Toward a Fault-Tolerant Control Allocation Evaluation Framework for eVTOL Aircraft

Garrett D. Asper* and Craig A. Woolsey†

Virginia Tech, Blacksburg, VA 24061

April 22, 2025

Abstract

Control allocation for electric vertical takeoff and landing (eVTOL) aircraft presents unique challenges, including increased actuator redundancy, evolving control effectiveness throughout the transition flight envelope, and a higher potential for control effector faults. This paper presents a fault-tolerant control allocation evaluation framework grounded in a geometric interpretation of aircraft capability. The framework leverages the attainable force and moment set, a convex representation of achievable six degree-of-freedom virtual control commands, to assess algorithm performance across the hover, transition, and cruise flight regimes. A collection of classical and optimization-based linear control allocation methods is evaluated using a high-fidelity model of the NASA Lift Plus Cruise vehicle. Algorithms are assessed by their ability to fill the attainable force and moment set volume under nominal and fault conditions, with additional metrics to evaluate control effort use and computational demand. Results reveal that the sign-preserving, direct allocation, and mixed L_1 optimization approaches consistently reproduce the set of attainable forces and moments throughout the transition envelope while exhibiting robust fault-reconfiguration capability. The evaluation process and findings establish the proposed fault-tolerant control allocation evaluation framework as a scalable, quantitative analysis tool for selecting control allocation algorithms suitable for real-time implementation in an overactuated eVTOL flight control system.

1 Introduction

The emergence of distributed electric propulsion (DEP) technology has introduced a variety of control effector placements and configurations for eVTOL aircraft. eVTOL aircraft experience highly nonlinear aeropropulsive interactions in flight when using control effectors to generate desired forces and moments. To characterize complicated effector interactions in fixed-wing aircraft, the attainable moment set (AMS) was introduced in [1, 2] as a convex hull of feasible moments that an aircraft can produce within actuator saturation constraints about a particular, linearized trim flight condition. eVTOL aircraft introduce novel complexities to the control allocation problem due to their ability to reconfigure control effector use and, in some cases, orientation based on the current flight condition–hover, transition, or cruise–within the broader transition envelope [3]. As a result, attainable moments alone do not capture the full performance of eVTOL aircraft, necessitating the inclusion of three degree-of-freedom (3DOF) forces in addition to moments when characterizing attainable virtual control commands. This paper introduces a complete 6DOF formulation that includes 3DOF force and 3DOF moment coupling called the attainable force and moment set (AFMS), later decoupled into longitudinal and lateral-directional 3DOF AFMS.

Control allocation is the process of determining inputs to control effectors, including motors and surfaces, to meet a desired virtual control command. With a highly redundant, overactuated system, there are potentially infinite feasible control input solutions that accurately generate the desired virtual control. The control allocation problem solves the underdetermined system of equations to calculate the best control solution. However, the control solution must also be physically attainable, that is, within the actuator saturation constraints [1, 2]. A control effector fault is defined as a partial or complete loss of control effectiveness. The presence of a fault shrinks the AFMS and provides further constraints on the control allocation problem. As a result, [4] coined the term "fault-hiding," which is when a control allocation algorithm successfully reconfigures the remaining available control authority to produce the virtual

*Graduate Student, Kevin T. Crofton Department of Aerospace and Ocean Engineering, garrettasper@vt.edu

†Professor, Kevin T. Crofton Department of Aerospace and Ocean Engineering, cwoolsey@vt.edu

control command accurately. Such reconfiguration in the allocation yields the nominal input-output response of the plant, allowing the baseline control algorithm to maintain the original performance of the system within the AFMS without relying on active fault-tolerant control methods [5].

Linear approaches to control allocation remain valid for small perturbations about a trim condition. The Moore-Penrose pseudoinverse approach [6] provides a minimum-norm solution but fails to respect actuator saturation constraints. The weighted pseudoinverse approach allows the definition of a weighting matrix that defines actuator use prioritization but also fails to account for saturation limits directly [7]. The redistributed pseudoinverse or cascaded generalized inverse approach addresses actuator saturation directly by reshaping control usage when an effector approaches its limits [7, 8]. The pseudoinverse, weighted pseudoinverse, and redistributed pseudoinverse all use a least-squares (LS) formulation. The direct allocation approach proposed searching the AMS for a feasible control solution [7]. However, the original formulation exceeded real-time computation capabilities. Breakthroughs in convex optimization, including the development of tools to implement linear programming (LP) [10] and quadratic programming (QP) [12] optimization problems to become feasible for real-time implementation. Bodson presented the direct allocation and mixed L_1 optimization algorithms in standard LP form to directly compare their performance on the Lockheed innovative control effectors (ICE) fixed-wing aircraft model [9]. Finally, solvers for QP problems enabled the development and demonstration of the sign-preserving and direction-preserving formulations evaluated during the X-33 program [11].

This paper builds on established methods for control allocation techniques and performance analysis to present a novel evaluation framework applicable to eVTOL aircraft. While prior work has explored the attainable moment set as a tool for visualizing control limits about a trim condition, this paper introduces the attainable force and moment set to analyze the complete 6DOF actuation capabilities of eVTOL configurations. The AFMS framework accounts for both force and moment generation in a coupled formulation and is subsequently decoupled into longitudinal and lateral-directional axes as is common practice in eVTOL linear model development [3] to enable three-dimensional visualization and analysis. These AFMS provide unique insight into the evolution of control authority across the hover, transition, and cruise flight regimes, known as the transition envelope, where actuator effectiveness varies significantly. This report provides an eVTOL-focused study of control allocation algorithms and their ability to optimize and reconfigure control throughout the entire transition envelope in the presence of control effector saturation and fault constraints. By computing the volume of the virtual control commands that each algorithm can feasibly reproduce within the AFMS, the framework provides a quantitative benchmark for algorithm performance in nominal and fault conditions. To the authors' knowledge, this is the first study to apply AFMS-based evaluation to directly assess how well control allocation algorithms reconfigure and hide faults across the entire eVTOL transition envelope—offering a scalable method for determining algorithm suitability for real-time implementation and integration with flight control systems. The proposed framework establishes a practical and quantitative baseline for determining whether a control allocation algorithm can be considered "fault-tolerant" throughout the entire transition envelope.

