Flight Control Design and Testing of a Tiltwing VTOL Aircraft

Leo Panish

St Hugh's College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Michaelmas 2025

To my grandfather, whose passion for flight and ceaseless ingenuity inspired my own love for aviation and engineering.

And to Wolfgang. You will forever be in our hearts.

Acknowledgments

First and foremost, I would like to give my utmost gratitude to my supervisor and mentor, Prof. Marko Bacic, who has supported my progression and growth as an engineer since my time as an undergrad here at Oxford. He has made this opportunity for me to pursue a DPhil possible, in every respect, and his guidance and support has helped shape who I am today. I am forever grateful for his contribution to my life.

I would also like to thank Dr Chris Nicholls and Dr Brian Tang. Chris contributed countless hours to helping me with flight testing and displayed great skill in holding the tether rope, and the chaos of the first untethered flight test remains a very special memory for me. Brian has always contributed invaluable insight, and I have very much enjoyed our discussions both in the lab and at the pub.

Thanks go to my thesis examiners, Prof. Mark Cannon and Prof. Sophie Armanini, for their valuable contributions and feedback on my work.

I would like to thank the Engineering and Physical Sciences Research Council (EPSRC) and Rolls-Royce plc for funding this research.

A big thanks to Jason McCluskey, who has helped me numerous times as far back as my 4YP, Gerard Shields at PrintPool for last-minute 3D-printing of tiltwing parts, and Steve Peart at the Cirencester & District Model Aircraft Club as well as Alan Smith at the Oxford Model Flying Club for making outdoor space available for me to carry out flight testing.

I would like to thank all of my friends, with special acknowledgment to Niko, Owen, Georgie and Oli, Arno, Boris, Linh and Máté, Joe, Max, and Alex for their support during my time as a DPhil student.

Lastly, I would like to thank my family: My grandmother, for her continuous support in my studies since the beginning of my undergrad. Shejeru, whom I am excited to soon call my sister. A special acknowledgment to my brother, Maxfield, for repeatedly telling me to "just get done already"... In all seriousness, his blunt honesty and unquestionable commitment to my success keeps me grounded, and I am beyond lucky to have him in my life. And the greatest acknowledgment to my parents, who have been the ultimate catalysts to my success. My father, for his wisdom about projects, people, and helping me stay aligned with my goals. And my mother, for her unconditional love and support, and for putting Oxford on my radar all those years ago.

A thesis made possible by...

Image courtesy of RS Components Ltd., reprinted with permission.

Abstract

Transition vertical take-off and landing (VTOL) aircraft configurations have risen in popularity as a result of the emerging "air taxi" market, i.e. the idea of point-to-point hybrid and electric aircraft transport solutions as a means to enable intra-city civilian air transport amidst rising global climate concerns. Such configurations benefit from the characteristics of both helicopters (high maneuverability and vertical flight) and conventional fixed-wing aircraft (efficient cruise) but remain a great challenge in the context of flight stability and control. This thesis explores the modeling, flight controller design, and testing of a unique tiltwing aircraft configuration: the Rolls-Royce eVTOL concept.

Minimum-energy transition trajectories are found that satisfy stall and dynamic constraints to gain an understanding of the aircraft's achievable flight envelope. By studying the effects of center-of-gravity (CG) perturbations in combination with active flow control (AFC) concepts, the theoretical limits of performance are determined, where it is found that constant-altitude transition can be achieved in the forward direction (hover-to-cruise) but not the reverse (cruise-to-hover). Comparative solutions reveal that a redesign of the tail wing could expand the achievable flight envelope of the aircraft during backward transition. In either case, the use of AFC serves to maximize aircraft performance and increase robustness to CG variations.

A full-envelope flight controller that employs feedback linearization in a cascade design combined with fault-tolerant control allocation is developed for a scaled model of this aircraft. The inner-loop attitude control is designed first, where nonlinear dynamic inversion (NDI) is paired with LQR to regulate the attitude rates, and a weighted pseudo-inversion control allocation scheme is implemented to map virtual controls to effector demands. It is found that the inner-loop bandwidth must be limited to avoid exciting a structural resonance mode of the 3D-printed airframe, but controller performance is nevertheless satisfactory in both simulation and hover flight. Good robustness is demonstrated in response to motor failure and wind gust disturbances.

The general problem of finding a full-envelope solution for the outer-loop velocity control design is investigated. A novel solution is proposed through the careful selection of thrust-aligned virtual control inputs and outputs, resulting in a globally-valid controller that is generalizable to a wide range of transition VTOL configurations. This outer-loop design is applied to the tiltwing aircraft of this thesis, and its completed full-envelope flight controller is evaluated in simulation and hover flight, demonstrating proof-of-concept of this novel outer-loop approach. A formal robustness analysis of such a controller design is carried out, elucidating the implications of physical design on the uncertainty bounds.

Contents

1 Introduction **1**
- 1.1 Motivation: Tiltwing VTOL Aircraft and their Challenges 3
- 1.2 Purpose of Research 7
- 1.3 Thesis Content 8

2 Key Concepts and Literature Review **10**
- 2.1 Introduction 10
- 2.2 General Flight Dynamics 11
  - 2.2.1 Coordinate Frames 11
  - 2.2.2 6-DoF Equations for a Rigid Body 13
- 2.3 Feedback Linearization Control 15
  - 2.3.1 Mathematical Concepts 18
    - 2.3.1.1 Lyapunov Stability Theory 18
    - 2.3.1.2 Lie Derivatives and Lie Brackets 21
    - 2.3.1.3 Diffeomorphisms and State Transformations 22
    - 2.3.1.4 The Frobenius Theorem and Involutivity 23
  - 2.3.2 Full-State Linearization 24
  - 2.3.3 Input-Output Linearization 27
- 2.4 Control Allocation for Over-Actuated Systems 30
  - 2.4.1 General Problem Formulation 30
  - 2.4.2 Techniques for Linear Control Effector Models 31
  - 2.4.3 Techniques for Nonlinear Control Effector Models 34
- 2.5 Trajectory Optimization Methods 35
- 2.6 Active Flow Control Applied to Aircraft 39
- 2.7 Uncertainty Modeling and Robustness Analysis 41
  - 2.7.1 Representing Uncertainty 42
  - 2.7.2 Robust Stability for Unstructured Uncertainty 44
  - 2.7.3 The Structured Singular Value 46
  - 2.7.4 H_∞ and μ-synthesis Control 47

3 Tiltwing Aircraft Modeling and Dynamics **49**
- 3.1 Introduction 49
- 3.2 Overview of the Aircraft 50
  - 3.2.1 Assumed Mission Profile 53
- 3.3 Mathematical Modeling 54

3.3.1 Gravity . . . 54
3.3.2 Propeller Forces and Moments . . . 55
3.3.2.1 Thrust Forces . . . 55
3.3.2.2 Thrust Moments . . . 56
3.3.2.3 Reaction Torques . . . 56
3.3.2.4 Prop Thrust and Torque Modeling . . . 56
3.3.2.5 Gyroscopic Moments . . . 58
3.3.3 Aerodynamic Forces and Moments . . . 58
3.3.3.1 Aerodynamic Coefficients and Stability Derivatives . . . 59
3.3.3.2 Propeller-Wing Interactions . . . 65
3.3.3.3 Stability-to-Body Frame Transformation . . . 67
3.4 Model Implementation . . . 68
3.5 Model Validation . . . 69
3.6 Analysis of the Dynamics . . . 74
3.6.1 Cruise . . . 74
3.6.2 Hover . . . 75
3.6.3 Transition . . . 76

4 Transition Trajectory Optimization 79
4.1 Introduction . . . 79
4.2 Simplified Model and Dynamics . . . 83
4.2.1 Three-Degree-of-Freedom Rigid Body Equations . . . 84
4.2.2 Thrust Forces and Moments . . . 85
4.2.3 Aerodynamics . . . 85
4.2.4 Active Flow Control Model . . . 86
4.3 Optimization Methodology and Implementation . . . 87
4.3.1 Optimization Problem Setup . . . 87
4.3.2 Computing a Guess Trajectory . . . 91
4.4 Results . . . 93
4.4.1 Forward Transition . . . 93
4.4.1.1 Without AFC . . . 93
4.4.1.2 With AFC . . . 95
4.4.2 Backward Transition . . . 96
4.4.2.1 Without AFC . . . 96
4.4.2.2 With AFC . . . 97
4.5 Conclusion . . . 105

5 Inner-Loop Design and Testing 107
5.1 Introduction . . . 107
5.2 Full-Envelope Flight Controller Overview . . . 109
5.3 Inner-Loop Design . . . 110
5.3.1 Controller Architecture . . . 110
5.3.2 NDI Attitude Rate Controller . . . 112
5.3.3 Euler Angle Tracking . . . 115
5.3.4 Control Allocation . . . 116

5.3.4.1 Integral Anti-Windup Strategy . . . 118
5.3.4.2 Vertical Take-Off Schedule . . . 120
5.4 Simulations . . . 122
5.4.1 Controller Tuning . . . 122
5.4.2 Results . . . 124
5.5 Flight Testing . . . 128
5.5.1 Experimental Setup . . . 128
5.5.2 Tethered Flight . . . 130
5.5.3 Untethered Flight . . . 135
5.6 Conclusion . . . 139

6 Generalized Full-Envelope Outer-Loop Control Strategies 140
6.1 Introduction . . . 140
6.2 General Problem Setup . . . 148
6.3 Selecting a Generalized Force Input and Compound Velocity Output 150
6.4 Strategy 1 (simplified case) . . . 154
6.4.1 Controller Definition . . . 154
6.4.2 Internal Dynamics . . . 155
6.4.3 Controlling the Internal Dynamics through Pitch Reference Commands . . . 156
6.5 Strategy 1 (general case) . . . 159
6.6 Strategy 2: Formal Design of $V_{c_\perp}$ Using a Pseudo-Control Allocation Scheme . . . 160
6.6.1 Motivation and Problem Setup . . . 160
6.6.2 Allocation Method . . . 163
6.6.3 Discussion . . . 163
6.7 Strategy 3: Coupled MIMO Approach . . . 164
6.7.1 MIMO Pseudo-Control Allocation Motivation . . . 165
6.7.2 Full-Envelope MIMO Design . . . 166
6.7.3 Discussion . . . 168
6.8 Conclusion . . . 171

7 Outer-Loop Design and Completed Flight Controller 173
7.1 Introduction . . . 173
7.2 Outer-Loop Design . . . 174
7.2.1 Decoupled Design (Strategy 1) . . . 174
7.2.2 Coupled MIMO Design (Strategy 3) . . . 177
7.2.3 Control Allocation (revisited) . . . 178
7.2.3.1 Transition Trajectory Schedule . . . 179
7.2.4 Lateral Velocity Controllers . . . 180
7.2.4.1 Hover Lateral Motion . . . 180
7.2.4.2 Cruise Coordinated Turn . . . 181
7.3 Completed Flight Controller Architecture . . . 182
7.4 Simulations . . . 183
7.4.1 Strategy 1 . . . 183

7.4.2 Strategy 3 . . . 187
7.4.3 Lateral Motion . . . 191
7.5 Flight Testing . . . 192
7.5.1 Experimental Setup . . . 192
7.5.2 Tethered Flight . . . 194
7.6 Robustness Analysis . . . 197
7.6.1 Analytical Uncertainty Description for a Tiltwing . . . 198
7.6.1.1 Error Modeling . . . 203
7.6.1.2 Robust Stability Condition . . . 204
7.6.2 Numerical Robustness Proof . . . 205
7.6.2.1 Results . . . 208
7.6.2.2 Validation . . . 209
7.6.3 Comparative Designs: H_∞ and μ-synthesis . . . 212
7.7 Conclusion . . . 215

8 Summary and Conclusions **216**

Bibliography **221**

List of Figures

1.1 Historical VTOL Aircraft 1
1.2 Examples of Transition VTOL Aircraft 2
1.3 Examples of Modern Tiltwings for Urban Air Mobility 3
1.4 Example thrust vs lift generation throughout transition for a tiltwing 5

