

Readiness Barrier Functions: Forward-Invariant Control Authority for Overactuated Multirotor Allocation

Giuseppe Silano¹, *Senior Member, IEEE*

Abstract—Allocation schemes that greedily maximize a readiness metric over the actuator fiber bundle of an overactuated multirotor produce commands that jump between disconnected optimal strata, demanding actuator rates no motor can deliver; effort-minimizing schemes are continuous but cannot guarantee that wrench-rate authority stays above any certified level. We reconcile the two by treating authority as a forward-invariant quantity: a control barrier function on the log-determinant of the drag-aware actuator-authority co-metric, enforced at torque level by a quadratic program in the allocation null space. A single design inequality renders the certified set compact and strictly interior to the actuator box, with the readiness cost of any rotor deactivation given in closed form as $\ln(n/(n-m))$ for symmetric designs. Tracking is sacrificed only through an explicit alignment ratio, with wrench error bounded by $\mathcal{O}(\rho^{-1/2})$ and a robust variant handles motor-parameter uncertainty with a closed-form floor shift independent of the airframe matrix. On a hexarotor and a fully-actuated octorotor the closed-form gap matches simulation to machine precision; in the authority-scarce regime greedy maximization violates the certified floor and commits wrench errors up to eighty times larger than the proposed filter, which holds invariance of the certified set at negligible tracking cost.

Index Terms—Aerial Systems: Mechanics and Control, Robot Safety, Optimization and Optimal Control, Redundant Robots.

I. INTRODUCTION

Overactuated multirotors—tilted-propeller hexarotors, fully actuated octorotors, omnidirectional platforms—carry more rotors than the wrench dimensions they control [1], [2]. The redundancy is deliberate, enabling fault tolerance, decoupled force-and-moment generation, and physical interaction, but it creates a persistent decision problem: at every instant a continuum of rotor-speed vectors produces the commanded wrench, and the allocator must select one [3], [4].

The classical answer minimizes actuator effort, treating allocation as constrained least-squares or a linear program [3], [4]. This template serves standard and fully-actuated multirotors [5], [6] and predictive controllers with actuator models [7], but handles control authority implicitly through the feasible set, never as a certified quantity. Such allocators are lightweight yet indifferent to a critical dynamic metric: *wrench-rate authority*, the symmetric acceleration capacity to generate rapid wrench variations. Since thrust and drag both scale quadratically with spin [7], [8], a slow rotor produces thrust with vanishing slope while a saturated rotor has spent its torque budget




Figure 1: Achievable wrench-rate set $\mathcal{R}$: sweet spot (blue), rotor 1 stopped (orange dashed), near saturation (red dotted). The commanded $\dot{w}_{\text{des}}$ must lie inside $\mathcal{R}$ to be realizable.

fighting drag; authority peaks at an interior *sweet spot* and collapses at both ends. The collapse takes two geometric forms: stopping one rotor squashes the achievable wrench-rate set *anisotropically*, saturating all rotors shrinks it *isotropically* (Figure 1)—neither visible to an effort-minimizing cost.

Determinant-based manipulability, originating with Yoshikawa [9], serves broadly as a singularity-avoidance penalty. Recent work extends it to multirotors via a drag-aware authority co-metric—*Drag-Aware Aerodynamic Manipulability* (DAAM)—whose log-determinant measures the remaining wrench-rate volume, maximized along the *task fiber* of redundant speeds satisfying the wrench command to repel both collapse modes [10]. That framework, however, has a critical open defect: the fiberwise maximizer is set-valued, with optimal selections on disconnected sheets [10], [11].

The obstruction is structural. The signed-quadratic thrust map partitions rotor-speed space into sign orthants whose boundaries are the loci where a rotor reverses direction; there the thrust slope vanishes, the allocation Jacobian degenerates, and the actuator rate sustaining a bounded wrench rate diverges [12]. Tracking a continuous task across orthants thus demands an acceleration no motor can deliver: the rotors saturate and inject the very wrench error the allocator exists to null. The standard remedy, low-pass filtering [3], [4], bounds the rate demand but drags the system *through* the low-authority crossing rather than around it (Section IV). Across the literature, then, authority is a byproduct of effort minimization or an unconstrained maximization objective, never a strict forward-invariant property—and a greedy maximizer violates the very level it was designed to protect the moment a crossing becomes necessary.

We therefore argue that control authority should be *certified* rather than greedily maximized. We treat the readiness

¹Ricerca sul Sistema Energetico, Milan, Italy, and Czech Technical University, Prague, Czech Republic (e-mail: silangu@fel.cvut.cz).

This work was partially funded by the research fund for the Italian Electrical System (decree n. 388, Nov. 6th, 2024) and the GACR project no. 26-22419S.

The author used Google Gemini 3 and Anthropic Claude 4.8 Opus for proofreading and exposition refinement, but independently verified all outputs and assumes full responsibility for the article's content.

log-determinant as a Control Barrier Function (CBF) [13], grounded in set-invariance theory [14] and its robust and discrete-time extensions [15], [16], and floor it at the torque level via a quadratic program acting in the allocation null space. Because the barrier is a determinant on a fibered actuator space rather than a workspace constraint, feasibility turns on the geometry of the fiber tangent space, and the single inequality $h(v) \geq 0$ fences both failure modes at once. The certified set has 2^n components, so any physical transit between them must cross the zero-spin boundary and pay an exact readiness penalty; enforcing the barrier confines the trajectory to its initial component and guarantees *forward invariance* of the certified set for the modeled dynamics. When tracking conflicts with the floor, the filter relaxes the wrench by a slack δ bounded by $\mathcal{O}(\rho^{-1/2})$, with ρ the barrier relaxation weight, which permits a continuous allocation where a greedy maximizer must jump.

This work makes four contributions. (i) *Readiness geometry*: a closed-form readiness gradient and an exact component-separation barrier (Theorem 2), the readiness lost when rotor k deactivates being $-\ln(1 - \lambda_k)$ for a capacity-weighted leverage score λ_k . (ii) *Safe set and invariance*: a single design inequality $\bar{\ell} > \ell^\dagger$ rendering the safe set compact and strictly interior to the actuator box, removing auxiliary saturation constraints (Proposition 5), with forward invariance for the modeled dynamics (Theorem 8). (iii) *Tracking*: an exact alignment ratio dictating when tracking stays exact (Theorem 10) and an $\mathcal{O}(\rho^{-1/2})$ wrench-error bound otherwise (Theorem 10). (iv) *Robustness*: a variant separating a worst-case metric bound from robust invariance (Theorem 11) via a closed-form drift term, without which the guarantee fails on 38% of plants at 30% mismatch. The certificates cover the allocation subproblem as a structural artifact of fiberwise optimization, independent of airframe or airspeed, and are validated on a hexarotor and a fully-actuated tilted octotoror over parametric uncertainty and transport delay.