2 The Control Allocation Problem

Consider a rigid-bodied, eVTOL aircraft. Let m denote the vehicle mass, $\boldsymbol{I} \in \mathbb{R}^{3\times3}$ the inertia matrix, and g the gravitational acceleration. The rotation matrix mapping from the inertial frame to the body frame is represented by $\boldsymbol{R}_{\text{IB}}$, and the third standard basis is defined as $\boldsymbol{e}_3 = [0\ 0\ 1]^\intercal$. We define $\boldsymbol{v} = [u\ v\ w]^\intercal$ as the translational velocity expressed in the body-frame, and $\boldsymbol{\omega} = [p\ q\ r]^\intercal$ as the rotational velocity expressed in the body-frame. The vehicle is subject to force $\boldsymbol{F} = [X\ Y\ Z]^\intercal$ and moment $\boldsymbol{M} = [L\ M\ N]^\intercal$. The translational and rotational dynamic equations of the vehicle in the body frame are

$$\dot{\boldsymbol{v}} = \boldsymbol{v} \times \boldsymbol{\omega} + g\boldsymbol{R}_{IB}^\intercal \boldsymbol{e}_3 + \frac{1}{m}\boldsymbol{F}$$
$$\dot{\boldsymbol{\omega}} = \boldsymbol{I}^{-1}(\boldsymbol{I}\boldsymbol{\omega} \times \boldsymbol{\omega}) + \boldsymbol{I}^{-1}\boldsymbol{M}$$

Let the mass matrix $\boldsymbol{\mathcal{M}}$ be given as

$$\boldsymbol{\mathcal{M}} = \begin{bmatrix} m\mathbb{I} & \boldsymbol{0} \\ \boldsymbol{0} & \boldsymbol{I} \end{bmatrix}$$

The vector $\boldsymbol{\tau}$ is defined as the virtual control input used for control design, where the control law computes a $\boldsymbol{\tau}_{\text{cmd}}$ to the vehicle from the current state and desired state. In reality, $\boldsymbol{\tau}$ captures the aerodynamics of the system as a function of the vehicle state $\boldsymbol{x} = [\boldsymbol{v}^\intercal\ \boldsymbol{\omega}^\intercal]^\intercal \in \mathbb{R}^6$ and control input $\boldsymbol{u} \in \mathbb{R}^m$ defined as

$$\boldsymbol{\tau} = \boldsymbol{g}(\boldsymbol{x}, \boldsymbol{u}) = \begin{bmatrix} \boldsymbol{F} \\ \boldsymbol{M} \end{bmatrix} \in \mathbb{R}^6$$

The system dynamics can be written as

$$\dot{\boldsymbol{x}} = \boldsymbol{f}(\boldsymbol{x}) + \boldsymbol{\mathcal{M}}^{-1}\boldsymbol{\tau} \tag{1}$$

where the function $\boldsymbol{f}(\boldsymbol{x})$ captures the rigid-body dynamics of the vehicle without direct control inputs. In general, the actuator model $\boldsymbol{\tau} = \boldsymbol{g}(\boldsymbol{x}, \boldsymbol{u})$ is nonlinear. A standard approach to linearizing a nonlinear system is to consider a particular trim state $\boldsymbol{x}_0$ with a corresponding trim control input $\boldsymbol{u}_0$. Let the perturbation from the trim state, control input, and virtual control be defined as $\Delta\boldsymbol{x} = \boldsymbol{x} - \boldsymbol{x}_0$, $\Delta\boldsymbol{u} = \boldsymbol{u} - \boldsymbol{u}_0$, and $\Delta\boldsymbol{\tau} = \boldsymbol{\tau} - \boldsymbol{\tau}_0$ respectively. A Taylor series expansion of the force and moment reveals the components of the linearized virtual control vector as

$$\begin{aligned}\boldsymbol{\tau} = \boldsymbol{g}(\boldsymbol{x}, \boldsymbol{u}) &= \begin{bmatrix} \boldsymbol{F} \\ \boldsymbol{M} \end{bmatrix} \\ &\approx \boldsymbol{g}(\boldsymbol{x}_0, \boldsymbol{u}_0) + \left.\frac{\partial \boldsymbol{g}}{\partial \boldsymbol{u}}\right|_{\boldsymbol{x}_0, \boldsymbol{u}_0} (\boldsymbol{u} - \boldsymbol{u}_0) \\ &\approx \underbrace{\boldsymbol{B}\boldsymbol{u}_0}_{\boldsymbol{\tau}_0} + \underbrace{\boldsymbol{B}_f \Delta\boldsymbol{u}}_{\Delta\boldsymbol{\tau}}\end{aligned}$$

where $\Delta\boldsymbol{\tau} = \boldsymbol{B}_f \Delta\boldsymbol{u}$, $\boldsymbol{B} \in \mathbb{R}^{6\times m}$ is the control effectiveness matrix, and $\boldsymbol{B}_f \in \mathbb{R}^{6\times m}$ is the control effectiveness matrix with fault information.

Consider a flight control algorithm that computes a desired virtual control command represented as

$$\boldsymbol{\tau}_{\text{cmd}} = \boldsymbol{\tau}_0 + \Delta\boldsymbol{\tau}_{\text{cmd}}$$

where $\boldsymbol{\tau}_0$ is the virtual control at the particular trim condition. Using the previously described small-perturbation framework, the control allocation problem can be expressed as

$$\Delta\boldsymbol{\tau}_{\text{cmd}} = \boldsymbol{B}_f \Delta\boldsymbol{u} \tag{2}$$

Because there are typically more actuators than controlled axes, particularly in an eVTOL aircraft, the control solution $\Delta\boldsymbol{u}$ that produces $\Delta\boldsymbol{\tau}_{\text{cmd}}$ is underdetermined. Such a system where $n < m$ has infinitely many control solutions $\Delta\boldsymbol{u}$ among available control effectors to satisfy Eq. (2).

