; max torque 1.0 mN m (allocator/failure/direction studies) or 7.0 mN m (portability/law-swap campaigns)
Orbit	circular, 7000 km radius (≈622 km altitude), inclination ≈45°, period ≈5800 s (nominal); orbital elements randomized per trial in the portability and law-swap campaigns, held at this reference in the allocator/failure/direction studies
Geomagnetic field	IGRF-13 (`ppigrf`) with Skyfield Earth-orientation frames
Disturbances	time-varying geomagnetic field is the only environmental driver in the allocation-campaign plant (gravity-gradient, drag, and SRP not attached)
Control rate Δt	2 s
Convergence threshold	5° final pointing error
Trials / horizon	100 paired; 1000 s (2000 s for the magnetic laws of §IV-C)
Canonical PD gains	$(k_p, k_d, k_c) = (5 \times 10^{-5}, 10^{-3}, 10^{-3})$

tial body rate of random direction and magnitude $\sim\mathcal{U}(0.1, 2.0)$°/s, and an *independent* random pointing goal (a random ECI direction for vector tasks, a random target orientation for full attitude); the resulting maneuvers are large (median initial error ≈60° for vector, ≈132° for full), the portability sweep varying configuration at fixed gains and the law-swap study varying the control law at fixed configuration. The allocator study (§IV-F) instead isolates allocation behavior near the target: it holds the goal at a fixed reference, perturbs the initial attitude by a small random-axis rotation $\sim\mathcal{U}(0.1, 0.5)$ rad, and additionally sweeps initial stored wheel momentum $h_0 \in \{0, 0.3, 0.6, 0.9\}\, h_{\max}$ on a deliberately harsh short-horizon testbed.

Table 6: One PD law across configurations and tasks, 100 paired trials per cell on identical orbits and initial conditions; each cell gives LP / QP convergence (fraction reaching 5°) and mean final error. QP improves the magnetorquer-only cases by accepting direction error, while the LP matches or outperforms once reaction-wheel authority is present (Section IV-F).

Config	**Task**	**Conv. %** LP / QP	**Mean (°)** LP / QP
3MTQ+0RW	vector	1 / 7	94.2 / 52.4
3MTQ+1RW	vector	80 / 77	8.8 / 7.3
3MTQ+3RW	vector	100 / 100	0.0 / 0.0
3MTQ+0RW	full	0 / 2	121.1 / 96.3
3MTQ+1RW	full	42 / 39	44.4 / 54.9
3MTQ+3RW	full	100 / 100	0.0 / 0.0

IV-B. Same Law, Different Configs

The framework's modularity claim is operational only if a single control law, dropped unchanged into the pipeline, behaves consistently across hardware. We test this directly by running the framework's quaternion PD with LP allocation (at canonical gains $(k_p, k_d, k_c) = (5\times10^{-5}, 10^{-3}, 10^{-3})$ shared with the controllability and allocation studies) on three actuator complements (3MTQ, 3MTQ+1RW, 3MTQ+3RW) and two task types (vector pointing and full-attitude regulation). One hundred paired trials per cell, with each scenario seen by every configuration.

Convergence rises monotonically with actuator authority, from 1% on 3MTQ-only to 80% (vector) and 42% (full) on 3MTQ+1RW, to 100% on 3MTQ+3RW, with the mean final error falling in step (Table 6, Figure 3). The two magnetic-only cells are interpretable without further work: 3MTQ is controllable for vector pointing and, over an orbit of field variation, for full attitude as well (Section III), so neither is uncontrollable in principle [3]. However, within this 1000 s horizon, magnetic-only authority is far too slow to drive a random initial error below the 5° threshold. Vector pointing performs better than full attitude, and both will converge over many orbits. This ceiling is the hardware's, not the allocator's: in a paired comparison on the same orbits and gains, swapping the direction-preserving LP for the direction-relaxing QP lifts convergence only to 7% (vector) and 2% (full), both on a timescale well short of one orbit (Table 6). The QP does, however, recover a large part of the

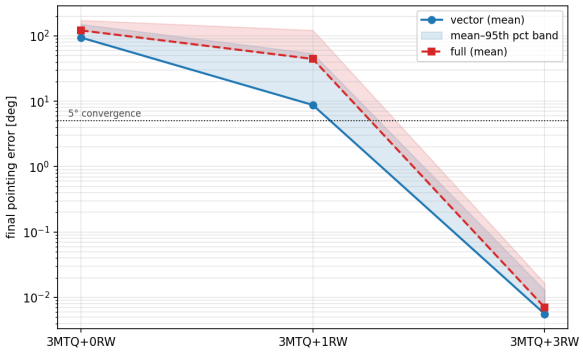


Figure 3: One PD law across actuator complements: mean final pointing error collapses by roughly four orders of magnitude as actuator authority grows from 3MTQ to 3MTQ+3RW, on both vector and full-attitude tasks. Dotted line marks the 5° convergence threshold.

residual (vector error nearly halves ($94^\circ \rightarrow 52^\circ$) and full attitude improves to 96° from 121°) so much of the LP's larger mean-error gap is its own all-or-nothing response to near-infeasible commands (Section IV-F), not a hardware floor. The framework's claim is therefore layered: the law ports cleanly across the pipeline, the one-orbit convergence ceiling is set by the actuator geometry, and how near a given allocator drives toward that ceiling is a separate, swappable choice.