2.1 Coordinate Frames 11
2.2 Feedback linearization control loop 18
2.3 Example circulation control airfoil with a blown flap that forms a cylindrical trailing-edge 39
2.4 Lift and drag curves for the circulation control airfoil of Figure 2.3 . . 40
2.5 Two types of uncertainty representations for an uncertain system G (represented at the system input) 42
2.6 General block diagram configurations with Δ uncertainty 43
2.7 Uncertain system G represented by a multiplicative uncertainty W_Δ in feedback with a controller K. This diagram can be used to derive the $M\Delta$-structure for this system. 46

3.1 Tiltwing Aircraft 51
3.2 Scaled Model Schematic (dimensions in mm) 51
3.3 Assumed Mission Profile 54
3.4 Tiltwing Moment Arms, Prop Directions, and Wing Angles 55
3.5 Propeller Coefficients vs Advance Ratio 57
3.6 Tiltwing Aerodynamic Modeling 59
3.7 Longitudinal Aero Coefficients vs Angle of Attack ($\delta = 0$) 62
3.8 Wing and Tail C_L vs Angle of Attack (range of δ) 62
3.9 Airspeed Vector Diagrams 66
3.10 Tiltwing aircraft wing and tail geometries set up for VLM 72
3.11 Tiltwing aircraft wing, tail, and fuselage geometries in OpenVSP set up for panel method 72
3.12 Cruise mode pole locations (three operating points) 75
3.13 Hover mode pole locations (with nominal cruise for comparison) . . . 76
3.14 Transition trajectory used for linear analysis 77
3.15 All longitudinal pole locations throughout transition with $\pm 2.6\%$ CG position variations. The color shift from orange to blue indicates moving from hover to cruise. 78

3.16 Longitudinal pole locations throughout transition – zoomed near origin. The color shift from orange to blue indicates moving from hover to cruise. 78

4.1 Full-Scale Aircraft Schematic 84
4.2 Forward trajectories without AFC – airspeed, altitude, thrusts, and wing angles. CG variation of ±2.6% the aircraft's length 100
4.3 Forward trajectories without AFC – pitch angle, elevator deflection, wetted AoAs, and AFC demands. CG variation of ±2.6% the aircraft's length 100
4.4 Forward trajectories with AFC – airspeed, altitude, thrusts, and wing angles. CG variation of ±2.6% the aircraft's length 101
4.5 Forward trajectories with AFC – pitch angle, elevator deflection, wetted AoAs, and AFC demands. CG variation of ±2.6% the aircraft's length 101
4.6 Backward trajectories without AFC – airspeed, altitude, thrusts, and wing angles. CG variation of ±2.6% the aircraft's length 102
4.7 Backward trajectories without AFC – pitch angle, elevator deflection, wetted AoAs, and AFC demands. CG variation of ±2.6% the aircraft's length 102
4.8 Backward trajectories with AFC – airspeed, altitude, thrusts, and wing angles. CG variation of ±2.6% the aircraft's length 103
4.9 Backward trajectories with AFC – pitch angle, elevator deflection, wetted AoAs, and AFC demands. CG variation of ±2.6% the aircraft's length 103
4.10 Backward trajectory with NACA 63(2)-615 tail airfoil, AFC, and aft-shifted CG – airspeed, altitude, thrusts, and wing angles 104
4.11 Backward trajectory with NACA 63(2)-615 tail airfoil, AFC, and aft-shifted CG – pitch angle, elevator deflection, wetted AoAs, and AFC demands 104

5.1 General full-envelope flight controller architecture 110
5.2 Inner-loop flight controller composed of inner-loop attitude controller subsystem (NDI attitude rate control + Euler-angle tracking) and control allocation. 111
5.3 NDI attitude rate controller in closed-loop with the tiltwing aircraft and control allocation. 114
5.4 Integral Anti-Windup Strategy 119
5.5 Constrained MIMO input vector scaling example 120
5.6 Simplified thrust and moment balance during take-off 121
5.7 Yaw rate, yaw angle, and aileron demand during structural vibrations excited by the NDI attitude rate controller with sufficiently-high bandwidth 124

5.8 Roll and pitch step responses of 0.1 rad and yaw rate step response of 0.1 rad s^{-1} in hover simulations. Horizontal dashed line indicates step input value. . . . 126
5.9 Prop and aileron demands corresponding to each attitude step response from Figure 5.8 . . . 126
5.10 Attitude responses to failure of left-inner wing prop at $t = 5$ s in hover simulations. Vertical dotted line indicates moment of failure. . . . 127
5.11 Prop and aileron demands during failure of left-inner wing motor at $t = 5$ s in hover simulations. Vertical dotted line indicates moment of failure. . . . 127
5.12 Indoor (tethered) testing setup . . . 128
5.13 Outdoor (untethered) testing setup . . . 129
5.14 Roll, pitch, and yaw rate responses during example flight . . . 132
5.15 Prop and aileron demands corresponding to each attitude response from Figure 5.14 . . . 132
5.16 Attitude responses during example take-off. The first vertical dotted line indicates first moment of lift-off while the second vertical dotted line indicates full hover. . . . 133
5.17 Prop and aileron demands during example take-off. The first vertical dotted line indicates first moment of lift-off while the second vertical dotted line indicates full hover. . . . 133
5.18 Attitude responses to failure of left-inner wing prop at flight time $t =$ 184.3 s. Vertical dotted line indicates moment of failure. . . . 134
5.19 Prop and aileron demands during failure of left-inner wing prop at flight time $t = 184.3$ s. Vertical dotted line indicates moment of failure. 134
5.20 Example roll, pitch, and yaw rate maneuvering during hover flight . . 137
5.21 Prop and aileron demands corresponding to each attitude maneuver from Figure 5.20 . . . 137
5.22 Attitude responses during take-off. The first vertical dotted line indicates first moment of lift-off while the second vertical dotted line indicates full hover. . . . 138
5.23 Prop and aileron demands during take-off. The first vertical dotted line indicates first moment of lift-off while the second vertical dotted line indicates full hover. . . . 138

6.1 2-dimensional thrust-vectored aircraft . . . 141
6.2 Fixed-wing tri-rotor configuration from Yayla et al. . . . 146
6.3 Generic 2-DoF pointmass free-body diagram . . . 149
6.4 Force and Velocity Diagrams . . . 153
6.5 Pseudo-control allocation scheme diagram showing flow of variable assignments for $V_{c_\perp}$ control using θ . . . 162

7.1 Linear approximation of internal dynamics regulation using θ_{ref} mapping (Strategy 1) . . . 177
7.2 Forward transition trajectory solution for the scaled model . . . 179

7.3 Backward transition trajectory solution for the scaled model 180
7.4 Aircraft free-body diagram during coordinated turn 181
7.5 Completed full-envelope flight controller layout 182
7.6 V_x step response in hover (Strategy 1) 184
7.7 V_z step response in hover (Strategy 1) 184
7.8 V_x step response in cruise (Strategy 1) 185
7.9 V_z step response in cruise (Strategy 1) 185
7.10 Controller performance in forward transition (Strategy 1) 185
7.11 Controller performance in backward transition (Strategy 1) 186
7.12 V_x step response in hover (Strategy 3) 189
7.13 V_z step response in hover (Strategy 3) 189
7.14 V_x step response in cruise (Strategy 3) 190
7.15 V_z step response in cruise (Strategy 3) 190
7.16 Controller performance in forward transition (Strategy 3) 190
7.17 Controller performance in backward transition (Strategy 3) 191
7.18 V_y step response in hover 192
7.19 Turn rate step response in cruise ($V_x = 20\ \mathrm{m\,s^{-1}}$) 192
7.20 V_x maneuvering in hover flight 196
7.21 V_z maneuvering in hover flight 196
7.22 Controller performance during wing rotation tests in hover flight . . . 197
7.23 Simplified longitudinal tiltwing model for pitch rate dynamics 199
7.24 Outer-loop θ-dynamics in the presence of uncertain q-loop dynamics . 207
7.25 RSM vs 3σ uncertainty bounds for tiltwing in transition 208
7.26 Forward transition simulations, $n_s = 10^3$ and $3\sigma = 3\%$. 211
7.27 Backward transition simulations, $n_s = 10^3$ and $3\sigma = 6\%$. 211
7.28 Forward transition simulations, $n_s = 10^3$ and $3\sigma = 9\%$. 211
7.29 Backward transition simulations, $n_s = 10^3$ and $3\sigma = 10\%$. 211
7.30 Interconnection structure for H_∞ and μ-synthesis control of feedback-linearized pitch rate. 212
7.31 Robust controllers compared with the current LQR design, $K_{dc}^v = 2000$. 215
7.32 Robust controllers compared with the current LQR design, $K_{dc}^v = 1$. . 215

List of Publications

The following is a list of conference and journal papers that were published as a result of the research carried out for this thesis:

- Panish, L., and Bacic, M., "Tiltwing eVTOL Transition Trajectory Optimization," *AIAA Journal of Aircraft*, Vol. 62, No. 1, pp. 81–93, Jan–Feb 2025. https://doi.org/10.2514/1.C037862.
- Panish, L., and Bacic, M., "A Generalized Full-Envelope Outer-Loop Feedback Linearization Control Strategy for Transition VTOL Aircraft," *AIAA Aviation Forum*, June 2023. https://doi.org/10.2514/6.2023-4511.
- Panish, L., Nicholls, C., and Bacic, M., "Nonlinear Dynamic Inversion Flight Control of a Tiltwing VTOL Aircraft," *AIAA SciTech Forum*, January 2023. https://doi.org/10.2514/6.2023-1910.
- Panish, L., and Bacic, M., "Transition Trajectory Optimization for a Tiltwing VTOL Aircraft with Leading-Edge Fluid Injection Active Flow Control," *AIAA SciTech Forum*, January 2022. https://doi.org/10.2514/6.2022-1082.

Chapter 1

Introduction

The idea of vertical flight has been around for thousands of years, the earliest known history being children's bamboo flying toys of 400 BC China and Japan known as *taketombo* [1]. However, it wasn't until the early 20$^{\text{th}}$ century that the first manned vertical take-off and landing (VTOL) aircraft was flown, and it would take until after World War II, with the invention of the turbine engine, for this technology to develop as we recognize it today [2].

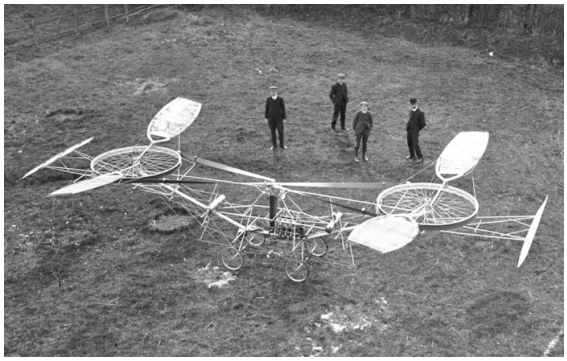

(a) The Cornu Helicopter. First manned VTOL, 1907 [3]

(b) The Kaman K-225. First turbine-powered helicopter, 1951 [4]

Figure 1.1: Historical VTOL Aircraft

VTOL aircraft can be categorized into broadly two groups: rotorcraft and powered-lift. The most widely-known example of rotorcraft is the helicopter, where all forces required to maintain flight are provided by an engine-driven rotor, though other con-

figurations exist including the autogyro and the gyrodyne. Conversely, powered-lift aircraft are a type of transition aircraft that get their VTOL capabilities from engine power but use a non-rotating aerodynamic surface to generate lift (i.e. a wing) during horizontal flight. Example configurations include tiltwings, tiltrotors, and tiltducts, all of which operate by shifting the direction of thrust depending on flight mode. Tiltwings achieve this by tilting the main wing with the engines/props attached, while tiltrotors have a fixed horizontal wing with engine nacelles that tilt. Tiltducts operate similarly but through the use of ducted fans. Two examples of such thrust-to wing-borne transition VTOL aircraft are the Canadair CL-84 tiltwing and the Bell-Boeing V-22 Osprey tiltrotor, shown in Figure 1.2.

