

Tilt as a Certified Resource: Preserving Motor Wrench-Rate Authority on Articulated Multirotors

Giuseppe Silano¹ and Martin Saska²

Abstract—Fully-actuated multirotor aerial vehicles must not only track nominal wrenches but retain the “readiness” to modulate them rapidly under disturbances. Classical effort-minimizing allocators ignore this dynamic limit, whereas maximizing readiness leads to topologically disconnected optimal sheets demanding physically impossible actuator rates. Enforcing a readiness safety floor on fixed-geometry symmetric platforms further encounters a zero-sum degeneracy: motor-speed redistribution cannot improve authority without conceding wrench tracking. This paper uses active morphology to break the degeneracy, treating servo tilt as a geometric resource supplying authority-recovery directions unavailable to static rotors. We construct a configuration-dependent, motor-only readiness certificate—the log-volume of the reachable wrench-rate set—that explicitly excludes servo capacity, preventing a “ghost capacity fallacy” in which the certificate would falsely credit slow mechanical kinematic limits instead of collapsing accurately at motor saturation. The certificate is enforced as a Control Barrier Function (CBF) within a Unified Physical-Command Quadratic Program acting on motor torques and servo setpoints. Closed-loop simulations of an articulated octorotor under severe gust disturbances show classical allocators diverging and uncertified articulated allocators violating the safety floor, while the proposed CBF filter bounds the system state and preserves vehicle authority.

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

I. INTRODUCTION

Fully-actuated Multi Rotor Aerial Vehicles (MRAVs) have expanded aerial robotics’ operational envelope [1], [2], decoupling translational and rotational dynamics through over-actuation to track arbitrary six-degree-of-freedom trajectories and sustain physical interaction with the environment [3]. This redundancy introduces the classically ill-posed control allocation problem: mapping a 6-dimensional target wrench to a higher-dimensional actuator-command space [4].

The classical answer minimizes actuator effort via the Moore–Penrose pseudo-inverse or a quadratic program [4]—computationally light, but blind to a distinct dynamic limit we call *wrench-rate readiness*: the rotors’ remaining capacity

to generate rapid wrench variations. Propeller thrust and drag both scale quadratically with spin rate, imposing competing limits: a slow rotor has a shallow thrust slope and cannot modulate force quickly, while a rotor near saturation has spent its torque budget fighting drag and has no headroom to accelerate. Readiness therefore peaks at an interior *sweet spot* and collapses at both extremes of the speed range.

Consider an articulated MRAV holding a contact force against a wall when a gust arrives: it may track its commanded wrench exactly, every rotor inside its speed limits, effort-minimization problem solved to optimality, and still fail—the rotors spinning fast enough that drag has consumed the torque budget, leaving no capacity to accelerate [5], [6]. Producing a nominal wrench and retaining the authority to change it are distinct capabilities.

Recent geometric work formalizes the latter through a drag-aware authority co-metric, Drag-Aware Aerodynamic Manipulability (DAAM), whose log-determinant measures the reachable wrench-rate volume [6], [7]. The co-metric $D(v)$, with $v \in \mathbb{R}^n$ the rotor spin rates, maps actuator bounds into task space; and the *readiness level* $L(v) = \ln \det D(v)$ is that set’s scalar log-volume. Maximizing L along the *task fiber*—the continuum of redundant rotor speeds satisfying the current wrench command—repels the allocation from both collapse modes; two obstacles remain.

The first is topological: the fiberwise maximizer is set-valued, with optimal selections on disconnected sheets. The signed-quadratic thrust map partitions rotor-speed space into *sign orthants*—convex regions where every rotor keeps a fixed spin direction, bounded by loci where it reverses. At such a zero crossing the thrust slope vanishes, the allocation Jacobian degenerates, and the actuator rate needed to sustain a bounded wrench rate diverges, so the fiber decomposes into components a bounded-rate allocator cannot cross (Figure 1). Low-pass filtering bounds the rate demand but introduces phase lag, dragging the trajectory through the low-authority crossing rather than avoiding it.

The second is geometric, and motivates this paper. Rather than maximize authority, one may certify it via a floor on L as a Control Barrier Function (CBF) [8], [9]. On symmetric fixed-geometry vehicles this meets a structural limitation [10]: at a regular constrained optimum—interior in rotor speed, wrench-rate Jacobian of full row rank—every motor-speed perturbation preserving the commanded wrench leaves L unchanged to first order, so the constant-wrench ascent available to motor redistribution is zero. Physically, the readiness gained by accelerating one rotor is cancelled by decelerating another; authority at the floor can then be recovered only by conceding wrench tracking. Section IV formalizes this; the configuration

¹Giuseppe Silano is with the Department of Power Generation Technologies and Materials, Ricerca sul Sistema Energetico (RSE) S.p.A., 20134 Milan, Italy, and also with the Department of Cybernetics, Czech Technical University in Prague, 12135 Prague, Czech Republic (e-mail: giuseppe.silano@fel.cvut.cz).

²Martin Saska is with the Department of Cybernetics, Czech Technical University, 12135 Prague, Czech Republic (e-mail: martin.saska@fel.cvut.cz).

This work was partially funded by the research fund for the Italian Electrical System (decree n. 388, Nov. 6th, 2024), by the GAČR project no. 26-22419S, by the CTU grant no. SGS26/077/OHK3/1T/13, and by the European Union under ROBOPROX reg. no. CZ.02.01.01/00/22_008/0004590.

The author used Google Gemini 3 and Anthropic Claude 5 Opus for proofreading and exposition refinement, but independently verified all outputs and assumes full responsibility for the article’s content.




Figure 1: Allocation section over commanded wrench w for rotor speed v_k . Greedy maximization jumps between optimal sheets $\Sigma_{\pm}$ (dashed), demanding physically impossible actuator rates. The proposed barrier ($h \geq 0$) strictly confines the allocation to a single certified band, resolving authority-tracking conflicts via an explicit wrench slack δ .

that efficient hover selects is exactly where motor-only certification becomes indefensible.

We break the degeneracy through active morphology [1], [2], treating servo tilt as geometric resource rather than a tracking variable: physically tilting the rotors alters the allocation matrix’s geometry, contributing wrench-rate directions unreachable by redistributing spin among fixed rotors. We are deliberate about scope—articulation is bidirectional, and an arbitrary tilt may raise or lower the metric—but a strictly favorable direction always *exists* where static motor redistribution has none, and a certified allocator will select it.

We therefore certify rather than maximize: a configuration-dependent, motor-only readiness certificate enforces a floor on the log-determinant, crediting no servo authority and remaining conservative in the Löwner order—not merely from caution, but because a servo-inclusive metric stays finite at motor saturation and would report a vehicle with no torque headroom as healthy. Restricting the metric to motor capacity preserves the divergence of L at saturation and, with it, the property that a single inequality fences both collapse modes (Section IV).