IV-C. Same Config, Different Laws

The pipeline's Stage 2 (Figure 1) is a control-law slot, not a fixed law. We demonstrate the flexibility of this architecture by routing three structurally different attitude laws through one identical actuator interface. Each law plugs into the framework at Stage 2 with no changes to the surrounding pipeline; Stages 1, 3, 4, and 5 are shared across all three runs. Specifically, we test a quaternion PD (a scalar-gain form of Wie's eigenaxis regulator [1]), Lovera and Astolfi's global magnetic inertial-pointing law [3], and Wisniewski's PD-on-quaternion magnetic controller [4]. Beyond swapping one law for another, the framework encourages exploration. It allows laws written for other actuator configurations to be tested on a bus, and it lets the operator change the goal type to suit the mission. The comparison is run on 3MTQ+1RW, 100 paired trials at $t_f = 2000$ s, on both vector-pointing and full-attitude tasks. Each trial seeds the three laws identically.

Extending a law past its native hardware. Two of the three laws were built for different hardware. Lovera's and Wisniewski's laws command magnetorquers only and never actuate a wheel, and magnetorquer torque is confined at each instant to the plane perpendicular to the field (Section III). On magnetorquers alone, full three-axis regulation is far too slow to converge within the campaign horizon, so in their published wheel-idle form the two laws tumble on this task, reaching 0 and 2 of 100 full-attitude trials (the last row of each block, Table 7). With the wheel available via the allocator to the same feedback law, an additional control axis is supplied, and the two laws regulate full attitude at 29% and 39% 4. This improvement comes from the allocator using the new actuator, as the laws are identical. A similar jump in performance when the wheel is made accessible appears on the vector-pointing tasks, from 41% to 80% and 16% to 72%. A magnetorquer-only law from the literature thus runs on a hybrid bus by reconfiguration alone, enabling a designer to benchmark laws that were never written for the actuators on hand.

Goal type as a design lever. This framework also allows a designer to convert full-attitude control laws to vector-pointing, as described in Section II-B. Relaxing full attitude to a vector goal lifts every law sharply, at both hardware levels, as shown in Table 7. With the wheel, convergence rises from the 29–39% range to 72–80%, and PD from 49% to 83%. Even with the wheel unused, the magnetic laws climb from 0–2% to 16–41%. The two levers also compound. A wheel-idle full-attitude law converging in 0–2% of trials reaches 72–80% once it has both the wheel and the relaxed goal, a far larger change than either change brings alone. For a designer weighing options, this is a valuable test, as a configuration that looks marginal on full attitude might be adequate for the pointing goal required. Which magnetic law is better changes with the goal: Lovera leads on vector and Wisniewski on full, at both hardware levels, demonstrating the value of exploring different options, as this framework allows.

IV-D. Gain Scaling Across Inertias

Portability is not solely about the actuators onboard a spacecraft. A controller tuned on one spacecraft should transfer to a differently-sized one once its gains are scaled with the inertia (Section II-C). Holding the framework's PD law and scenario set fixed on a 3MTQ+3RW full-attitude task (100 trials) and scaling the canonical gains by $\mathrm{tr}J/\mathrm{tr}J_1$, convergence is invariant across a twentyfold inertia range. The convergence is 100% on all three buses (Table 8), whereas the unscaled canonical gains hold near their tuning inertia but collapse to 1% (mean error 27°) on the 20× bus. Using inertia-

Table 7: Three control laws through the framework's actuator interface (3MTQ+1RW, 100 paired trials, $t_f = 2000$ s). Convergence is the fraction of trials below 5° final error; Δ vs. PD is the paired mean-error delta. The last two rows run Lovera and Wisniewski in their *published* magnetic-only form (wheel idle), serving as a negative control isolating the framework's contribution.

Task	Law	Conv. %	Mean (°)	Med (°)	p95 (°)	Δ (°)
Vector	PD	83	5.50	1.06	39.54	+0.00
	Lovera	80	7.76	0.14	52.76	+2.26
	Wisniewski	72	6.01	2.28	20.89	+0.52
	Lovera (RW idle)	41	31.9	15.1	102.9	+26.4
	Wisniewski (RW idle)	16	39.0	37.3	104.9	+33.5
Full	PD	49	35.13	5.37	108.73	+0.00
	Lovera	29	54.40	51.38	130.95	+19.27
	Wisniewski	39	28.70	7.61	101.81	−6.42
	Lovera (RW idle)	0	116.09	118.94	171.57	+80.96
	Wisniewski (RW idle)	2	110.97	115.79	174.82	+75.84