II. MODELING AND NOTATION

An overactuated multirotor with n rotors produces a control wrench $w \in \mathbb{R}^m$ ($n > m$) via $f(v) = A\phi(v)$, where $v \in \mathbb{R}^n$ collects the signed spin rates—we assume bidirectional propulsion with torque-controlled rotors capable of reversal through $v_i = 0$ — $\phi(v) = v \odot |v|$ is the signed-quadratic thrust map, and $A \in \mathbb{R}^{m \times n}$ is the constant, full-row-rank allocation matrix with columns $A_{\bullet i}$. The *allocation fiber* $\mathcal{F}_w := f^{-1}(w)$ is the continuum of redundant configurations producing w , from which the allocator selects one point. Per rotor,

$$m_i \dot{v}_i = -b_i v_i |v_i| + \tau_i, \quad |\tau_i| \leq \bar{\tau}_i, \quad (1)$$

where m_i is the motor-propeller inertia, $b_i > 0$ the lumped drag coefficient, and $\bar{\tau}_i > 0$ the torque limit. The quadratic form $b_i v_i |v_i|$ is the quasi-static rotor-drag model [7], [8], valid at hover and low advance ratio; in fast forward flight inflow modulates b_i , treated as bounded parametric variation in Section III-D2. Stacking (1) gives the vector dynamics $\dot{v} = d(v) + M^{-1}\tau$, with $M = \text{diag}(m_i)$, and the drag vector $d(v) = -M^{-1} \text{diag}(b)\phi(v)$. The *Symmetric Acceleration*

Capacity (SAC) is the largest torque-limited acceleration available symmetrically at spin rate v_i :

$$\bar{a}_i(v_i) = \frac{\bar{\tau}_i - b_i v_i^2}{m_i}, \quad \mathcal{E} = \{v : |v_i| < v_i^{\text{sat}} := \sqrt{\bar{\tau}_i/b_i} \forall i\}, \quad (2)$$

where v_i^{sat} is the *saturation speed* and $\mathcal{E}$ the open feasible box on which $\bar{a}_i > 0$. With allocation Jacobian $J(v) = 2A \text{diag}(|v|)$ and SAC weighting $W(v) = \text{diag}(\bar{a}_i^2)$, the deliverable wrench rates form the ellipsoid $\mathcal{R}(v) = \{JW^{1/2}u : \|u\| \leq 1\}$, whose Gram matrix is the *drag-aware authority co-metric* [10]:

$$D(v) = J(v)W(v)J(v)^\top = 4A \Psi(v) A^\top, \quad (3)$$

$$\psi_i(v_i) = v_i^2 \bar{a}_i(v_i)^2, \quad (4)$$

where $\Psi = \text{diag}(\psi_i)$. Equivalently $\mathcal{R}(v) = \{x : x^\top D^{-1}x \leq 1\}$, with semi-axes $\sqrt{\lambda_i(D)}$ and m -dimensional volume $\omega_m \sqrt{\det D}$ (ω_m the unit-ball volume). The *readiness level*

$$L(v) := \ln \det D(v) = 2 \ln(\text{vol}_m \mathcal{R}(v)/\omega_m). \quad (5)$$

is therefore a log-volume of achievable wrench rate.

A floor on L is informative on three counts. As a *volume*, $e^{L/2}$ is proportional to omnidirectional wrench-rate authority; as a *worst case*, since $\lambda_{\min}(D) \geq e^{\bar{\ell}}/\bar{\Lambda}^{m-1}$ (Proposition 5(iv)), the floor guarantees an authority radius along *every* wrench-rate axis, not merely on average; and as a *conditioning measure*, the determinant simultaneously penalizes the loss of any axis while retaining sensitivity in the remaining directions [9], [10]—a trace criterion misses a collapsed axis and $\lambda_{\min}$ ignores the retained ones. Figure 1 makes the two collapse modes geometrically distinct.

We reserve h for the *shifted barrier*, where $\bar{\ell} \in (\ell^\dagger, L^{\text{op}})$ is the chosen authority floor (specified in Section III):

$$h(v) = L(v) - \bar{\ell}, \quad \mathcal{C} = \{v \in \mathcal{E} : h(v) \geq 0\}, \quad (6)$$

$\mathcal{C}$ being the *certified-authority set*. Throughout, A satisfies the following: **A1**, $\text{rank } A = m$, $n > m$, and $A_{\bullet i} \neq 0$ for every i .

We use $i \in \{1, \dots, n\}$ as a generic rotor index in sums and stacked expressions, and k to denote a specific rotor singled out for analysis, as in the drop-out of Theorem 2.

Remark 1 (The column condition is not implied by rank). *Full row rank does not force $A_{\bullet i} \neq 0$: deleting a column of a 4×6 hexarotor matrix still leaves rank 4. A rotor with $A_{\bullet k} = 0$ contributes nothing to the wrench and would not appear in any physical design, yet $\lambda_k = 0$, $\Delta_k = 0$, and the floor window (11) collapses, so the condition must be stated.*

III. THEORY AND ALLOCATION ARCHITECTURE

The rotor-speed space $\mathcal{E}$ is partitioned by the hyperplanes $v_k = 0$ into 2^n *sign orthants*, within each of which L is smooth with global maxima at the sweet spots $\pm v^*$. The fiberwise $\arg \max L$ is therefore set-valued: a greedy allocator can undergo a discontinuous branch change when the optimal strata $\pm v^*$ exchange dominance as the commanded wrench reverses. Any such change passes rotor k through $v_k = 0$, where its thrust slope vanishes, its column drops out of D , and the Jacobian loses rank (Figure 2); our first result quantifies the cost.




Figure 2: Allocation fiber bundle. $\mathcal{C}$ confines the trajectory to a certified arc of $f^{-1}(w)$, isolated from dropout ($v_k = 0$) and saturation.

Theorem 2 (Dropout Gap). *The largest readiness level attainable with rotor k inactive ($\psi_k = 0$) is $L^{\max} - \Delta_k$, where*

$$\begin{aligned} \Delta_k &= -\ln(1 - \lambda_k) \in (0, +\infty], \\ \lambda_k &= \psi_k^* A_{\bullet k}^\top S^{-1} A_{\bullet k} \in (0, 1], \end{aligned} \quad (7)$$

with λ_k the ψ^* -weighted leverage of rotor k . Here, $S := D(v^*)/4$ is the co-metric evaluated at the sweet spot, so $S = \sum_i \psi_i^* A_{\bullet i} A_{\bullet i}^\top$ with $\psi_k^* := \psi_k(v_k^*)$. Moreover: (i) $\Delta_k = +\infty$ iff rotor k is essential (i.e., the submatrix A_{-k} missing column k has rank $A_{-k} < m$); (ii) $\sum_k \lambda_k = m$; (iii) for symmetric designs $\lambda_k = m/n$ and

$$\Delta = \ln\left(\frac{n}{n-m}\right), \quad (8)$$

which is a function of the redundancy ratio alone.

Proof. With $\psi_k = 0$, Loewner monotonicity and the matrix determinant lemma applied to the rank-1 update of S yield $\det D \leq 4^m \det S(1 - \lambda_k)$, with equality when all other rotors are at v^* . Since λ_k is a diagonal entry of an orthogonal projector in the S -norm, $\lambda_k \in [0, 1]$; $\lambda_k = 1$ iff $\text{rank } A_{-k} < m$, giving (i). Part (ii) follows because the projector has trace m ; (iii) holds when equal ψ_i^* make $S = \psi^* A A^\top$. $\square$

The leverage score λ_k measures how *irreplaceable* rotor k is: since $\sum_k \lambda_k = m$, a large λ_k means rotor k owns a disproportionate share of the m wrench dimensions, and jumping across $v_k = 0$ costs the exact penalty Δ_k , infinite for an essential rotor. On a symmetric hexarotor ($n = 6$, $m = 4$), each crossing costs $\Delta = \ln 3 \approx 1.10$ nats regardless of size, propeller, or payload. The design goal follows: *confine the rotor-speed state to a single sign orthant and never pay this penalty*. We thus define the *dropout level*—the highest readiness retained if any rotor approaches zero spin—as

$$\ell^\dagger := L^{\max} - \min_k \Delta_k. \quad (9)$$

A floor $\bar{\ell} > \ell^\dagger$ then fences every zero-spin configuration with one scalar inequality, locking the vehicle inside its initial orthant. This gap depends on the physical capacities ψ^* , not A alone: for a hexarotor with $\pm 30\%$ spread in $(\bar{\tau}_i, b_i)$ the true gaps scatter to $\{1.066, 1.994, 0.486\}$ nats against a uniform geometric 1.099—a 0.9 nat error that can void the guarantee.