However, constraints arise as each control effector has physical limits. In this paper, the term saturation constraints will be used to refer to the range of control deflections δ a surface can achieve or the range of speeds Ω a motor can achieve. Some formulations of the control allocation problem consider control rate limits. Rate limits define the speed at which an actuator can change its commanded value, $\dot{\delta}$ for control deflection speed, and $\dot{\Omega}$ for motor acceleration. Only control saturation constraints are considered in this report. If the total control saturation constraints are defined by $\boldsymbol{u}_{\min} \leq \boldsymbol{u} \leq \boldsymbol{u}_{\max}$, then one can shift these limits by the trim control $\boldsymbol{u}_0$. The shift yields the control increment saturation constraints $\Delta\boldsymbol{u}_{\min} = \boldsymbol{u}_{\min} - \boldsymbol{u}_0$ and $\Delta\boldsymbol{u}_{\max} = \boldsymbol{u}_{\max} - \boldsymbol{u}_0$. As a result, the control solution is feasible if it satisfies the saturation constraints

$$\Delta\boldsymbol{u}_{\min} \leq \Delta\boldsymbol{u} \leq \Delta\boldsymbol{u}_{\max} \tag{3}$$

Such physical control effector limits presented in Eq. (3) may render Eq. (2) unable to be solved if $\Delta\boldsymbol{\tau}_{\text{cmd}}$ is outside of the physical capabilities of the aircraft.

Additional constraints on the control allocation problem arise in the presence of actuator faults. With each actuator fault, the AFMS shrink, imposing further physical aircraft constraints on the control allocation problem. The control effector health matrix $\boldsymbol{\Lambda} = \text{diag}(\lambda_1, \ldots, \lambda_m)$ allows fault information to inform a control allocation algorithm of new constraints. Each element $\lambda_i \in [0, 1]$ represents the effectiveness, where $\lambda_i = 1$ denotes a fully functional effector and $\lambda_i = 0$ indicates a complete failure. As a result, the control effectiveness matrix with fault information becomes

$$\boldsymbol{B}_f = \boldsymbol{B}\boldsymbol{\Lambda} \tag{4}$$

Algorithms that solve the control allocation problem in Eq. (2) subject to physical actuator constraints in Eq. (3) are considered fault-tolerant control allocation (FTCA) algorithms. Once the control allocation calculates $\Delta\boldsymbol{u}$, it is summed with the trim control input such that the control input $\boldsymbol{u} = \Delta\boldsymbol{u} + \boldsymbol{u}_0$ is commanded to the plant as shown in Figure 1.

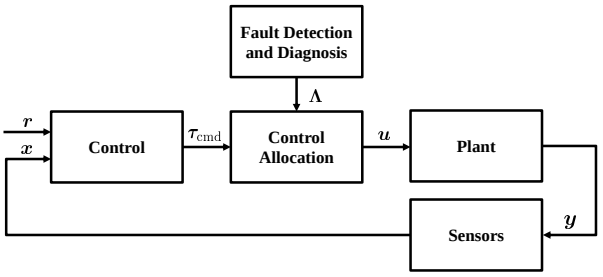

Figure 1: Control system block diagram.

3 Control Allocation Algorithms

Approaches to solving the linear control allocation problem vary considerably. Linear model-based methods include least-squares, quadratic programming, and linear programming problem formulations. Although control allocation methods are comprised of various formulations, there are parallels in the cost functions J that the algorithm seeks to minimize with the solution $\Delta\boldsymbol{u}$. The primary goal of control allocation algorithms is to minimize the error between the commanded virtual control perturbation $\Delta\boldsymbol{\tau}_{\text{cmd}}$ and

the virtual control perturbation realized by the vehicle $\Delta\boldsymbol{\tau}$. The cost function for error minimization is represented by

$$J = \|\Delta\boldsymbol{\tau} - \Delta\boldsymbol{\tau}_{\text{cmd}}\| \tag{5}$$

Additionally, the control effort minimization problem seeks to lower the percentage of available control used and is given by

$$J = \|\Delta\boldsymbol{u} - \Delta\boldsymbol{u}_{\text{pref}}\| \tag{6}$$

where $\Delta\boldsymbol{u}_{\text{pref}}$ is a preferred actuator perturbation. Additional constraints on control effector use can be considered, including prioritizing the use of effectors that are least prone to failure or use the least power. However, these formulations are not considered in this report.

The pseudoinverse and weighted pseudoinverse methods serve as a baseline for comparison, as they are the most commonly used control allocation approaches due to their ease of implementation. These methods use an LS formulation and do not account for control saturation constraints. The redistributed pseudoinverse method also uses an LS formulation and attempts to distribute control effort among healthy actuators within the saturation constraints. The sign-preserving QP approach is the first of three optimization methods evaluated. Direct allocation and mixed L_1 optimization are formulated as a standard LP problem. These three optimization approaches all account for saturation and fault constraints. Table 1 presents the control allocation methods evaluated in this report, including their problem formulation and the primary references used for implementation.

Table 1: Comparison of control allocation approaches and baseline problem formulations.

Method	Problem	Ref.
Pseudoinverse	LS	[6]
Weighted Pseudoinverse	LS	[7]
Redistributed Pseudoinverse	LS	[7, 8]
Sign-Preserving	QP	[11, 12]
Direct Allocation	LP	[7, 9, 10]
Mixed L_1 Optimization	LP	[9, 10]

4 Control Allocation Evaluation

While all control allocation algorithms share the objective of computing a feasible $\Delta\boldsymbol{u}$ that achieves $\Delta\boldsymbol{\tau}_{\text{cmd}}$, their ability to do so robustly, efficiently, and with minimal actuator strain under saturation and fault conditions varies. This section presents a comprehensive benchmark methodology designed to expose these differences. Using the NASA Lift Plus Cruise (LPC) eVTOL platform as a simulation testbed, the performance of each algorithm is evaluated based on geometric control authority coverage of the AFMS, control effort required to track the commanded virtual control inputs, and computational efficiency. The AFMS coverage metric quantifies the percentage of the theoretical control envelope of an aircraft that an algorithm can use, while the control effort and timing metrics provide practical insight into control effector strain and real-time onboard implementation feasibility. Collectively, these metrics enable a rigorous assessment of FTCA algorithm performance across the transition flight envelope.

4.1 eVTOL Platform

The eVTOL aircraft used to evaluate the control allocation algorithms in this report is the NASA LPC vehicle. The authors of [3, 13] provide an overview of the NASA LPC model and the generic urban air mobility (GUAM) simulation environment created in MATLAB® and Simulink®. The NASA LPC vehicle is shown in Figure 2, with control effector labels and a table with corresponding effector names and saturation constraints. Applicable units are given in degrees and RPM for surfaces and motors, respectively.

