IAC-26,B4,6A,2,x109468

One Wheel Is Enough: Control Strategies for the 3+1 Actuator Architecture Across Small Satellite Mission Profiles

Niclas Scheuer[a], Patrick McKeen[b*], Kerri Cahoy[b]

[a] *Department of Mechanical and Process Engineering, ETH Zürich, Zürich, Switzerland*
[b] *Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, MA, USA*
[*] *Corresponding author, pmckeen@mit.edu*

Abstract

Small satellites commonly employ either magnetorquers alone or a full reaction-wheel set for attitude control, leaving intermediate actuator architectures comparatively unexplored. This paper investigates one such architecture, the "3+1" architecture, consisting of three orthogonal magnetorquers and a single variable-speed reaction wheel. We ask how much capability of a conventional three-wheel system can be recovered with only one wheel. Geometric analysis shows that the additional wheel substantially expands the instantaneous control authority of the magnetorquer system: in most attitudes, the combined actuators can generate three-axis torque while the wheel contributes to disturbance rejection. The 3+0, 3+1, and 3+3 architectures are compared in disturbance-inclusive 6U simulations using both vector and quaternion attitude objectives. Across the disturbance-inclusive 6U simulations, median final attitude errors ranged from 79°–131° for 3+0, 0.51°–0.74° for 3+1, and 0.47°–0.61° for 3+3, demonstrating that one reaction wheel recovers most of the full configuration's nominal pointing performance. First-order sizing studies relate actuator authority to disturbance-rejection and momentum-management requirements, providing practical design guidance for the 3+1 architecture. The principal limitation of 3+1 is not nominal pointing authority but momentum management. We compare discrete pointing/desaturation mode switching, continuous pointing-and-unloading co-allocation, and a finite-horizon planner that exploits predicted magnetic geometry to jointly manage attitude error, wheel momentum, and actuator constraints. These approaches expose a trade between pointing accuracy, pointing availability, momentum regulation, and onboard compute requirements. The results show that the progression from 3+0 to 3+3 is strongly nonlinear: under the conditions studied, a single reaction wheel moves system performance much closer to 3+3 than to 3+0, making 3+1 a potentially attractive SWaP-C architecture for missions whose pointing requirements are compatible with active momentum management.

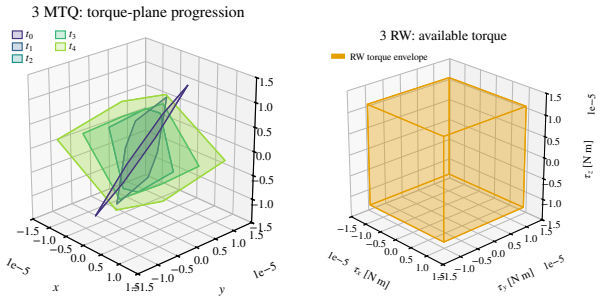
1. Introduction

Small satellite attitude control is normally achieved one of two ways: magnetorquers only or a full set of reaction wheels. A triad of orthogonal magnetorquers (MTQs) is lightweight and has no moving mechanisms, but can only produce torque perpendicular to the local magnetic field. Three or more reaction wheels (RWs) provide fast, precise three-axis control, while adding precision mechanisms and stored momentum that must be unloaded. Between these choices is a system with three magnetorquers and one variable-speed reaction wheel, written 3+1 here.

We evaluate 3+1 as a primary attitude-control architecture and ask what pointing capabilities it can provide and what ultimately limits its performance. Adding a single reaction wheel to three magnetorquers substantially expands the available torque directions, providing instantaneous three-axis authority over much of the operating geometry. This capability comes with a corresponding challenge: the wheel has finite momentum capacity and must eventually be unloaded, potentially interrupting or degrading pointing. This paper therefore examines both sides of the 3+1 architecture: the pointing capability recovered by the first reaction wheel and the mission-level penalty associated with managing its momentum.

1.1 Magnetorquers and Reaction Wheels

Magnetorquers are electromagnets that generate control torque by interacting with an external magnetic field. They are lightweight, inexpensive, power-efficient, and contain no moving parts, making them particularly attractive for small satellites in low Earth orbit [1, 2]. A set of three orthogonal magnetorquers can generate a bounded magnetic



(a) Magnetorquer triad torque envelope.

(b) Reaction-wheel triad torque envelope.

Figure 1: Time-varying torque authority of the conventional actuator architectures: (a) three magnetorquers and (b) three reaction wheels.

dipole $\mathbf{m} \in \mathbb{R}^3$. The resulting control torque is

$$\tau_m = \mathbf{m} \times \mathbf{B}, \tag{1}$$

where $\mathbf{B} \in \mathbb{R}^3$ is the local geomagnetic field. The cross product constrains the instantaneous torque to a plane perpendicular to $\mathbf{B}$, giving only two-dimensional instantaneous torque authority. As the spacecraft progresses along its orbit, however, the magnetic-field direction changes and the available torque plane rotates with it, as illustrated in Figure 1(a).

Reaction wheels generate control torque by accelerating or decelerating a rotating mass, exchanging angular momentum with the spacecraft. A conventional three-wheel set provides direct three-axis torque authority within the torque and momentum limits of the wheels, enabling rapid and precise attitude control. Unlike magnetorquers, however, reaction wheels do not exchange angular momentum with the external environment during normal operation. Consequently, sustained external disturbances cause momentum to accumulate in the wheels until one approaches its maximum speed. This accumulated momentum must then be unloaded using an external actuator, typically magnetorquers or thrusters. Figure 1(b) illustrates the three-dimensional reaction-wheel torque envelope.

1.2 Hardware Choice

Relative to a 3+3 bus, 3+1 removes two reaction wheels and retains the same magnetorquers. Table 1 gives representative per-unit values from published catalog data [3]. At this scale, removing two wheels saves roughly 0.3–0.4 kg and up to tens of watts of combined peak demand. These reductions in size, weight, power, and cost (SWaP-C) repeat across a constellation. It also removes the mounting, balancing, jitter, thermal, acceptance, and lifetime analyses required from those two wheels.

Table 1: Class-level figures for a 15 mN m s CubeSat reaction wheel and a 0.6 A m^2 magnetorquer rod, from published catalog data for flight-heritage units [3]. Per unit.

	Reaction wheel	Magnetorquer rod
Mass	145–200 g	∼40 g
Volume	∼68 cm^3	∼18 cm^3
Peak power	up to 15 W	∼0.45 W
Rotating hardware	rotor, bearings	none

1.3 Relation to prior work

Magnetorquer-only and reaction-wheel systems are both supported by mature control methods. Magnetic control includes B-dot detumbling [4], periodic feedback systems [5, 6], nonlinear and projected magnetic feedback laws [7], sliding-mode [8], and planning-based approaches [9]; see [1] for an overview.

Conventional reaction-wheel systems use direct three-axis feedback control, while magnetorquers or other external actuators are typically used for momentum unloading. Established desaturation approaches include optimal magnetic unloading [10], control allocation that separates pointing and unloading objectives [11], continuous dumping through wheel-set nullspaces [12], and combined globally stabilizing laws [13].

One-wheel systems also have substantial flight heritage, primarily in momentum-bias architectures. SNAP-1 used a pitch-axis momentum wheel and three magnetic rods for approximately 3° nadir pointing [14, 15], while CanX-2 flew the same actuator complement on a three-kilogram nanosatellite [16]. In these systems, the wheel primarily supplied stored angular momentum and gyroscopic stiffness, with the magnetorquers providing attitude correction and momentum regulation. Following reaction-wheel failures, FUSE also operated for years using one remaining wheel and magnetic torquer bars; the wheel controlled one axis while the magnetic actuators controlled the remaining axes, with science observations restricted to attitudes compatible with the available control authority [17, 18].

Direct use of a single variable-speed wheel combined with magnetorquers as a control actuator has received less attention. Predictive-control studies of magnetically actuated dual-spin spacecraft command a wheel about a nominal bias momentum [19, 20], while related CubeSat work has investigated fine pointing with one wheel and magnetic actuation under limited initial errors and disturbance models [21, 22]. In contrast, this work considers 3+1 as a primary, low-bias attitude-control architecture for general slew, acquisition, and pointing tasks, with the wheel and magnetorquers allocated jointly under actuator and momentum constraints.

1.4 Contributions

This work builds on generalized attitude-control [23], estimation [24], and trajectory-planning methods [25, 26] previously developed by the authors. The methods are applied here to characterize 3+1 as an actuator architecture rather than to develop a new architecture-specific control law.