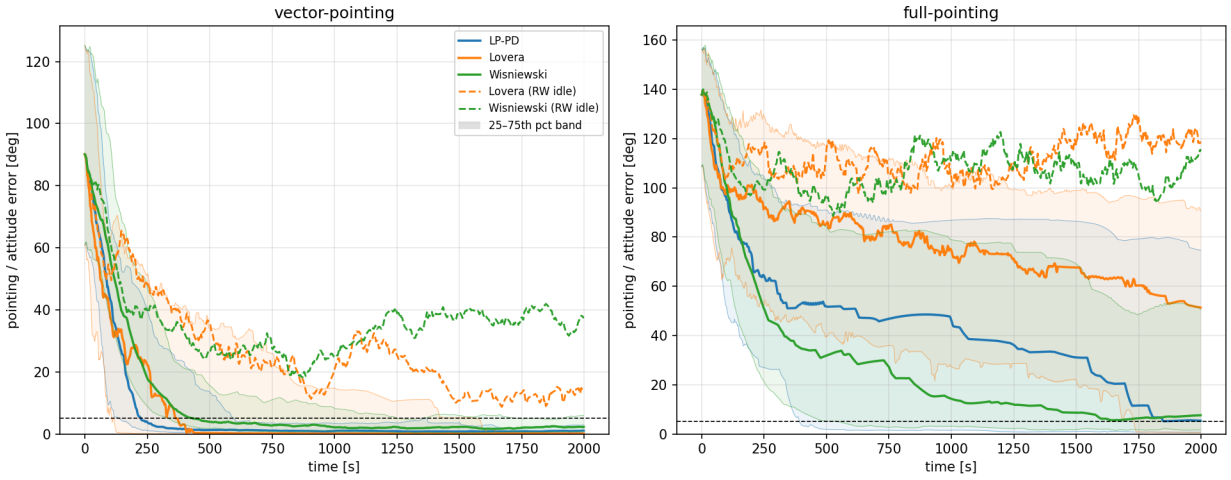


Figure 4: Pointing-error trajectories for three laws through the same pipeline on 3MTQ+1RW (100 paired trials), on the vector (left) and full-attitude (right) tasks. All three are comparable on the vector task; on full attitude the framework's PD leads and the RW-enabled magnetic laws trail but converge, which they almost never do in their published magnetic-only form.

proportionality for gains can enable simple and effective tuning across different scales of spacecraft.

IV-E. Automatic Reallocation under Failure

A separate consequence of putting allocation in its own pipeline stage is that an actuator loss is a matter for the allocator alone: the control law and compensation see no change, and the LP simply solves the same problem over a reduced polytope. To show both that the framework absorbs a failure transparently and what graceful degradation looks like when the survivors cannot span the task, we fail one reaction wheel at $t = 500$ s during nadir-locked (full-LVLH) tracking on a 600 km circular orbit and run a full orbit ($t_f = 6000$ s) under IGRF-13 and gravity-gradient torque, with no artificial body disturbance. Two configurations see the same failure: an over-actuated 3MTQ+3RW and a controllability-marginal 3MTQ+1RW.

Figure 5 shows the contrast. Losing a wheel takes 3MTQ+3RW to 3MTQ+2RW, whose two wheels and three magnetorquers still span $\mathbb{R}^3$. The allocator redistributes the lost wheel's contribution and the full three-axis LVLH error stays at 0.0° with no measurable transient. Losing the wheel on 3MTQ+1RW leaves a magnetorquer-only platform, no longer controllable

Table 8: Gain scaling across inertias (framework PD, 3MTQ+3RW, full attitude, 100 trials). Gains scaled by $\mathrm{tr}J/\mathrm{tr}J_1$ hold convergence at 100% over a $20\times$ inertia range; the unscaled canonical gains collapse on the largest bus.

			Conv. %	
Inertia	$\mathrm{tr}J$	gain×	scaled	unscaled
J_1 (reference)	0.048	1.0	100	100
J_2 (BeaverCube)	0.076	1.6	100	100
J_3 ($20\times J_1$)	0.960	20.0	100	1

Table 9: Recovery after a reaction-wheel failure drops a 3MTQ+1RW bus to magnetorquers only (3MTQ+0RW): relaxing the goal (Stage 1) versus swapping the allocator (Stage 5), over one orbit. "Boresight" is the peak boresight-to-nadir error; "Full LVLH" is the three-axis error, left free (—) on the reduced-goal runs.