4.2 Attainable Moment Set

To benchmark control allocation algorithms within realistic operating conditions, we must first define the bounds of physically achievable control commands. This section introduces the geometric construct used to quantify these bounds. Durham proposed the attainable moment set (AMS) to evaluate the performance capabilities of a fixed-wing aircraft [2, 7]. The AMS, represented by Φ, is a three-dimensional convex hull that characterizes the attainable moments a fixed-wing aircraft can physically achieve with feasible control inputs that satisfy Eq. (3). The AMS can be constructed as a perturbation in moment $\Delta\boldsymbol{M}$ from the trim moment $\boldsymbol{M}_0$ about a particular trim flight condition. The AMS can be plotted, and the physical limits of the aircraft can be subsequently determined using a facet search method in which control effectors are swept between their minimum and maximum saturation limits as described in [7].

Since control allocation seeks to determine a feasible $\Delta\boldsymbol{u}$ such that a $\boldsymbol{\tau}_{\text{cmd}}$ is realized by the system, it is necessary to determine the maximum realizable $\boldsymbol{\tau}$ of the system using the AMS. The AMS computation method presented in [7] only considers the ability of an aircraft to achieve three-axis $\Delta\boldsymbol{M}$ from a partic-

Symbol	Control Effector	Min	Max
n_1	Lift Rotor 1	550	1550
n_2	Lift Rotor 2	550	1550
n_3	Lift Rotor 3	550	1550
n_4	Lift Rotor 4	550	1550
n_5	Lift Rotor 5	550	1550
n_6	Lift Rotor 6	550	1550
n_7	Lift Rotor 7	550	1550
n_8	Lift Rotor 8	550	1550
n_9	Pusher Rotor	0	1750
δ_{a_L}	Left Aileron	-30	+30
δ_{a_R}	Right Aileron	-30	+30
δ_{e_L}	Left Elevator	-30	+30
δ_{e_R}	Right Elevator	-30	+30
δ_r	Rudder	-30	+30

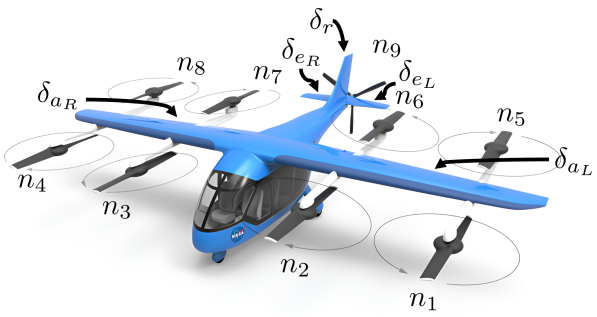

Figure 2: Control effector labels and saturation constraints for the NASA LPC vehicle adapted from [3].

ular trim condition. As demonstrated in [7], control effector faults decrease the volume of the AMS, as fewer moments are able to be physically achieved by the aircraft.

4.3 Attainable Force and Moment Set

The traditional AMS approach developed by Durham [7] is suitable for fixed-wing aircraft as it neglects the attainable forces due to lift L and thrust T along the longitudinal axis by considering them as sufficient to sustain the current trim condition. However, eVTOL aircraft have the ability to change the effective direction of T, particularly throughout the transition flight envelope. For this reason, the complete 6DOF virtual control perturbation $\Delta\boldsymbol{\tau}$ comprised of 3DOF $\Delta\boldsymbol{F}$ and 3DOF $\Delta\boldsymbol{M}$ must be used to generate the AFMS of the vehicle about a particular trim condition. To simplify the 6DOF AFMS for visualization and analysis, the assumption of longitudinal and lateral-directional decoupling is employed to yield two three-dimensional AFMS computed using the facet search approach described in [7]. Table 2 presents the force and moment components of each AFMS as defined throughout this report.

Table 2: AFMS longitudinal and lateral-directional axes force and moment components.

Axis	Forces and Moments
Longitudinal	ΔX, ΔZ, ΔM
Lateral-Directional	ΔY, ΔL, ΔN

The AFMS in the nominal, fault-free flight condition changes as the eVTOL aircraft moves through the transition flight envelope. In the case of the NASA LPC vehicle, both AFMS in the longitudinal and lateral-directional axes are parameterized by forward airspeed V as described in [3]. Such changes in the AFMS in the longitudinal and lateral-directional axes are presented in Figure 3, where three flight conditions are shown to represent hover, mid-transition, and cruise flight in nominal, fault-free conditions.

As faults occur in the aircraft, the AFMS decrease in volume. As with the nominal flight condition, any force and moment combination inside the modified AFMS is feasible and can still be physically achieved by the aircraft. Figure 4 demonstrates how the AFMS changes in the presence of control effector faults. The red convex hull is the reduced AFMS due to a complete fault in lift motor n_1 and left aileron δ_{a_L}, while the exterior hull is the nominal AFMS.

The longitudinal and lateral-direction AFMS have been established as a method to understand the physical limits in generating $\Delta\boldsymbol{\tau}$ throughout the flight envelope in the presence of saturation constraints and control effector faults. Durham observed that an ideal control allocation algorithm can solve for a feasible $\Delta\boldsymbol{u}$ that satisfies Eq. (2) subject to Eq. (3) for each feasible moment in the AMS [7]. For the case of an eVTOL vehicle, an ideal control allocation algorithm can solve for a feasible $\Delta\boldsymbol{u}$ that satisfies Eq. (2) subject to Eq. (3) for the longitudinal and lateral-directional AFMS throughout the transition envelope from hover to cruise flight.

The $\Delta\boldsymbol{\tau}_{\text{cmd}}$ that are feasible for a particular control allocation algorithm form a convex hull Π [7]. Ideally, the AFMS produced by a control allocation algorithm Π matches the AFMS of the vehicle Φ, meaning the algorithm can fully leverage the available control authority. To quantify this performance, the volume ratio is defined as $\text{vol}(\Pi)/\text{vol}(\Phi)$, expressed as a percentage. A volume ratio of 100% indicates that the algorithm successfully reproduces the entire attainable set at a given trim condition using admissible control effector commands.