Contributions. The core contributions of this manuscript are fourfold:

- A first-order geometric characterization of the motor-only readiness degeneracy at a regular constrained optimum, with explicit regularity hypotheses (Section IV).
- A configuration-dependent readiness certificate, conservative in the Löwner order, whose conservatism is what preserves the saturation indication that a servo-inclusive metric destroys (Sections III–IV).
- Enforcement of that certificate on motor and servo commands through a memoryless Unified Physical-Command Quadratic Program (UPC-QP) with relative-degree-one barrier dynamics and an online feasibility margin, structurally free of integrator wind-up (Section V).
- A controlled three-way evaluation isolating the causal contribution of the readiness constraint from articulation alone under severe disturbance, in which an articulated allocator without the certificate still violates the floor (Section VI).

II. RELATED WORK

Allocation for overactuated systems is mature [4]; MRAV specialization is well developed: tilted-rotor platforms decouple translational and rotational dynamics [2], [11]; actuation properties are characterized via force–moment decoupling [5]; surveys cover the platform class [1] and applications [3]. The dominant allocator—pseudo-inverse or effort-minimizing program—ignores actuator bounds and saturates afterwards. Constrained dynamic allocation embeds those limits directly [12], reference governors enforce admissibility [13], full-pose tracking extends to limited actuation [14] and bounded lateral force [15], and null-space methods exploit redundancy instantaneously [16] or over a horizon [17]. All certify that the commanded wrench is currently *attainable*; none floors the capacity to *change* it, which is what decides whether a gust can be answered.

Measuring capability by reachable-set volume originates with Yoshikawa’s manipulability index [18]. DAAM [6] adapts this to MRAVs by weighting each rotor by its *aerodynamic* state, so the metric collapses as drag consumes torque margin, trading off against energy when maximized along the task fiber [7]. We adopt this co-metric but not its use: a floor rather than a maximum sidesteps the set-valued, disconnected fiberwise maximizer.

Such a floor is naturally a CBF [19]. Determinant-type capability barriers appear in manipulator control: [8] certify singularity avoidance with an exponential CBF on the manipulability index. Closest to this work, [20] certify a margin from rank-deficient regions of a state-dependent input–output map, using one barrier per eigenvalue—a construction requiring *simple* eigenvalues. This is where our setting diverges: a symmetric airframe carries exactly repeated eigenvalues by construction, so per-eigenvalue barriers lose differentiability precisely where we need them, whereas $\ln \det D$ stays smooth wherever $D \succ 0$. Their map is also weighted uniformly, while ours carries the drag-dependent capacity producing the collapse in the first place.

Feasibility of the enforcing program is a known concern, addressed through conditions on the optimal control problem, input-constrained constructions, soft-minimum barriers, and composite certificates for actuator lag [19]; we inherit rather than resolve it, since our barrier acts on physical commands whose admissible set is a box, giving a closed-form per-step margin that we monitor instead of assume.

Rotor tilting is itself established for reshaping the feasible wrench set [2], [11], with recent omnidirectional platforms exploring actuation-aware allocation experimentally [17], [21]: there, tilt *reaches* wrenches; we use it to *preserve* the ability to change them. Fault-tolerant control is the complement in time, accommodating rotor loss after it occurs [22]; readiness acts before.

III. SYSTEM MODEL AND READINESS CERTIFICATE

This section fixes the vehicle model, constructs the readiness co-metric, and separates three objects the paper keeps distinct. Every quantity here is computable from the current state by a single $m \times m$ factorization.

A. Physical actuator dynamics

The vehicle carries $n > m = 6$ actively articulable rotor groups. Rotor i spins at a signed rate v_i , axis tilted by servo angle α_i . Physical inputs are motor torques $\tau \in \mathbb{R}^n$ and the servo setpoints $\alpha_c \in \mathbb{R}^n$:

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

$$\dot{\alpha}_i = \frac{1}{\tau_s} (\alpha_{c,i} - \alpha_i), \quad |\dot{\alpha}_i| \leq \bar{u}, \quad \alpha_i \in [\alpha, \bar{\alpha}], \quad (2)$$

with lumped motor-propeller inertia $J_{m,i}$, drag coefficient b_i , torque limit $\bar{\tau}_i$, servo time constant τ_s , and servo rate limit $\bar{u}$. The first-order rotor model with quadratic drag (1) is standard and experimentally identified [5]; the servo is a symmetric first-order lag, chosen to keep physical commands in the barrier's first derivative—whether that suffices to avoid high-order constructions [23] is established in Section V, not assumed.

Equation (1) determines the certificate's basis. For $v_i > 0$ the admissible spin acceleration satisfies $|\dot{v}_i| \leq \bar{a}_i(v_i)$, where

$$\bar{a}_i(v_i) = (\bar{\tau}_i - b_i v_i^2) / J_{m,i}, \quad v_i^{\text{sat}} = \sqrt{\bar{\tau}_i / b_i}, \quad (3)$$

is the *symmetric acceleration capacity*: maximal at rest, decreasing monotonically to zero at the drag-limited terminal speed v_i^{sat} , where drag consumes the entire torque budget—formalizing that a rotor may be far from any speed limit yet have no capacity to accelerate.

B. Wrench map and wrench-rate jacobians

The body wrench follows the signed-quadratic thrust map

$$w = A(\alpha) \phi(v) \in \mathbb{R}^m, \quad \phi_i(v_i) = v_i |v_i|, \quad (4)$$

where $A_{\bullet,i} \in \mathbb{R}^m$ collects the thrust direction and induced drag torque of rotor i at cant α_i . With $A'_{\bullet,i} = \partial A_{\bullet,i} / \partial \alpha_i$ the tilt direction of column i , differentiating (4) gives

$$\dot{w} = J_v(v, \alpha) \dot{v} + J_\alpha(v, \alpha) \dot{\alpha}, \quad (5)$$

with $J_v = 2A(\alpha) \text{diag}(|v|)$ and $[J_\alpha]_{\bullet,i} = A'_{\bullet,i} \phi_i(v)$. Substituting (1)–(2) renders the wrench rate affine in the physical command $u = [\tau^\top, \alpha_c^\top]^\top$,

$$\dot{w} = M(v, \alpha) u + d(v, \alpha), \quad (6)$$

with input map and drift

$$M = [J_v D_m \quad \frac{1}{\tau_s} J_\alpha], \quad d = -J_v D_m b \odot \phi(v) - \frac{1}{\tau_s} J_\alpha \alpha, \quad (7)$$

where $D_m = \text{diag}(J_{m,i}^{-1})$, $b = (b_1, \dots, b_n)$, and $\odot$ is the elementwise product. Drift d collects command-independent terms: drag torque to overcome and the servo lag on current tilt. Equation (6) is the form the Section V allocator acts on.