We first analyze how the addition of a single reaction wheel to a magnetorquer triad changes the instantaneous torque geometry relative to 3+0 and 3+3, and use this geometry to predict the resulting slew, disturbance-rejection, and momentum-unloading capabilities in Section 2. We then test these predictions through closed-loop slew and pointing simulations of the three architectures under common conditions in Section 3. Analytical disturbance-rejection and actuator-sizing relationships are developed to identify the regimes in which 3+1 is practical are developed in Section 4. Section 5 concerns itself with the principal operational limitation of the 3+1 architecture, reaction-wheel momentum management. We use two approaches for this problem: discrete mode switching between pointing and desaturation, and continuous co-allocation of pointing and unloading objectives. We quantify the resulting trade between momentum regulation, pointing accuracy, and pointing availability. As a third approach, we evaluate in Section 6 a feasibility-aware trajectory planner that uses knowledge of future magnetic geometry to jointly manage pointing and wheel momentum. Finally, representative mission-level studies in Section 7 evaluate the resulting pointing performance and operational availability of the 3+1 architecture. The resulting architectural trade space and its implications for small-satellite attitude control are discussed in Section 8.

2. Geometric Analysis

2.1 Achievable Torque Sets

Throughout this paper, all vectors and torques are expressed in the spacecraft body frame.

Let n_m and n_w denote the number of magnetorquers and reaction wheels, respectively. Their actuator-axis matrices are defined as

$$\mathbf{A}_m = [\hat{\mathbf{a}}_{m,1} \quad \cdots \quad \hat{\mathbf{a}}_{m,n_m}] \in \mathbb{R}^{3\times n_m},$$
$$\mathbf{A}_w = [\hat{\mathbf{a}}_{w,1} \quad \cdots \quad \hat{\mathbf{a}}_{w,n_w}] \in \mathbb{R}^{3\times n_w},$$

where each column is an actuator axis.

For the magnetorquers, let $\mathbf{u}_m \in \mathbb{R}^{n_m}$ contain the commanded magnetic dipole of each actuator. The total magnetic dipole is

$$\mathbf{m} = \mathbf{A}_m \mathbf{u}_m.$$

The resulting spacecraft torque is

$$\tau_m = \mathbf{m} \times \mathbf{B} = -[\mathbf{B}]_\times \mathbf{A}_m \mathbf{u}_m, \tag{2}$$

where $\mathbf{B} \in \mathbb{R}^3$ is the local geomagnetic field and $[\cdot]_\times$ denotes the skew-symmetric cross-product matrix.

Let the admissible magnetorquer command set be

$$\mathscr{U}_m = \{\mathbf{u}_m \in \mathbb{R}^{n_m} \mid \underline{\mathbf{u}}_m \le \mathbf{u}_m \le \overline{\mathbf{u}}_m\},$$

where the inequalities are applied elementwise. The instantaneous magnetorquer torque envelope is therefore

$$\mathscr{T}_m(\mathbf{B}) = \{-[\mathbf{B}]_\times \mathbf{A}_m \mathbf{u}_m \mid \mathbf{u}_m \in \mathscr{U}_m\}.$$

For the reaction wheels, let $\mathbf{u}_w \in \mathbb{R}^{n_w}$ contain the commanded wheel torques. The corresponding torque applied to the spacecraft is

$$\tau_w = -\mathbf{A}_w \mathbf{u}_w. \tag{3}$$

With the admissible command set

$$\mathscr{U}_w = \{\mathbf{u}_w \in \mathbb{R}^{n_w} \mid \underline{\mathbf{u}}_w \le \mathbf{u}_w \le \overline{\mathbf{u}}_w\},$$

the reaction-wheel torque envelope is

$$\mathscr{T}_w = \{-\mathbf{A}_w \mathbf{u}_w \mid \mathbf{u}_w \in \mathscr{U}_w\}.$$

Because the torques produced by the two actuator classes add, the instantaneous torque envelope of the combined actuator architecture is the Minkowski sum

$$\boxed{\mathscr{T}(\mathbf{B}) = \mathscr{T}_m(\mathbf{B}) \oplus \mathscr{T}_w.} \tag{4}$$

Thus, for any instantaneous magnetic-field vector, the complete set of body torques available to the spacecraft is obtained by adding every feasible magnetorquer torque to every feasible reaction-wheel torque. Figure 2 visualizes these sets for the 3+0, 3+1, and 3+3 architectures.