In creating Π for each control allocation algorithm, the vertices of Φ were commanded as $\Delta\boldsymbol{\tau}_{\text{cmd}}$. If the control allocation approach was able to determine a

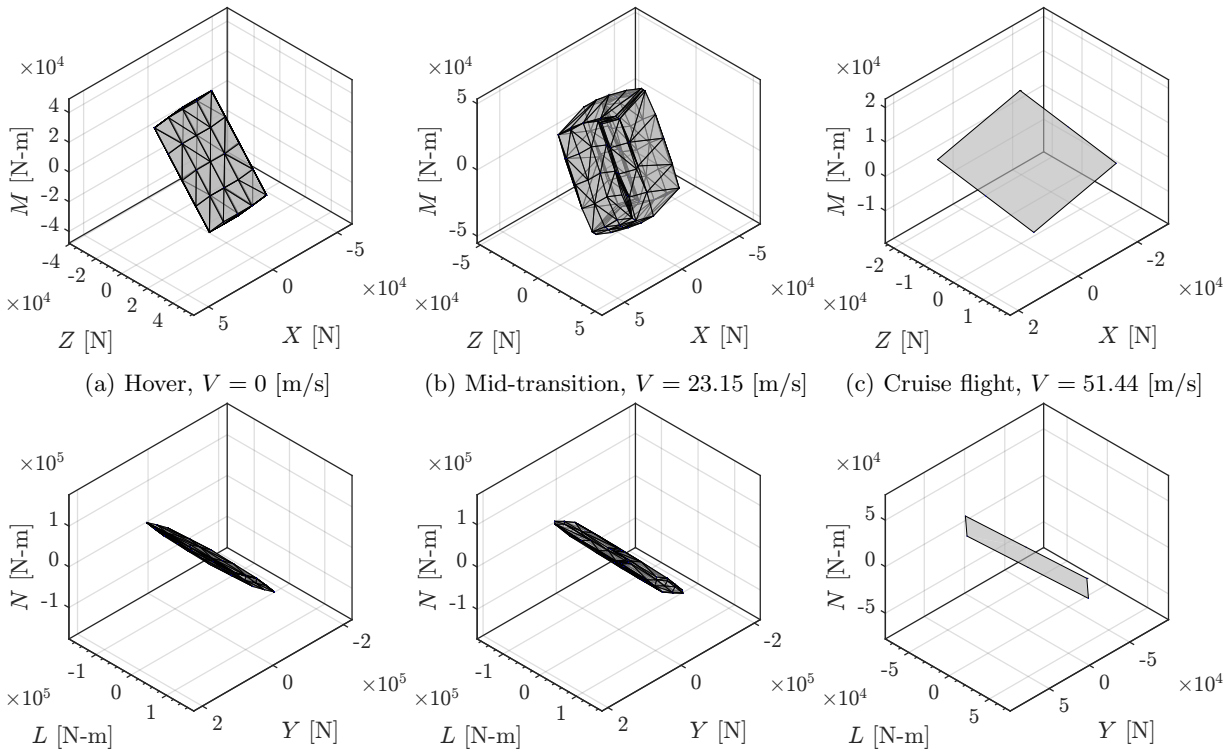

(a) Hover, $V = 0$ [m/s] (b) Mid-transition, $V = 23.15$ [m/s] (c) Cruise flight, $V = 51.44$ [m/s]

(d) Hover, $V = 0$ [m/s] (e) Mid-transition, $V = 23.15$ [m/s] (f) Cruise flight, $V = 51.44$ [m/s]

Figure 3: NASA LPC nominal longitudinal (top row) and lateral-directional (bottom row) AFMS.

solution for the control increment $\Delta \boldsymbol{u}$ within the saturation limits, and the resultant $\Delta \boldsymbol{\tau} \approx \Delta \boldsymbol{\tau}_{\text{cmd}}$, the test point was kept and added as a vertex of Π. If the constraints for any vertex of Φ were violated, the evaluation algorithm added uniformly distributed candidate vertices interior to Φ to asses the feasible volume of Π.

The nominal, fault-free condition was used to determine the volume ratio and compare each control allocation approach. Additionally, several fault scenario benchmarks were used to examine how the volume ratio of each algorithm changed. However, as the volume of Φ decreases, the vertices of the AFMS become less extreme. As a result, satisfactory control allocation algorithms with a volume ratio of 100% in the nominal case continued to do so in the presence of faults. However, the volume ratios of algorithms that underperformed in the nominal case typically improved in the fault scenarios. Consequently, only the largest volume of the AFMS at each flight condition was considered in the control allocation algorithm benchmark, as viable FTCA approaches successfully reconfigured the control increment solution when provided with accurate fault information.

As seen in Figure 5, the performance of each algorithm and its ability to produce feasible control input solutions is plotted throughout the transition envelope. For the remainder of this report, control allocation algorithms that have a volume ratio of 100% throughout the transition flight envelope in both the longitudinal and lateral-directional axes are considered FTCA algorithms.

In the longitudinal axis, the pseudoinverse algorithm performs well in hover, degrades in coverage early into the transition regime, and performs well starting when the lift motor effectiveness is set to zero in the linearized model at $V = 48.87$ [m/s], decreasing the number of candidate effectors. The weighted pseudoinverse similarly performs better in early transition and degrades through transition, improving only when the lift motors are removed. In the lateral-directional axes, the pseudoinverse fails to find feasible solutions until late transition, while the weighted pseudoinverse has limited coverage in hover, none in mid-transition, and improves in late-transition. Due to the poor performance of the pseudoinverse and weighted pseudoinverse approaches, they will not be considered FTCA algorithms. The redistributed pseudoinverse algorithm achieves near-complete coverage of Φ in the longitudinal axis but exhibits mild degradation in lateral-directional performance during early transition. Although the redistributed pseu-