C. The motor-only readiness co-metric

Weighting each rotor by its capacity (3), the wrench rates achievable by the motors at fixed articulation contain the ellipsoid $\mathcal{E} = \{J_v W^{1/2} \zeta : \|\zeta\|_2 \leq 1\}$ with $W = \text{diag}(\bar{a}_i^2)$, whose squared volume is proportional to the determinant of the DAAM co-metric

$$D(v, \alpha) = J_v W J_v^\top = 4 A(\alpha) \Psi(v) A(\alpha)^\top, \quad (8)$$




Figure 2: Readiness certificate geometry. (a) Actuator-rate box bounded by symmetric capacity $\bar{a}_i(v_i)$. (b) Reachable wrench-rate zonotope under-approximated by inscribed ellipsoid $\mathcal{E}$. (c) Readiness level $L = \ln \det D$ measures the log-volume of this inner approximation.

with $\Psi = \text{diag}(\psi_i)$ and the per-rotor readiness weight is

$$\psi_i(v_i) = v_i^2 \bar{a}_i(v_i)^2, \quad (9)$$

where a prime denotes differentiation with respect to the rotor's own spin rate, $\psi'_i = d\psi_i/dv_i$. The scalar readiness level and the shifted safety barrier are

$$L(v, \alpha) = \ln \det D(v, \alpha), \quad h(v, \alpha) = L(v, \alpha) - \ell_{\min}, \quad (10)$$

where $\ell_{\min} \in \mathbb{R}$ is the designer-chosen *readiness floor*. The certified set is $\{(v, \alpha) : h(v, \alpha) \geq 0\}$, and Section VI sets $\ell_{\min}$ relative to nominal hover readiness.

Two properties of (9) explain why a single scalar inequality suffices: ψ_i vanishes at $v_i = 0$ (zero thrust slope, no wrench rate contribution) and at $v_i = v_i^{\text{sat}}$ ($\bar{a}_i = 0$, no acceleration capacity)—in either case $\det D$ loses that rotor's contribution, so a floor on L fences *both* failure modes, dropout and saturation, without auxiliary logic. Moreover $\mathcal{E}$ is the image of the unit 2-ball, contained in the actuator-rate box, so (8) under-approximates the true reachable zonotope (Figure 2) and L never overstates authority; we avoid the maximum-volume (Löwner–John) ellipsoid, since the 2-ball image is closed-form and differentiable in (v, α) , and the resulting conservatism is one-sided, hence safe.

Remark 1 (Three distinct objects). *The motor-only co-metric $D(v, \alpha)$ in (8), depending only on current state, is the sole basis of barrier h ; the augmented metric $D_{\text{art}} = D + \sum_i \rho_i A'_{\bullet,i} A'_{\bullet,i}^\top$, with $\rho_i = \bar{u}^2 \phi_i(v)^2$, accounting for wrench rate from servo motion, is used in Section IV only to quantify the authority the certificate declines to credit and never enters a control constraint; and the plant (1)–(2) supplies the actuator rates, with u the only quantity the allocator commands. Proposition 6 shows that restricting the certificate to motor capacity is the right choice.*

D. Gradients in closed form

Both gradients needed to enforce (10) are available analytically at every control instant.

Lemma 2 (Readiness gradients). *Let $D(v, \alpha) \succ 0$ and define the rotor leverage s_i and tilt cross-leverage c_i ,*

$$s_i = A_{\bullet,i}^\top D^{-1} A_{\bullet,i}, \quad c_i = A_{\bullet,i}^\top D^{-1} A'_{\bullet,i}. \quad (11)$$

Then

$$\frac{\partial L}{\partial v_i} = 4 \psi'_i(v_i) s_i, \quad \frac{\partial L}{\partial \alpha_i} = 8 \psi_i(v_i) c_i. \quad (12)$$

Proof. By Jacobi's formula, $\partial \ln \det D = \text{tr}(D^{-1} \partial D)$. Only the i -th term of (8) depends on v_i or α_i : $\partial D / \partial v_i = 4\psi_i' A_{\bullet i} A_{\bullet i}^\top$ and $\partial D / \partial \alpha_i = 4\psi_i (A_{\bullet i}' A_{\bullet i}^\top + A_{\bullet i} A_{\bullet i}'^\top)$. Taking traces against D^{-1} gives (12), using symmetry of D^{-1} . $\square$

Both expressions share one factorization of D , so the cost per step is a single $m \times m$ inverse and $2n$ quadratic forms—negligible for embedded flight control.

Define the capacity-weighted *leverage score* $\sigma_i = 4\psi_i s_i$. The trace identity gives

$$\sum_i \sigma_i = \text{tr}(D^{-1} D) = m, \quad (13)$$

so the m wrench dimensions are shared among n rotors, σ_i the fraction of readiness volume rotor i supports.

This also quantifies the cost of losing a rotor: setting $\psi_i = 0$ removes the rank-one term $4\psi_i A_{\bullet i} A_{\bullet i}^\top$ from (8), the matrix determinant lemma gives $\det(D - 4\psi_i A_{\bullet i} A_{\bullet i}^\top) = \det(D) (1 - \sigma_i)$, so the readiness drop is exactly

$$\Delta_i = -\ln(1 - \sigma_i). \quad (14)$$

The fraction $\sigma_i \rightarrow 1$ marks an essential rotor whose loss makes D singular, $\Delta_i \rightarrow \infty$. On a symmetric platform at uniform spin, $\sigma_i = m/n$, and (14) reduces to the finite *dropout gap* $\Delta = \ln(n/(n - m))$ —a bounded, predictable degradation.

Equation (13) has one further consequence: when every σ_i is equal, $\nabla_v L$ is proportional to the vector of ones, so a constant-wrench motor redistribution trades readiness between rotors at no net gain—stated precisely with regularity conditions in Section IV.

IV. ARTICULATION AS A READINESS RESOURCE

We now show that rotor-speed redistribution alone cannot defend the certificate at the configuration efficient flight selects, that articulation supplies what is missing, and that the certificate must nonetheless exclude servo authority.

A. The zero-sum degeneracy

Assumption 1. v^* is interior to $\{v : 0 < v < v^{\text{sat}}\}$, $J_v(v^*)$ has full row rank m , and v^* is a regular local maximizer of $L(\cdot, \alpha)$ subject to $w(v, \alpha) = w^*$, so a unique Lagrange multiplier $\lambda \in \mathbb{R}^m$ exists.