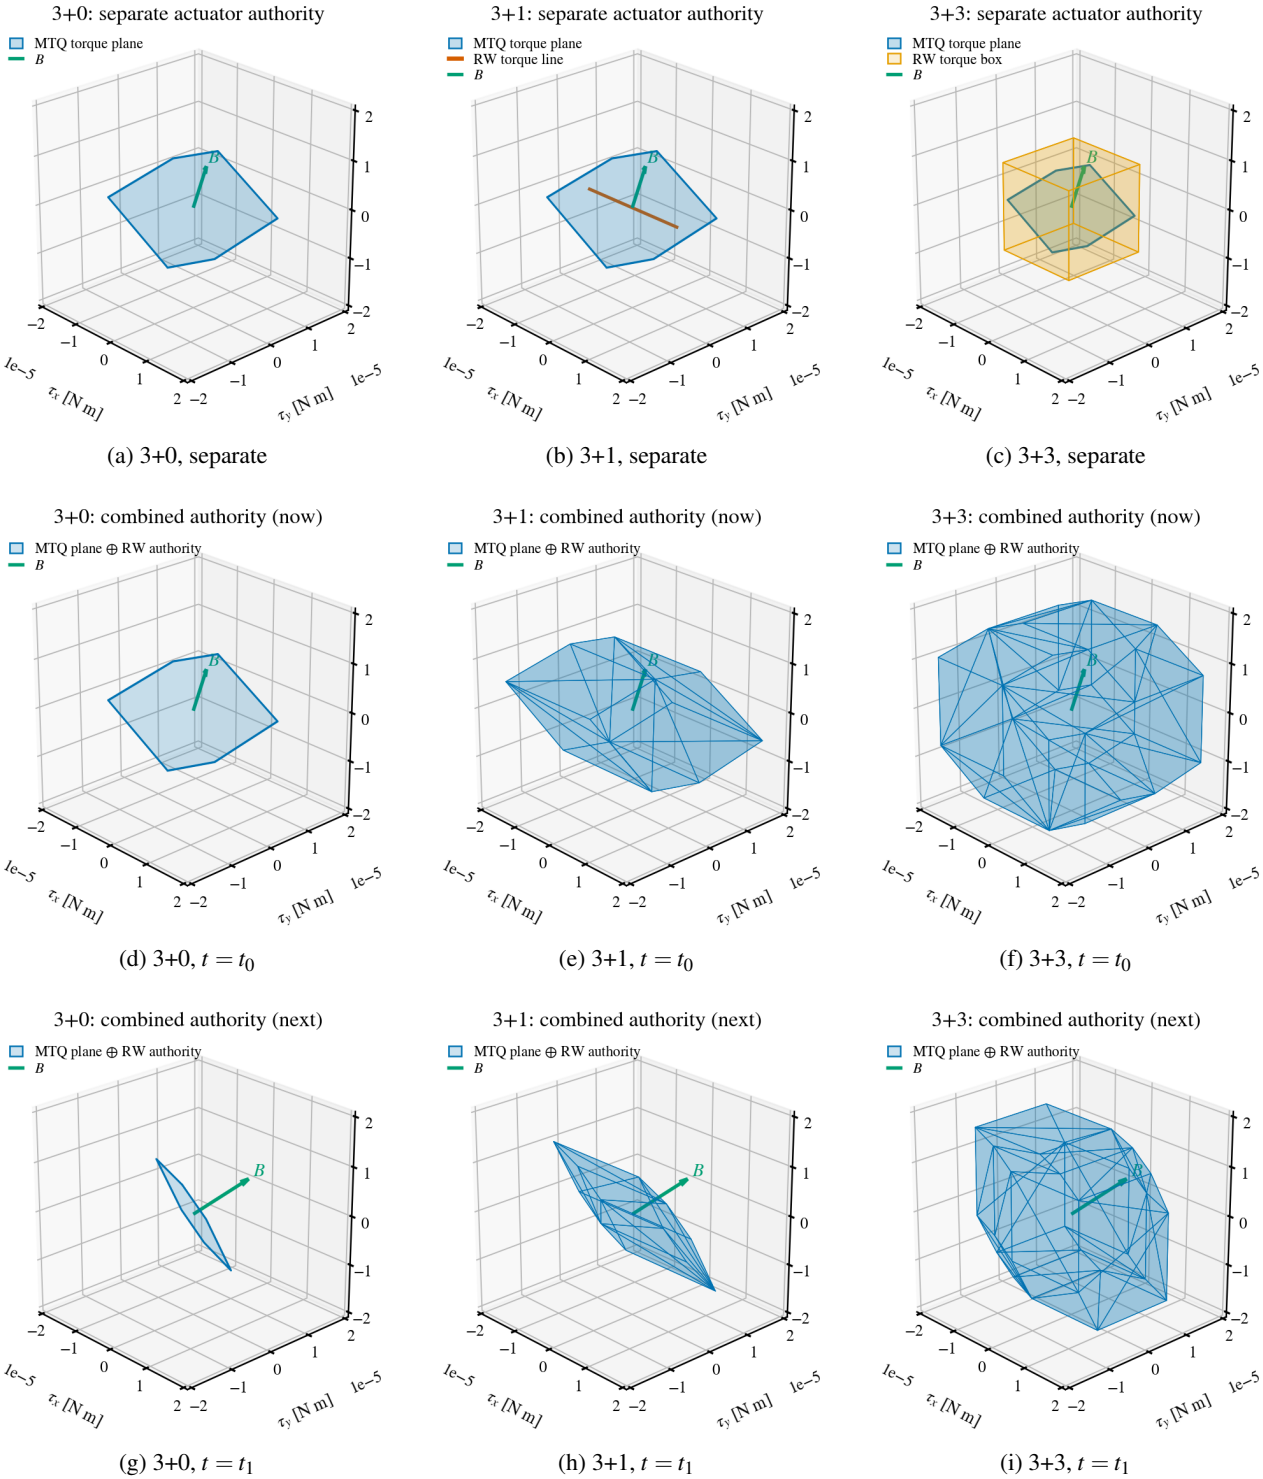
2.2 Torque Generation

The dimensionality of the achievable torque set provides a direct geometric measure of instantaneous torque authority. For 3+0, $\mathscr{T} = \mathscr{T}_m$ is confined to the plane perpendicular to $\mathbf{B}$ and therefore has only two-dimensional instantaneous torque authority, as shown in Figs. 2d and 2g. As the spacecraft progresses along its orbit, however, the direction of $\mathbf{B}$ changes and the magnetic torque plane rotates. Consequently, torque directions unavailable at one instant can become available later.

Adding a single reaction wheel fundamentally changes this geometry. The wheel contributes torque along its axis $\hat{\mathbf{a}}_w$, such that the combined 3+1 torque set spans $\mathbb{R}^3$ whenever

$$\hat{\mathbf{a}}_w^T \mathbf{B} \neq 0.$$

Thus, 3+1 provides instantaneous three-dimensional torque authority for almost all wheel–field geometries, as illustrated in Figs. 2e and 2h. Only when $\hat{\mathbf{a}}_w^T \mathbf{B} = 0$, with the wheel axis lying in the magnetic torque plane, does the set become two-dimensional. Near this singular geometry, three-dimensional authority is retained but weakens in the added direction. Unlike 3+0, which relies on temporal variation of the magnetic field for three-axis control,



(a) 3+0, separate (b) 3+1, separate (c) 3+3, separate

(d) 3+0, $t = t_0$ (e) 3+1, $t = t_0$ (f) 3+3, $t = t_0$

(g) 3+0, $t = t_1$ (h) 3+1, $t = t_1$ (i) 3+3, $t = t_1$

Figure 2: Achievable torque sets for the 3+0, 3+1, and 3+3 architectures with the instantaneous magnetic field direction B shown. The first row shows the magnetorquer and reaction-wheel torque sets separately. The second row shows their combined achievable torque sets at the initial magnetic-field orientation, and the third row shows the combined sets after the magnetic-field orientation has changed. The combined torque set is the Minkowski sum of the individual actuator torque sets.

3+1 therefore provides direct three-axis authority almost everywhere.

2.3 Disturbance Rejection

Torque generation only considers whether a desired torque direction is available; disturbance rejection additionally requires sufficient authority in that direction. For an external disturbance τ_d, instantaneous rejection is possible if

$$-\tau_d \in \mathscr{T}(\mathbf{B}).$$

For a spacecraft expected to operate over arbitrary attitudes, disturbances may act in any body-frame direction. A useful measure of disturbance rejection is therefore the largest torque magnitude that can be opposed in every direction.

For 3+0, the instantaneous torque set is planar and consequently always contains an uncontrolled direction parallel to $\mathbf{B}$. Thus, no nonzero disturbance magnitude can be rejected in every direction instantaneously. The addition of one reaction wheel removes this null direction for almost all magnetic geometries, giving 3+1 finite three-dimensional disturbance-rejection authority. However, as $\hat{\mathbf{a}}_w$ approaches perpendicularity with $\mathbf{B}$, the combined torque volume becomes increasingly thin and the minimum available torque over all directions decreases, approaching zero at the singular geometry $\hat{\mathbf{a}}_w^T\mathbf{B} = 0$. A nonsingular 3+3 wheel set, in contrast, maintains finite torque authority in all three dimensions independently of the magnetic-field direction. Section 4 quantifies these disturbance-rejection margins relative to representative environmental torques and actuator sizes.

2.4 Momentum Management

Reaction wheels accumulate momentum and must eventually be unloaded. $\mathscr{T}_{\text{desat}}$ denotes the torque space where reaction wheel and magnetorquer torque can cancel, enabling desaturation of the wheels without torquing the satellite. It is defined as

$$\boxed{\mathscr{T}_{\text{desat}} = \mathscr{T}_m \ominus (-\mathscr{T}_w)} \tag{5}$$

where $\ominus$ denotes the Pontryagin difference. For 3+1, a nonzero wheel-unloading torque can be cancelled exactly by the magnetorquers only when the wheel axis lies in the magnetic torque plane, i.e.,

$$\hat{\mathbf{a}}_w^T\mathbf{B} = 0.$$

At all other instantaneous geometries, torque-free unloading is not possible.

The simultaneous pointing and unloading problem can be written directly as a combined desired torque τ_{ref}

$$\tau_{\text{ref}}^{\star} = \tau_p + \tau_w,$$

with feasibility requiring

$$\tau_p + \tau_w \in \mathscr{T}_m(\mathbf{B}(t)).$$

This is the same time-varying feasibility problem as 3+0 pointing: a desired torque must lie within the instantaneous magnetic torque set, whose orientation changes with $\mathbf{B}(t)$. For 3+1 desaturation, the desired magnetic torque simply contains both the pointing and unloading objectives. Thus, the geometric constraints governing 3+0 pointing reappear in 3+1 as constraints on simultaneous pointing and momentum unloading.

2.5 Geometric Summary

The geometric analysis predicts three key properties of the 3+1 architecture. First, adding a single reaction wheel expands the instantaneous torque set from a plane to a three-dimensional volume for almost all magnetic geometries. Second, disturbance rejection is similarly improved but remains dependent on the relative wheel–field geometry, with reduced authority near the singular configuration. Finally, momentum management emerges as the principal limitation: simultaneous pointing and torque-free desaturation reduce to the same time-varying magnetic feasibility problem as classical 3+0 control. The following sections test these predictions through closed-loop simulation and mission-level evaluation.

3. Slewing Performance

3.1 Control and Allocation

Slewing performance is compared for the 3+0, 3+1, and 3+3 architectures using the common 6U spacecraft and simulation environment summarized in Table 2. The same high-level attitude controller is used in each case, so differences in performance arise primarily from the available actuator authority and its allocation. Purposefully-tuned controllers for each architecture (especially 3+0) may have small improvements in performance, especially at longer time horizons. The shared controller is retained for comparability.

Following the generalized actuator-allocation framework of [23], the attitude controller first generates a reference body torque τ_{ref}. For a quaternion attitude objective, the attitude error is formed from the current attitude q and desired attitude q_d as

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

and a proportional-derivative control law generates

$$\tau_p = -k_p \, \text{sgn}(q_{e,0}) \, \mathbf{q}_{e,v} - k_d \, \omega_e,$$

where $q_{e,0}$ and $\mathbf{q}_{e,v}$ are the scalar and vector components of the error quaternion and ω_e is the angular-rate error.