Goal	Alloc.	Boresight (°)	Full LVLH (°)	Rate (°/s)
Full LVLH	LP	71.3	179.9	1.01
Full LVLH	QP	0.25	0.47	0.06
Reduced	LP	13.5	—	0.07
Reduced	QP	1.66	—	0.04

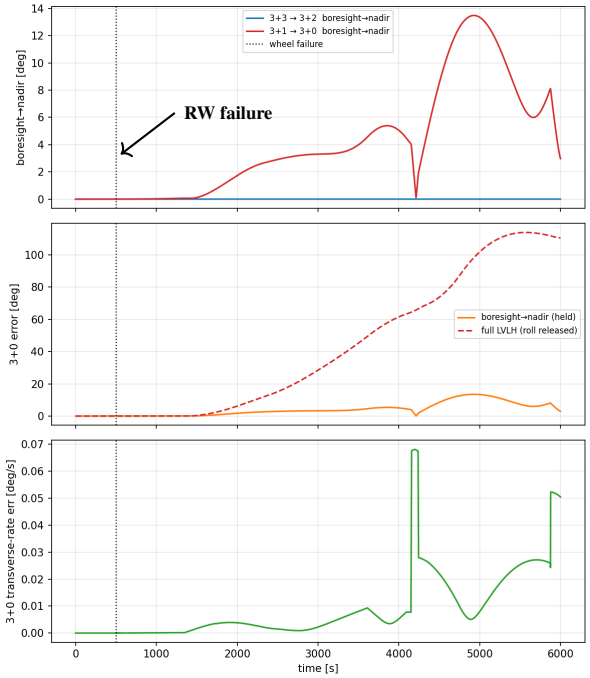


Figure 5: Reaction-wheel failure at $t = 500$ s during full-orbit nadir (LVLH) tracking (600 km orbit, IGRF-13 + gravity gradient, no artificial torque; direction-preserving LP). *Top:* boresight-to-nadir error: 3MTQ+3RW→3MTQ+2RW holds 0°; 3MTQ+1RW→3MTQ+0RW, run on a reduced-attitude goal, stays bounded (peak 13.5°, settling ∼3°). *Middle* (3MTQ+0RW): held boresight (solid) bounded, released roll (dashed) drifts to ∼110°. *Bottom:* transverse rate <0.07°/s is bounded, and not tumbling.

for full three-axis LVLH (Section III). With a reduced-attitude goal (boresight on nadir), the platform keeps the pointing intact. The boresight transient peaks at 13.5° and settles toward 3° over the orbit. The transverse body rate stays below 0.07°/s throughout, so there is no tumbling.

Both recoveries are configurations the operator sets, not automatic responses, and neither touches the control law. The reduced-attitude hold above is one lever (Stage 1), dropping the roll demand the magnetorquers cannot meet. Another, stronger, option is swapping the LP for QP (Stage 5). The QP holds full three-axis LVLH where the LP tumbles, so the goal need not be relaxed at all (Table 9). The LP tumbles on a full-attitude goal is a property of its direction-preserving geometry, analyzed in Section IV-F. The bounded outcome also leans on the environment: disabling gravity gradient takes the same run to 33.8° from 3.0°, so the passive torque is doing real work, not the allocator alone.

IV-F. Swapping the Allocator

The same modularity that makes Stage 2 a control-law slot makes Stage 5 an allocator slot. Inside the achievable polytope every allocator returns the requested torque, so the LP and QP differ only near the boundary, in the underactuated regime. There the LP preserves the wrapped law's certificate (Section V-A) at the cost of magnitude, while the QP recovers pointing at the cost of torque direction. Across the configuration grid, LP performs better wherever a wheel is present. The QP wins only on the pure-magnetorquer bus, where the LP zeroes the along-field request but QP lifts vector convergence $1\% \to 7\%$ and nearly halves the residual ($94° \to 52°$) as shown in Table 6, though the slew still does not converge. Adding a reaction wheel changes this. On 3MTQ+1RW full attitude the LP's direction-preservation beats the QP (42% vs. 39%, 44° vs. 55°), because the wheel supplies the along-field axis that the QP's least-squares tilt wastes. With full actuation the two coincide. The three allocators are the LP of (10), a two-stage priority-restoring variant on LP (LP2, introduced below), and the QP of (11). Indicative single desktop-core solve times are LP1≈350 μs, LP2≈700 μs, QP≈40 μs, all comfortably

real-time.

LP preserves torque direction; QP does not. Figure 6 quantifies this difference. Across the achievable envelope the LP holds the achieved torque parallel to the request to a mean of $\sim$0.001°. It returns a scaled-down torque, or zero when the direction is unachievable, but never a tilted one. The QP instead tilts the toward the achievable plane when the request has a component along the unachievable direction. This tilt is near zero for interior requests and grows toward 90° near the boundary, averaging 12–33° per configuration and peaking near 88° on the magnetorquer-only bus. This worst case falls along the instantaneous field direction, the axis the controllability analysis of (Section III) marks as unachievable. Geometrically, QP's largest direction error coincides with the thinnest dimension of the achievable torque polytope (Figure 2). Because a tilted torque can flip the sign of the $\dot{V}$ contribution in a Lyapunov analysis, this matters for stability, not just accuracy (Section V-A).

The default LP saturates ungracefully, and a priority-restoring variant does not fix it. Direction versus magnitude is one allocator metric. Saturation behavior is a second, internal to the LP. The LP combines the direct gyroscopic (or disturbance) feedforward compensation and the PD feedback into a single scale factor α. Under saturation, these are not distinguished, and it cannot protect the feedforward. As stored wheel momentum approaches saturation, α falls to $\approx$0.03, the allocator delivers only a few percent of the request, and pointing degrades. The body rate is still held throughout ($\tau_{\text{gyro}} \perp \omega$ does no work, confirmed in closed loop with wheels at 0.95–0.99 $h_{\max}$), so the failure mode is lost pointing authority, not instability.