Proposition 3 (Zero-sum authority degeneracy). *Under Assumption 1, $\nabla_v L(v^*) = J_v(v^*)^\top \lambda$, and the constant-wrench authority ascent available through rotor-speed redistribution,*

$$\gamma = \max \{ \nabla_v L^\top \xi : J_v \xi = 0, |\xi| \leq \bar{a} \}, \quad (15)$$

is zero. Here $\xi \in \mathbb{R}^n$ is a candidate spin-acceleration perturbation.

Proof. Interiority removes the box constraints; full row rank of $J_v(v^*)$ is the linear independence constraint qualification under which λ exists uniquely. The first-order necessary condition for the constrained maximizer is therefore $\nabla_v L(v^*) = J_v(v^*)^\top \lambda$. Any feasible ξ in (15) lies in $\ker J_v(v^*)$, whence

$$\nabla_v L^\top \xi = (J_v^\top \lambda)^\top \xi = \lambda^\top (J_v \xi) = 0. \quad (16)$$




Figure 3: Authority ascent geometry. (a) At v^* , the readiness gradient is orthogonal to $\ker J_v$, yielding zero constant-wrench ascent ($\gamma = 0$). (b) Active tilt supplies a recovery direction outside $\ker J_v$, restoring positive ascent ($\gamma_{\text{art}} > 0$) at the same commanded wrench.

The linear objective therefore vanishes on the entire feasible set, which is nonempty since $\xi = 0$ is admissible, so its maximum is zero. $\square$

Two qualifications: the result is *first-order*—no admissible constant-wrench perturbation increases L to first order, not that the state cannot move—and *local*, holding at v^* under Assumption 1. What it establishes is that the recovery direction available to rotor-speed redistribution vanishes exactly where readiness is optimal, so a motor-only allocator meeting the floor there must concede wrench tracking to restore authority.

The mechanism is concrete (Figure 3a): on a symmetric octorotor in hover, $n = 8$, $m = 6$ leave a genuine two-dimensional null space, and since all rotors share an operating point, (13) gives every rotor $\sigma_i = m/n$ and $\nabla_v L \propto \mathbf{1}$. Accelerating one rotor group while holding the wrench forces compensating deceleration elsewhere, and since every rotor has identical thrust sensitivity and drag headroom, the readiness gained is cancelled exactly by the readiness lost—the exchange is zero-sum, which is (16) read backwards.

B. Restoring the ascent through articulation

Admitting the servo tilt rates $\mu = \dot{\alpha} \in \mathbb{R}^n$ as a second resource enlarges (15) to

$$\gamma_{\text{art}} = \max_{\xi, \mu} \nabla_v L^\top \xi + \nabla_\alpha L^\top \mu \quad \text{s.t. } J_v \xi + J_\alpha \mu = 0, \quad |\xi| \leq \bar{a}, \quad |\mu| \leq \bar{u}, \quad (17)$$

where the equality constraint holds the commanded wrench—tilt perturbs it, motors absorb the perturbation (Figure 3b). This *actuator-envelope* quantity bounds what the actuators could do at the current state; Section V decides what they actually do.

Proposition 4 (Certified articulated ascent). *Let J_v^+ be the Moore–Penrose pseudoinverse; define the reduced tilt gradient and the induced motor-correction direction*

$$\tilde{g} = \nabla_\alpha L - J_\alpha^\top (J_v^+)^T \nabla_v L, \quad z = J_v^+ J_\alpha \text{sign}(\tilde{g}). \quad (18)$$

Then $\gamma_{\text{art}} \geq \beta \|\tilde{g}\|_1 \bar{u}$, with the dimensionless feasibility factor $\beta = \min\{1, \min_i \bar{a}_i / (\bar{u} |z_i|)\}$.

Proof. Take the tilt rate $\mu^* = \bar{u} \text{sign}(\tilde{g})$, which saturates the servo box, and the minimum-norm motor correction $\xi^* = -J_v^+ J_\alpha \mu^* = -\bar{u} z$ holding the wrench, so the pair satisfies the equality constraint of (17) by construction. Substituting into the objective and using $\text{sign}(\tilde{g})^\top \tilde{g} = \|\tilde{g}\|_1$ gives $\nabla_v L^\top \xi^* + \nabla_\alpha L^\top \mu^* = \|\tilde{g}\|_1 \bar{u}$. The pair need not be admissible, since ξ^*

may exceed the motor box; but scaling both components by $\beta \in (0, 1]$ preserves the equality constraint, scales the objective linearly, and by the definition of β yields $|\beta \xi_i^*| = \beta \bar{u} |z_i| \leq \bar{a}_i$ for every i . The scaled pair is thus feasible and attains $\beta \|\tilde{g}\|_1 \bar{u}$, which lower-bounds the maximum. $\square$

The factor β is the price of holding the wrench: tilting perturbs it, motors must undo the perturbation, and $\beta < 1$ when they lack headroom to do so at full servo rate.

At v^* the construction simplifies: since $\nabla_v L$ vanishes there, the second term of (18) drops and $\tilde{g}$ becomes purely geometric, $\tilde{g}_i = 2A_{\bullet i}^\top (AA^\top)^{-1} A_{\bullet i}'$, independent of spin, alternating in sign around a symmetric airframe—opposite-sign tilt on adjacent arms inflates readiness volume while induced wrench perturbations cancel in pairs. A gradient proportional to $\mathbf{1}$ cannot do this in the rotor-speed null space, where the same pairing cancels the readiness gain as well; articulation escapes the trade because it changes the geometry of $A(\alpha)$ rather than redistributing effort within a fixed one, avoiding the drag bottleneck entirely.

Remark 5 (Bidirectionality). *Articulation is not uniformly beneficial: an arbitrary tilt perturbation may lower L as easily as raise it. Proposition 4 is an existence statement: a favorable direction exists where rotor-speed redistribution has none.*

C. Why the certificate excludes servo authority

Servo motion produces wrench rate, so a metric crediting it exists: $D_{\text{art}} = D + \sum_i \rho_i A_{\bullet i}' A_{\bullet i}'^\top$ with $\rho_i = \bar{u}^2 \phi_i(v)^2$. We deliberately do not certify against it.

Proposition 6 (Conservatism and saturation indication). (i) $D \preceq D_{\text{art}}$ in the Löwner order, hence $L \leq L_{\text{art}}$ pointwise and $\{L \geq \ell_{\min}\} \subseteq \{L_{\text{art}} \geq \ell_{\min}\}$: certifying $h \geq 0$ certifies the articulated authority as well. (ii) If $\text{rank}[A_{\bullet 1}' \phi_1, \dots, A_{\bullet n}' \phi_n] = m$ at saturation, then L_{art} remains finite as $v \rightarrow v^{\text{sat}}$, whereas $L \rightarrow -\infty$.