The allocator then maps the commanded body torque to feasible magnetorquer and reaction-wheel commands. Using the actuator mappings defined in Equation 2 & 3, the instantaneous torque mapping is

Table 2: Reference 6U estimator-in-the-loop Monte Carlo setup. Each case uses 100 randomized runs; architecture, allocator, and pointing-goal type are varied.

Spacecraft	
Bus	6U, 12.0 kg
Inertia	$\mathrm{diag}(0.13, 0.10, 0.05)$ kg m^2
Actuators	3 body-axis MTQs; 0, 1, or 3 reaction wheels
Actuator limits	0.6 A m^2 MTQ dipole; 2 mN m wheel torque; 15 mN m s capacity
Environment and estimation	
Orbit	Circular, 400 km, $i = 97^\circ$, period 5554 s
Disturbances	Gravity-gradient, drag, SRP, and residual dipole
Estimator update	1 s; rotation-vector quaternion error representation
Sensors	3 gyros, 3 magnetometers, and 2 star trackers
Simulation and control	
Time step	$\Delta t = 1$ s
Duration	One orbit, 5554 s by default
Monte Carlo	100 randomized runs per case
Gains	$k_p = 2.9 \times 10^{-4}$, $k_d = 8.7 \times 10^{-3}$, $k_h = 1.0 \times 10^{-3}$
Allocators	LP and QP
Pointing goals	Random quaternion and vector pointing
Architectures	3+0, 3+1, and 3+3

$$\tau = \underbrace{\begin{bmatrix} -[\mathbf{B}]_\times \mathbf{A}_m & -\mathbf{A}_w \end{bmatrix}}_{\mathbf{G}(\mathbf{B})} \begin{bmatrix} \mathbf{u}_m \\ \mathbf{u}_w \end{bmatrix} = \mathbf{G}(\mathbf{B})\,\mathbf{u},$$

subject to the actuator bounds $\underline{\mathbf{u}} \le \mathbf{u} \le \overline{\mathbf{u}}$. Two generalized allocators are considered: A linear-programming allocator (LP), and a quadratic-programming allocator (QP).

The linear-programming allocator preserves the direction of the requested torque and finds the largest feasible magnitude in that direction:

$$\begin{aligned} \max_{\mathbf{u},\alpha} \quad & \alpha \\ \text{s.t.} \quad & \mathbf{G}(\mathbf{B})\mathbf{u} = \alpha \hat{\tau}_{\mathrm{ref}}, \\ & \underline{\mathbf{u}} \le \mathbf{u} \le \overline{\mathbf{u}}, \\ & 0 \le \alpha \le \|\tau_{\mathrm{ref}}\|, \end{aligned} \tag{6}$$

where $\hat{\tau}_{\mathrm{ref}} = \tau_{\mathrm{ref}}/\|\tau_{\mathrm{ref}}\|$. The LP therefore returns either the requested torque or the largest feasible scaled torque parallel to it.

The quadratic-programming allocator instead minimizes the Euclidean error between the requested and achievable torque:

$$\begin{aligned} \min_{\mathbf{u}} \quad & \|\mathbf{G}(\mathbf{B})\mathbf{u} - \tau_{\mathrm{ref}}\|_2^2 \\ \text{s.t.} \quad & \underline{\mathbf{u}} \le \mathbf{u} \le \overline{\mathbf{u}}. \end{aligned} \tag{7}$$

Unlike the LP, the QP may rotate the achieved torque away from the requested direction when the command lies outside the achievable torque set. Figure 4 illustrates this geometric distinction.

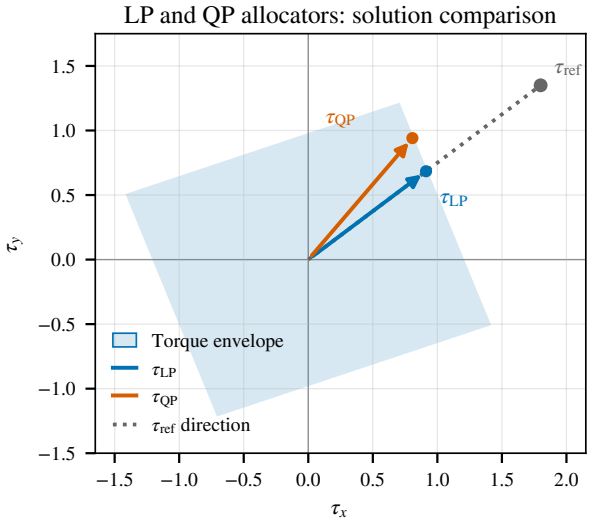


Figure 3: LP

Figure 4: Geometric interpretation of generalized torque allocation. The LP preserves the requested torque direction and selects the largest feasible magnitude along that direction. The QP minimizes Euclidean torque error and may therefore change the torque direction when the requested torque lies outside the achievable set.

Both full-quaternion and vector-pointing objectives are considered. The former specifies the complete spacecraft attitude, whereas the latter only matches a body-fixed boresight to an inertial direction and leaves rotation about that direction unconstrained.

3.2 Pointing Performance

The disturbance-inclusive pointing results are shown in Figure 5, while Figure 6 shows the clearest distinction: 3+0 does not consistently converge to the commanded pointing objective, whereas both 3+1 and 3+3 rapidly reduce the initial attitude error.

The final-error results show the same trend (Figure 5). For vector pointing, median errors are 0.54° (LP) and 0.51° (QP) for 3+1, compared with 86.39° and 79.31° for 3+0 and approximately 0.47° for 3+3. For quaternion pointing, the corresponding medians are 0.74° and 0.68° for 3+1, 131.37° and 112.89° for 3+0, and approximately 0.61° for 3+3. The architecture-level median trend is consistent across both pointing objectives and both LP and QP allocation, although the 3+1 cases retain occasional high-error outliers.

These results support the geometric analysis of Sec-

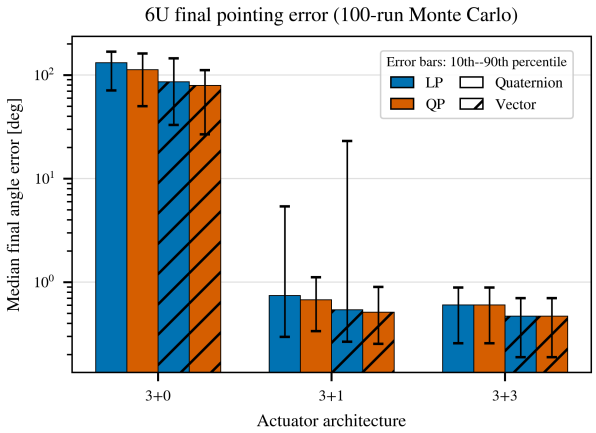


Figure 5: Disturbance-inclusive final vector and quaternion pointing error for the 3+0, 3+1, and 3+3 architectures using LP and QP allocation. Error bars denote the median and 10th–90th percentile range over 100 Monte Carlo runs.

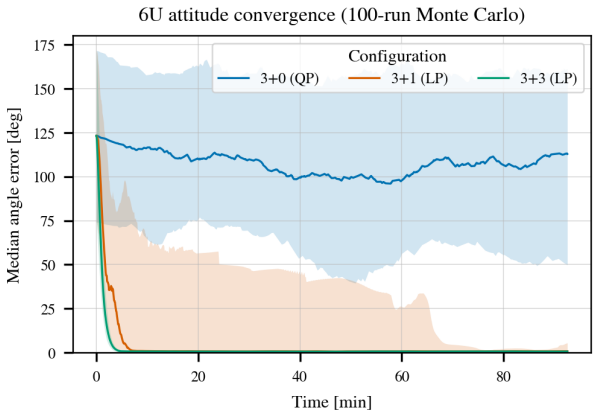


Figure 6: Disturbance-inclusive vector-pointing convergence for the 3+0, 3+1, and 3+3 architectures. Lines denote the median and shaded regions the 10th–90th percentile range over 100 Monte Carlo runs.

tion 2.2: the first reaction wheel provides a qualitative increase in instantaneous torque authority. For the cases considered here, this moves 3+1 pointing performance substantially closer to 3+3 than to 3+0.