A two-stage variant (LP2) that delivers the feedforward first and spends the remaining authority on feedback was considered. The first stage allocates as much of the compensation torque as possible:

$$\begin{aligned} \max_{\mathbf{u}_{\text{comp}},\alpha} \quad & \alpha \\ \text{s.t.} \quad & B_\tau(t)\mathbf{u}_{comp} = \alpha\left(\tau_{\text{gyro}} + \tau_{\text{frame}} + \tau_{\text{dist}} + \tau_{\text{damp}}\right), \\ & \mathbf{u}_{\text{comp}} \in \mathscr{U}, \quad 0 \le \alpha \le 1. \end{aligned} \tag{18}$$

If $\alpha = 1$, the compensation terms are fully realized. If $\alpha < 1$, the allocator degrades gracefully by applying the largest feasible compensation torque in the required direction. The feasibility of the first stage with $\alpha = 1$ also provides a useful actuator-sizing heuristic: it tests whether the selected actuators can counter the worst-case compensation torque in arbitrary alignment before any authority is reserved for feedback stabilization.

Table 10: Each pipeline stage is exercised by a distinct experiment: the modularity is a tested claim, not only an architectural one.

Component	Isolated by	What it shows
Goal	full-LVLH $\rightarrow$ reduced after wheel failure (§IV-E)	drops the infeasible roll demand
Law	PD, Lovera, Wisniewski through one interface (§IV-C)	different laws swap in unchanged
Allocation	LP vs. QP, paired (§IV-F)	certificate-preserving vs. recovery
Config.	one PD law across actuator sets (§IV-B)	one law ports across configurations
Gains	inertia sweep (§IV-D)	one law transfers across scale

The remaining actuator authority is then assigned to the stabilizing reference torque. Let $\mathbf{u}^{\star}_{\text{comp}}$ and $\alpha^{\star}$ denote the solution of (18). The residual feasible set is

$$\mathscr{U}_{rem} = \left\{ \mathbf{u}_{ref} \;:\; \mathbf{u}^{\star}_{comp} + \mathbf{u}_{ref} \in \mathscr{U} \right\}. \tag{19}$$

The second stage solves

$$\begin{aligned} \max_{\mathbf{u}_{ref},\beta} \quad & \beta \\ \text{s.t.} \quad & B_\tau(t)\mathbf{u}_{ref} = \beta\,\tau_{ref}, \\ & \mathbf{u}_{ref} \in \mathscr{U}_{rem}, \quad 0 \le \beta \le 1. \end{aligned} \tag{20}$$

In testing, LP2 helped only when the feedforward is itself feasible. LP2 matches LP exactly on the fully actuated 3MTQ+3RW, but on the underactuated 3MTQ+1RW the gyroscopic term is already largely unreachable ($1 - \bar{\alpha} \approx 0.85$ even at zero stored momentum), so LP2 uses scarce magnetorquer authority to an unreachable target and trails LP (55.8° vs. 47.9° mean, in LP's favor). We keep LP as the default. Making LP2 sound would need a momentum/disturbance-aware governor that keeps $(\omega, \mathbf{h}_{\text{rw}})$ inside the reliably actuable core of the time-varying torque polytope (Section III), which we leave to future work.

V. DISCUSSION

V-A. Stability of the Wrapped Law

The stability of a wrapped control law depends on how the allocator preserves the stabilizing torque delivered to