(a) The Canadair CL-84 tiltwing [5]

(b) The Bell-Boeing V-22 Osprey tiltrotor [6] (Image licensed under CC BY-SA 4.0)

Figure 1.2: Examples of Transition VTOL Aircraft

Recent improvements in hybrid and electric technology have made civil transition VTOL aircraft a more present reality than they were back in the late 20$^{\text{th}}$ century. Hence, in recent years, there has been a rising interest in Urban Air Mobility (e.g. "air taxis"), the idea of transport aircraft with VTOL capabilities designed for intra- and/or inter-city travel [7, 8]. Many of the concepts studied adopt either hybrid [9, 10] or all-electric powertrains [11, 12] due to considerations of emissions and ease of flight control and redundancy offered by multiple propulsors. However, the usage of electrical storage and propulsion reduces the overall payload and range capability of flight vehicles relative to those powered by kerosene. Consequently, transition

aircraft, which combine VTOL capabilities required for operation in dense, urban environments with the efficiency of fixed-wing operation during cruise, tend to be a popular design choice for air taxi concepts [9, 12]. Two examples of such air taxi concepts are shown in Figure 1.3. The Airbus A^3 Vahana is a tandem tiltwing design that is popular in current literature regarding tiltwing modeling and controller design, and the Rolls-Royce eVTOL concept is the configuration of interest in this thesis.

(a) The Airbus A^3 Vahana [13]
(Image licensed under CC BY-SA 4.0)

(b) The Rolls-Royce eVTOL concept [14]
(©Rolls-Royce)

Figure 1.3: Examples of Modern Tiltwings for Urban Air Mobility

1.1 Motivation: Tiltwing VTOL Aircraft and their Challenges

There was an immense amount of research and development in the 1960s–1980s of various transition aircraft configurations [15, 16, 17, 18, 19]. Tiltwings [16], tiltrotors [17, 18], and tiltducts [19] were three major classes that were investigated. However, only one aircraft of this type (Bell-Boeing V-22 Osprey) has reached operational service, and one nears introduction (Leonardo AW609), both being tiltrotors. The relative merits of tiltwings vs tiltrotors are a subject of much debate [20, 21]. Each have notable pros and cons, ranging from mechanical complexity to the flight control design, but both present significant flight control challenges during conversion between thrust- and wing-borne flight as the aircraft are often pushed to their operational limits within their feasible flight envelopes.

Tiltwings can be considered an attractive choice from a systems-level point of view. Comparing the thrust generation capabilities of tiltwings and tiltrotors, the latter experience significant down-loading from the rotor wake in hover due to its fixed wing, which effectively behaves as a flat plate interfering with the airflow and thus requires extra thrust and power to compensate [22]. Conversely, for a tiltwing, the wing chord is always aligned with the prop wake, rendering this effect negligible in comparison. As a further consequence, tiltrotors tend to have much larger rotor blades than a tiltwing's prop blades. This in turn reduces the disk loading, which can improve efficiency in hover but makes cruise less efficient [22]. Hence, tiltwings are better suited for commercial corporate and commuter markets because their generally smaller prop radius means much faster and more efficient cruise, and hover times are not significant in such operations [23]. Mechanically, a tiltwing will also tend to be less complex than a tiltrotor, rendering tiltwings more reliable and easier/cheaper to manufacture. The tilting mechanism itself for a tiltrotor generally requires intricate drive systems due to the nacelles being widely separated, while everything moves synchronously on a tiltwing. As an example, the Bell-Boeing V-22 Osprey tiltrotor drive system consists of five gearboxes and an interconnected drive shaft to ensure one-engine-inoperative operations [23]. Furthermore, there are weight drawbacks for a tiltrotor due to aeroelasticity effects. Specifically, tiltrotors are particularly susceptible to a whirl flutter instability involving a coupling of the rotors and the fuselage-wing bending modes as a result of typically lower engine speeds than for a tiltwing [22]. To compensate, significant stiffening of the wing structure is required, which adds a weight penalty [23].

However, tiltwings present some unique challenges, a particular one being the transition maneuver itself from hover to cruise flight (and back). This maneuver for tiltwing aircraft, visualized in Figure 1.4, is particularly challenging due to the shifting wing angle while the aircraft rapidly changes airspeed [24]. In hover-to-cruise

(forward) transition, tiltwing aircraft must generate the necessary thrust to balance the aircraft's weight while also accelerating to a speed where sufficient lift can be generated by the wings. In cruise-to-hover (backward) transition, the same balancing act must be achieved but in reverse. In both cases, the shifting wing angle and rapidly changing airspeed often pushes the local angle of attack (AoA) of the wings to their aerodynamic stall regime, which is problematic for aircraft controllability, performance, and safety. In a report by McKinney et al. [25], the phenomenon of wing aerodynamic stall in transition for tiltwing aircraft is discussed in detail. While the propeller slipstream blowing over the wing tends to help keep the local AoA on the wing within the limits of stall, stalling is still likely to occur unless the aircraft climbs throughout the maneuver (as illustrated in Figure 1.4), which helps reduce the angle between the wing chord and the aircraft's flight path, and if high-lift devices (such as flaps or slats) are used. This phenomenon was validated by flight testing of the Canadair CL-84 tiltwing [16], and while increasing altitude helps avoid stall throughout transition, this is not preferred for civil applications (e.g. air taxis), where minimal altitude variation may be necessary for operation within narrow urban flight corridors and smoother maneuvers may be desired for passenger comfort.

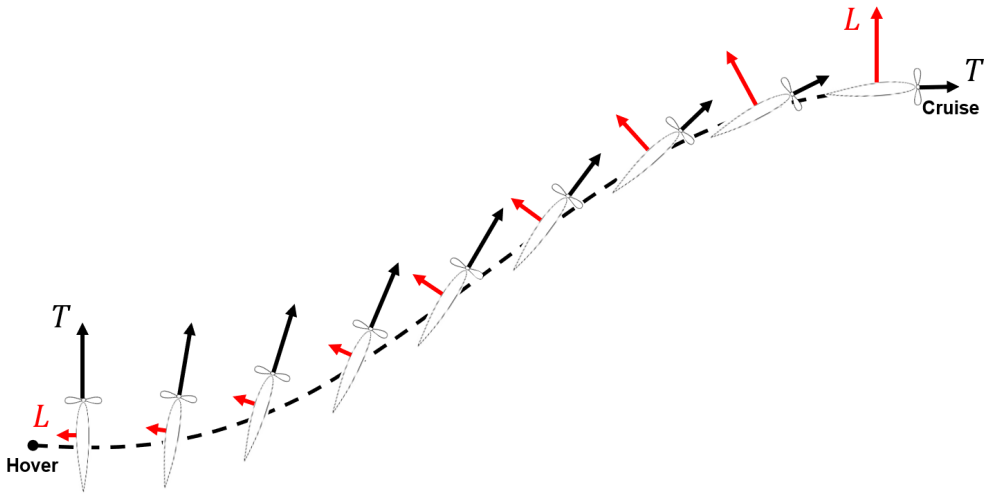

Figure 1.4: Example thrust vs lift generation throughout transition for a tiltwing

Another challenge, which applies to all transition VTOL aircraft, is in the area of the flight control design. The flight control is a particularly challenging problem due to the significant variation of these aircraft's flight dynamics (including complex prop-wing aerodynamic coupling/interactions) as they move through their complete operating regime (hover, cruise, and transition). These configurations, which are fundamentally characterized by relying on thrust generation for hover mode and lift generation for cruise, generally exhibit a lack of independent thrust actuation in both the vertical and horizontal directions. Dynamically, this necessitates that the thrust demand be coupled with the angle at which it will act in order to maintain full controllability of the longitudinal velocity states, or otherwise, the vertical airspeed dynamics exhibits changing relative degrees throughout the flight envelope [26]. This fundamental difference in operation (thrust- vs wing-borne flight), which is exemplified in Figure 1.4, thus poses a unique control challenge during the transition maneuver and proves to be a key constraint in finding full-envelope flight control solutions for the design of the outer-loop (i.e. velocities) dynamics for these aircraft, rendering such control design to be a complex task without either assuming independent thrust actuation in both the vertical and horizontal directions and/or implementing a gain schedule or other switching behavior. Hence, typical control approaches require designing many flight controllers for various modes of flight and finding ways to smoothly switch between them, a laborious task that has potential downsides. Furthermore, the current lack of generalized, full-envelope control strategies emphasizes the complexity of controlling such aircraft and indicates a potential lack of understanding of the underlying dynamics of these vehicles.

1.2 Purpose of Research

This thesis investigates full-envelope flight control design for a tiltwing VTOL aircraft through the application of a nonlinear control technique called feedback linearization. Typically, the lack of independent thrust actuation in both the vertical and horizontal directions limits the applicability of a feedback linearization controller in the context of the outer-loop (velocities) design. However, this research proposes a novel solution through the careful selection of thrust-aligned virtual control inputs and outputs, resulting in a globally valid controller that is generalizable to a wide range of transition VTOL configurations. Combined with a full-envelope inner-loop (attitudes) solution, such a control concept tackles one of the core difficulties in controlling these aircraft (the lack of independent thrust actuation) and provides a deeper understanding of the underlying dynamics governing how these aircraft perform.

To demonstrate the effectiveness of such a control strategy, this controller design is evaluated in both simulation and flight testing using a dynamically-scaled tiltwing demonstrator aircraft based on the Rolls-Royce eVTOL concept (shown in Figure 1.3b). This configuration serves as an ideal platform to test the validity and effectiveness of this control strategy while additionally solving/overcoming its unique complexities, such as the challenge of avoiding wing aerodynamic stall throughout the transition maneuver. Hence, the contributions of this thesis are:

- A complete 6-degree-of-freedom (DoF) model of this unique tiltwing aircraft
- A novel full-envelope feedback linearization-based controller design
- A full-envelope control allocation methodology designed for fault-tolerance (an important safety consideration in the context of air taxis)
- Trajectory optimization solutions for this tiltwing aircraft to determine feasible minimum-energy transition trajectories

- Simulation and flight testing results as well as a formal robustness analysis to demonstrate the validity and effectiveness of this controller design

The design, implementation, and demonstration of this system through the research presented in this thesis helps widen the design space for tiltwing aircraft (and transition aircraft in general) and contributes toward minimizing energy use and maximizing passenger comfort and safety in newly-emerging Urban Air Mobility markets.

1.3 Thesis Content

Chapter 2 reviews key concepts pertinent to the content of this thesis. This includes general aircraft modeling and dynamics, feedback linearization control and control allocation, numerical methods for solving trajectory optimization/general optimal control problems, and uncertainty modeling and robustness analysis. Chapter 3 then presents the complete 6-DoF model and dynamics of the tiltwing aircraft of interest in this work, which will be used throughout the thesis, including a brief analysis of its dynamics across its complete flight-envelope. Before considering the controller design, feasible transition trajectories are explored in Chapter 4. This will not only help build a further understanding of the dynamics and performance of this aircraft, but this methodology will also be later used to generate outer-loop reference inputs for the full-envelope flight controller. The full-envelope flight controller design and all associated results are then presented across the next three chapters. Chapter 5 begins this trilogy by discussing the inner-loop attitude controller design, simulation predictions, and initial flight testing results. Focus is then shifted in Chapter 6 to the general problem of the outer-loop velocity control design for thrust- to wing-borne transition VTOLs, where a novel approach is presented that demonstrates how feedback linearization can be applied in this context through the careful selection of thrust-aligned virtual control inputs and outputs. While the focus of this thesis is

on tiltwings, the concepts in this chapter are presented in this more general context to emphasize key points in this novel approach and to build an understanding of the fundamental dynamics governing how these aircraft operate and behave. Chapter 7 then shifts the focus back to the specific tiltwing configuration of interest in this thesis. By applying the generalized outer-loop control approach that was introduced in the previous chapter, the outer-loop velocity controller design and completed flight controller for this aircraft are presented, and the effectiveness of this completed design is demonstrated through simulation, flight testing, and a formal robustness proof and analysis. Lastly, conclusions for the thesis are drawn in Chapter 8.