4. Disturbance Rejection

Disturbance rejection can be evaluated directly using the LP allocator of Equation 6. For a disturbance torque τ_d, the requested torque is simply set to

$$\tau_{\text{ref}} = -\tau_d. \tag{8}$$

The disturbance is fully rejected when the allocator can realize the complete requested torque, i.e.,

$$\alpha^\star = \|\tau_d\|. \tag{9}$$

As shown in Section 2, the instantaneous authority of the 3+1 architecture is geometry-dependent and can become singular. Consequently, no finite actuator sizing can guarantee rejection of an arbitrary disturbance for all geometries. We therefore characterize disturbance rejection statistically over randomized disturbance and magnetic-field directions.

Figure 7 provides a first-order actuator-sizing map for this trade. For several disturbance magnitudes τ_d, the curves show combinations of reaction-wheel torque and magnetorquer dipole for which 90% of the sampled geometries permit full disturbance rejection. The characteristic knee of each curve identifies a balanced sizing region: increasing either actuator substantially beyond this point provides little additional benefit unless the other is also increased. More generally, the architecture presents a trade between disturbance magnitude and geometric coverage: larger disturbances require greater actuator authority, while requiring rejection over a larger fraction of attitudes and magnetic-field directions further increases the required sizing. A designer may therefore estimate τ_d using a mission-specific disturbance model and use the figure to select an appropriate 3+1 actuator combination. Representative flown actuator combinations are included for comparison.

5. Momentum Management

As discussed in Section 2.4, momentum management is the principal operational challenge of the 3+1 architecture. We therefore consider two simple strategies that trade pointing performance against momentum regulation. The first uses explicit switching between pointing and desaturation modes, while the second blends the two objectives continuously.

5.1 Mode-Switching Desaturation

The mode-switching strategy separates pointing and momentum management into two operating modes, as illustrated in Figure 8. During normal operation, the spacecraft performs vector pointing using the LP allocator and the reaction wheel is used without an explicit unloading objective. When the wheel momentum exceeds $0.75\,h_{\max}$, the controller enters a dedicated desaturation mode.

During desaturation, pointing is temporarily released and the magnetorquers are used to reduce wheel momentum while damping spacecraft rates. Let

$$\hat{\mathbf{B}} = \frac{\mathbf{B}}{\|\mathbf{B}\|}, \qquad \omega_\perp = \left(\mathbf{I} - \hat{\mathbf{B}}\hat{\mathbf{B}}^T\right)\omega.$$

A magnetic rate-damping term is combined with a wheel-momentum dumping command and projected into the

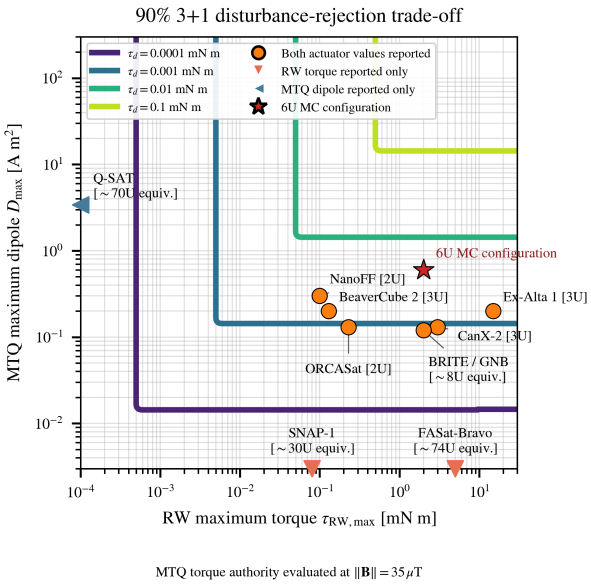


Figure 7: First-order 3+1 disturbance-rejection sizing trade-off. Each curve gives combinations of reaction-wheel torque and magnetorquer dipole that fully reject a disturbance of magnitude τ_d for 90% of sampled disturbance and magnetic-field directions. MTQ authority is evaluated at $\|\mathbf{B}\| = 35\ \mu$T. Representative spacecraft actuator combinations are shown where published values are available.



Figure 8: Mode-switching momentum-management logic for the 3+1 architecture. The spacecraft enters desaturation mode at 75% wheel momentum capacity and returns to pointing mode once the momentum is reduced below 25%.

achievable magnetic torque plane,

$$\tau_{\mathrm{ref}}^{\star} = -k_d \boldsymbol{\omega}_{\perp} - \left(\mathbf{I} - \hat{\mathbf{B}}\hat{\mathbf{B}}^T\right) k_h \left(\mathbf{h}_{\mathrm{rw}} - \mathbf{h}_{\mathrm{target}}\right).$$

The corresponding actuators are obtained using the allocators from Section 3. Because pointing is not maintained during this mode, the available magnetic authority can be devoted primarily to reducing wheel momentum and body rates. Once the wheel momentum falls below $0.25\,h_{\max}$, the controller returns to normal pointing mode.

Figure 9 shows the mode-switching strategy for the 6U spacecraft using the configuration from Table 2, including a representative pointing and desaturation time history together with the resulting Monte Carlo mode statistics. The fraction of time spent in desaturation provides a direct measure of the operational cost of momentum management, since this time is unavailable for useful pointing. The pointing uptime is initially low because the preceding slew drives the reaction wheel toward saturation, requiring frequent desaturation early in the run. As the simulation progresses, this transient effect diminishes and the uptime approaches approximately 90%, with subsequent saturation driven primarily by the accumulation of disturbance-induced momentum.

5.2 Co-Allocation

Mode switching preserves pointing performance while in the pointing mode, but requires dedicated desaturation intervals during which useful pointing is lost. An alternative is to unload the reaction wheel continuously while maintaining attitude control, accepting some degradation in pointing accuracy.

We implement this by blending the actuator commands generated by the nominal pointing controller and the desaturation controller,

$$\boxed{\mathbf{u} = (1-\gamma)\mathbf{u}_{\mathrm{point}} + \gamma \mathbf{u}_{\mathrm{desat}}, \qquad 0 \le \gamma \le 1.} \tag{10}$$

The pointing command $\mathbf{u}_{\mathrm{point}}$ is obtained from the same LP allocation used in Equation 6, while $\mathbf{u}_{\mathrm{desat}}$ is generated using the momentum-dumping and rate-damping law of Section 5.1.

Figure 10 shows the effect of varying γ for the 6U spacecraft. At low values of γ, insufficient unloading causes the wheel to saturate rapidly, degrading pointing performance. At high values, excessive authority is devoted to desaturation, again increasing pointing error. Between these extremes, a broad intermediate region maintains wheel momentum while retaining high pointing accuracy. Co-allocation therefore replaces the discrete pointing downtime of mode switching with a tunable continuous trade between pointing and momentum management.

6. Finite-Horizon Planning

The momentum-management strategies of Section 5 are reactive: they respond to the instantaneous attitude error and wheel state without considering how the magnetic control geometry will evolve. A finite-horizon planner can instead use this future information when selecting the attitude and actuator trajectory.