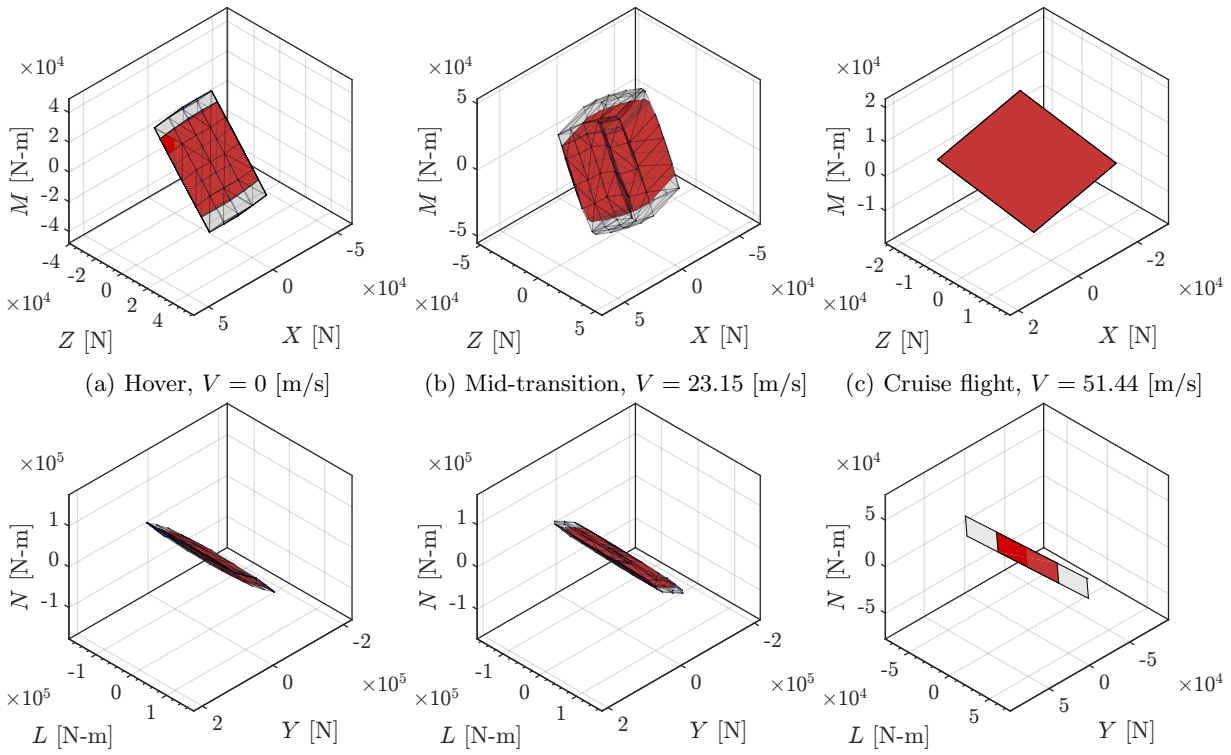

(a) Hover, $V = 0$ [m/s] (b) Mid-transition, $V = 23.15$ [m/s] (c) Cruise flight, $V = 51.44$ [m/s]

(d) Hover, $V = 0$ [m/s] (e) Mid-transition, $V = 23.15$ [m/s] (f) Cruise flight, $V = 51.44$ [m/s]

Figure 4: NASA LPC longitudinal (top row) and lateral-directional (bottom row) axes with the AFMS of the full loss of effectiveness in n_1 and δ_{a_L} in red, and the nominal AFMS in gray.

doinverse performs well overall, full coverage of Φ is desired, so it will not be considered as an FTCA algorithm. Finally, the sign-preserving, direct allocation, and mixed L_1 optimization approaches yield a volume ratio of 100% and continue to cover the full AFMS of the vehicle throughout the transition envelope in the longitudinal and lateral-directional axes in both nominal and fault conditions. As a result, these three algorithms will be considered FTCA algorithms. The performance results of the pseudoinverse, weighted pseudoinverse, and redistributed pseudoinverse algorithms will be shown throughout the rest of this report for a baseline comparison but will not be discussed as they are not FTCA algorithms.

4.4 Control Effort

Control effort is a valuable secondary evaluation metric for FTCA algorithms. Inefficient use of the available control pushes control effectors closer to their saturation limits. Additionally, unnecessary overuse of actuators may lead to a higher likelihood of faults and a subsequent decrease in the AFMS. In this study, an average of $\Delta \boldsymbol{u}$ solutions across the hull was taken at each flight condition. This average was then normalized by the saturation limits to compute a percent control effort used to achieve $\Delta \boldsymbol{\tau}_{\text{cmd}}$. The normalized control perturbation is defined by

$$\Delta \boldsymbol{u}_{\text{norm}} = \frac{\Delta \overline{\boldsymbol{u}} - \Delta \boldsymbol{u}_{\text{min}}}{\Delta \boldsymbol{u}_{\text{max}} - \Delta \boldsymbol{u}_{\text{min}}}$$

In the nominal, fault-free case, very little difference is noticeable in the control effort used by each algorithm. Only minor differences are noted in late transition where the sign-preserving and direct allocation approaches use slightly more effort than mixed L_1 optimization.

Figure 6 shows that a fault condition can provide additional insight for control effort evaluation. The fully failed effectors in the shown fault condition include n_1, n_4, n_8, and δ_{a_R}. Each FTCA algorithm successfully redistributed $\Delta \boldsymbol{u}$ among the remaining healthy control effectors. However, direct allocation uses significantly less control effort when accounting for the effector failures. As a result, the control effort used by the direct allocation algorithm provides more margin before violating saturation constraints than the sign-preserving or mixed L_1 optimization algorithms.