The content of Chapter 4 was published in the AIAA Journal of Aircraft [27] following preliminary ideas for the work that were presented as paper 2022-1082 at the AIAA SciTech 2022 conference [28]. The inner-loop controller design, control allocation, and hover simulation and (tethered) flight testing results of Chapter 5 were originally presented as paper 2023-1910 at the AIAA SciTech 2023 conference [29]. Preliminary work on Chapter 6 and its associated content in Chapter 7 was originally presented as paper 2023-4511 at the AIAA Aviation 2023 conference [26], and the completed work (Chapters 6 and 7) is planned to be submitted as a journal paper(s) to the AIAA Journal of Guidance, Control, and Dynamics.

Chapter 2

Key Concepts and Literature Review

2.1 Introduction

This chapter provides a background of the key dynamics and control concepts relevant to the content of this thesis. Section 2.2 gives an overview of general flight dynamics and how a generic aircraft can be modeled as a 6-DoF rigid body. A thorough exploration of feedback linearization control is then covered in Section 2.3, as this is the core nonlinear control technique used in this work, followed by a discussion of control allocation as a method for allocating inputs in an over-actuated system such as a tiltwing. Focus is then shifted in Section 2.5 to a review of trajectory optimization and numerical methods for solving general optimal control problems. Section 2.6 gives a brief overview of active flow control (AFC) and current literature on relevant aircraft applications. Lastly, Section 2.7 reviews system uncertainty and robustness in the context of analysis and control design.

2.2 General Flight Dynamics

2.2.1 Coordinate Frames

To derive the aircraft equations of motion, four coordinate frames are required: Earth-fixed, body, wind, and stability. The Earth-fixed and body frames are the primary coordinate systems used to define the relevant state variables of the aircraft (position, attitude, velocity, rotation rate), while the wind and stability frames are primarily used for aerodynamic modeling.

The Earth-fixed frame, also known as the topodetic reference system [30] or the inertial North-East-Down (NED) frame [31], is fixed in space with its x-axis directed North, y-axis East, and z-axis down into the Earth. The aircraft's body axis system is also a NED-oriented coordinate frame but is instead aligned with the aircraft's position and orientation. The origin of this coordinate system lies at the aircraft's center of gravity (CG), and its x-axis (denoted x_b, etc.) points toward the nose of the aircraft, the y-axis toward the aircraft's right wing, and the z-axis perpendicularly downward, as shown in Figure 2.1a.

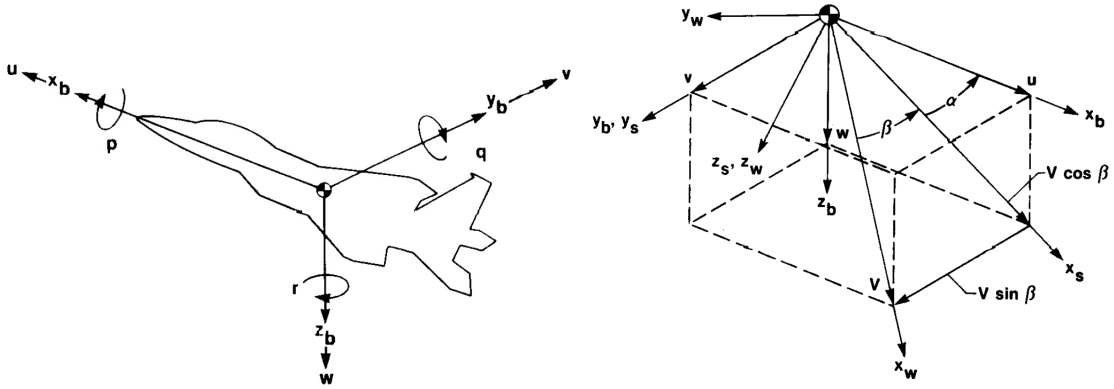

(a) Aircraft Body Axes [30] (b) Wind Axes relative Body and Stability [30]

Figure 2.1: Coordinate Frames

The aircraft's body velocities are denoted u, v, and w for components in the x_b, y_b, and z_b directions, respectively, while its rotation rates about those axes are denoted p,

q, and r. The attitude of the aircraft is described by the three Euler angles (ϕ, θ, ψ) representing roll, pitch, and yaw, respectively. These angles are measured as counter-clockwise positive rotations of the aircraft's body axes relative the Earth-fixed frame. Note that because the aircraft's attitude is measured relative the Earth-fixed frame, the Euler angle time derivatives ($\dot{\phi}$, $\dot{\theta}$, $\dot{\psi}$) are not equal to the body rotation rates (p, q, r) [32]. In a similar fashion to the aircraft's attitude, the position of the aircraft (x, y, z) is measured by the location of its body frame relative the Earth-fixed frame.

The wind coordinate frame is defined by a radial vector (the freestream airspeed V) and two angles (the freestream angle of attack α and the sideslip angle β), as shown in Figure 2.1b. Note that the sideslip angle β is taken as clockwise positive, and subscripts w and s indicate wind and stability, respectively. The freestream angle of attack α is defined as the angle between the horizontal body axis (x_b) and the freestream airspeed vector in the longitudinal direction, while the sideslip angle is defined as the angle between the longitudinal plane and the airspeed vector in the lateral direction.

Mathematically, the freestream airspeed V, freestream angle of attack α, and sideslip β can be expressed in terms of the body-axis velocities [32]:

$$\begin{aligned} V = \|\mathbf{V}\| &= \sqrt{u^2 + v^2 + w^2} \\ \alpha &= \arctan\left(\frac{w}{u}\right) \\ \beta &= \arcsin\left(\frac{v}{V}\right) \end{aligned} \tag{2.1}$$

Note that for conventional fixed-wing aircraft, α typically represents the incidence angle of the main wing (i.e. its angle of attack), but this assumes that the wing is fixed and its chord is aligned with the aircraft's longitudinal axis. In the context of a tiltwing aircraft, it will be necessary to define this differently to account for both the wing angle ξ and any prop-wing aerodynamic interactions [33, 34].

Lastly, the stability frame is closely related to the wind axis system but is oriented relative the wind frame through a rotation of β and relative the body frame through a rotation $-\alpha$. This coordinate frame is commonly used to define aerodynamic coefficients that are employed in standard empirical aerodynamic models [30, 31, 32, 35, 36].

2.2.2 6-DoF Equations for a Rigid Body

A generic aircraft can be modeled as a rigid body with three translational and three rotational degrees of freedom, resulting in a set of 12 simultaneous nonlinear differential equations describing the aircraft's position, attitude, velocity, and rotation rate dynamics. The equations for the aircraft's position and attitude (Earth-fixed frame) can be derived purely from kinematics, while its velocity and rotation rate (body frame) are derived using the Newton-Euler equations for a rigid body.

Beginning with the latter, we apply the Newton-Euler equations to the aircraft's CG in the body reference frame [30]:

$$\begin{aligned} \mathbf{F}_{net} &= m(\dot{\mathbf{V}} + \boldsymbol{\Omega}_b \times \mathbf{V}) \\ \mathbf{M}_{net} &= \mathbf{J}\dot{\boldsymbol{\Omega}}_b + \boldsymbol{\Omega}_b \times (\mathbf{J}\boldsymbol{\Omega}_b) \end{aligned} \tag{2.2}$$

where $\mathbf{F}_{net}$ and $\mathbf{M}_{net}$ are 3×1 column vectors containing the net external forces and moments on the body, respectively, m is the mass of the aircraft, $\mathbf{V}$ is the body velocity vector, $\boldsymbol{\Omega}_b$ contains the body rotation rates, and $\mathbf{J}$ is the symmetric inertia tensor with elements I_{ij}. To obtain a set of equations in state-space form, (2.2) above is rearranged such that all first-order derivatives are on the left-hand side:

$$\begin{aligned} \dot{\mathbf{V}} &= -\boldsymbol{\Omega}_b \times \mathbf{V} + \frac{1}{m}\mathbf{F}_{net} \\ \dot{\boldsymbol{\Omega}}_b &= \mathbf{J}^{-1}\left(-\boldsymbol{\Omega}_b \times (\mathbf{J}\boldsymbol{\Omega}_b) + \mathbf{M}_{net}\right) \end{aligned} \tag{2.3}$$

To derive the equations for the aircraft's position and attitude states, we make use of the following transformation matrices [30, 31, 32, 36, 37]:

$$\mathbf{T}_{be} = \begin{bmatrix} c_\theta c_\psi & c_\theta s_\psi & -s_\theta \\ s_\phi s_\theta c_\psi - c_\phi s_\psi & s_\phi s_\theta s_\psi + c_\phi c_\psi & s_\phi c_\theta \\ c_\phi s_\theta c_\psi + s_\phi s_\psi & c_\phi s_\theta s_\psi - s_\phi c_\psi & c_\phi c_\theta \end{bmatrix} \quad \mathbf{R}_{be} = \begin{bmatrix} 1 & 0 & -s_\theta \\ 0 & c_\phi & s_\phi c_\theta \\ 0 & -s_\phi & c_\phi c_\theta \end{bmatrix} \tag{2.4}$$

where c_ϕ and s_ϕ are shorthand for $\cos(\phi)$ and $\sin(\phi)$, etc. These matrices transform translations and rotations from the inertial Earth-fixed frame to the local body-axis system. Hence,

$$\begin{aligned} \begin{bmatrix} \dot{x} & \dot{y} & \dot{z} \end{bmatrix}^T &= \mathbf{T}_{be}^{-1}\mathbf{V} \\ \begin{bmatrix} \dot{\phi} & \dot{\theta} & \dot{\psi} \end{bmatrix}^T &= \mathbf{R}_{be}^{-1}\mathbf{\Omega}_b \end{aligned} \tag{2.5}$$

Expanding (2.3) and (2.5), and assuming symmetry in the aircraft's xz-plane (i.e. $I_{xy} = I_{yz} = 0$), gives the set of 12 simultaneous nonlinear differential equations of motion for a generic aircraft:

$$\begin{bmatrix} \dot{x} \\ \dot{y} \\ \dot{z} \\ \dot{\phi} \\ \dot{\theta} \\ \dot{\psi} \\ \dot{u} \\ \dot{v} \\ \dot{w} \\ \dot{p} \\ \dot{q} \\ \dot{r} \end{bmatrix} = \begin{bmatrix} u(c_\theta c_\psi) + v(s_\phi s_\theta c_\psi - c_\phi s_\psi) + w(c_\phi s_\theta c_\psi + s_\phi s_\psi) \\ u(c_\theta s_\psi) + v(c_\phi c_\psi + s_\phi s_\theta s_\psi) + w(c_\phi s_\theta s_\psi - s_\phi c_\psi) \\ -u(s_\theta) + v(s_\phi c_\theta) + w(c_\phi c_\theta) \\ p + q(s_\phi t_\theta) + r(c_\phi t_\theta) \\ q(c_\phi) - r(s_\phi) \\ q(s_\phi/c_\theta) + r(c_\phi/c_\theta) \\ rv - qw + F_{x_{net}}/m \\ pw - ru + F_{y_{net}}/m \\ qu - pv + F_{z_{net}}/m \\ I_1\left(pqI_2 + qrI_3 + I_{zz}M_{x_{net}} + I_{xz}M_{z_{net}}\right) \\ pr(I_{zz} - I_{xx})/I_{yy} + (r^2 - p^2)I_{xz}/I_{yy} + M_{y_{net}}/I_{yy} \\ I_1\left(pqI_4 - qrI_2 + I_{xz}M_{x_{net}} + I_{xx}M_{z_{net}}\right) \end{bmatrix} \tag{2.6}$$

where,

$$\begin{aligned} I_1 &= 1/(I_{xx}I_{zz} - I_{xz}^2) & I_3 &= -I_{zz}^2 + I_{yy}I_{zz} - I_{xz}^2 \\ I_2 &= I_{xz}(I_{xx} - I_{yy} + I_{zz}) & I_4 &= I_{xx}^2 - I_{xx}I_{yy} + I_{xz}^2 \end{aligned} \tag{2.7}$$

Hence, to derive the dynamic model for a specific aircraft, the net external forces $\mathbf{F}_{net} = [F_{x_{net}} \quad F_{y_{net}} \quad F_{z_{net}}]^T$ and moments $\mathbf{M}_{net} = [M_{x_{net}} \quad M_{y_{net}} \quad M_{z_{net}}]^T$ appearing in (2.6) must be defined for the particular aircraft's configuration and effector model.