The floor's upper limit is the mission-dependent worst-case achievable readiness

$$L^{\text{op}} := \min_{w \in \mathcal{W}} \max_{v \in f^{-1}(w) \cap \mathcal{E}} L(v), \quad (10)$$

so a certifiable floor must lie in the window

$$\ell^\dagger < \bar{\ell} \leq L^{\text{op}}, \quad (11)$$

non-empty iff $L^{\text{op}} > \ell^\dagger$; we set $\bar{\ell} = \ell^\dagger + \kappa(L^{\text{op}} - \ell^\dagger)$, $\kappa \in (0, 1)$.

Proposition 3 (Certifiable Floor). *If (11) is violated ($\bar{\ell} > L^{\text{op}}$), the certified set intersects some commanded fiber $f^{-1}(w)$ in the empty set, so no allocator can hold $h \geq 0$ while tracking. Certifiability is thus a property of the task-vehicle pair, not the filter: a drone commanded near its thrust ceiling may have no certifiable floor.*

A. Readiness geometry and physical sweet spot

To avoid the dropout penalty we must understand how L varies across $\mathcal{E}$. The co-metric enters each rotor only through the scalar $\psi_i = (v_i \bar{a}_i)^2$ —the squared product of *thrust sensitivity* v_i (a fast rotor converts spin changes into larger wrench changes) and *torque headroom* $\bar{a}_i$ (drag consumes part of the budget just to hold speed). Since $D = 4A\Psi A^\top$, $\Psi = \text{diag}(\psi_i)$, is polynomial in v_i^2 on $\mathcal{E}$, the barrier h inherits no kink from the modulus in ϕ and is C^∞ wherever $D \succ 0$.

Lemma 4 (Readiness Gradient). *Wherever $D(v) \succ 0$,*

$$\begin{aligned} \frac{\partial h}{\partial v_i}(v) &= \frac{8 v_i \bar{a}_i(v_i)(\bar{\tau}_i - 3b_i v_i^2)}{m_i} s_i(v), \\ s_i &:= A_{\bullet i}^\top D^{-1} A_{\bullet i} > 0, \end{aligned} \quad (12)$$

with $4 \sum_i \psi_i s_i = m$. Each component vanishes exactly at $v_i \in \{0, \pm v_i^*\}$ with $v_i^* = v_i^{\text{sat}}/\sqrt{3}$, and L attains its global maximum $L^{\max} = \ln \det(4S)$ at the 2^n sweet-spot configurations $|v_i| = v_i^*$.

Proof. Jacobi's formula applied to (3) gives $\partial h / \partial v_i = 4\psi_i'(v_i) s_i$; by standard differentiation $\psi_i' = 2v_i \bar{a}_i(\bar{\tau}_i - 3b_i v_i^2)/m_i$, which is (12). The trace identity $4 \sum_i \psi_i s_i = \text{tr}(D^{-1}D) = m$ is immediate. Roots of ψ_i' are as stated. Global maximality follows from Loewner monotonicity of D in each ψ_i and monotonicity of $\det$ on the positive-definite cone; the 2^n degeneracy holds because ψ_i is even in v_i . $\square$

The gradient vanishes at a *physical sweet spot*: $v_i^* \approx 58\%$ of saturation speed, where drag $b_i(v_i^*)^2 = \bar{\tau}_i/3$ consumes one third of the torque budget and leaves two thirds as reserve. The 2^n sign degeneracy is physical—a rotor at $-v^*$ is as ready as one at $+v^*$, so the greedy selector has no tiebreaker and may jump when the strata exchange dominance, whereas the barrier certifies that authority can be maintained inside a single sweet-spot component.

B. Certified safe set and forward invariance

The key insight—illustrated in Figure 3—is that a *single* scalar inequality $h(v) \geq 0$ simultaneously fences both physical failure modes: the zero-spin locus (where $\psi_k \rightarrow 0$ kills the thrust slope) and the saturation boundary (where drag has




Figure 3: Certified set $\mathcal{C}$ for $n = 2$ ($\bar{\tau} = b = m = 1$): the four sign-orthant components (blue) sit strictly between the zero-spin hyperplanes $v_i = 0$ (orange) and saturation boundary $\partial\mathcal{E}$ (red dashed), arrows mark the uniform margin on both sides (Proposition 5(i),(iv)).

exhausted the torque budget). No auxiliary saturation logic is needed. We now prove this and establish the set's geometric properties. We assume henceforth that $\bar{\ell} > \ell^\dagger$ (A2, verified offline from (7)–(9)) and that $\bar{\ell}$ is a regular value of h (A3, holds generically by Sard's theorem).

Proposition 5 (Structure of the Certified-Authority Set). *Under A1–A3, there exists $\varepsilon > 0$ with $\psi_i(v_i) \geq \varepsilon$ for all i on $\mathcal{C} = \{v \in \mathcal{E} : h(v) \geq 0\}$. Consequently: (i) $\mathcal{C} \subset \text{int } \mathcal{E}$ and $|v_i| > 0$ on $\mathcal{C}$; (ii) $\text{rank } J(v) = m$ and $D(v) \succ 0$ on $\mathcal{C}$; (iii) $\mathcal{C}$ is compact; (iv) $\lambda_{\min}(D) \geq \underline{\lambda} := e^{\bar{\ell}}/\bar{\Lambda}^{m-1}$ on $\mathcal{C}$, with $\bar{\Lambda} := 4 \sum_i \psi_i^* \|A_{\bullet i}\|^2$; (v) $\partial\mathcal{C}$ is a smooth compact hypersurface with $\nabla h \neq 0$.*

Proof. For each i , Loewner monotonicity gives $\ln \det D \leq g_i(\psi_i)$ where $g_i(0) = L^{\max} - \Delta_i \leq \ell^\dagger < \bar{\ell}$ by A2; by continuity $\psi_i \geq \varepsilon_i$ on $\mathcal{C}$. Set $\varepsilon = \min_i \varepsilon_i$; then $|v_i| > 0$ and $\bar{a}_i > 0$, giving (i)–(ii). The set $K_\varepsilon := \{v : \psi_i \geq \varepsilon \forall i\}$ is closed and bounded in $\bar{\mathcal{E}}$, hence compact; $\mathcal{C} \subseteq K_\varepsilon$ gives (iii). For (iv): $\det D \geq e^{\bar{\ell}}$ and $\lambda_k(D) \leq \bar{\Lambda}$. Item (v) is A3 with the regular value theorem. $\square$

Proposition 5(i) guarantees that any trajectory in $\mathcal{C}$ automatically respects the actuator limits—speeds stay away from zero and below saturation with no explicit saturation check—while part (iv) bounds $\lambda_{\min}(D)$ from below, so the wrench-rate ellipsoid keeps a guaranteed volume in every direction. Figure 3 shows this for $n = 2$: the four certified components sit strictly between the zero-spin hyperplanes and the saturation boundary, separated by a uniform margin on both sides.