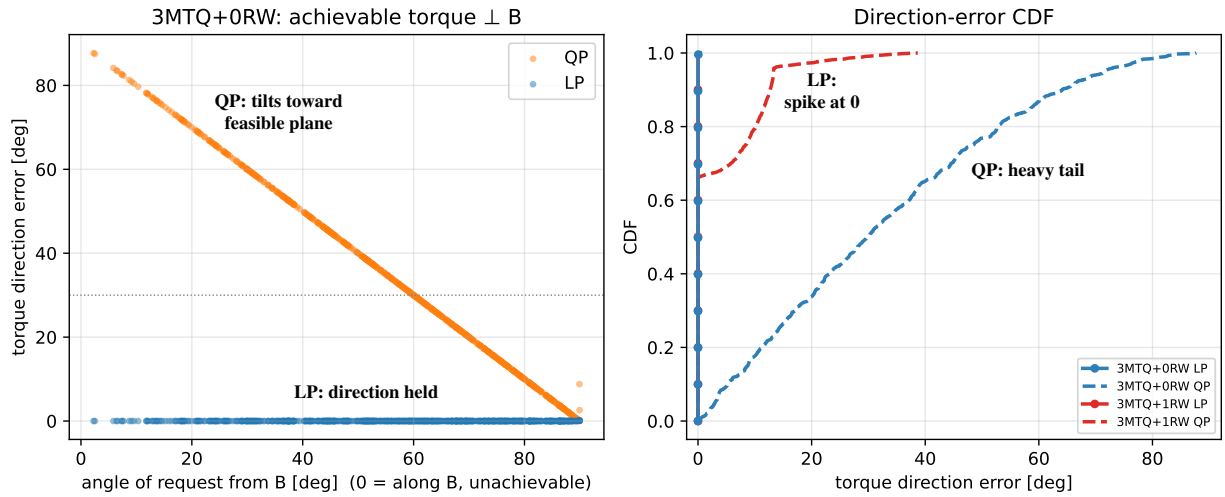


Figure 6: Torque-direction error of LP versus QP allocation across requested directions, underactuated configurations. *Left* (3MTQ+0RW): QP error grows toward 90° as the request aligns with the field (the unachievable direction), while LP holds direction to ~ 0. *Right:* direction-error CDFs: LP is a spike at zero, QP a heavy tail.

the system. The following discussion provides an initial stability interpretation of the proposed allocation strategies. A complete proof for the time-varying, saturated, and underactuated case remains open and is the subject of ongoing investigation. In addition, the experimental campaign presented in this work did not allow an exhaustive comparison between LP and LP2 under all operating conditions, so the conclusions should be interpreted as preliminary.

Assume the nominal control law admits a Lyapunov certificate $V \geq 0$ whose derivative depends on the applied body torque τ_a through

$$\dot{V} = \omega_e^\top \tau_a + g(x), \qquad \dot{V} \leq 0 \quad \text{when} \quad \tau_a = \tau_{ref}. \tag{21}$$

This structure covers common attitude controllers such as quaternion PD, rate damping, LQR, and magnetic control laws like those of [2,4] (whose certificates share this structure, closed via an orbit-averaging argument). The compensation framework augments the stabilizing reference torque with model-based terms, as described in 8. The additional terms represent known plant dynamics or disturbances that should ideally be canceled before applying the stabilizing control action.

LP: Direction preservation of the full compensated command LP treats the complete compensated command as a single torque direction and seeks the largest feasible scaling (see 10). The resulting applied torque is

$$\tau_a = \alpha \left(\tau_{\text{ref}} + \tau_{\text{gyro}} + \tau_{\text{frame}} + \tau_{\text{dist}} + \tau_{\text{damp}} \right). \tag{22}$$

This preserves the direction of the total compensated command and avoids the perpendicular torque components introduced by a distance-minimizing allocator. However, it also scales the compensation terms $\tau_{\text{gyro}} + \tau_{\text{frame}} + \tau_{\text{dist}} + \tau_{\text{damp}}$. Thus, under saturation, the stabilizing torque and the compensation torques compete for the same actuator authority.

The main practical concern is not necessarily that the compensation torque is infeasible by itself. Rather, a large reference command τ_{ref}, for example from aggressive PD gains or large attitude errors, can dominate the combined command in (22), force $\alpha < 1$, and thereby reduce the delivered compensation torque even when the actuator set could have produced the required compensation terms alone. Consequently, LP has an attractive direction-preserving property, but its performance can degrade when the feedback command is large relative to the compensation torque.

LP2: Prioritizing feedforward compensation The LP2 allocator in (18) and (20) was designed to address this problem. In that case, the applied torque is therefore

$$\tau_a = \alpha^\star \left(\tau_{\text{gyro}} + \tau_{\text{frame}} + \tau_{\text{dist}} + \tau_{\text{damp}} \right) + \beta \tau_{\text{ref}}. \tag{23}$$

Consider a quaternion PD controller like $\tau_{\text{ref}} = -k_p q_{ev} - k_d \omega_e$ with the Lyapunov function

$$V = \tfrac{1}{2} \omega_e^\top J \omega_e + 2k_p (1 - q_{e0}). \tag{24}$$

(Note this meets (21) with $g(x) = k_p \omega_e^\top q_{ev}$ and $\dot{V} = -k_d \|\omega_e\|^2$ at full torque.) For such a controller, if $\alpha^* =$

1, the applied LP2 torque of (23) gives the same scalar degradation of the stabilizing part of the Lyapunov argument:

$$\dot{V} = -\beta k_d \|\omega_e\|^2 + (1-\beta) k_p \omega_e^\top q_{ev}. \tag{25}$$

Thus, mild saturation or rate-dominated motion remains benign, and the nominal proof is recovered as $\beta \to 1$. If $\alpha^\star < 1$, the compensation itself is saturated and the same full stability guarantee no longer applies; nevertheless, the first stage still preserves the compensation direction and provides a controlled degradation mode.

QP: Euclidean minimization For comparison, the distance-minimizing QP solves the problem in (11). Under saturation, the delivered torque may contain a component not aligned with the compensated command, such that $\tau_a = \gamma\tau_{\text{comp}} + \tau_\perp$ with $\tau_\perp \perp \tau_{\text{comp}}$. The Lyapunov derivative then contains $\omega_e^\top \tau_\perp$ which has arbitrary sign. Therefore, although the QP minimizes instantaneous Euclidean torque error, it does not generally preserve the directional structure used by the Lyapunov certificate.