2.3 Feedback Linearization Control

Feedback linearization is a powerful nonlinear control technique that involves canceling the nonlinearities of a system model through feedback, allowing one to apply familiar linear control techniques such as PID or LQR [38]. Where a local Jacobian linearization linearizes a system about a specific operating point (a specific point in the system domain), feedback linearization achieves linearization over an entire region of the domain. This region is often a subset of the domain, but in special cases it can be the entire domain itself [39]. However, due to relying on inverting the system model, feedback linearization can be quite sensitive to disturbances and modeling uncertainties. Nevertheless, due to its effectiveness at handling nonlinear systems, feedback linearization has been widely studied and used in aerospace control applications. Before diving into the details, we will begin with two simple examples to build intuition around how this nonlinear control technique works.

Example 2.1. Full-state linearization.

Consider the following nonlinear single-input-single-output (SISO) input-affine system:

$$\begin{aligned} \begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \end{bmatrix} &= \begin{bmatrix} x_2^2 - x_1 \\ x_1^3 + x_2 \end{bmatrix} + \begin{bmatrix} 0 \\ 1 \end{bmatrix} u \\ y &= x_1 \end{aligned} \tag{2.8}$$

To find a relationship between the input u and the output y, we can differentiate the output until the input appears in its dynamics, which in this case requires differentiating twice:

$$\ddot{y} = x_1 + 2x_1^3 x_2 + x_2^2 + 2x_2 u \tag{2.9}$$

Equating $\ddot{y}$ to a new input v allows us to construct the following inversion control law:

$$u = \frac{1}{2x_2}\left(v - x_1 - 2x_1^3 x_2 - x_2^2\right) \tag{2.10}$$

and a new state vector $\mathbf{z}$ can be constructed as $\mathbf{z} = [y \quad \dot{y}]^T$, resulting in the following linear system from v to y:

$$\begin{aligned} \dot{\mathbf{z}} &= \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \mathbf{z} + \begin{bmatrix} 0 \\ 1 \end{bmatrix} v \\ y &= \begin{bmatrix} 1 & 0 \end{bmatrix} \mathbf{z} \end{aligned} \tag{2.11}$$

Hence, this system is fully controllable, and it is therefore possible to design a state-feedback control law for v that stabilizes this system. This is known as full-state linearization. ■

Example 2.2. Input-output linearization.

Consider the same nonlinear system from (2.8), but we will now take the output to be $y = x_2$. Once again, we will differentiate the output y until the input u appears in its dynamics, which in this case only requires differentiating once:

$$\dot{y} = x_1^3 + x_2 + u \tag{2.12}$$

Thus, by equating $\dot{y}$ to a new input v, this system will now have the following inversion control law:

$$u = v - x_1^3 - x_2 \tag{2.13}$$

and a new state vector $\mathbf{z}$ can be constructed such that $z_1 = y$. However, the choice of z_2 is now (somewhat) arbitrary because z_2 will not appear in the linear system from v to y:

$$\begin{aligned} \dot{z}_1 &= v \\ y &= z_1 \end{aligned} \tag{2.14}$$

To complete the new state vector $\mathbf{z}$, we will select z_2 to ensure invertibility of this state transformation. The simplest choice for this is $z_2 = x_1$, which results in the following dynamics for $\mathbf{z}$:

$$\dot{\mathbf{z}} = \begin{bmatrix} v \\ z_1^2 - z_2 \end{bmatrix} \tag{2.15}$$

Hence, only part of the system's dynamics have been linearized, and the state z_2 remains unobservable by this state transformation. This is known as input-output linearization, and in order for this system to be stable, it is necessary that the unobservable state z_2 be stable on its own. ■

For a general discussion of the key principles of feedback linearization, we consider nonlinear multi-input-multi-output (MIMO) input-affine[1] square[2] systems of the following form:

$$\begin{aligned} \dot{\mathbf{x}} &= \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{u} \\ \mathbf{y} &= \mathbf{h}(\mathbf{x}) \end{aligned} \tag{2.16}$$

Feedback linearization seeks to apply a change of variables $\mathbf{z} = \boldsymbol{\phi}(\mathbf{x})$ and a state feedback control law $\mathbf{u} = \mathbf{E}^{-1}(\mathbf{x})\left(\mathbf{v} - \mathbf{b}(\mathbf{x})\right)$ in a new input $\mathbf{v}$ that together transform the nonlinear system given in (2.16) into an equivalent decoupled linear system in Brunovsky form (a chain of integrators) [39]. Figure 2.2 shows a block diagram of such a feedback linearization control loop.

[1]Note that non-affine systems of the form $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{u})$ may be considered as well, but these require finding the inverse mapping for $\mathbf{u}$, such as via an appropriate variable substitution or numerical methods [40].

[2]I.e. the number of inputs and outputs are equal.

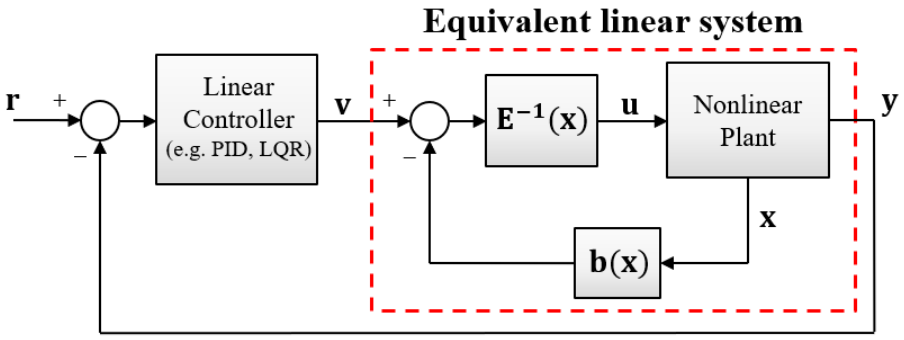

Figure 2.2: Feedback linearization control loop

As was shown in Examples 2.1 and 2.2, there are two possibilities with feedback linearization: full-state linearization, where the system is completely linearized, and input-output linearization, where only the mapping from $\mathbf{v}$ to $\mathbf{y}$ is linearized while the system itself may be only partially linearized [39]. It can be shown that a full-state linearized system, or an input-output linearized system that is minimum phase, is locally asymptotically stabilizable [40]. We will first introduce a few necessary mathematical concepts before arriving at the aforementioned change of variables and state feedback control law. Note that the bold notation for vectors and matrices will be omitted for the rest of this section.

2.3.1 Mathematical Concepts

2.3.1.1 Lyapunov Stability Theory

Consider the nonlinear autonomous closed-loop system $\dot{x} = f(x)$. A state x^* is an equilibrium state of this system if it satisfies:

$$f(x^*) = 0 \tag{2.17}$$

Determining the stability of a nonlinear system is not as obvious as it is for linear (time-invariant) systems, where the locations of the system's poles provide information about whether the time-domain state trajectories will grow or decay. For

nonlinear systems, a more general definition of stability is required:

Theorem 2.1. *If the equilibrium state of the nonlinear system is $x^* = 0$, then this equilibrium state is said to be stable if, for any $R > 0$, there exists some $r > 0$ such that if $\|x(0)\| < r$, then $\|x(t)\| < R$ for all $t \geq 0$. Otherwise, the equilibrium point is unstable.* [40]

This definition of stability, known as stability in the sense of Lyapunov, means that the system trajectory can be kept arbitrarily close to the origin provided it starts sufficiently close to it. Note that the equilibrium point need not be at the origin; if instead x^* is some non-zero value, then this same definition of stability applies simply by shifting the origin by $x - x^*$. For the remainder of this discussion on stability, only the stability of equilibrium points at the origin will be considered without loss of generality.

Considering the above definition of stability, notions of local vs global stability can be inferred. *Local stability* refers to the system trajectory remaining within some finite neighborhood around the origin (i.e. $R < \infty$). In contrast, if this neighborhood can be made arbitrarily large (i.e. $R \to \infty$) and stability in the sense of Lyapunov is still satisfied, then this is known as *global stability*. Additionally, notions of stability can be extended to describe whether the system trajectory will approach the origin, and if so, how quickly it does: An equilibrium point is said to be *asymptotically stable* if, in addition to being stable, $x(t) \to 0$ as $t \to \infty$; furthermore, an equilibrium point is said to be *exponentially stable* if the rate of convergence is bounded by an exponential function of the initial state, i.e. $\|x(t)\| \leq \alpha\|x(0)\|e^{-\lambda t}$ for all $t > 0$, where α and λ are strictly positive numbers [40]. Note that these definitions are not needed for linear (time-invariant) systems as, in contrast to general (nonlinear) systems, asymptotic stability is always global and exponential.

To formally determine stability for a general (nonlinear) system, we will use Lyapunov's direct method, which is an extension of the concept of energy dissipation in

a physical system [41]: If the total energy of a system continuously decreases with time, then the system must eventually settle down to an equilibrium point. This can be formulated mathematically through the use of positive definite functions. A scalar function $V(x)$ is defined as *locally positive definite* if $V(0) = 0$ and, in some region around the origin, $V(x) > 0$. Furthermore, if this region can be extended to span the whole state space, then $V(x)$ is said to be *globally positive definite* [40]. In contrast, a function is said to be *negative definite* if $-V(x)$ is positive definite. Additionally, the notion of a *semi-definite* function can be defined if the inequality instead satisfies $V(x) \geq 0$. Note that these definitions imply that $V(x)$ will have a minimum/maximum, and in the case of a definite function, this minimum/maximum is unique. Lastly, if the vector x denotes the state of the system, then the scalar function $V(x)$ is implicitly a function of time; hence, its time derivative can be written out using the chain rule [40]:

$$\dot{V}(x) = \frac{\partial V}{\partial x}\dot{x} = \frac{\partial V}{\partial x}f(x) \tag{2.18}$$

Lyapunov's theorem for stability can now be stated:

Theorem 2.2. *If, in some region around the origin, there exists a scalar function $V(x)$ that is positive definite and its derivative $\dot{V}(x)$ is negative semi-definite, then the equilibrium point at the origin is locally stable. If, in addition, $\dot{V}(x)$ is negative definite, then this point is locally asymptotically stable. Furthermore, if this region can be arbitrarily extended such that $V(x) \to \infty$ as $\|x\| \to \infty$ (and $\dot{V}(x)$ remains strictly negative definite), then this point is globally asymptotically stable.* [40]

If a scalar function $V(x)$ can be found that satisfies the above, then $V(x)$ is called a *Lyapunov function* for the system $\dot{x} = f(x)$. Note that this theorem is a sufficient condition for stability but not a necessary one; hence, finding a Lyapunov function can be a non-trivial task, and struggling to find one does not mean a system is necessarily

unstable. That said, in practice a common Lyapunov function to form is simply a quadratic in the state vector x, i.e. $V(x) = x^T P x$.