Proof. (i) The difference $D_{\text{art}} - D = \sum_i \rho_i A_{\bullet i}' A_{\bullet i}'^\top$ is a sum of rank-one positive semidefinite terms, hence $\succeq 0$. Monotonicity of $\ln \det$ on the positive-definite cone gives $L \leq L_{\text{art}}$, and the set inclusion is immediate. (ii) As $v \rightarrow v^{\text{sat}}$ we have $\bar{a} \rightarrow 0$ by (3), so $\Psi \rightarrow 0$, $D \rightarrow 0$ and $L \rightarrow -\infty$. The servo weights behave oppositely: $\rho_i \rightarrow \bar{u}^2 \phi_i(v^{\text{sat}})^2 > 0$, since ϕ_i is evaluated at the *largest* admissible spin. Under the rank hypothesis $\det D_{\text{art}}$ tends to a finite positive limit. $\square$

Part (ii) forces the choice: under a gust the response rests on motors (high electrical bandwidth), while servos are mechanically slow. A barrier built on D_{art} would nevertheless read those slow kinematic limits as authority—the *ghost capacity* fallacy—reporting a vehicle with no remaining torque margin as safe. Restricting the metric to motor capacity avoids this and, by part (i), costs only conservatism.

D. The admissible tilt set

Tilt-dependent allocation matrices lose rank at flat configurations [11], usually handled by a dedicated constraint; here the floor supplies it.

Proposition 7 (Singularity exclusion). *Let $L^{\max}(\alpha) = \max_v L(v, \alpha)$ and $\mathcal{A}(\ell_{\min}) = \{\alpha : L^{\max}(\alpha) > \ell_{\min}\}$. Then $\mathcal{A}(\ell_{\min})$ contains no configuration with $\text{rank } A(\alpha) < m$, and on $\mathcal{A}(\ell_{\min})$ the certified set $\{h \geq 0\}$ has nonempty interior in v .*

Proof. If $\text{rank } A(\alpha) < m$ then $D(v, \alpha) = 4A\Psi A^\top$ is rank-deficient for every v by (8), so $\det D = 0$, $L^{\max}(\alpha) = -\infty < \ell_{\min}$, and $\alpha \notin \mathcal{A}(\ell_{\min})$. For the second claim, some v attains $L(v, \alpha) > \ell_{\min}$; continuity of L in v gives a neighborhood with $h > 0$. $\square$

The floor defending authority thus also fences the singularity, needing no separate avoidance logic. The two effects pull oppositely on cant angle: the available ascent $\|\tilde{g}\|_1$ is largest at shallow cant, while $\mathcal{A}(\ell_{\min})$ excludes the shallowest configurations precisely because they approach rank deficiency.

V. ENFORCEMENT THROUGH PHYSICAL COMMANDS

What remains is enforcing the floor using hardware commands—motor torques and servo setpoints—rather than the idealized actuator rates of (17).

A. Barrier dynamics and relative degree

Write $x = (v, \alpha)$, differentiating h along (1)–(2) gives

$$\begin{aligned} \dot{h} &= \sum_i \frac{\partial L}{\partial v_i} \frac{\tau_i - b_i v_i |v_i|}{J_{m,i}} + \sum_i \frac{\partial L}{\partial \alpha_i} \frac{\alpha_{c,i} - \alpha_i}{\tau_s} \\ &= a_{\text{CBF}}(x)^\top u - \delta_h(x), \end{aligned} \quad (19)$$

which is affine in the physical command $u = [\tau^\top, \alpha_c^\top]^\top$. Here

$$a_{\text{CBF}} = \begin{bmatrix} D_m \nabla_v L \\ D_s \nabla_\alpha L \end{bmatrix}, \quad \delta_h = \nabla_v L^\top D_m b \odot \phi(v) + \nabla_\alpha L^\top D_s \alpha, \quad (20)$$

with $D_m = \text{diag}(J_{m,i}^{-1})$ and $D_s = \tau_s^{-1} I$. The scalar δ_h is the *barrier drift*, distinct from the wrench-rate drift d of (7). Both gradients come from Lemma 2 in closed form, so (20) costs one factorization of D per step.

Assumption 2. $a_{\text{CBF}}(x) \neq 0$, equivalently $(\nabla_v L, \nabla_\alpha L) \neq (0, 0)$, on the operating set.

Under Assumption 2 the physical commands appear in h 's first derivative with a nonvanishing coefficient, so the barrier has relative degree one and high-order constructions [23]—with the added conservatism they impose—are not required.

Assumption 2 cannot be strengthened to a sign or magnitude condition, since the coefficient multiplying τ_i is $(\partial L / \partial v_i) / J_{m,i}$, zero wherever $\nabla_v L$ is—exactly the configuration of Proposition 3, where the motor block of a_{CBF} vanishes and the servo block carries the entire row. We therefore monitor Assumption 2 online rather than assert it globally.

B. The unified physical-command quadratic program

Let w_{des} be the wrench requested by the SE(3) outer loop and $w = A(\alpha)\phi(v)$ the wrench currently produced; the target wrench rate

$$\dot{w}_{\text{tar}} = K_w (w_{\text{des}} - w), \quad K_w > 0, \quad (21)$$




Figure 4: (a) The certified set $\{h \geq 0\}$ in the actuator state $x = (v, \alpha)$, with outward gradient ∇h near the boundary. (b) The same requirement in physical-command space: the barrier row (23) cuts a half-space from the actuator box $\mathcal{U}$, and u^* projects the unconstrained minimizer of (22) onto it. The margin ρ_{CBF} of (26) is positive exactly when the row is enforceable.

is a proportional feedback evaluated at every control instant.

Definition 1 (UPC-QP). *With M, d from (6)–(7), weights $W \succ 0$ and $R \succ 0$, trim $u_{\text{ref}} = [b \odot \phi(v)^\top, \alpha^\top]^\top$ and barrier gain $\chi > 0$,*

$$u^* = \arg \min_u \frac{1}{2} \|Mu + d - \dot{w}_{\text{tar}}\|_W^2 + \frac{1}{2} \|u - u_{\text{ref}}\|_R^2 \quad (22)$$

$$\text{s.t. } a_{\text{CBF}}^\top u \geq -\chi h + \delta_h, \quad (23)$$

$$|\tau| \leq \bar{\tau}, \quad (24)$$

$$\alpha_c \in [\alpha - \tau_s \bar{u}, \alpha + \tau_s \bar{u}] \cap [\underline{\alpha}, \bar{\alpha}]. \quad (25)$$

Two features: (23) is $\dot{h} \geq -\chi h$ rewritten through (19), a single barrier inequality (Figure 4b); and (25) encodes the servo rate limit *exactly* rather than as a penalty, so the returned command is always executable.