It remains to prove that a deliverable torque always exists to keep the state inside $\mathcal{C}$, which requires care because drag makes acceleration and deceleration profoundly unequal. Decelerating is easy—drag assists, with bound $(\bar{\tau}_i + b_i v_i^2)/m_i$ growing with spin—while accelerating is hard, with bound $\bar{a}_i = (\bar{\tau}_i - b_i v_i^2)/m_i$ shrinking to zero at saturation (nearly 50 : 1 at $0.98 v_i^{\text{sat}}$). Since the allocator cannot anticipate the next demand direction, authority means capability *both* ways, so the metric must be weighted by the binding symmetric direction.

Lemma 6 (Viability via Symmetric Acceleration Capacity). *For every $v \in \text{int } \mathcal{E}$, the interval $[-\bar{a}_i, \bar{a}_i]$ is the largest symmetric box of admissible accelerations within $[a_i^-, a_i^+]$. Consequently,*

$$\sup_{|\tau| \leq \bar{\tau}} \nabla h(v)^\top (d(v) + M^{-1}\tau) \geq \sum_i |\partial_i h| \bar{a}_i \geq 0, \quad (13)$$

with strict inequality whenever $\nabla h(v) \neq 0$.

Proof. Evaluating $a_i^\pm = (\pm \bar{\tau}_i - b_i v_i |v_i|)/m_i$ for all signs of v_i confirms $[-\bar{a}_i, \bar{a}_i] \subseteq [a_i^-, a_i^+]$, tight at one endpoint. Choosing $a_i = \bar{a}_i \text{sgn}(\partial_i h)$ gives (13). $\square$

Lemma 6 makes the CBF admissible set non-empty: a physically deliverable torque always exists to prevent readiness from falling. The result hinges on weighting D by the *symmetric* capacity $\bar{a}_i$ rather than the asymmetric braking margin; otherwise the barrier would be useless at high spin, where the 50 : 1 acceleration/deceleration asymmetry is most extreme.

C. Control barrier allocation filter

With the certified safe set established, we now present the controller that enforces it. The filter acts as a *post-processor*: at every control step it receives a nominal torque command τ_{ref} from a nominal allocator (e.g., an effort-minimizing). It then computes the applied torques τ closest to τ_{ref} that keep the state inside $\mathcal{C}$. When no conflict exists the filter is inactive and reproduces the nominal command exactly. Let $\mu := \dot{w}_{\text{des}} + K_w \tilde{w}$ be the task-acceleration demand, with wrench error $\tilde{w} := w_{\text{des}} - A\phi(v)$ and diagonal gain $K_w > 0$. The allocation filter solves:

$$\begin{aligned} \min_{\tau, \delta} \quad & \|\tau - \tau_{\text{ref}}\|_R^2 + \rho \|\delta\|^2 \\ \text{s.t.} \quad & \nabla h(v)^\top a \geq -\alpha(h(v)), \\ & J(v)a = \mu + \delta, \\ & |\tau| \leq \bar{\tau}, \\ & a = d(v) + M^{-1}\tau. \end{aligned} \quad (14)$$

Here $R \succ 0$ is a diagonal weight matrix, and the slack variable δ absorbs the wrench deficit when the safety barrier and the tracking demand are in genuine conflict, with ρ setting the exchange penalty. Crucially, the barrier gradient $\nabla h(v)$ naturally encodes the hardware limits: a motor with a small acceleration capacity $\bar{a}_i$ exerts proportionally less weight on the gradient, ensuring the filter actively protects the most vulnerable rotors first.

Lemma 7 (Uniform Strict Feasibility). *Under A1–A3, $\exists \sigma > 0 : \max_{|\tau| \leq \bar{\tau}} [\nabla h^\top (d + M^{-1}\tau) + \alpha(h)] \geq \sigma$ uniformly on $\mathcal{C}$. The minimizer of (14) is unique and locally Lipschitz on $\mathcal{C}$.*

Proof. If $\nabla h \neq 0$, Lemma 6 supplies a strictly interior torque; if $\nabla h = 0$ then $v \in \text{int } \mathcal{C}$ (Proposition 5(v)) and the constraint holds trivially. Uniformity follows from compactness of $\mathcal{C}$. The minimizer is Lipschitz by strong convexity and Hoffman's bound under the uniform Slater condition [17]. $\square$

Theorem 8 (Valid CBF and Forward Invariance). *Under A1–A3, h is a valid CBF on $\mathcal{C}$ and (14) is feasible for all $\rho > 0$. For any locally Lipschitz selection $\tau(t, v)$ satisfying*

the barrier row with $v(0) \in \mathcal{C}$, the closed-loop trajectory is unique, complete, and satisfies $v(t) \in \mathcal{C} \subset \text{int } \mathcal{E}$ for all $t \geq 0$, remaining in the connected sign-orthant component containing $v(0)$.

Proof. The admissible set is non-empty by Lemma 6 (when $\nabla h \neq 0$) or trivially (when $\nabla h = 0 \in \text{int } \mathcal{C}$). Feasibility of (14) holds because δ is unconstrained. The minimizer is locally Lipschitz (Lemma 7), guaranteeing local existence and uniqueness. Along any solution $\dot{h} \geq -\alpha(h)$; applying the standard comparison lemma yields $h(v(t)) \geq 0$ [13], [14]. Compactness precludes finite escape, and interior containment follows from Proposition 5(i). $\square$

The CBF barrier row $\nabla h^\top a \geq -\alpha(h)$ acts dynamically as a *margin-dependent speed limit*: deep inside $\mathcal{C}$ where authority is abundant, the constraint is inactive and the filter passes the nominal command through unchanged; as the state approaches $\partial\mathcal{C}$, the permitted rate of authority expenditure shrinks continuously to zero, so the boundary is approached asymptotically and never crossed. Crucially, the filter never evaluates the set-valued $\arg \max_{f^{-1}(w)} L$ —it solves a strongly convex Quadratic Program (QP) to produce a single continuous selection, eliminating the sign jumps that force greedy allocators to saturate the commanded motor acceleration.

Remark 9 (Sampled-data implementation). *Theorem 8 certifies invariance in continuous time. Under zero-order hold with step Δt , the barrier row holds only at samples, so between them h may dip by at most $\alpha(h)\Delta t + \frac{1}{2}L_h\Delta t^2$ —the first-order term being the permitted rate $-\dot{h}(t_k) \leq \alpha(h(t_k))$ and L_h bounding $\ddot{h}$ on $\mathcal{C}$. This $\mathcal{O}(\Delta t)$ margin is absorbed by the strict interior exactly as transport delay is; at 1–2 kHz it is negligible, and a discrete-time CBF [16] removes the Δt dependence at the cost of a nonconvex per-step check.*

D. Task relaxation and robustness

1) *Tracking bounds and the zero-sum exchange*: Whether the filter can simultaneously maintain the floor and track the task exactly depends on the available null-space geometry. Define the maximum authority ascent achievable at constant wrench,

$$\gamma(v) := \max\{\nabla h^\top \xi : \xi \in \ker J(v), |\xi_i| \leq \bar{a}_i\} \geq 0, \quad (15)$$

and the *alignment ratio*

$$\vartheta(v) := \gamma(v) / \|W^{1/2} \nabla h(v)\|_2 \in [0, \sqrt{n}]. \quad (16)$$

On a symmetric platform at near-uniform spin, $\nabla h \propto \mathbf{1} = A^\top e_{\text{coll}} \in \text{range } J^\top$ (where e_{coll} is the collective thrust axis), so $\gamma = 0$ exactly: the null space of J is orthogonal to ∇h by construction, and any redistribution that speeds some rotors must slow others by the same amount—because all rotors are equally sensitive, the gains and losses cancel identically. The alignment ratio measures precisely how far from this zero-sum regime the current configuration lies.