2.3.1.2 Lie Derivatives and Lie Brackets

Consider the vector fields[3] $f(x)$ and $h(x)$ on the domain $D \subset \mathbb{R}^n$. We define the Lie derivative of element i in h with respect to f as:

$$L_f h_i(x) = \frac{\partial h_i}{\partial x} f(x) \tag{2.19}$$

Lie derivatives can be defined recursively using the following notation [39]:

$$\begin{aligned} L_f^0 h_i(x) &= h_i(x) \\ L_f^k h_i(x) = L_f L_f^{k-1} h_i(x) &= \frac{\partial L_f^{k-1} h_i}{\partial x} f(x) \end{aligned} \tag{2.20}$$

The Lie derivative is a scalar function that represents the directional derivative of h_i along the direction of the vector field f, and its relevance to dynamical systems can be seen by considering the nonlinear system of (2.16). Assuming $u = 0$ for this example, the time derivatives of the ith output of y can be written using Lie derivatives as:

$$\begin{aligned} \dot{y}_i &= \frac{\partial h_i}{\partial x} \dot{x} = L_f h_i(x) \\ \ddot{y}_i &= \frac{\partial L_f h_i}{\partial x} \dot{x} = L_f^2 h_i(x) \end{aligned} \tag{2.21}$$

and so on [40].

Considering once again the vector fields $f(x)$ and $h(x)$, the Lie bracket of f and h is a third vector field defined as:

$$[f, h] = \frac{\partial h}{\partial x} f(x) - \frac{\partial f}{\partial x} h(x) \tag{2.22}$$

[3]For consistency with terminology from differential geometry, of which many of the following mathematical concepts are based, a vector function $f : \mathbb{R}^n \to \mathbb{R}^n$ will be denoted a *vector field* in $\mathbb{R}^n$.

The Lie bracket $[f, h]$ can also be expressed as $ad_f h$ (where ad stands for "adjoint"), which is useful for defining Lie brackets recursively [40]:

$$\begin{aligned} ad_f^0 h &= h(x) \\ ad_f^k h &= [f, ad_f^{k-1} h] \end{aligned} \tag{2.23}$$

In contrast to the Lie derivative, the Lie bracket describes the directional derivative of the entire vector field h with respect to the vector field f.

2.3.1.3 Diffeomorphisms and State Transformations

A diffeomorphism is a generalization of the concept of a coordinate transformation. Specifically, a vector field ϕ in $\mathbb{R}^n$ that is defined within a region Ω is a diffeomorphism if it is smooth, and its inverse ϕ^{-1} exists and is also smooth. If the region Ω spans all of $\mathbb{R}^n$, then ϕ is a global diffeomorphism, though these are rare [40].

To determine whether a nonlinear function $\phi(x)$ is a (local) diffeomorphism, it suffices to show that its Jacobian $\partial\phi/\partial x$ is non-singular at a point $x = x_0$ in Ω. That is, if all of the gradients of the elements in ϕ (i.e. $\partial\phi_i/\partial x$) are linearly independent at the point $x = x_0$, then $\phi(x)$ defines a (local) diffeomorphism in a neighborhood of x_0 in Ω [40]. (Hence, if this neighborhood, and thus Ω, can be extended to include all $x \in \mathbb{R}^n$, then $\phi(x)$ defines a global diffeomorphism.)

A diffeomorphism can be used to carry out a state transformation of a nonlinear system. Consider the diffeomorphism $\phi(x)$ and the nonlinear system of (2.16). A change of variables $z = \phi(x)$ can be achieved by differentiating z with respect to time:

$$\dot{z} = \frac{\partial \phi}{\partial x}\dot{x} = \frac{\partial \phi}{\partial x}\left(f(x) + g(x)u\right) \tag{2.24}$$

and substituting $x = \phi^{-1}(z)$ [40].

2.3.1.4 The Frobenius Theorem and Involutivity

The Frobenius theorem is an important tool that provides necessary and sufficient conditions for the solvability of a specific class of partial differential equations (PDEs) [40]. The relevance of the Frobenius theorem will become apparent later when discussing the internal dynamics of an input-output linearized system, but a full treatment of this theorem won't be necessary. Instead, we will focus on the so-called *involutivity condition*.

Consider the following set of first-order PDEs:

$$\begin{aligned} \frac{\partial \psi}{\partial x} f(x) &= 0 \\ \frac{\partial \psi}{\partial x} g(x) &= 0 \end{aligned} \tag{2.25}$$

where f and g are known vector fields in $\mathbb{R}^n$ and $\psi(x)$ is an unknown scalar function. The Frobenius theorem states that a solution $\psi(x)$ exists if, and only if, there exist scalar functions $\alpha_1(x)$ and $\alpha_2(x)$ such that the Lie bracket of f and g can be expressed as a linear combination of the vector fields f and g [40], that is:

$$[f, g] = \alpha_1 f + \alpha_2 g \tag{2.26}$$

This is known as the *involutivity condition* on these vector fields. Note that this can be generalized to more than two equations/vector fields, in which case involutivity is defined as being able to express all possible Lie brackets formed from the set of vector fields as a linear combination of the original set of vector fields [40].

As an example, involutivity can be used to determine whether a system is full-state linearizable. Consider the system given in (2.16). For simplicity in this example, we will assume this system is SISO such that $f(x)$ and $g(x)$ are both column vectors[4]. It can be shown [40] that this system is full-state linearizable if, and only if, it is

[4]The MIMO conditions for full-state linearizability can be found in [42].

controllable[5] and the set $\{g, ad_f g, ..., ad_f^{n-2} g\}$ is involutive.

2.3.2 Full-State Linearization

Consider the MIMO system given in (2.16), where x is an n-dimensional state vector, u is an m-dimensional vector of inputs, y is an m-dimensional vector of outputs, g is an $n \times m$ matrix, and f and h are vector fields on the domain $D \subset \mathbb{R}^n$. Using the notation of Lie derivatives from Section 2.3.1.2, the kth derivative of output y_i is defined as:

$$y_i^{(k)} = L_f^k h_i(x) + \sum_{j=1}^{m} L_{g_j} L_f^{k-1} h_i(x) u_j \tag{2.27}$$

where g_j is the jth column of g corresponding to input u_j. If $L_{g_j} L_f^{k-1} h_i(x) \neq 0$ for at least one j, then the input u will be present in the above equation. When u first appears in the kth derivative of y_i, we call $\rho_i = k$ the individual relative degree of output i, and the total relative degree of the MIMO system is given by $\rho = \sum_{i=1}^{m} \rho_i$, $0 \leq \rho \leq n$ [40]. The relative degree characterizes the lowest-order derivative in which the input u appears in output y_i, and it is analogous to the difference between the number of poles and zeros for linear systems [40].

Applying (2.27) to each y_i gives the following input-output dynamics for this system [40]:

$$\begin{bmatrix} y_1^{(\rho_1)} & y_2^{(\rho_2)} & \cdots & y_m^{(\rho_m)} \end{bmatrix}^T = b(x) + E(x)u \tag{2.28}$$

where,

$$b(x) = \begin{bmatrix} L_f^{\rho_1} h_1(x) & L_f^{\rho_2} h_2(x) & \cdots & L_f^{\rho_m} h_m(x) \end{bmatrix}^T \tag{2.29}$$

[5] A generalized controllability condition for nonlinear (SISO) systems is that the vector fields $\{g, ad_f g, ..., ad_f^{n-1} g\}$ are linearly independent [40].

and,

$$E(x) = \begin{bmatrix} L_{g_1}L_f^{\rho_1-1}h_1(x) & L_{g_2}L_f^{\rho_1-1}h_1(x) & \cdots & L_{g_m}L_f^{\rho_1-1}h_1(x) \\ L_{g_1}L_f^{\rho_2-1}h_2(x) & L_{g_2}L_f^{\rho_2-1}h_2(x) & \cdots & L_{g_m}L_f^{\rho_2-1}h_2(x) \\ \vdots & \vdots & \ddots & \vdots \\ L_{g_1}L_f^{\rho_m-1}h_m(x) & L_{g_2}L_f^{\rho_m-1}h_m(x) & \cdots & L_{g_m}L_f^{\rho_m-1}h_m(x) \end{bmatrix} \tag{2.30}$$

is known as the decoupling matrix. Hence, to cancel the nonlinearities of the system in (2.16), we simply invert (2.28) to get the following state feedback control law:

$$u = E^{-1}(x)\big(v - b(x)\big) \tag{2.31}$$

where v is a new input vector generated by a linear controller acting on the following (decoupled) linearized system:

$$\begin{bmatrix} y_1^{(\rho_1)} & y_2^{(\rho_2)} & \cdots & y_m^{(\rho_m)} \end{bmatrix}^T = v \tag{2.32}$$

Note that, considering (2.31), it is therefore necessary that $E(x)$ be invertible.

Thus far, no assumptions have been made about the total relative degree of the system ρ. If we now consider the case in which $\rho = n$, the system is said to be full-state linearizable, which means the above input-output mapping allows for full control of all of the system's states and there will be no internal dynamics (i.e. dynamics that are unobservable by this state transformation). Hence, (2.32) implies that full-state linearizable systems will have the following state transformation $z = \phi(x)$:

$$z = \begin{bmatrix} \begin{bmatrix} y_1 & \dot{y}_1 & \cdots & y_1^{(\rho_1-1)} \end{bmatrix}^T \\ \vdots \\ \begin{bmatrix} y_m & \dot{y}_m & \cdots & y_m^{(\rho_m-1)} \end{bmatrix}^T \end{bmatrix} \tag{2.33}$$

which constitutes a diffeomorphism and results in the following linear state-space in Brunovsky form [39]:

$$\begin{aligned} \dot{z} &= Az + Bv \\ y &= Cz \end{aligned} \tag{2.34}$$

where,

$$A = \begin{bmatrix} A_1 & 0 & \cdots & 0 \\ 0 & A_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & A_m \end{bmatrix} \quad B = \begin{bmatrix} B_1 & 0 & \cdots & 0 \\ 0 & B_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & B_m \end{bmatrix} \quad C = \begin{bmatrix} C_1 & 0 & \cdots & 0 \\ 0 & C_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & C_m \end{bmatrix} \tag{2.35}$$

and,

$$A_i = \begin{bmatrix} 0_{(\rho_i-1)\times 1} & I_{(\rho_i-1)\times(\rho_i-1)} \\ 0 & 0_{1\times(\rho_i-1)} \end{bmatrix} \quad B_i = \begin{bmatrix} 0_{(\rho_i-1)\times 1} \\ 1 \end{bmatrix} \quad C_i = \begin{bmatrix} 1 & 0_{1\times(\rho_i-1)} \end{bmatrix} \tag{2.36}$$

A special case exists for systems that are full-state linearizable and $y = x$ (which also implies the individual relative degrees are all equal to 1). Frequently referred to as nonlinear dynamic inversion (NDI) within the aerospace community, these conditions simplify $b(x) = f(x)$ and $E(x) = g(x)$ in (2.29) and (2.30). For example, aircraft attitude rate dynamics often satisfy these criteria when arranged into the following input-affine form with virtual controls (generalized moments):

$$\begin{bmatrix} \dot{p} \\ \dot{q} \\ \dot{r} \end{bmatrix} = \begin{bmatrix} \dot{p}_{passive} \\ \dot{q}_{passive} \\ \dot{r}_{passive} \end{bmatrix} + \begin{bmatrix} \frac{-I_{zz}}{I_{xz}^2 - I_{xx}I_{zz}} & 0 & \frac{-I_{xz}}{I_{xz}^2 - I_{xx}I_{zz}} \\ 0 & \frac{1}{I_{yy}} & 0 \\ \frac{-I_{xz}}{I_{xz}^2 - I_{xx}I_{zz}} & 0 & \frac{-I_{xx}}{I_{xz}^2 - I_{xx}I_{zz}} \end{bmatrix} \begin{bmatrix} M_x \\ M_y \\ M_z \end{bmatrix} \tag{2.37}$$

where the subscript "passive" simply indicates the parts of the attitude rate dynamics that do not depend on the effector model (M_x, M_y, and M_z). Hence, this system can

be directly inverted to form the following NDI control law:

$$u = g^{-1}(x)\big(v - f(x)\big) \tag{2.38}$$

and the change of variables $z = \phi(x)$ simplifies to $z = x$ for this NDI system.

One variant of NDI is incremental nonlinear dynamic inversion (INDI). INDI exploits time-scale separation principles between the system's states and inputs while utilizing incremental control commands to replace $f(x)$ in (2.38) with current measurements of $\dot{x}$. This method can help increase robustness of the feedback linearization approach due to less reliance on the system model. However, it in turn has the drawback of relying on sufficiently accurate measurements of, typically, acceleration states.