Proposition 8 (Well-posedness and redundancy resolution). *The objective (22) is strictly convex, the box is nonempty compact, and whenever (23) is compatible with that box the minimizer u^* is unique.*

Proof. The Hessian of (22) is $H = M^\top W M + R$. Since $M \in \mathbb{R}^{m \times 2n}$ with $2n > m$, the tracking term $M^\top W M$ is singular: it is positive semidefinite with a kernel of dimension at least $2n - m$. Adding $R \succ 0$ gives $z^\top H z \geq z^\top R z > 0$ for all $z \neq 0$, so $H \succ 0$ and the objective is strictly convex. The box is nonempty because the current α satisfies (25) by construction and $\tau = 0$ satisfies (24), and it is compact as a finite product of closed bounded intervals. Intersecting with the half-space (23) preserves convexity and compactness, and a strictly convex function attains a unique minimum on a nonempty compact convex set. $\square$

This also resolves the redundancy without an explicit null-space parameterization: on the $(2n - m)$ -dimensional kernel of $M^\top W M$ the tracking term is blind, and the R -weighted term selects the command closest to trim u_{ref} . When w_{des} is not instantaneously reachable, the solver degrades tracking optimally in the metric W while (23)–(25) continue to hold exactly—the certificate is never traded against the task.

C. Feasibility margin and conditional invariance

The guarantee below is conditional on (23) being satisfiable within the box, a condition that is checkable in closed form.

Definition 2 (Feasibility margin). *With $\mathcal{U}$ the box (24)–(25),*

$$\rho_{\text{CBF}}(x) = \max_{u \in \mathcal{U}} a_{\text{CBF}}^\top u + \chi h - \delta_h. \quad (26)$$

Because $a_{\text{CBF}}^\top u$ is linear and $\mathcal{U}$ is a Cartesian product of scalar intervals, the maximizer is the vertex aligned with $\text{sign}(a_{\text{CBF}})$ and (26) evaluates in $\mathcal{O}(n)$ operations, available *only* because the admissible command set is a box; coupled actuator constraints would need a linear program instead.

Geometrically, ρ_{CBF} is the signed distance by which the best available command clears the barrier row (Figure 4b), evaluated at every instant; a positive value certifies the row is enforceable at that state. On $\rho_{\text{CBF}} \leq 0$ the controller terminates and reports—(23) is never relaxed into a soft penalty, since a silently dropped constraint would void the certificate precisely when it binds.

Theorem 9 (Conditional forward invariance). *Let h be continuously differentiable on an open set containing the closed-loop trajectory, let Assumption 2 hold along it, and suppose $\rho_{\text{CBF}}(x(t)) \geq 0$ on $[0, T]$ with a locally Lipschitz solution map $x \mapsto u^*(x)$. If $h(x(0)) \geq 0$ then $h(x(t)) \geq 0$ for all $t \in [0, T]$.*

Proof. Feasibility means the applied u^* satisfies (23), which by (19) reads $\dot{h} \geq -\chi h$ pointwise along the closed-loop vector field. Local Lipschitz continuity of u^* gives existence and uniqueness of the closed-loop solution on $[0, T]$. Comparing h with $y(t) = h(x(0))e^{-\chi t}$ (solving $\dot{y} = -\chi y$) via the comparison lemma yields $h(x(t)) \geq h(x(0))e^{-\chi t} \geq 0$ iff $h(x(0)) \geq 0$. $\square$

Theorem 9 states the trajectory never leaves the certified set (Figure 4a). Its hypotheses are monitored, not established once. Feasibility is certified online through ρ_{CBF} and Assumption 2 is checked each step; we do not establish the Lipschitz property of the solution map, and strict convexity gives uniqueness of u^* but not continuity in x . A stronger, unconditional statement would require proving feasibility over the whole operating envelope, which we have not done.

VI. SIMULATION STUDY

Validation uses a reconstructed actively articulated octorotor ($n = 8, m = 6$). Structural parameters, rotor placement, and the wrench map matrix derive from the Omnidirectional Octorotor architecture [21], augmented here with servo-actuated tilt; the physical parameters (Table I) combine literature measurements with synthesized capacity bounds. The torque limit $\bar{\tau}$ is set so $v^{\text{sat}} = \sqrt{\bar{\tau}/b}$ coincides with the manufacturer’s maximum rotor speed, coupling the capacity model to empirical actuator envelopes. The operating point $v^* = v^{\text{sat}}/\sqrt{3} = 577.4 \text{ rad s}^{-1}$ follows from (9), and the readiness floor is $\ell_{\text{min}} = L(v^*, \alpha_0) - 2 = 26.098$ nats. The lumped drag coefficient b is identified with the static propeller torque coefficient c_T ; the resulting non-dimensional coefficients ($C_T = 0.101, C_Q = 0.0071, C_P = 0.044$ at $\rho = 1.225 \text{ kg m}^{-3}, D = 0.229 \text{ m}$) fall within the UIUC Propeller Database’s [24] thin-electric spectrum.

Table I: Platform parameters.

Symbol	Value	Unit	Symbol	Value	Unit
n, m	8, 6	—	J_m	5×10^{-5}	kg m^2
arm length	0.246	m	τ_s	0.05	s
mass	2.0	kg	b	$= c_\tau$	—
inertia	diag(0.0217, 0.0217, 0.04)	kg m^2	$\bar{\tau}$	0.137	N m
c_f	8.59×10^{-6}	N s^2	v^{sat}	1000	rad s^{-1}
c_τ	1.37×10^{-7}	N m s^2	v^*	577.4	rad s^{-1}
$\bar{u}$	180–462	$^\circ \text{s}^{-1}$	$\ell_{\min}$	26.098	nats
$\alpha, \bar{\alpha}$	± 30	$^\circ$			

Table II: Closed-loop metrics ($\bar{u} = 276^\circ \text{s}^{-1}$). $h_{\min} < 0$ denotes strict violation of the authority floor.

Task	Allocator	RMS _p [m]	$h_{\min}$	sat. [%]	tilt [°]	RMS _w
Hover	All	0.0000	+1.940	0.00	0.00	0.000
Aggressive	Static P-Inv	1.0263	+1.681	0.00	0.00	0.011
	Fixed-Tilt DAAM	1.0250	+1.567	0.00	0.00	0.075
	Uncertified DAAM	1.0250	+1.625	0.00	0.08	0.075
	Robust CBF Filter	1.0250	+1.625	0.00	0.08	0.075
Mild gust	Static P-Inv	1.0738	−0.440	0.00	0.00	0.016
	Fixed-Tilt DAAM	1.0737	−0.749	0.00	0.00	0.101
	Uncertified DAAM	1.0737	−0.156	0.00	0.23	0.101
	Robust CBF Filter	1.0737	+0.251	0.00	1.06	0.101
Strong gust	Static P-Inv	diverged at $t = 6.005 \text{ s}$				
	Fixed-Tilt DAAM	1.5461	−63.25	20.99	0.00	1.320
	Uncertified DAAM	1.4444	−15.30	0.18	11.80	0.214
	Robust CBF Filter	1.4444	+0.012	0.21	14.98	0.212



To isolate the contributions of structural articulation and the proposed CBF, four allocation schemes are evaluated under an identical SE(3) tracking loop:

- 1) **Static Pseudo-Inverse:** standard minimum-norm baseline [4] $\phi_{\text{des}} = A^+ w_{\text{des}}$, fixed tilt $\alpha_{0,i} = \pm 15^\circ$, ex-post actuator clipping.
- 2) **Fixed-Tilt DAAM:** Definition 1 with static tilt and the barrier constraint (23) omitted.
- 3) **Uncertified Articulated DAAM:** Definition 1 with active tilt, but without the barrier constraint (23).
- 4) **Robust CBF Filter:** Definition 1, incorporating both active articulation and the barrier constraint (23).