Theorem 10 (Strict Feasibility and Bounded Task Relaxation). *Let $\underline{\lambda}$ be as in Proposition 5(iv) and suppose $\|\mu\| \leq (1-\beta)\sqrt{\underline{\lambda}}$ for some capacity split ratio $\beta \in (0, 1)$. If*

$$\beta\gamma(v) + \alpha(h(v)) \geq (1-\beta)\|W^{1/2} \nabla h(v)\|_2, \quad (17)$$

then (14) has a feasible point with $\delta = 0$. On $\partial\mathcal{C}$ this reduces to $\vartheta(v) \geq (1-\beta)/\beta$. When $\delta^* \neq 0$ the minimizer satisfies

$$\|\delta^*\| \leq c_\tau / \sqrt{\rho}, \quad c_\tau := 2\|R\|^{1/2}\|\bar{\tau}\|, \quad (18)$$

and the wrench error obeys $\dot{\tilde{w}} = -K_w \tilde{w} - \delta^*$, yielding

$$\|\tilde{w}(t)\| \leq e^{-K_w t} \|\tilde{w}(0)\| + c_\tau / (K_w \sqrt{\rho}); \quad (19)$$

the ultimate wrench error is $\mathcal{O}(\rho^{-1/2})$.

Proof. Strict feasibility: set $a = a_0 + \xi$ with $a_0 = WJ^\top D^{-1}\mu$ the horizontal lift and ξ feasible for the program (15) scaled by β . Then $Ja = \mu$ and $\delta = 0$. Admissibility of a follows from $\|W^{-1/2}a_0\| \leq (1-\beta)$ (Proposition 5(iv)) and $|\xi_i| \leq \beta\bar{a}_i$; the barrier row follows from (17) via Cauchy–Schwarz. When $\delta^* \neq 0$: comparing the optimal value with the $(\tau^{(0)}, 0)$ feasible point from the strict-feasibility case and using strong convexity of the QP objective gives (18). The error dynamics $\dot{\tilde{w}} = -K_w \tilde{w} - \delta^*$ yield (19) by the standard ISS bound. $\square$

On symmetric platforms the null-space exchange is zero-sum ($\vartheta \rightarrow 0$), so spin redistribution cannot recover readiness; the margin comes instead from the strict interior $\alpha(h) > 0$, which Proposition 3 guarantees by keeping $\bar{\ell}$ below L^{op} . The $\mathcal{O}(\rho^{-1/2})$ bound is a design rule—quadruple ρ to halve the tracking error—and the slack δ enters the outer loops as a small bounded disturbance, trading a persistent tracking error for the elimination of an impulsive one at the moment authority collapses.

2) *Parametric robustness*: The results above are nominal; motor parameters vary with temperature, battery state, and propeller wear. Two effects must be handled. The first is the *metric*: the true $L(v; \zeta)$ at the true parameters ζ may lie below the modelled $h + \bar{\ell}$. The second, and more dangerous, is the *drift*: the barrier is enforced using the modelled dynamics $d(v; \hat{\zeta})$, where $\hat{\zeta} = (\bar{\tau}, b)$ are the nominal parameter values, while the true plant runs $d(v; \zeta)$ with $\zeta \in \mathcal{Z}_p$, and the residual is not small. Bounding a metric is not the same as controlling its derivative; the two must be treated separately.

Acceleration headroom $\bar{a}_i = (\bar{\tau}_i - b_i v_i^2) / m_i$ is affine in both the motor torque limit $\bar{\tau}_i$ and the drag coefficient b_i : a weaker motor ($\bar{\tau}_i$ low) and a draggier propeller (b_i high) compound the headroom loss through the same fraction. Bounding both by a relative half-width $p \in (0, 1)$ —so $\bar{\tau}'_i \in [\bar{\tau}_i(1 \pm p)]$ and $b'_i \in [b_i(1 \pm p)]$ —the combined worst case is attained at a single vertex of the parameter box, the doubly degraded corner $\zeta^- = (\bar{\tau}(1-p), b(1+p))$, collapsing an otherwise 2^{2n} vertex enumeration to a single evaluation.

Theorem 11 (Worst-Case Metric and Robust Forward Invariance). *Let $\mathcal{Z}_p = \{(\bar{\tau}', b') : \bar{\tau}'_i \in [\bar{\tau}_i(1 \pm p)], b'_i \in [b_i(1 \pm p)]\}$, $\bar{a}_i^- := (\bar{\tau}_i(1-p) - b_i(1+p)v_i^2) / m_i$, $D^- := 4A \text{diag}(v_i^2(\bar{a}_i^-)^2)A^\top$, and $\mathcal{E}^- := \{v : |v_i| < \sqrt{\bar{\tau}_i(1-p)/(b_i(1+p))}\} \subset \mathcal{E}$. For all $\zeta \in \mathcal{Z}_p$ and $v \in \mathcal{E}^-$: $D(v; \zeta) \succeq D^-(v)$, so $L(v; \zeta) \geq L^-(v) := \ln \det D^-(v)$. Moreover, if τ satisfies the tightened barrier row*

$$\nabla h_p^\top M^{-1} \tau \geq -\alpha(h_p) - \min_{\zeta \in \mathcal{Z}_p} \nabla h_p^\top d(v; \zeta), \quad (20)$$

where $h_p := L^- - \bar{\ell}$, the row being enforced with the worst-case available torque $|\tau_i| \leq \bar{\tau}_i(1-p)$, and

$$\min_{\zeta \in \mathcal{Z}_p} \nabla h_p^\top d(v; \zeta) = \nabla h_p^\top d(v; b) - p\|c\|_1, \quad (21)$$

where $c_i := -\partial_i h_p \cdot b_i v_i |v_i| / m_i$. Hence, $\mathcal{C}_p := \{v \in \mathcal{E}^- : h_p(v) \geq 0\}$ is forward invariant for every $\zeta \in \mathcal{Z}_p$.