2.3.3 Input-Output Linearization

We will now consider the case in which $\rho < n$. As no assumptions were made about the system relative degree when deriving the feedback linearization control law in Section 2.3.2, (2.28)–(2.32) remain valid. However, the state transformation $z = \phi(x)$ now becomes:

$$z = \begin{bmatrix} \mu \\ \psi \end{bmatrix} \tag{2.39}$$

where the vector μ represents the external (controllable) dynamics of the original system with corresponding linear state-space in Brunovsky form (see (2.35) and (2.36)):

$$\begin{aligned} \dot{\mu} &= A\mu + Bv \\ y &= C\mu \end{aligned} \tag{2.40}$$

and there will remain $n - \rho$ states ψ_i that are unobservable by this input-output transformation. These states, known as the internal dynamics, are significant as their

lack of controllability/observability means they must remain stable on their own in order for the overall input-output linearized system to be stable.

To find the internal dynamics, we consider the fact that the state transformation $\phi(x)$ must be a valid diffeomorphism, which means that ϕ^{-1} must exist (see Section 2.3.1.3). Hence, it is necessary that ψ be chosen such that $\phi(x)$ is invertible. Fortunately, it can be shown [42] that it is always possible to find $n - \rho$ functions ψ_i that ensure the mapping $\phi(x)$ has a Jacobian matrix that is non-singular at some $x = x_0$, thus making $\phi(x)$ a local diffeomorphism in a neighborhood of x_0 [40]. Note that this means that ψ in itself does not have any prescribed form but simply indicates that the gradients of the elements of ψ and μ must be linearly independent. However, if it turns out that the set of all g_j is involutive near x_0, then it is possible to choose ψ_i such that $L_{g_j}\psi_i = 0$ (for all $1 \leq i \leq n - \rho$ and $1 \leq j \leq m$), which physically means that the input u will not appear in the internal dynamics. Note that for MIMO systems, it is not uncommon that the set of all g_j is *not* involutive. In contrast for SISO systems, since there is only one g_j, involutivity of the set g_j (i.e. a single column vector) is *always* satisfied, so $L_{g_j}\psi_i = 0$ is expected [42]. For the general case (i.e. without assuming involutivity), the internal dynamics can be expressed generically in the following way [42]:

$$\dot{\psi} = q(\mu, \psi) + p(\mu, \psi)u \tag{2.41}$$

(2.41) above and (2.40) with the substitution of $v = b(x) + E(x)u$ and $x = \phi^{-1}(\mu, \psi)$ together compose the *normal form* of the system's dynamics in the so-called "normal coordinates" μ and ψ.

To determine the stability of the internal dynamics, it suffices to show that the system's so-called "zero-dynamics", i.e. the internal dynamics when the system's outputs are held constant at zero, are stable [40]. Analyzing the zero-dynamics is generally more tractable than analyzing the internal dynamics because the zero-dynamics define an intrinsic property of the system (a differential equation) that only depends on the

internal states ψ, while the internal dynamics generally also depend on the external states μ and possibly the inputs u [40]. To find the zero-dynamics, we consider the constraint that the system's outputs are all held constant at zero, which implies that all of their time derivatives are zero (this equivalently implies $\mu = \dot{\mu} = 0$); hence, for a system to operate in the zero-dynamics, the initial external states must be $\mu_0 = 0$ and the input u must be such that it maintains the condition $y_j^{(\rho_j)} = 0$ (for all j). From (2.28), this means u must be [42]:

$$u_0 = -E^{-1}(0, \psi)b(0, \psi) \tag{2.42}$$

where $x = \phi^{-1}(\mu, \psi)$ has been used and $E(\mu, \psi) := E(\phi^{-1}(\mu, \psi))$, etc. Substituting $\mu = 0$ and $u = u_0$ into (2.41) thus gives the system's zero-dynamics [42]:

$$\dot{\psi} = q(0, \psi) - p(0, \psi)E^{-1}(0, \psi)b(0, \psi) \tag{2.43}$$

which can be seen to only depend on the internal states ψ. Note that in the case in which the set of all g_j is involutive and ψ is found that satisfies $L_{g_j}\psi_i = 0$ (for all i and j), then the zero-dynamics reduce to $\dot{\psi} = q(0, \psi)$. However, this can be interpreted as simply working in a different set of coordinates and is not strictly necessary.

It is important to note that while stability of the zero-dynamics indicates stability of the internal dynamics, it only does so *locally*; that is, even if the zero-dynamics are found to be globally stable, the internal dynamics are only guaranteed to be locally stable. Lastly, general nonlinear systems with (asymptotically) stable zero-dynamics are known as (asymptotically) minimum phase. This distinction follows from the linear analog, where it can be shown [40] that the poles of the zero-dynamics of a linear system are exactly the zeros of that system. Hence, if an input-output linearized system is asymptotically minimum phase, then it is said to be locally asymptotically stable. In contrast, full-state linearizable systems are always locally asymptotically

stabilizable by virtue of being fully controllable.

2.4 Control Allocation for Over-Actuated Systems

Many transition VTOL configurations have multiple propulsors and control surfaces, all of which can generate moments and forces on the aircraft. This typically results in an over-actuated mechanical system. This is of course desirable, as it allows for flexibility in the control design and can help maintain full operability in light of actuator failure, an important safety consideration for VTOL aircraft. The problem of distributing the control effort across such a set of redundant actuators is called control allocation. Control allocation works by connecting high-level motion controllers producing virtual control inputs (typically generalized forces and moments) to low-level physical actuator inputs [43]; as such, the high-level motion controllers need not know the details of which actuators are to be used given a particular flight mode or failure condition.

2.4.1 General Problem Formulation

The primary objective of the control allocation is to determine the real actuator inputs $\mathbf{u}$ which jointly produce the commanded virtual controls $\boldsymbol{\tau}_c$ [43]. It is possible that this objective cannot be met if the commanded virtual controls are greater than the physically producible forces and moments $\boldsymbol{\tau} = \mathbf{f}_B(\mathbf{x}, \mathbf{u})$ of the real actuators. If a feasible $\mathbf{u}$ cannot be found such that $\boldsymbol{\tau} = \boldsymbol{\tau}_c$, then the control allocation must resort to minimizing the allocation error $\boldsymbol{\tau}_c - \boldsymbol{\tau}$ instead, which usually results in some loss of performance. This can be set up as the following general optimization problem

[43]:

$$\begin{aligned} \min_{\mathbf{u}} \quad & \|\mathbf{Qs}\| \\ \text{s.t.} \quad & \mathbf{s} = \boldsymbol{\tau}_c - \mathbf{f}_B(\mathbf{x}, \mathbf{u}) \\ & \mathbf{u}_{min} \leq \mathbf{u} \leq \mathbf{u}_{max} \\ & \dot{\mathbf{u}}_{min} \leq \dot{\mathbf{u}} \leq \dot{\mathbf{u}}_{max} \end{aligned} \tag{2.44}$$

where $\mathbf{s}$ is some slack variable (representing the allocation error) and $\mathbf{Q}$ is a weighting matrix. For over-actuated systems, $\dim(\mathbf{u}) > \dim(\boldsymbol{\tau})$ and so (2.44) does not have a unique solution. Therefore, it is possible to introduce secondary objectives $J(\mathbf{x}, \mathbf{u})$, such as minimizing energy or redistributing the weighting applied to particular actuators (e.g. in the case of failure), into the above cost function.

2.4.2 Techniques for Linear Control Effector Models

Most control allocation algorithms assume a linear control effector model of the form $\boldsymbol{\tau} = \mathbf{Bu}$, where $\mathbf{B}$ is a matrix (typically called the control effectiveness matrix) that maps the inputs $\mathbf{u}$ to the physically producible forces and moments $\boldsymbol{\tau}$ [43]. Hence, if $\boldsymbol{\tau} = \boldsymbol{\tau}_c$ is given as an input to the control allocation from a high-level motion controller, the control effectiveness matrix $\mathbf{B}$ can be simply inverted to determine the real actuator inputs required to produce the demanded generalized forces and moments. The main challenge, however, is that $\mathbf{B}$ is not a square matrix (for over-actuated systems), so directly inverting the linear model is not possible.

The most common way to deal with this is to use generalized inverses (also called pseudo-inverses) [44]. These can be formulated in the context of a quadratic program (QP) [43]:

$$\begin{aligned} \min_{\mathbf{u}} \quad & \frac{1}{2}(\mathbf{u} - \mathbf{u}_p)^T \mathbf{W} (\mathbf{u} - \mathbf{u}_p) \\ \text{s.t.} \quad & \boldsymbol{\tau}_c = \mathbf{Bu} \end{aligned} \tag{2.45}$$

where $\mathbf{W}$ is a positive definite weighting matrix and $\mathbf{u}_p$ is the preferred value of $\mathbf{u}$. In

the case that $\mathbf{u}$ is unconstrained and $\mathbf{B}$ is full rank, (2.45) has the following explicit solution:

$$\begin{aligned} \mathbf{u} &= (\mathbf{I} - \mathbf{CB})\mathbf{u}_p + \mathbf{C}\boldsymbol{\tau}_c \\ \mathbf{C} &= \mathbf{W}^{-1}\mathbf{B}^T(\mathbf{B}\mathbf{W}^{-1}\mathbf{B}^T)^{-1} \end{aligned} \tag{2.46}$$

where $\mathbf{C}$ is a generalized inverse of $\mathbf{B}$. This solution can be derived from optimality conditions using Lagrange multipliers [45, 46], and for the case that $\mathbf{W} = \mathbf{I}$ and $\mathbf{u}_p = 0$, the solution reduces to the Moore-Penrose pseudo-inverse [44, 47, 48].

While computationally convenient that there is an explicit solution, the unconstrained nature of this problem does not guarantee that $\mathbf{u}$ will lie within the physical limits of the actuators even if the demanded generalized forces and moments are physically attainable [43]. Specifically, if $\mathbb{A}$ represents the set of attainable forces and moments of the control effectors (known as the attainable set), then the unconstrained pseudo-inverse could demand values of $\mathbf{u}$ outside its physical bounds even if $\boldsymbol{\tau}_c \in \mathbb{A}$. (It is worth mentioning that if $\boldsymbol{\tau}_c \notin \mathbb{A}$, then $\mathbf{u}$ will always lie outside its physical bounds.) A few methods are available to deal with such cases, though others exist: on-line constrained optimization, direct allocation, the redistributed pseudo-inverse, and daisy chaining.

An on-line constrained optimization can be set up in many ways, but in a linear context it can be formulated as the same QP as in (2.45) but with the addition of actuator limits as in (2.44). Hence, provided $\boldsymbol{\tau}_c \in \mathbb{A}$, a solution will be found. However, this on-line constrained QP will of course still break down if this is not the case (i.e. if $\boldsymbol{\tau}_c \notin \mathbb{A}$). Said another way, the problem must be feasible.

To ensure a solution is still found when $\boldsymbol{\tau}_c \notin \mathbb{A}$, one can resort to minimizing the allocation error $\boldsymbol{\tau}_c - \boldsymbol{\tau}$ instead. One such method to do this is direct allocation. Proposed by Durham [48], this method searches for the largest feasible solution $\mathbf{Bu}$ such that the direction of $\boldsymbol{\tau}_c$ is preserved (the importance of preserving MIMO input

vector directionality is discussed in [49]):

$$\begin{aligned} \max_{\alpha} \quad & \alpha \\ \text{s.t.} \quad & \alpha \boldsymbol{\tau}_c = \mathbf{B}\mathbf{u} \\ & \alpha \boldsymbol{\tau}_c \in \mathbb{A} \end{aligned} \tag{2.47}$$

where α is a scalar between 0 and 1. Unfortunately, solving this problem can be quite challenging for systems in which $\dim(\mathbf{u})$ is large due to the significant amount of facets and vertices composing the boundary of $\mathbb{A}$ [43]. Many numerical algorithms have been proposed as an improvement over (2.47), all with varying computational complexities [43, 47, 50, 51, 52, 53, 54].