We use a feasibility-aware trajectory planner [25] based

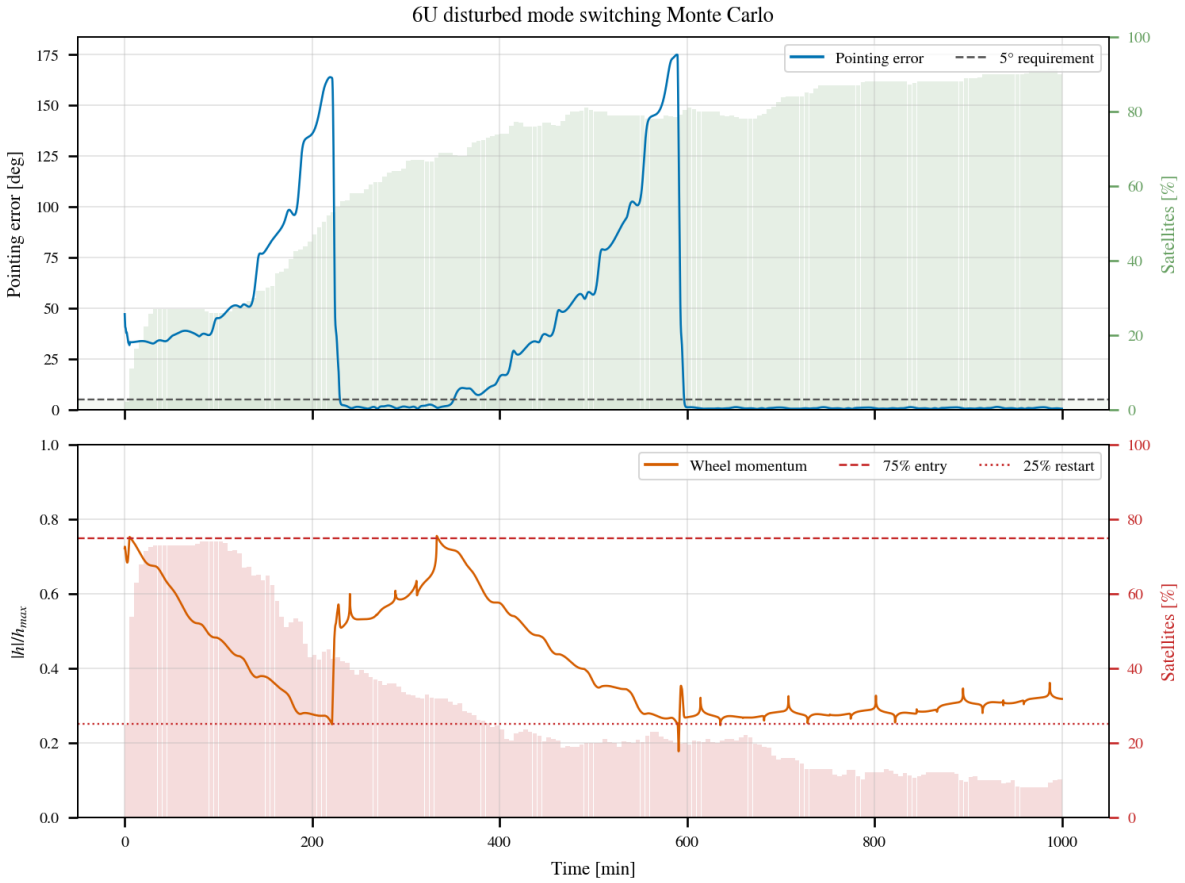


Figure 9: Mode-switching momentum-management performance for the 6U spacecraft over 100 Monte Carlo simulations. The bars show, as a function of time, the fraction of Monte Carlo runs operating in pointing and desaturation modes. The line shows a single representative Monte Carlo run to illustrate the mode-switching behavior.

on the Augmented Lagrangian Trajectory Optimizer (ALTRO) [27]. The planner propagates a model of the spacecraft, orbit, and geomagnetic field over a finite horizon and optimizes the complete actuator trajectory subject to the spacecraft dynamics and actuator constraints. Pointing error, reaction-wheel momentum, and control effort are included simultaneously in the objective. The resulting trajectory can therefore sacrifice locally optimal pointing or actuation when doing so improves future magnetic authority or reduces wheel momentum. This foresight comes at the cost of substantially greater onboard computation than the reactive methods considered above.

Figure 11 compares the wheel momentum deposited during a representative slew. Mode switching and co-allocation leave approximately 6.98 and 5.05 mN m s of peak momentum change, respectively, whereas the feasibility-aware planner reduces this to 0.54 mN m s. Rather than correcting momentum accumulation after it occurs, the planner selects a slew trajectory that avoids depositing most of the momentum in the first place.

The resulting pointing–momentum trade is shown in Figure 12. Reactive mode switching and the co-allocation sweep expose the expected compromise between final pointing accuracy and remaining wheel momentum. The planner moves the solution toward the lower-left region of this trade space, simultaneously achieving low pointing error and low stored momentum.

7. Mission Capability and Uptime

7.1 BeaverCube II Case Study

To evaluate the 3+1 architecture in a representative mission scenario, we consider BeaverCube II, a 3U CubeSat requiring ground-station tracking [28]. Relevant spacecraft and mission parameters are given in Table 3. Unlike inertial

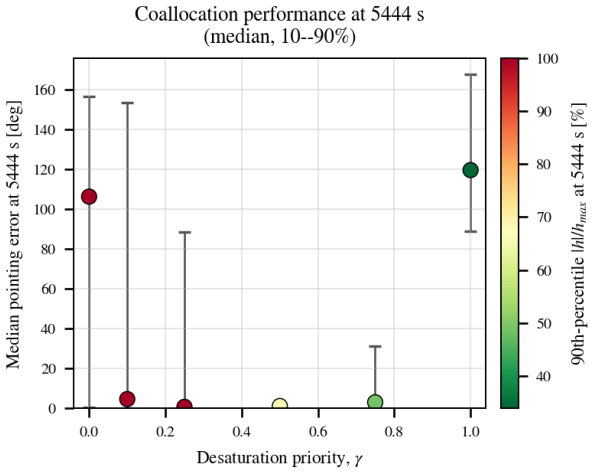


Figure 10: Co-allocation performance for the 6U spacecraft as a function of desaturation priority γ. Markers show the median final pointing error, with error bars spanning the 10th–90th percentiles, while marker color indicates the 90th-percentile reaction-wheel momentum as a percentage of maximum capacity.

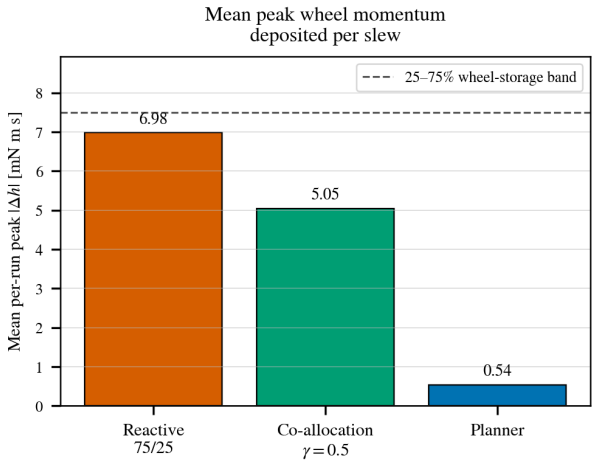


Figure 11: Peak reaction-wheel momentum deposited during a representative 6U slew for reactive mode switching, continuous co-allocation with $\gamma = 0.5$, and feasibility-aware finite-horizon planning. The planner substantially reduces momentum accumulation by accounting for future actuator and magnetic-field geometry when constructing the slew.

pointing, this task requires continuous tracking of a moving ground target throughout the pass.

Using 3+1 with LP allocation, the spacecraft successfully acquires the ground station from randomized initial attitudes and tracks it through the pass (Figure 13). The median pointing error falls below the 5° connection requirement shortly after acquisition. Across the 100 Monte Carlo runs, the median usable contact time is 10.12 min, with a 10–90 percentile range of 7.25–11.33 min, compared with a geometric maximum of 11.33 min, demonstrating effective mission-level tracking with the 3+1 architecture.

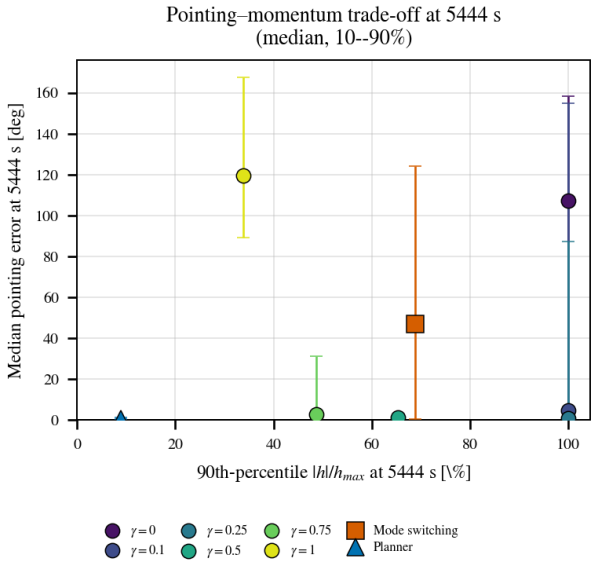


Figure 12: Pointing–momentum trade-off for the 6U spacecraft. Co-allocation results are shown across desaturation priorities γ, together with mode switching and the feasibility-aware planner. The planner achieves both low final pointing error and low reaction-wheel momentum by optimizing the maneuver over a finite horizon.

7.2 Uptime Performance

The preceding sections establish that 3+1 can achieve useful pointing with LP allocation, but sustained mission performance depends on how reaction-wheel momentum is managed. We therefore combine the 6U pointing results with the momentum accumulation and desaturation behavior of the mode-switching, co-allocation, and finite-horizon planning strategies.