Overall, LP, LP2, and QP represent different compromises between torque accuracy, compensation priority, and stability preservation. LP preserves the direction of the full compensated command, but large feedback torques can cause the compensation terms to be unnecessarily scaled. LP2 prioritizes compensation first and then allocates the remaining authority to the stabilizing reference torque, which is often preferable when the model-based terms in (8) are significant. The QP minimizes torque error, but may rotate the applied torque away from the stabilizing direction. These conclusions are useful design guidance, but the rigorous stability theory for saturated, time-varying, underactuated allocation remains open.

V-B. Practical Considerations

The allocation layer adds an LP or QP solve to each control cycle. For the small problem sizes here (six decision variables with box constraints), the solve is inexpensive. On a single desktop core the LP runs in $\approx 350\,\mu$s and the QP in $\approx 40\,\mu$s (Section IV-F), so at typical 1–10 Hz ADCS rates the allocation is a small fraction of the control cycle. A flight-representative embedded timing is left to implementation.

The compensation terms (gyroscopic, frame rotation, disturbance feedforward) require information from the onboard estimator, specifically the estimated angular velocity, reaction wheel momentum, and disturbance torques. The framework is compatible with standard EKF/UKF estimation pipelines; it consumes their outputs but does not modify the estimator itself. If the estimator provides poor disturbance estimates, the disturbance compensation can be disabled individually without affecting the rest of the pipeline—the control law simply operates without that feedforward term.

The framework does not fix a bad control law. It wraps the law as a black box: a law that commands unstable torques will have those torques faithfully allocated to the actuators. The contribution is in the interface, not the control design—rarely a problem in practice, since the literature laws tested in Section IV-C are well-characterized.

V-C. Limitations

The geometric limit on magnetorquer-only attitude is the hardware's, not the framework's. Any law applied directly to the same actuators hits the same issue. How much of the feasible set an allocator extracts is a swappable choice (Section IV-F), but choosing when a constrained QP should relax direction rather than preserve it remains open.

Section V-A gives the stability argument and the regimes in which it holds; a fully general proof for the time-varying underactuated case, and for laws stabilized by an environmental torque, remains open. The framework also assumes torque-based control laws (or laws whose actuator commands can be mapped back to torque), and gain scaling between differently-sized spacecraft is approximate. This is sufficient for trade studies, but flight implementations should expect some per-mission adjustment.

V-D. Future Work

Completing the stability argument of Section V-A for the general underactuated case is a key theoretical need. Beyond that, the goal conversion between full-attitude and reduced-attitude representations offers significant room for improvement. The current approach selects the nearest quaternion in the goal set, but this does not exploit the freedom that reduced-attitude goals provide. The work in [10] identifies several alternatives: selecting the free rotation parameter γ to minimize the predicted Lyapunov rate of change, to minimize the distance to a sliding surface, to maximize alignment between desired and achievable torque given the current magnetic field geometry, or to enforce keep-out zone avoidance. A switching or blending approach across these strategies may yield the best results.

The allocator analysis of Section IV-F points to a momentum-aware governor (LP3) as a natural next step: a feasibility-restoring layer that keeps the $(\boldsymbol{\omega}, \mathbf{h}_{\text{rw}})$ state inside the reliably actuable core of the time-varying torque polytope, so that the priority-restoring LP2 becomes unconditionally sound. This connects the allocator to the controllability geometry of Section III.

Automatic gain scaling based on the controllability analysis would further reduce per-mission tuning. Additionally, hardware validation through a ground testbed or flight demonstration remains the essential next step.

VI. CONCLUSION

We have presented a modular framework that decouples attitude control law design from actuator configuration. Any torque-producing control law can be inserted at the framework's Stage 2 slot (Figure 1); the surrounding stages (goal formulation, optional torque mapping, feedforward compensation, and bound-respecting allocation) adapt the law to the actual hardware. The framework was validated across three structurally distinct control laws (PD and the magnetic laws of Lovera and Wisniewski), three actuator configurations (3MTQ+0RW, 3MTQ+1RW, 3MTQ+3RW), failure scenarios, and a paired allocator comparison spanning two allocators (LP, QP). On the controllability-appropriate diagonal cells, a single PD law unchanged across actuator complements converges in 80% (3MTQ+1RW vector) to 100% (3MTQ+3RW) of trials, and the law-swap experiment finds three literature control laws comparable on the vector task once they all share the framework's allocation interface.

The compact controllability analysis (Section III) provides the prior-check: a mission designer can tell from Table 4 whether the chosen actuator configuration meets the task's required-controllable-subspace dimension before running any closed-loop simulation. The same geometry that decides attitude controllability also decides reaction-wheel desaturation, via the torque-polytope intersection condition. The allocation campaign exposed a defect in the standard direction-preserving LP, identified the regime in which a priority-restoring two-stage LP (LP2) helps, and pointed to a momentum-aware governor as the natural next step. Across these results the pipeline does more than port a law. On the magnetorquer-only bus the framework separates three findings a single controller would combine: controllability analysis shows convergence over multiple orbits based on hardware, the one-orbit convergence numbers show it is too slow under either allocator, and the allocator comparison shows where measurable progress is unused.