System dynamics are integrated explicitly at 5 ms over a 10 s horizon across five environments: a benign hover, a 0.5 m position step, an aggressive lateral maneuver, and that maneuver under mild (3 N) and strong (8 N) raised-cosine lateral gusts. The complete pipeline—plant models, allocator parameters, seeds, and data-generation scripts—is at <https://anonymous.4open.science/r/tilt-certified-readiness-418A>; the allocator is implemented in Python, with (22) solved as a bounded least-squares program.

The gusts enter as external additive forces on the airframe; rotor coefficients c_f, c_τ and the allocation map $A(\alpha)$ remain at static-hover values throughout, isolating the allocation problem from the fluid dynamics of high-advance-ratio crosswinds. Modeling the inflow-dependent degradation a true 8 N gust would induce is left to future work.

Authority preservation and closed-loop tracking. When authority is abundant (hover, aggressive), Table II shows all allocators tracking nearly identically and the barrier strictly inactive (0 of 2000 instants, $h \geq 1.567$): the CBF costs nothing outside the critical boundary. The Static Pseudo-Inverse is most accurate in wrench tracking (0.011 vs. 0.075 Root Mean Square (RMS)) since it inverts the static map directly rather than through the rate-based DAAM formulation.

Figure 5: Readiness trajectories under the strong gust: Static Pseudo-Inverse (green), Fixed-Tilt DAAM (red), Uncertified Articulated DAAM (orange), Robust CBF Filter (blue). Readiness floor $h = 0$ (dashed) and gust onset at $t = 2 \text{ s}$ (dash-dotted), both gray.Figure 6: Physical actuator commands under the strong gust, Robust CBF Filter (disturbance onset $t = 2 \text{ s}$, dashed gray). Top: rotor speeds, avoiding saturation and zero-crossings. Bottom: tilt angles, using the mechanical range to absorb the commanded wrench.

Under the mild gust, tracking stays uniform across allocators (isolating readiness as the only quantity that moves) while every uncertified allocator breaches the floor; only the Robust CBF Filter holds it, at 1.06° of tilt. Under the strong gust the Static Pseudo-Inverse diverges at $t = 6.005 \text{ s}$, having reached 49.2% saturation and $h_{\min} = -54.53$. Active articulation alone (Uncertified DAAM) recovers most of the actuator-usage benefit relative to fixed geometry—saturation and wrench error both fall by roughly an order of magnitude—yet still violates the floor by 15.30 nats: articulation is necessary but not sufficient. Re-solving with and without (23) at every instant shows the barrier row active on only 5.8% of steps (median command displacement 11.4%): it converts the Uncertified DAAM’s recovered capacity into a certified floor, bounding the state at the envelope’s edge ($h_{\min} = +0.012$) without added saturation (Figure 5).

The commands issued are physically executable (Figure 6): rotor speeds stay within $[0.20, 0.89] v^{\text{sat}}$, reaching neither saturation nor a zero-crossing—the two loci at which the certificate collapses—and tilt stays inside $\pm 30^\circ$, touching the limit only transiently. The vehicle rejects the disturbance through articulation, as Section IV predicts, rather than forcing motors past capacity.

Robustness to parametric uncertainty. We first perturb the plant parameters while leaving the model class intact. On the strong-gust trajectory, mass ($\pm 10\%$), torque limit ($\pm 15\%$), and thrust coefficient ($\pm 10\%$) are randomized and the servo-rate limit is drawn from $\{180, 276, 318, 462\}^\circ \text{s}^{-1}$, over eight trials with shared initial states and seeds across allocators.

The Robust CBF Filter holds $h_{\min} > 0$ in all eight (median $+0.018$, worst $+0.004$) against 0/8 for the Fixed-Tilt DAAM (median -58.24), at lower median RMS position error (1.565 vs. 1.689 m). Since ρ_{CBF} stays positive throughout (minimum 81.85), Theorem 9’s hypotheses hold, and the $+0.004$ worst case shows the filter tracking the true capacity boundary rather than resting on slack.

Robustness to unmodeled aerodynamics. We next introduce a *structural* mismatch. The allocator is unchanged—barrier, gradients, and allocation follow the static thrust map $A(\alpha)\phi(v)$ of (4)—but the integrator applies a higher-fidelity wrench it never sees, computed from the air-relative velocity $v_{\text{air}} = v_{\text{body}} - R^{\top} v_{\text{wind}}$: advance-ratio thrust loss, an in-plane rotor H -force with induced torque, and quadratic fuselage drag. The Robust CBF Filter still holds the floor ($h_{\min} = +0.012$), articulating to absorb the discrepancy as bounded wrench-tracking error via the W -weighted cost of (22) while the barrier stays hard. The uncertified allocators fail more severely than before: Fixed-Tilt DAAM collapses to $h_{\min} = -63.25$ at 21.0% saturation, the Static Pseudo-Inverse diverges at $t = 6.01$ s. Surviving a plant it does not model, the certificate enforces *structural*, not model-matched, safety.

VII. CRITICAL DISCUSSION AND FUTURE WORK

The readiness level is a log-volume, which can be large while the reachable set is thin in one direction, so $h \geq 0$ does not imply authority is available in whatever direction a disturbance demands. To test this we defined, for analysis only, the worst-direction authority

$$\eta(v, \alpha) = \min_{\|y\|=1} \max \{y^{\top} \dot{w} : \dot{w} \in J_v(v, \alpha) \mathcal{B}_{\bar{a}}\}, \quad (27)$$

with $\mathcal{B}_{\bar{a}} = \{\xi : |\xi| \leq \bar{a}\}$ the motor-rate box: the smallest, over unit wrench-rate directions, of the largest response that box can deliver in that direction, one linear program per sampled direction. Over 40 states sampled along the gust trajectories with 60 directions each, states with $h < 0$ have median $\eta = 2.74$ against 31.25 for $h \geq 0$ —an order of magnitude, precisely the collapse the floor exists to prevent. Above the floor, however, correlation between h and η is $+0.04$, and between h and the median-direction response -0.34 . Consequently, $\ln \det D$ is used strictly as a threshold, avoiding a directional linear program in the control loop. Restricting the certificate to motor authority also introduces measurable conservatism: at the operating point v^* , the discarded margin $L_{\text{art}} - L$ is 4.20, 6.38, 7.19, and 9.49 nats at the four commercial servo rates of Table I (180 – 462° s^{-1})—excluding servo capacity to preserve the saturation divergence (Proposition 6) leaves this margin uncredited, particularly for high-speed actuators—a trade-off that preserves single-inequality enforceability, leaving a unified motor-servo readiness barrier an open problem.