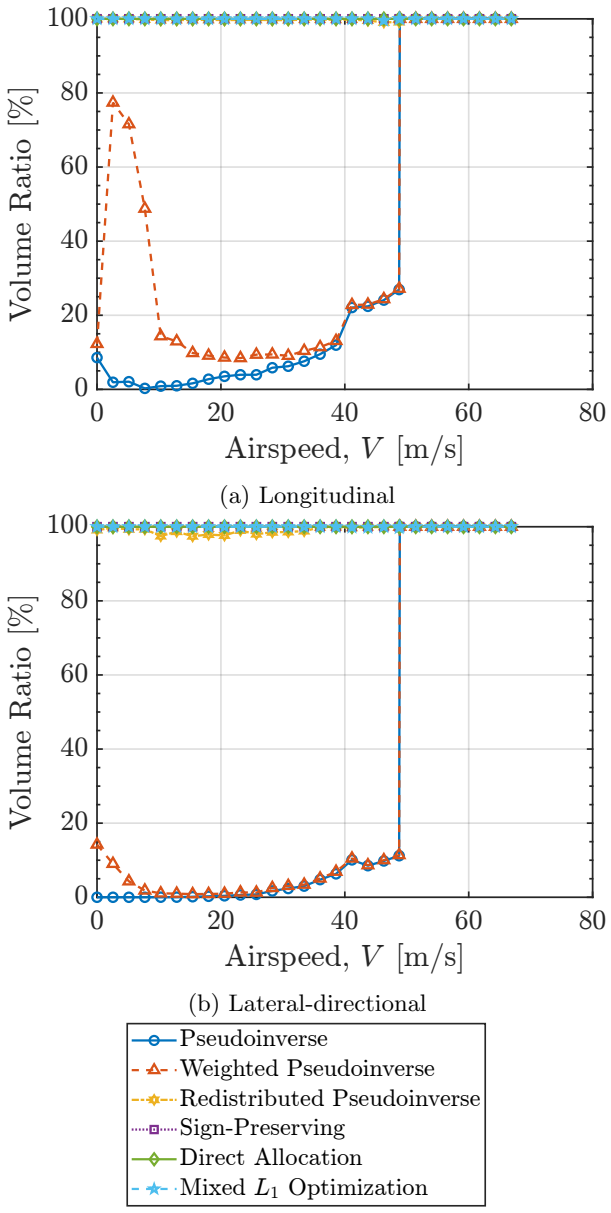

Figure 5: Volume ratio of Φ filled by Π throughout the transition flight envelope.

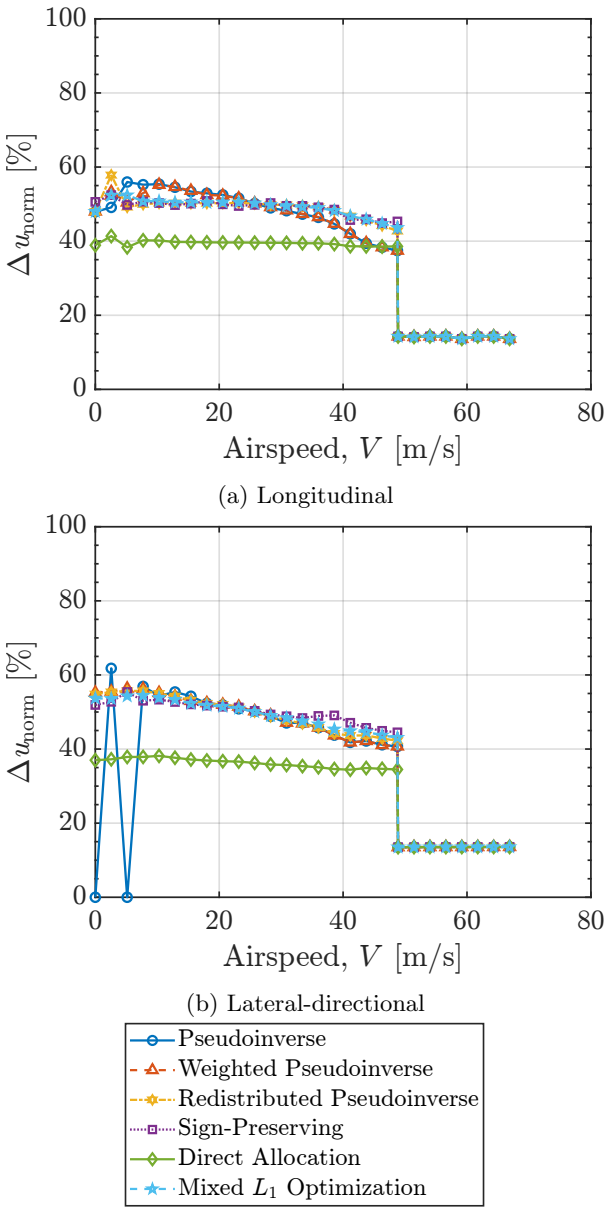

Figure 6: Average control effort use with full loss of effectiveness in n_1, n_4, n_8, and δ_{a_R}.

4.5 Computational Efficiency

The final metric used to evaluate the FTCA algorithms is the average time each approach requires to compute a solution to the control allocation problem. Each FTCA algorithm was coded as an m-file in MATLAB® 2024a and timed using an evaluation script on a mobile workstation with an Intel® Core™ i9-13950HX processor running at 2.20 GHz with 32 GB of RAM. Figure 7 presents the average time in seconds to compute $\Delta \boldsymbol{u}$ for each point in the AFMS plotted by flight condition through the transition envelope. In the nominal, fault-free case, the mixed L_1 optimization algorithm requires significantly more time despite using the same computationally efficient simplex LP solver used by direct allocation. The mixed L_1 optimization formulation contains slack variables for error that increase the dimension of the cost function and, thus, the computational burden of the algorithm. The sign-preserving approach is much more efficient than mixed L_1 optimization but is less efficient than the direct allocation

algorithm. The algorithms show a similar trend when evaluated at the fault condition described in Sec. 4.4, but take slightly less time to solve the control allocation problem as the number of effectors to optimize has decreased.