The intended contribution is operational: turning "compare to literature" from a week-long reimplementation task into a configuration change. Lovera's and Wisniewski's magnetic-only laws, which on a 3MTQ+1RW spacecraft would normally leave the wheel idle, were exercised through the framework with no modification to their feedback structures. In such a case, they converged on the same Monte Carlo testbed as the framework's own PD, demonstrating that the allocation layer supplies the third-axis authority their original publications could not. For small-satellite teams whose ADCS development is bottlenecked by per-mission control-law reimplementation, the framework is the practical implementation: ground a controller in literature, port it as configuration, and spend the saved time exploring more options and tuning your solution rather than rebuilding one.

REFERENCES

[1] B. Wie, *Space Vehicle Dynamics and Control*, 2nd ed. Reston, VA, USA: AIAA Education Series, 2008.

[2] M. Lovera and A. Astolfi, "Global magnetic attitude control of spacecraft in the presence of gravity gradient," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 42, no. 3, pp. 796–805, Jul. 2006.

[3] M. Lovera and A. Astolfi, "Global magnetic attitude control of inertially pointing spacecraft," *J. Guid. Control Dyn.*, vol. 28, no. 5, pp. 1065–1072, 2005.

[4] R. Wisniewski, "Sliding mode attitude control for magnetic actuated satellite," *IFAC Proc. Volumes*, vol. 31, no. 22, pp. 179–184, 1998.

[5] A. Loría, E. Panteley, D. Popović, and A. R. Teel, "A nested Matrosov theorem and persistency of excitation for uniform convergence in stable nonautonomous systems," *IEEE Trans. Autom. Control*, vol. 50, no. 2, pp. 183–198, Feb. 2005.

[6] G. Sechi, M. Buonocore, F. Cometto, M. Saponara, A. Tramutola, B. Vinai, G. André, and M. Fehringer, "In-flight results from the drag-free and attitude control of GOCE satellite," in *Proc. 18th IFAC World Congr.*, Milan, Italy, 2011.

[7] C. M. Pong and D. W. Miller, "Reduced-attitude boresight guidance and control on spacecraft for pointing, tracking, and searching," *J. Guid. Control Dyn.*, vol. 38, no. 6, pp. 1027–1035, 2015.

[8] E. A. Hogan and H. Schaub, "Three-axis attitude control using redundant reaction wheels with continuous momentum dumping," *J. Guid. Control Dyn.*, vol. 38, no. 10, pp. 1865–1871, 2015.

[9] P. McKeen and K. Cahoy, "Method for pointing spacecraft while minimizing change in orientation," *J. Guid. Control Dyn.*, vol. 45, no. 8, pp. 1554–1557, 2022.

[10] P. McKeen, "Computational methods to improve satellite attitude determination and control with a focus on autonomy, generalizability, and underactuation," Ph.D. dissertation, Massachusetts Inst. Technol., Cambridge, MA, USA, 2024.

[11] P. McKeen, N. Scheuer, and K. Cahoy, "A modular open-source Python framework for small satellite attitude determination and control," presented at the Small Satellites, Systems & Services Symp. (4S/SmallSat Europe), 2026.

[12] P. do Vale Pereira *et al.*, "BeaverCube: Coastal imaging with VIS/LWIR CubeSats," in *Proc. 34th Annu. Small Satellite Conf.*, Logan, UT, USA, 2020.

[13] E. Silani and M. Lovera, "Magnetic spacecraft attitude control: A survey and some new results," *Control Eng. Practice*, vol. 13, no. 3, pp. 357–371, 2005.

[14] M. Yu. Ovchinnikov and D. S. Roldugin, "A survey on active magnetic attitude control algorithms for small satellites," *Prog. Aerosp. Sci.*, vol. 109, art. 100546, 2019.

[15] F. L. Markley and J. L. Crassidis, *Fundamentals of Spacecraft Attitude Determination and Control*, ser. Space Technology Library, vol. 33. New York, NY, USA: Springer, 2014.

[16] H. Schaub and J. L. Junkins, *Analytical Mechanics of Space Systems*, 4th ed. Reston, VA, USA: AIAA Education Series, 2018.

[17] J. R. Wertz, D. F. Everett, and J. J. Puschell, Eds., *Space Mission Engineering: The New SMAD*. Hawthorne, CA, USA: Microcosm Press, 2011.

[18] P. Fortescue, G. Swinerd, and J. Stark, Eds., *Spacecraft Systems Engineering*, 4th ed. Chichester, UK: Wiley, 2011.

[19] C. Cappelletti, S. Battistini, and B. K. Malphrus, Eds., *CubeSat Handbook: From Mission Design to Operations*. London, UK: Academic Press, 2021.

[20] J. Solà, "Quaternion kinematics for the error-state Kalman filter," arXiv:1711.02508, 2017.