Figure 14 shows the fraction of time spent within 5° and 1° of the commanded pointing direction. At the 5° level, co-allocation outperforms reactive mode switching. By continuously unloading the wheel, it reduces the frequency of large momentum excursions and therefore avoids some of the explicit pointing downtime associated with mode switching. The trade reverses for fine pointing: mode switching spends more time within 1°, since pointing and desaturation are separated rather than simultaneously competing for actuator authority. The planner provides the highest 5° availability and comparable fine-pointing performance,

Table 3: BeaverCube II controller-in-the-loop evaluates a 3-MTQ, 1-reaction-wheel spacecraft during a simulated ground-station pass.

Spacecraft	
Bus	BeaverCube II, 4 kg
Inertia	$\begin{bmatrix} 0.03136 & 0.00006 & -0.00671 \\ 0.00006 & 0.03409 & -0.00012 \\ -0.00671 & -0.00012 & 0.01004 \end{bmatrix}$ kg m^2
Actuators	Three orthogonal MTQs and one $+\hat{z}$ reaction wheel
Limits	0.2 A m^2 MTQ dipole; 2.3 mN m wheel torque; 3.6 mN m s capacity
Boresight	$+\hat{y}$
Environment and simulation	
Orbit	Circular, 635 km altitude
Disturbances	Gravity-gradient, drag, and solar-radiation pressure
Ground pass	Boston, 42.36° N, 71.06° W
Time step	$\Delta t = 2$ s
Duration	1000 s horizon-to-horizon pass
Monte Carlo	100 randomized runs
Control	
Controller	LP MTQ–reaction-wheel co-allocation
Gains	$k_p = 5.0 \times 10^{-5}$, $k_d = 2.0 \times 10^{-3}$
Pointing	Ground-station line-of-sight tracking with 5° link criterion

although its 1° result is slightly below the reactive case.

The distinction becomes more pronounced as slew demand increases. Figure 15 estimates daily pointing availability using the measured momentum deposited per slew together with the corresponding desaturation burden. Reactive mode switching degrades first as repeated slews drive increasingly frequent desaturation episodes. Co-allocation extends the useful operating range by reducing momentum accumulation, but eventually also loses pointing availability as slew demand rises. The planner remains substantially less sensitive because it shapes each maneuver to limit momentum buildup before desaturation is required.

These results highlight that the principal benefit of improved momentum management is not necessarily lower instantaneous pointing error, but greater mission availability. Mode switching favors fine-pointing quality at the cost of explicit downtime, co-allocation trades some accuracy for fewer desaturation interruptions, and finite-horizon planning achieves the highest overall uptime by reducing the momentum burden at its source.

8. Discussion

The results consistently show that the transition from 3+0 to 3+1 is qualitatively more important than the actuator count alone suggests. The geometric analysis of Section 2 shows that adding a single reaction wheel expands the instantaneous torque set from a plane to a three-dimensional

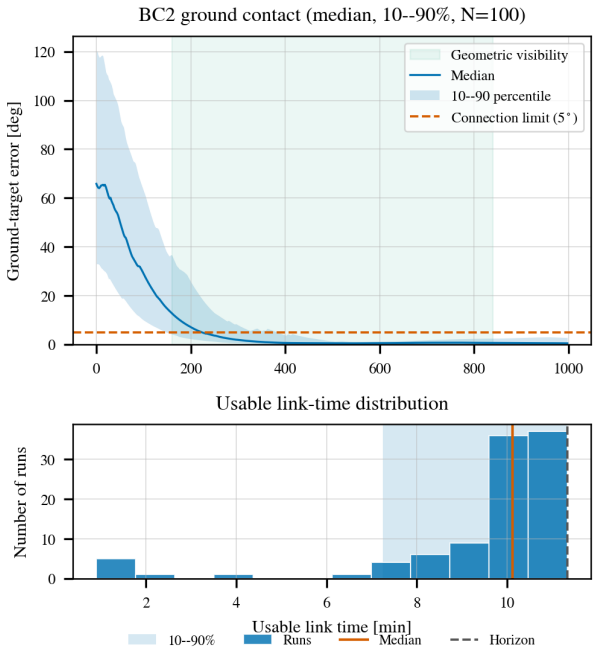


Figure 13: BeaverCube II ground-station tracking with 3+1 and LP allocation. Top: median pointing error with the 10–90 percentile band relative to the 5° connection limit. Bottom: histogram of usable contact time across 100 Monte Carlo runs; the median is 10.12 min and the 10–90 percentile range is 7.25–11.33 min.

volume for almost all magnetic geometries (Figure 2). The closed-loop results of Section 3 confirm the practical consequence: 3+1 pointing performance lies much closer to 3+3 than to 3+0 across both vector and quaternion objectives (Figure 5).

This improvement does not remove the geometry dependence of magnetic actuation. The disturbance-rejection sizing map of Figure 7 shows that useful 3+1 performance depends on the balance between reaction-wheel torque and magnetorquer authority, with diminishing benefit from oversizing either actuator independently. Together with the SWaP-C comparison of Table 1, this defines the main hardware trade: 3+1 can recover much of the control capability of 3+3 while retaining a substantially reduced wheel complement.

The principal limitation is reaction-wheel momentum management. As predicted by Section 2.4, arbitrary pointing and torque-free unloading are generally incompatible. The mode-switching, co-allocation, and finite-horizon approaches of Sections 5 and 6 therefore trade pointing accuracy, downtime, momentum regulation, and computational complexity. The resulting mission-level availability trends

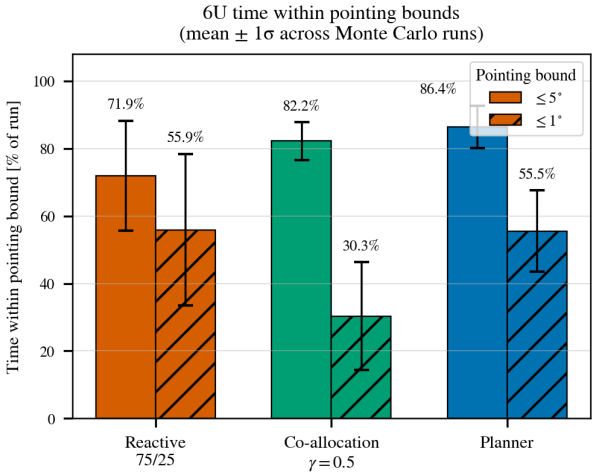


Figure 14: Fraction of 6U simulation time spent within 5° and 1° of the commanded pointing direction for reactive mode switching, co-allocation with $\gamma = 0.5$, and the feasibility-aware planner. Co-allocation improves coarse-pointing availability by reducing explicit desaturation downtime, while mode switching retains better fine-pointing accuracy.

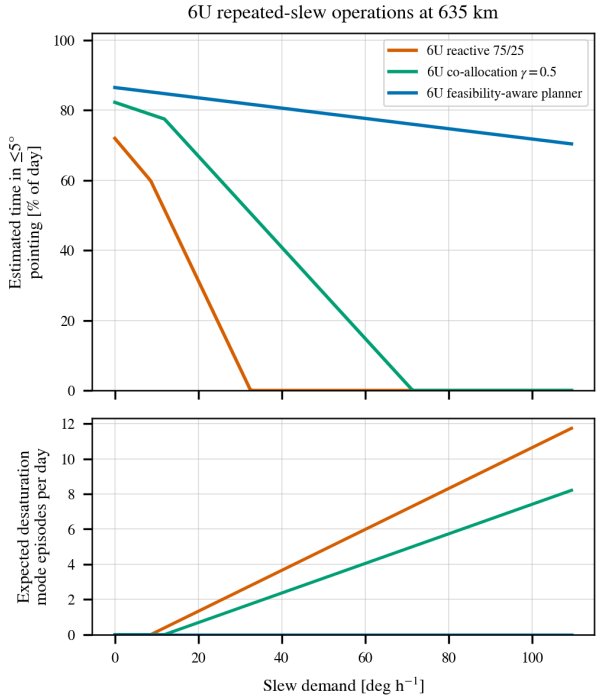


Figure 15: Estimated 6U operational uptime as a function of repeated-slew demand. Top: fraction of the day spent within 5° pointing. Bottom: expected number of desaturation-mode episodes per day. Reactive mode switching and co-allocation lose availability as slew demand increases, whereas finite-horizon planning remains effective over a much broader operating range by limiting wheel-momentum accumulation during each slew.

in Figs. 14 and 15 indicate that momentum management, rather than nominal pointing authority, is the main operational constraint on 3+1.