Proof. On $\mathcal{E}^-$, $\bar{a}_i(\zeta) \geq \bar{a}_i^- \geq 0$; squaring (both non-negative) and applying Loewner monotonicity and log-det monotonicity gives the metric bound. For invariance: $\dot{h}_p = \nabla h_p^\top (d(\zeta) + M^{-1}\tau) \geq \nabla h_p^\top M^{-1}\tau + \min_{\zeta'} \nabla h_p^\top d(\zeta') \geq -\alpha(h_p)$ by (20); the comparison lemma gives $h_p \geq 0$. Feasibility of (20) under $|\tau_i| \leq \bar{\tau}_i(1-p)$ follows from Lemma 6 applied to $\bar{a}_i^-$, so the reduced box still admits an authority-preserving input. The minimum in (21) is attained at a sign vertex because the drift component $d_i = -b_i v_i |v_i| / m_i$ is affine in b_i . $\square$

Corollary 12 (Exact Price of Robustness). *At the sweet-spot maximizer, $\psi_i^* = 4\bar{\tau}_i^3 / (27b_i m_i^2)$ scales by $(1-p)^3 / (1+p)$ at ζ^- , so the floor shift is independent of A :*

$$L_p^{\max} - L^{\max} = m \ln \left(\frac{(1-p)^3}{1+p} \right). \quad (22)$$

The certified authority volume $\det D^{1/2}$ scales by $((1-p)^3 / (1+p))^{m/2}$ and the per-direction radius by its m -th root. On an $m = 4$ hexarotor the retained volume is 44%, 18%, 7% at $p = 0.1, 0.2, 0.3$ (radius 81%, 65%, 51%).

The minimization in (21) costs a single $\mathcal{O}(n)$ pass—a weighted sum $p\|c\|_1$ over quantities the nominal filter already computes—with no extra QP constraints or vertex enumerations. The price (22) is exact and unavoidable, attained at $\zeta^- \in \mathcal{Z}_p$, so no method certifying the full box can recover more. Torque appears cubed in the capacity metric, making torque calibration roughly $3\times$ more valuable per unit relative error than drag calibration.

3) *Unmodeled input delay:* Under transport delay τ_d , the QP answers “what torque is safe?” about a state the vehicle has already left. Two mismatch terms accumulate: the drift at the stale state differs from the current drift, and the gradient has rotated since the QP was solved. The barrier may therefore continue to fall for an entire delay window after the QP has commanded it to stop.

Proposition 13 (Delay Undershoot). *Let the applied torque be $\tau^*(v(t - \tau_d))$ and define $K_1 := \sup_{v \in \mathcal{C}} \|\nabla(\nabla h^\top d)(v)\|$, $K_2 := \sup \|\nabla^2 h\|_\infty$, $V := \max_i \sup(|d_i| + \bar{\tau}_i/m_i)$, $\bar{H} := L^{\max} - \bar{\ell}$, and $K := K_1 + K_2 \max_i(\bar{\tau}_i/m_i)$. For a linear class- $\mathcal{K}$ function α ,*

$$h(v(t)) \geq -\eta(\tau_d), \quad (23)$$

where $\eta(\tau_d) := (\alpha\bar{H} + KV\tau_d)\tau_d + \frac{KV\tau_d}{\alpha}$ is the worst-case undershoot bound.

Proof. Decompose $\dot{h}(t)$ around the delayed state v^- . The stale barrier row contributes $\geq -\alpha h(v^-)$; the two mismatch brackets are bounded by $K_1 V \tau_d$ and $K_2 V \tau_d \max_i(\bar{\tau}_i/m_i)$. Using $h \leq \bar{H}$ and the standard comparison lemma yields (23). $\square$

Algorithm 1: Readiness-barrier allocation filter

Input: $A, (\bar{\tau}, b, m), \bar{\ell}, \alpha, K_w, \rho, R$
1 Offline: $\lambda_k \leftarrow \psi_k^* A_{\bullet,k}^\top S^{-1} A_{\bullet,k}$;
 $\ell^\dagger \leftarrow L^{\max} - \min_k(-\ln(1 - \lambda_k))$; verify (11)
2 while running do
 3 measure v ; compute $D, J, W, \nabla h$ (12), ϑ (16)
 4 $\tilde{w} \leftarrow w_{\text{des}} - A\phi(v)$; $\mu \leftarrow \dot{w}_{\text{des}} + K_w \tilde{w}$; $\tau_{\text{ref}} \leftarrow$
 effort-min QP
 5 solve (14) for (τ^*, δ^*) ; apply τ^*
6 end

Corollary 14 (Certified Delay Ceiling). *Raising the floor to $\bar{\ell} + \eta(\tau_d)$ certifies the original floor $\bar{\ell}$ under delay τ_d iff $\eta(\tau_d) \leq L^{\text{op}} - \bar{\ell}$. The certified ceiling τ_d^* is the unique root of $\eta(\tau_d^*) = L^{\text{op}} - \bar{\ell}$. Beyond τ_d^* , the tightened constraint set is empty.*

Battery sag (shrinking $\mathcal{Z}_p$) is handled by re-evaluating the window (11) at the current state; a rotor failure requires recomputing Δ_k on A_{-k} with two guards—clamp $\Delta_j = +\infty$ for any newly-essential rotor, and refuse to enforce an empty window. Algorithm 1 summarizes the offline verification and online loop.

IV. SIMULATION STUDY

We simulate two platforms spanning distinct redundancy ratios and task dimensions: a planar hexarotor ($n=6, m=4, \Delta = \ln 3$) and a fully-actuated tilted octorotor ($n=8, m=6, \Delta = \ln 4$), both with identical per-rotor parameters so $\lambda_k = m/n$ and Theorem 2 applies exactly. A fully-actuated hexarotor ($n=m=6$) is excluded: every rotor is essential ($\lambda_k = 1, \Delta_k = +\infty$) and the stratum-jump problem does not arise. Motor inertia and $\bar{\tau}_i/b_i$ are anchored to fixed-RPM wind-tunnel data from the UIUC Propeller Database¹ (APC 10 $\times$ 4.7 propeller profile), grounding saturation speeds in true aerodynamic limits; (1) is integrated at 1–2 kHz and (14) solved with OSQP².

We first verify that the geometric objects lie where the theory places them. Theorem 2 predicts $\Delta = \ln(n/(n-m))$ independently of motor parameters, propeller, or payload. Numerically, the hexarotor gives $L^{\max} = -10.128$, $\ell^\dagger = -11.226$ ($\Delta = 1.098 \approx \ln 3$) and the octorotor $L^{\max} = +1.448$, $\ell^\dagger = +0.062$ ($\Delta = 1.386 \approx \ln 4$)—both matching (8) to machine precision. The certifiable window $(\ell^\dagger, L^{\max})$ exceeds one nat on both, so (11) admits a non-degenerate range of floors rather than a knife-edge.

1) *Singularity avoidance and closed-loop tracking:* To expose the necessity of the barrier, four allocators drive the hexarotor through the same sinusoidal moment reversal at fixed collective: the effort-minimizing QP [3] (standard baseline), the greedy readiness maximizer of [10] (fiberwise $\arg \max h$), a low-pass filtered greedy variant (50 ms cutoff, the natural engineering patch), and the proposed filter (Algorithm 1, floor via Proposition 3 at $\kappa = 1/2$), summarized in Table I and Figures 4 and 5.