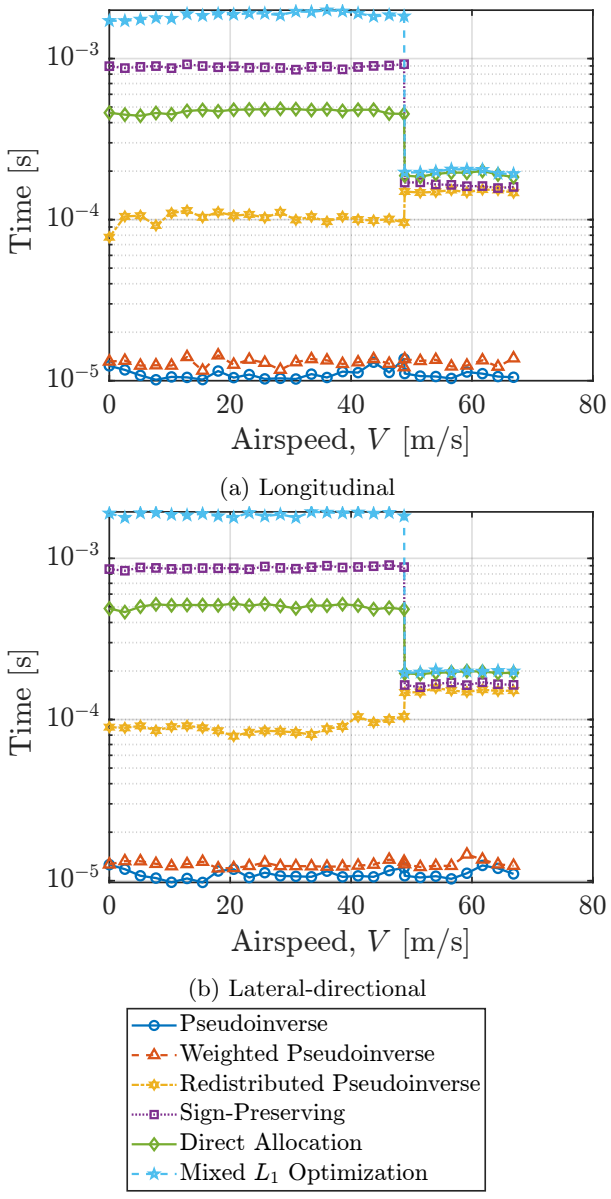

(a) Longitudinal

(b) Lateral-directional

Figure 7: Average time per computation.

5 Conclusions

eVTOL aircraft challenge conventional control allocation approaches due to their actuator redundancy, evolving control effectiveness across the transition envelope, and susceptibility to control effector faults. This paper introduced a novel, scalable evaluation framework grounded in the AFMS to benchmark the performance of control allocation algorithms across the full eVTOL transition envelope. Using a high-fidelity model of the NASA Lift Plus Cruise vehicle, a suite of classical and optimization-based linear allocation strategies was evaluated under nominal and fault conditions. The AFMS-based evaluation framework quantitatively assessed the ability of each algorithm to fully exploit available control authority and reconfigure in the presence of faults, establishing a baseline qualification for a control algorithm to be considered "fault-tolerant."

Among the algorithms tested, the sign-preserving, direct allocation, and mixed L_1 optimization approaches consistently produced feasible solutions throughout the entire AFMS at each flight condition in the transition envelope. These algorithms also demonstrated fault tolerance by maintaining full AFMS coverage in the presence of multiple effector failures. The direct allocation algorithm emerged as the most computationally- and control-efficient method under fault conditions, making it particularly well-suited for real-time implementation.

The proposed AFMS-based evaluation framework establishes a clear, quantitative benchmark for FTCA algorithm performance, offering a practical tool for selecting real-time capable allocation strategies for overactuated eVTOL flight control systems. The authors will extend the AFMS evaluation framework to a closed-loop control architecture implemented in a high-fidelity eVTOL simulation. This effort will quantify transient performance, tracking accuracy, and fault-hiding response latency under time-varying fault conditions across the transition flight envelope.

Acknowledgments

The first author acknowledges funding provided by the Virginia Space Grant Consortium through the graduate research fellowship and the National Science Foundation Graduate Research Fellowship Program.

References

[1] W. C. Durham, "Constrained control allocation," *Journal of Guidance, Control, and Dynamics*, vol. 16, no. 4, pp. 717–725, Jul. 1993. [Online]. Available: https://arc.aiaa.org/doi/10.2514/3.21072

[2] ——, "Constrained control allocation - Three-moment problem," *Journal of Guidance, Control, and Dynamics*, vol. 17, no. 2, pp. 330–336, Mar. 1994. [Online]. Available: https://arc.aiaa.org/doi/10.2514/3.21201

[3] B. M. Simmons, P. G. Buning, and P. C. Murphy, "Full-envelope aero-propulsive model identification for lift+cruise aircraft using computational experiments," in *Proc. AIAA AVIATION 2021 Forum*, July 2021. [Online]. Available: https://arc.aiaa.org/doi/abs/10.2514/6.2021-3170

[4] A. Marks, "A fault hiding approach for fault tolerant control of a class of vtol vehicles," Ph.D. dissertation, Cranfield University, Cranfield, UK, February 2015, [Online]. Available: http://dspace.lib.cranfield.ac.uk/handle/1826/9250.

[5] I. Bessa, V. Puig, and R. M. Palhares, "Reconfiguration blocks and fault hiding: Design, applications, and challenges," *Annual Reviews in Control*, vol. 56, p. 100896, 2023. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/S1367578823000330

[6] J. C. A. Barata and M. S. Hussein, "The Moore–Penrose Pseudoinverse: A Tutorial Review of the Theory," *Brazilian Journal of Physics*, vol. 42, no. 1-2, pp. 146–165, Apr. 2012. [Online]. Available: http://link.springer.com/10.1007/s13538-011-0052-z

[7] W. Durham, K. A. Bordignon, and R. Beck, *Aircraft control allocation*. Chichester, West Sussex, United Kingdom: Wiley, 2017.

[8] T. A. Johansen and T. I. Fossen, "Control allocation—A survey," *Automatica*, vol. 49, no. 5, pp. 1087–1103, May 2013. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0005109813000368

[9] M. Bodson, "Evaluation of Optimization Methods for Control Allocation," *Journal of Guidance, Control, and Dynamics*, vol. 25, no. 4, pp. 703–711, Jul. 2002, publisher: American Institute of Aeronautics and Astronautics. [Online]. Available: https://arc.aiaa.org/doi/10.2514/2.4937

[10] D. G. Luenberger and Y. Ye, *Linear and Nonlinear Programming*, third edition ed., ser. International Series in Operations Research & Management Science. Boston, MA: Springer US, 2008, no. 116.

[11] A. Simmons and A. Hodel, "Control allocation for the X-33 using existing and novel quadratic programming techniques," in *Proceedings of the 2004 American Control Conference*, vol. 2, Jun. 2004, pp. 1701–1706 vol.2, iSSN: 0743-1619. [Online]. Available: https://ieeexplore.ieee.org/document/1386824

[12] T. F. Coleman and Y. Li, "A reflective newton method for minimizing a quadratic function subject to bounds on some of the variables," *SIAM Journal on Optimization*, vol. 6, no. 4, pp. 1040–1058, 1996. [Online]. Available: https://doi.org/10.1137/S1052623494240456

[13] M. J. Acheson, I. M. Gregory, and A. Patterson, "Generic urban air mobility simulation," in *Proc. AIAA SCITECH 2025 Forum*, January 2025. [Online]. Available: https://arc.aiaa.org/doi/abs/10.2514/6.2025-0929