These results suggest that 3+1 is most attractive for SWaP-C-constrained missions that require substantially greater pointing capability than 3+0, but can tolerate geometry-dependent authority and active momentum management. Missions requiring uninterrupted high-precision three-axis control or frequent aggressive slewing remain better suited to conventional multiwheel architectures.

9. Conclusion

This paper evaluates the 3+1 actuator architecture as an intermediate option between magnetorquer-only and conventional multiwheel attitude control. The addition of a single reaction wheel fundamentally changes the available control geometry, providing three-dimensional instantaneous torque authority for almost all magnetic-field orientations. Under the conditions studied, this moves 3+1 pointing performance substantially closer to 3+3 than to 3+0.

The remaining limitation is not nominal pointing capability, but the finite momentum capacity of the single wheel. Disturbance rejection therefore depends on balanced actuator sizing, while sustained operation requires explicit momentum-management strategies. Reactive mode switching, continuous co-allocation, and finite-horizon planning provide progressively different trade-offs between pointing accuracy, availability, momentum regulation, and computational cost.

Overall, 3+1 should not be viewed as a replacement for 3+3, but as a distinct intermediate architecture. For missions whose pointing and availability requirements are compatible with active momentum management, one reaction wheel can recover a disproportionate fraction of the capability of a full wheel set while preserving much of the hardware simplicity of magnetic actuation.

References

[1] M. Y. Ovchinnikov and D. S. Roldugin, "A survey on active magnetic attitude control algorithms for small satellites," *Progress in Aerospace Sciences*, vol. 109, p. 100546, 2019.

[2] F. Miranda, "Guidance stabilization of satellites using the geomagnetic field," *International Journal of Aerospace Engineering*, vol. 2012, p. 231935, 2012.

[3] "Product datasheets: AAC Clyde Space RW400-series reaction wheels; CubeSpace CubeWheel series; NewSpace Systems NCTR-M-series magnetorquer rods," Product datasheets aggregated at satsearch.co, 2026, accessed September 2026.

[4] A. C. Stickler and K. T. Alfriend, "Elementary magnetic attitude control system," *Journal of Spacecraft and Rockets*, vol. 13, no. 5, pp. 282–287, 1976.

[5] M. Lovera, E. De Marchi, and S. Bittanti, "Periodic attitude control techniques for small satellites with magnetic actuators," *IEEE Transactions on Control Systems Technology*, vol. 10, no. 1, pp. 90–95, 2002.

[6] M. E. Pittelkau, "Optimal periodic control for spacecraft pointing and attitude determination," *Journal of Guidance, Control, and Dynamics*, vol. 16, no. 6, pp. 1078–1084, 1993.

[7] M. Lovera and A. Astolfi, "Spacecraft attitude control using magnetic actuators," *Automatica*, vol. 40, no. 8, pp. 1405–1414, 2004.

[8] R. Wisniewski, "Linear time-varying approach to satellite attitude control using only electromagnetic actuation," *Journal of Guidance, Control, and Dynamics*, vol. 23, no. 4, pp. 640–647, 2000.

[9] A. Gatherer and Z. Manchester, "Magnetorquer-only attitude control of small satellites using trajectory optimization," in *Proceedings of the AAS/AIAA Astrodynamics Specialist Conference*, Aug. 2019.

[10] F. Giulietti, A. A. Quarta, and P. Tortora, "Optimal control laws for momentum-wheel desaturation using magnetorquers," *Journal of Guidance, Control, and Dynamics*, vol. 29, no. 6, pp. 1464–1468, 2006.

[11] J.-F. Trégouët, D. Arzelier, D. Peaucelle, C. Pittet, and L. Zaccarian, "Reaction wheels desaturation using magnetorquers and static input allocation," *IEEE Transactions on Control Systems Technology*, vol. 23, no. 2, pp. 525–539, 2015.

[12] E. A. Hogan and H. Schaub, "Three-axis attitude control using redundant reaction wheels with continuous momentum dumping," *Journal of Guidance, Control, and Dynamics*, vol. 38, no. 10, pp. 1865–1871, 2015.

[13] G. Avanzini, E. L. de Angelis, F. Giulietti, and N. Serrano, "Attitude control of Low Earth Orbit satellites by reaction wheels and magnetic torquers," *Acta Astronautica*, vol. 160, pp. 625–634, 2019.

[14] W. H. Steyn, Y. Hashida, and V. Lappas, "An attitude control system and commissioning results of the SNAP-1 nanosatellite," in *Proceedings of the 14th Annual AIAA/USU Conference on Small Satellites*, Logan, UT, USA, 2000, paper SSC00-VIII-8.

[15] W. H. Steyn and Y. Hashida, "In-orbit attitude performance of the 3-axis stabilised SNAP-1 nanosatellite," in *Proceedings of the 15th Annual AIAA/USU Conference on Small Satellites*, Logan, UT, USA, 2001, paper SSC01-V-1.

[16] K. Sarda, S. Eagleson, E. Caillibot, C. Grant, D. Kekez, F. Pranajaya, and R. E. Zee, "Canadian advanced nanospace experiment 2: Scientific and technological innovation on a three-kilogram satellite," *Acta Astronautica*, vol. 59, pp. 236–245, 2006.

[17] D. J. Sahnow, J. W. Kruk, T. B. Ake, B.-G. Andersson, A. Berman, W. P. Blair, K. Boyer, M. Caplinger, M. Calvani, T. Civeit, V. V. D. Dixon, M. N. England, M. E. Kaiser, M. Kochte, H. W. Moos, and B. A. Roberts, "Operations with the new FUSE observatory: three-axis control with one reaction wheel," *Proceedings of SPIE*, vol. 6266, p. 626602, 2006.

[18] B. A. Roberts, J. W. Kruk, T. B. Ake, T. S. Englar, B. F. Class, and D. M. Rovner, "Three-axis attitude control with two reaction wheels and magnetic torquer bars," in *AIAA Guidance, Navigation, and Control Conference*, Providence, RI, Aug. 2004.

[19] R. D. Halverson, D. Gebre-Egziabher, and R. J. Caverly, "Attitude control of dual-spin satellites in low-Earth orbit via predictive control and magnetic actuation," in *AIAA SciTech Forum*, Orlando, FL, 2024, AIAA 2024-2278.

[20] R. D. Halverson and R. J. Caverly, "Time-varying model predictive attitude control for magnetically actuated dual-spin satellites," 2025, preprint submitted to Journal of Guidance, Control, and Dynamics.

[21] J. Li, M. Post, T. Wright, and R. Lee, "Design of attitude control systems for CubeSat-class nanosatellite," *Journal of Control Science and Engineering*, vol. 2013, p. 657182, 2013.

[22] X. Zhu, M. Sutherland, A. R. Girard, and B. E. Gilchrist, "Attitude control of a 3U CubeSat with combination of magnetorquers and reaction wheels," in *AIAA SciTech Forum*, 2023, AIAA 2023-0935.

[23] P. McKeen, N. Scheuer, and K. Cahoy, "Generalized attitude control for small spacecraft," in *Proceedings of the 40th Annual Small Satellite Conference*, Logan, UT, USA, 2026, paper SSC26-P2-54.

[24] P. McKeen and K. Cahoy, "Generalized attitude estimation for spacecraft," in *Proceedings of the 38th Annual Small Satellite Conference*, Logan, UT, USA, 2024, paper SSC24-WV-08.

[25] P. McKeen, N. Scheuer, and K. Cahoy, "A feasibility-aware attitude trajectory planner for small satellites," in *Proceedings of the 40th Annual Small Satellite Conference*, Logan, UT, USA, 2026, paper SSC26-FT-34.

[26] P. McKeen, A. Meredith, and K. Cahoy, "Trajectory-planning attitude control system for satellites with magnetic attitude control and one reaction wheel," in *AAS/AIAA Astrodynamics Specialist Conference*, 2022, paper AAS 22-810.

[27] T. A. Howell, B. E. Jackson, and Z. Manchester, "ALTRO: A fast solver for constrained trajectory optimization," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Macau, China, Nov. 2019, pp. 7674–7679.

[28] A. Bahlous-Boldi, C. Lisy, N. Mondal, B. Ferro, S. Kacker, M. Dahl, M. Anderson, and K. Cahoy, "BeaverCube II: Using AI-optimized processors on earth-observing CubeSats for autonomous image analysis and intelligent data handling," in *Proceedings of the 38th Annual Small Satellite Conference*, Logan, UT, USA, 2024.