The central pathology appears in the authority-scarce rows of Table I (collective 0.7–0.8), where greedy DAAM violates

¹<https://m-selig.ae.illinois.edu/props/propDB.html>

²<https://github.com/osqp>

Table I: Collective-thrust sweep (hexarotor, moment amplitude $1.5 \times \text{ref}$, $\bar{\ell}$ per Proposition 3). “thr.”: collective thrust as a fraction of the nominal hover value; “TV”: total variation of the commanded allocation; “wErr”: RMS wrench error; $h_{\min} = \min_t h$. Bold: floor violated.

thr.	$\bar{\ell}$	effort-min			greedy DAAM			low-pass DAAM			Algorithm 1		
		$h_{\min}$	TV	wErr	$h_{\min}$	TV	wErr	$h_{\min}$	TV	wErr	$h_{\min}$	TV	wErr
0.6		not certifiable: $L^{\text{op}} = -11.49 < \bar{\ell}^* = -11.23$											
0.7	-11.14	+0.08	2.8	0.0003	-0.54	31.5	0.145	-0.54	31.5	0.147	+0.12	2.8	0.0018
0.8	-11.00	+0.22	2.6	0.0003	-0.48	8.4	0.113	-0.48	8.4	0.114	+0.22	2.6	0.0003
0.9	-10.91	+0.31	2.5	0.0003	+0.31	2.4	0.0069	+0.31	2.4	0.0076	+0.31	2.5	0.0003
1.0	-10.87	+0.35	2.3	0.0003	+0.35	2.3	0.0069	+0.35	2.3	0.0076	+0.35	2.3	0.0003
1.1	-10.87	+0.35	2.2	0.0003	+0.36	2.2	0.0069	+0.36	2.2	0.0076	+0.35	2.2	0.0003
1.2	-10.90	+0.33	2.1	0.0003	+0.33	2.1	0.0069	+0.33	2.1	0.0076	+0.33	2.1	0.0003




Figure 4: Collective-thrust sweep: worst-case authority $\min_t h$ (left) and Root Mean Square (RMS) wrench error (right).

the floor it maximizes ($h_{\min} = -0.54$ at 0.7, -0.48 at 0.8). The mechanism is in the total-variation column: as the moment reverses, the two symmetric optima $\pm v^*$ trade dominance, each handover is an orthant crossing (Section III) that drives rotors through the zero-spin dead zone, and the commanded allocation jumps. Its “TV” spikes to 31.5 at collective 0.7, an order of magnitude above the 2.8 of both continuous allocators, because the motors are slammed against their limits at every handover. The finite motor time constant converts this discontinuity into a sustained tracking error: RMS wrench error reaches 0.145 at 0.7 against 0.0018 for the proposed filter (eighty-fold), widening to 0.113 vs. 0.0003 at 0.8. Violation and error are one defect—the signature of an allocation optimal pointwise yet discontinuous in time.

The full sweep (Figure 4) makes the regime structure explicit. The greedy trace dives below the floor only in the scarce band where the optima are farthest apart; from collective 0.9 the optima stop exchanging dominance, no crossing is demanded, and all allocators coincide, with wrench error time-aligned to the floor excursions. Where the barrier is inactive (≥ 0.9) the filter reproduces its nominal allocator exactly, incurring no cost; where active (0.7) it improves worst-case authority over the effort-minimizing baseline (+0.12 vs. +0.08) at a wrench-error price of 1.8×10^{-3} , matching the $\mathcal{O}(\rho^{-1/2})$ bound of Theorem 10. In actuator space, greedy DAAM chords across the orthant boundary through $v_i = 0$, while the filtered trajectory stays in one orthant with $\min_{i,t} |v_i| = 0.335$, the uniform clearance of Proposition 5(i), absorbing the reversal in the null space rather than jumping.

Low-pass filtering deserves a fair reading because its failure is instructive. It succeeds at its stated job—Figure 5 (left) shows the peak commanded rate falling from 731 to 14.6 s^{-2} —but a filtered reversal traverses the *same distance* through actuator space, merely slowly. The phase lag then




Figure 5: Low-pass filtering failure: peak commanded actuator rate (left) and barrier value $h(t)$ (right).

stretches the dwell below the floor to 0.756 s, *longer* than the 0.629 s of the unfiltered baseline at nearly the same depth (Figure 5, right): a rate hazard is traded for an exposure hazard. The proposed filter is a different object, with peak rate 0.5 s^{-2} and zero floor-violation time. Finally, along the closed loop the alignment ratio (16) satisfies $\vartheta \leq 7 \times 10^{-3}$: the null space is orthogonal to ∇h to numerical precision, the exchange is zero-sum, and the entire margin is supplied by the strict-interior term $\alpha(h) > 0$, not by null-space ascent.

2) *Monte-Carlo mission sweeps*: To confirm the single-maneuver result is not an adversarial artifact, we draw random certifiable missions with collective $\sim \mathcal{U}[0.7, 1.2]$ and moment amplitude $\sim \mathcal{U}[1.0, 2.0]$. The certifiability frontier runs diagonally because $\psi_i = v_i^2 \bar{a}_i^2$ vanishes as $v_i \rightarrow 0$: large moments at low collective leave no authority to certify. Over the full envelope the barrier is largely inactive and the filter reproduces its nominal allocator at no performance cost. In the scarce corner ($\mathcal{U}[0.68, 0.82]$) the separation is absolute: greedy DAAM violates the floor in 80% of missions with wrench error $(85.9 \pm 18.6) \times 10^{-3}$, whereas the proposed filter holds strictly positive authority $(+0.148 \pm 0.020)$ at error $(1.9 \pm 0.9) \times 10^{-3}$ —a forty-five-fold gap achieved *together* with the authority guarantee greedy maximization was meant to provide. Greedy and low-pass are dominated on both axes at once, so uncertified maximization is not a tradeoff against the filter but strictly worse.

3) *Resilience to unmodeled delay*: Transport delay in the torque path is a failure channel no allocation certificate can ignore: Electronic Speed Controller (ESC), bus scheduling, and rate-loop filtering place 1–10 ms between the QP solution and the applied torque. Figure 6 sweeps τ_d to failure across three regimes. Below the certified ceiling $\tau_d^* = 1.96 \text{ ms}$ (Corollary 14) the margined floor $\bar{\ell} + \eta(\tau_d)$ deterministically protects the state; between τ_d^* and $\approx 100 \text{ ms}$ the certificate is silent yet the floor degrades only mildly ($\min_t h = +0.140$ at 20 ms, $+0.097$ at 50 ms), because the strict interior keeps the stale gradient descent-blocking well past expiry. The $30\text{--}50\times$ gap reflects the worst-case constant $KV \approx 400 \text{ s}^{-1}$ (the product of the Lipschitz constant of $\nabla h^\top d$ and the maximum actuator speed), attained only when every rotor accelerates most damagingly at full torque. Real ESC delays thus straddle τ_d^* —the lower end covered by the theorem, the upper by an order-of-magnitude margin—and Corollary 14 states exactly how to certify 10 ms (raise the floor by $\eta(\tau_d)$) and when it is impossible (beyond τ_d^* the tightened set is empty).

4) *Verification of worst-case robust invariance*: Motor torque limits and drag coefficients are the two quantities a real




Figure 6: Delay sweep: certified (green), tolerant (orange), and failure (red) regimes; hatched: typical ESC range. $\min_t h$ (red), bound $-\eta(\tau_d)$ (dashed, Proposition 13), margined floor enforced for $\tau_d \leq \tau_d^*$ (blue squares, dotted).