Theoretical boundaries and algorithmic scope. Theorem 9’s guarantee rests on three conditions: continuous online feasibility of the allocation quadratic program, monitored via ρ_{CBF} ; a locally Lipschitz solution map, though strict convexity gives only a unique u^* , not continuity in the state, since the Lipschitz property can fail at active-set transitions; and the degeneracy of Proposition 3, proven for symmetric and near-symmetric

configurations only, leaving its persistence under manufacturing asymmetry untested. Empirically, the filter bounds the state precisely at the envelope’s edge ($h_{\min} = +0.012$), confirming the barrier formulation—not mere availability of articulation—preserves the vehicle under severe disturbance.

Methodological trade-offs and future directions. Integrating a rate-based CBF into the allocation framework trades a marginal loss of wrench-tracking precision in benign flight (Section VI) for strict authority preservation under disturbances that make classical, uncertified allocators diverge; articulation supplies the physical authority, the barrier enforces its correct expenditure. These conclusions are bounded by the reconstructed plant’s modeling assumptions: static drag approximations, an estimated motor inertia J_m , and unmodeled aeroelastic and servo-loading effects introduce parametric uncertainty in absolute readiness values, though the underlying geometric and thresholding properties remain invariant. Future work: hardware validation, broader robustness testing, and a less conservative barrier.

REFERENCES

- [1] M. Hamandi *et al.*, “Design of Multi-Rotor Aerial Vehicles: A Taxonomy Based on Input Allocation,” *Int. J. Robot. Res.*, vol. 40, no. 9, pp. 1015–1044, 2021.
- [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] A. Ollero *et al.*, “Past, Present, and Future of Aerial Robotic Manipulators,” *IEEE Trans. Robot.*, vol. 38, no. 1, pp. 626–645, 2022.
- [4] T. A. Johansen *et al.*, “Control allocation—A survey,” *Automatica*, vol. 49, no. 5, pp. 1087–1103, 2013.
- [5] 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.
- [6] A. Franchi, “Aero-Promptness: Drag-Aware Aerodynamic Manipulability for Propeller-driven Vehicles,” *arXiv:2603.07998*, 2026.
- [7] —, “Variable Aerodynamic Damping Actuation via Co-Contraction: A Structural Analogy with Variable Stiffness Actuation,” *IEEE Robot. Autom. Lett.*, vol. 11, no. 10, pp. 11 362–11 369, 2026.
- [8] V. Kurtz *et al.*, “Control Barrier Functions for Singularity Avoidance in Passivity-Based Manipulator Control,” in *IEEE Conf. Decis. Control*, 2021, pp. 6125–6130.
- [9] Z. Zhang *et al.*, “Constrained Passive Interaction Control: Leveraging Passivity and Safety for Robot Manipulators,” in *IEEE Int. Conf. Robot. Autom.*, 2024, pp. 13 418–13 424.
- [10] G. Silano, “Readiness Barrier Functions: Forward-Invariant Control Authority for Overactuated Multirotor Allocation,” *arXiv:2608.16335*, 2026.
- [11] M. Allenspach *et al.*, “Design and Optimal Control of a Tiltrotor Micro-Aerial Vehicle for Efficient Omnidirectional Flight,” *Int. J. Robot. Res.*, vol. 39, no. 10–11, pp. 1305–1325, 2020.
- [12] O. Härkegård, “Dynamic Control Allocation Using Constrained Quadratic Programming,” *J. Guid. Control Dyn.*, vol. 27, no. 6, pp. 1028–1034, 2004.
- [13] B. Convens *et al.*, “Control of Fully Actuated Unmanned Aerial Vehicles with Actuator Saturation,” *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 12 715–12 720, 2017.
- [14] M. Hamandi *et al.*, “Full-Pose Trajectory Tracking of Overactuated Multi-Rotor Aerial Vehicles With Limited Actuation Abilities,” *IEEE Robot. Autom. Lett.*, vol. 8, no. 8, pp. 4951–4958, 2023.
- [15] A. Franchi *et al.*, “Full-Pose Tracking Control for Aerial Robotic Systems With Laterally Bounded Input Force,” *IEEE Trans. Robot.*, vol. 34, no. 2, pp. 534–541, 2018.
- [16] Y. Su *et al.*, “Nullspace-Based Control Allocation of Overactuated UAV Platforms,” *IEEE Robot. Autom. Lett.*, vol. 6, no. 4, pp. 8094–8101, 2021.
- [17] R. Pretto *et al.*, “Receding-Horizon Nullspace Optimization for Actuation-Aware Control Allocation in Omnidirectional UAVs,” in *Int. Conf. Unmanned Aircr. Syst.*, 2026, pp. 80–87.
- [18] T. Yoshikawa, “Manipulability of Robotic Mechanisms,” *Int. J. Robot. Res.*, vol. 4, no. 2, pp. 3–9, 1985.

- [19] A. D. Ames *et al.*, “Control Barrier Functions: Theory and Applications,” in *Eur. Control Conf.*, 2019, pp. 3420–3431.
- [20] K. Forghani *et al.*, “Task-Space Singularity Avoidance for Control Affine Systems Using Control Barrier Functions,” *arXiv:2603.23753*, 2026.
- [21] M. Hamandi *et al.*, “An Omnidirectional Non-Tethered Aerial Prototype with Fixed Uni-Directional Thrusters,” in *IEEE Int. Conf. Robot. Autom.*, 2025, pp. 8649–8655.
- [22] M. W. Mueller *et al.*, “Stability and control of a quadcopter despite the complete loss of one, two, or three propellers,” *IEEE Int. Conf. Robot. Autom.*, pp. 45–52, 2014.
- [23] J. Breeden *et al.*, “High Relative Degree Control Barrier Functions Under Input Constraints,” in *IEEE Conf. Decis. Control*, 2021, pp. 6119–6124.
- [24] J. B. Brandt *et al.*, “UIUC propeller database, volumes 1–4,” University of Illinois at Urbana-Champaign, <https://m-selig.ae.illinois.edu/props/propDB.html>, 2022.