Table II: Robustness campaign (8 plants per level). “viol.”: fraction below the certified floor. Lower: aggressive floor ($\kappa = 0.9$), isolating the drift term of Theorem 11.

p	controller	mean	worst	viol.	wErr	$\bar{\ell}$
0%	nominal barrier	+0.164	+0.164	0.00	0.0003	-11.06
10%	effort-min QP	-0.196	-0.504	0.88	0.0035	-11.06
	nominal barrier	+0.014	-0.436	0.38	0.0047	-11.06
	robust (Thm. 11)	+1.82	+1.18	0.00	0.0193	-12.71
20%	effort-min QP	-0.610	-1.22	0.88	0.0062	-11.06
	nominal barrier	-0.103	-0.940	0.50	0.0127	-11.06
	robust	+3.56	+2.25	0.00	0.0347	-14.47
30%	effort-min QP	-1.08	-2.00	0.88	0.0083	-11.06
	nominal barrier	-0.231	-1.55	0.50	0.0392	-11.06
	robust	+5.44	+3.40	0.00	0.0466	-16.39
$\kappa=0.9$ ablation: metric bound alone vs. full robust filter						
10%	metric bound only	+0.114	+0.099	0.00	0.0165	-
	robust	+0.127	+0.104	0.00	0.0045	-
20%	metric bound only	+0.099	+0.066	0.00	0.0284	-
	robust	+0.156	+0.121	0.00	0.0068	-
30%	metric bound only	+0.025	-0.089	0.38	0.0367	-
	robust	+0.175	+0.131	0.00	0.0093	-

vehicle least holds constant—torque sags with voltage, drag drifts with wear—so we subject $(\bar{\tau}_i, b_i)$ to $\pm p$ mismatches over eight random plants per level (Table II). The nominal certificate is predictably brittle: enforced on the wrong plant it fails on 38% at $p = 10\%$ (worst case -0.436) and 50% at $p = 20\text{--}30\%$, while the uncertified effort-minimizing QP fails on 88% everywhere. This is not a flaw in the construction but the expected result of evaluating an exact object at the wrong point of parameter space—the failure Theorem 11 eliminates. The robust formulation restores zero violations at a price known before flight: Corollary 12 gives the floor shift $L_p^{\max} - L^{\max} = m \ln \frac{(1-p)^3}{1+p}$, independent of A and exact at the corner $\zeta^- \in \mathcal{Z}_p$, so the enforced floors descend from -11.06 to -12.71 , -14.47 , -16.39 as p grows to 30% (retained authority volume 44%, 18%, 7%).

The lower block of Table II isolates the ablation: with the metric bound alone (barrier at ζ^- , drift nominal) the certificate still fails on 38% at $p = 30\%$, since bounding a function is not bounding its derivative. Restoring the closed-form drift—one $\mathcal{O}(n)$ pass—returns violations to zero and cuts wrench error fourfold ($0.0367 \rightarrow 0.0093$); metric and drift bounds are separately necessary, sufficient only together.

V. CONCLUSION

Treating multirotor control authority as a certified, forward-invariant quantity rather than an objective to be maximized yields an allocation filter that protects the certified set from the two collapse modes to which effort- and readiness-based allocators are blind. For the modeled dynamics it gives exact confinement, an explicit alignment ratio dictating when tracking stays exact, and an $\mathcal{O}(\rho^{-1/2})$ bound otherwise, all from the closed-form inequality $\bar{\ell} > \ell^\dagger$ with no vertex enumeration. Because torque enters the metric cubically, compensating battery sag is roughly three times more valuable than refining drag estimates, and robust invariance is recovered in a single $\mathcal{O}(n)$ pass at a closed-form authority cost (Corollary 12).

Scope and limitations. Three boundaries delimit the guarantees. The certificates use the quasi-static drag model $-b_i v_i |v_i|/m_i$; in fast forward flight inflow makes b_i state-dependent, which Theorem 11 absorbs as bounded $\pm p$ variation though large excursions warrant online estimation of b_i . Invariance is proved in continuous time, with Remark 9 bounding the zero-order-hold error absorbed by the strict interior at kilohertz rates. Finally, ∇h uses measured speeds, so feedback noise enters the constraint row, tolerated by the same margin. These motivate flight validation under aerodynamic effects that deform D and online adaptation of $\mathcal{Z}_p$.

REFERENCES

- [1] A. Ollero *et al.*, “Past, Present, and Future of Aerial Robotic Manipulators,” *IEEE Trans. Rob.*, vol. 38, no. 1, pp. 626–645, 2022.
- [2] R. Rashad *et al.*, “Fully Actuated Multirotor UAVs: A Literature Review,” *IEEE Robot. Autom. Mag.*, vol. 27, no. 3, pp. 97–107, 2020.
- [3] T. A. Johansen *et al.*, “Control Allocation—A Survey,” *Automatica*, vol. 49, no. 5, pp. 1087–1103, 2013.
- [4] M. Bodson, “Evaluation of Optimization Methods for Control Allocation,” *J. Guid. Control Dyn.*, vol. 25, no. 4, pp. 703–711, 2002.
- [5] M. Ryll *et al.*, “A Novel Overactuated Quadrotor Unmanned Aerial Vehicle: Modeling, Control, and Experimental Validation,” *IEEE Trans. Control Syst. Technol.*, vol. 23, no. 2, pp. 540–556, 2015.
- [6] G. Michieletto *et al.*, “Fundamental Actuation Properties of Multirotors: Force–Moment Decoupling and Fail–Safe Robustness,” *IEEE Trans. Rob.*, vol. 34, no. 3, pp. 702–715, 2018.
- [7] D. Bicego *et al.*, “Nonlinear Model Predictive Control with Enhanced Actuator Model for Multi-Rotor Aerial Vehicles with Generic Designs,” *J. Intell. Robot. Syst.*, vol. 100, no. 3–4, pp. 1213–1247, 2020.
- [8] P.-J. Bristeau *et al.*, “The role of propeller aerodynamics in the model of a quadrotor UAV,” in *ECC*, 2009, pp. 683–688.
- [9] T. Yoshikawa, “Manipulability of Robotic Mechanisms,” *Int. J. of Robot. Res.*, vol. 4, no. 2, pp. 3–9, 1985.
- [10] A. Franchi, “Aero-Promptness: Drag-Aware Aerodynamic Manipulability for Propeller-driven Vehicles,” *arXiv:2603.07998*, 2026.
- [11] —, “Muscle Coactivation in the Sky: Geometry and Pareto Optimality of Energy vs. Aerodynamic Promptness and Multirotors as Variable Stiffness Actuators,” in *ICUAS*, 2026, pp. 235–242.
- [12] —, “Global Geometry of Orthogonal Foliations of Signed-Quadratic Systems,” *arXiv:2604.01912*, 2026.
- [13] A. D. Ames *et al.*, “Control Barrier Function Based Quadratic Programs for Safety Critical Systems,” *IEEE Trans. on Autom. Control*, vol. 62, no. 8, pp. 3861–3876, 2017.
- [14] F. Blanchini, “Set Invariance in Control,” *Automatica*, vol. 35, no. 11, pp. 1747–1767, 1999.
- [15] X. Xu *et al.*, “Robustness of Control Barrier Functions for Safety Critical Control,” in *IFAC ADHS*, vol. 48, no. 27, 2015, pp. 54–61.
- [16] A. Agrawal *et al.*, “Discrete Control Barrier Functions for Safety-Critical Control of Discrete Systems with Application to Bipedal Robot Navigation,” in *RSS*, 2017.
- [17] S. M. Robinson, “Strongly Regular Generalized Equations,” *Math. Oper. Res.*, vol. 5, no. 1, pp. 43–62, 1980.