To avoid any on-line optimization altogether (which could be desirable for computational robustness), a different modification from the original unconstrained pseudo-inverse is the redistributed pseudo-inverse. The redistributed pseudo-inverse, discussed by Virnig and Bodden [55] and Shi et al. [56], first solves the original unconstrained control allocation problem, and if the solution satisfies the actuator constraints then no further steps are required. However, if the solution $\mathbf{u}$ lies outside the physical bounds of the system, then the elements of $\mathbf{u}$ that violate the constraints must be saturated to form the admissible control vector $\overline{\mathbf{u}}$. To minimize the allocation error, a reduced control allocation is then solved involving only the elements of $\overline{\mathbf{u}}$ that are unsaturated and the remaining allocation error. This procedure is repeated until either a feasible solution is found or no improvement to the allocation can be made. While often effective, this method does not guarantee that a feasible solution will be found nor that the allocation error will be minimized [43].

A simpler alternative to the redistributed pseudo-inverse is the daisy chaining method, though it is often less effective [43]. This method works by grouping actuators into ranked groups and solving the control allocation problem separately. If any actuators in a higher-priority grouping saturates, that group freezes and the allocation

error is distributed among the second group, and so on [57]. Unfortunately, this may lead to sub-optimal solutions, in which certain actuators are not fully utilized [43].

2.4.3 Techniques for Nonlinear Control Effector Models

Nonlinear control allocation methods can give more exact solutions and may be necessary to achieve desired performance depending on the complexity of the effector model itself [43]. However, with increased modeling complexity comes increased computational complexity, and some nonlinear control allocation methods can prove too computationally inefficient depending on the system. Hence, a balance must be found between accuracy and performance.

The most general approach to control allocation for a nonlinear control effector model is to solve the general control allocation optimization problem of (2.44) [58, 59]. However, for any on-line optimization solution, a non-convex control effector model $\mathbf{f}_B(\mathbf{x}, \mathbf{u})$ could render the computation time too slow for real-time applications. One method to speed up computation time is to remove the actuator constraints in (2.44) and solve an unconstrained version of the problem. With actuator constraints removed, this optimization becomes effectively equivalent to finding the roots of the function $\mathbf{s}(\mathbf{u}) = \boldsymbol{\tau}_c - \mathbf{f}_B(\mathbf{x}_0, \mathbf{u})$, where $\mathbf{x}_0$ indicates the system's states at the current time step. A nonlinear root finding method can then be applied, such as Newton-Raphson or Levenberg-Marquardt within Matlab's `fsolve()` [60, 61, 62], which will typically find a solution faster than a constrained non-convex optimization. The drawback, of course, is that this is an unconstrained solution, so it is not guaranteed that $\mathbf{u}$ will be admissible.

Johansen et al. [58] proposed that a nonlinear program (NLP) as in (2.44) could be solved more efficiently by locally approximating the cost function by a quadratic cost function and linearizing the constraints. Johansen and Fossen [43] describe this as similar to sequential quadratic programming (SQP), though only one linear/quadratic

approximation is made per iteration. This method was found to have decent results with the example given in [58]. However, it cannot be generalized to arbitrary nonlinear systems, as several sequential linear/quadratic approximations may be required per iteration to achieve the necessary accuracy for systems that see large changes in allocated forces at each time step [59], and this could have implications on computation time in a real-time application.

As a simpler and more computationally-efficient alternative, local Jacobian linearizations can be applied to the nonlinear control effector model $\mathbf{f}_B(\mathbf{x}, \mathbf{u})$ such that:

$$\mathbf{B} = \left.\frac{\partial \mathbf{f}_B(\mathbf{x}, \mathbf{u})}{\partial \mathbf{u}}\right|_{(\mathbf{x}_0, \mathbf{u}_0)} \tag{2.48}$$

Various linear techniques, such as those described in Section 2.4.2, can then be applied as long as the operating point of the system is within the region in which the local linear approximation of $\mathbf{B}$ is sufficiently accurate. This is the approach that will be taken in this thesis, where the chosen linear technique will be a weighted pseudo-inverse algorithm due to its ease of implementation, computational efficiency, pre-determined closed-form solution, and ability to easily deal with actuator failures/faults.

2.5 Trajectory Optimization Methods

Trajectory optimization is a type of optimal control problem that looks for state and input trajectories which minimize a given cost (performance index) subject to constraints. In the context of tiltwing aircraft, this is not only useful for *optimality* (e.g. minimizing energy usage) but also for *feasibility* due to the narrow flight corridor in which such aircraft are able to operate during transition. A trajectory optimization problem can be formulated mathematically as the following general optimal control

problem:

$$\begin{aligned} \min_{t,\mathbf{u}(t)} \quad & J\left(t, \mathbf{x}(t), \mathbf{u}(t)\right) = \int_{t_0}^{t_F} L\left(t, \mathbf{x}(t), \mathbf{u}(t)\right) dt \\ \text{s.t.} \quad & \dot{\mathbf{x}} = \mathbf{f}\left(t, \mathbf{x}(t), \mathbf{u}(t)\right) \\ & \mathbf{x}_{min} \leq \mathbf{x}(t) \leq \mathbf{x}_{max} \\ & \mathbf{u}_{min} \leq \mathbf{u}(t) \leq \mathbf{u}_{max} \\ & \mathbf{C}_{P,min} \leq \mathbf{C}_P\left(t, \mathbf{x}(t), \mathbf{u}(t)\right) \leq \mathbf{C}_{P,max} \\ & \mathbf{C}_{B,min} \leq \mathbf{C}_B\left(t_0, t_F, \mathbf{x}(t_0), \mathbf{x}(t_F)\right) \leq \mathbf{C}_{B,max} \end{aligned} \tag{2.49}$$

where L is a quantity to be minimized in some sense, J is the cost functional involving L, and $\mathbf{C}_P$ and $\mathbf{C}_B$ are the path and boundary constraints, respectively.

While analytical methods may be applied in simple problems (e.g. infinite-horizon LQR [63]), optimal control problems almost always must be solved numerically due to the complexity of most applications [64]. Numerical methods for solving such problems can be divided into two major classes: indirect methods and direct methods [65]. Indirect methods involve first employing the calculus of variations to determine first-order optimality conditions, which then results in a multi-point boundary value problem to be solved. In contrast, direct methods solve the original problem directly (as the name suggests) by applying a process known as transcription, in which the above optimal control problem is discretized and converted to a finite-dimensional NLP to be solved using well-known optimization techniques [64]. As the latter method can be applied without explicitly deriving the necessary optimality conditions, this makes such an approach appealing for general purpose applications [65].

Transcription for optimal control problems is commonly achieved through a technique known as direct collocation. Direct collocation involves discretizing the problem and approximating the state and input trajectories as piecewise polynomials (i.e. splines) [64]. The constraints in the continuous optimal control problem (e.g. the differential equation defining the system's dynamics) are then enforced at specified

points known as collocation points [66]. These collocation points as well as the order of the approximating polynomials are determined based on a chosen quadrature rule (e.g. trapezoidal, Hermite-Simpson, etc.) employed by the collocation [67, 68]. A subset of direct collocation is orthogonal collocation, where the collocation points are chosen to be the roots of a family of orthogonal basis functions, such as Chebyshev or Legendre polynomials [64]. The advantage of orthogonal collocation methods is that the quadrature approximation of the definite integral becomes significantly more accurate [64].

Collocation methods historically have been implemented in two common ways: local collocation and global collocation [64, 69]. In local collocation, the time interval is split into multiple segments (a mesh), and within each segment the state is approximated using a polynomial of the same fixed degree. This so-called h method achieves convergence by varying the number of segments in the mesh, and direct collocation methods for optimal control problems were originally developed in this way [66, 70]. In contrast, global collocation (known as a p method [69]) uses a fixed mesh (i.e. a single segment) and instead achieves convergence by varying the degree of the polynomial. When orthogonal collocation is used globally, it is known as *pseudospectral* and has the desirable characteristic of converging at an exponential rate as a function of the number of collocation points/the order of the polynomial (for problems whose solutions are smooth and well-behaved) [64, 68].

More recently, attention has shifted to so-called hp methods, a hybrid between h and p methods that operates by varying both the mesh size and the polynomial degree [66, 69, 70, 71]. In an hp method, the mesh is adjusted adaptively, allowing potential discontinuities to be dealt with at the boundary of the mesh intervals (h part), while smoothness of the solution is exploited between the mesh intervals to take advantage of exponential convergence in these regions (p part) [70]. Hence, such a method achieves convergence with a significantly smaller finite-dimensional approximation

than h or p methods individually and thus improves the computational efficiency of solving the NLP [71]. An example trajectory optimization software that makes use of this kind of method is GPOPS-II [70]. GPOPS-II is a general-purpose optimal control problem solver that employs hp-adaptive Legendre-Gauss-Radau quadrature orthogonal collocation [70, 71] to approximate the continuous optimal control problem as a sparse NLP, which is then solved using the NLP solver IPOPT.

An alternative approach to solving optimal control problems that foregoes the need for general-purpose solvers is successive convexification [72, 73, 74, 75]. This method works by representing the nonlinear dynamics with difference-of-convex (DC) functions, which is possible to do when the dynamics are sufficiently regular [75]. These dynamics can then be successively linearized around predicted trajectories, while linearization errors can be treated as bounded disturbances [76], allowing for robust tube-based model predictive control (MPC) methods to be used for trajectory generation. Such an approach offers faster convergence and guaranteed convergence to a locally optimal solution, which in turn can enable real-time operation and additionally make it possible to implement directly as an MPC control law, if desired. However, this comes at the cost of a significant increase in the complexity of the initial problem setup.

For this thesis, trajectory optimization will be used to understand feasible minimum-energy trajectories for a tiltwing aircraft and to provide outer-loop reference inputs for a feedback linearization-based controller design. Hence, the computational efficiency and ability to integrate an MPC-based control law provided by successive convexification methods is not necessary in this context, and so the general-purpose solver GPOPS-II will be the main tool for trajectory optimization used in this research.

2.6 Active Flow Control Applied to Aircraft

Active flow control (AFC) is a rapidly advancing field of fluid mechanics concerned with favorably manipulating the fluid flow field of engineering fluid systems to improve efficiency and performance. Flow control strategies (both active and passive) are generally employed to delay/advance flow transition, suppress/enhance turbulence, or prevent/provoke flow separation in order to reduce drag, augment lift, and/or suppress flow-induced noise [77]. AFC involves introducing energy into the flow field and can be used to control complex dynamical processes (e.g. turbulence production in turbulent boundary layers) by controlling natural stabilities of the flow; because flow control is most effective when the control input is introduced locally at a high receptivity region, AFC methods can attain large effects using small, localized energy inputs [78]. Example AFC methods include steady/unsteady suction or blowing and the use of synthetic jets.

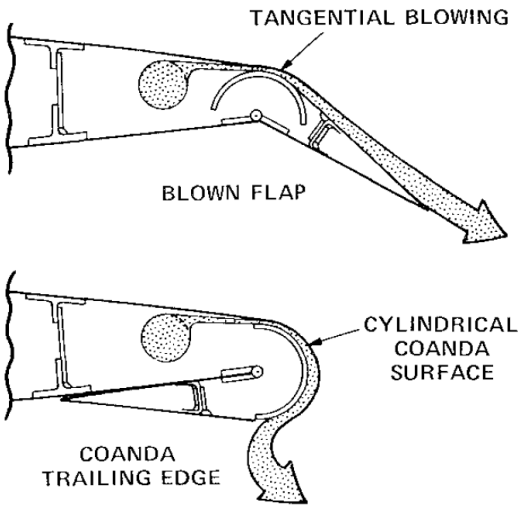

Figure 2.3: Example circulation control airfoil with a blown flap that forms a cylindrical trailing-edge. From [79]; reprinted by permission of the American Institute of Aeronautics and Astronautics, Inc.

Two useful aircraft applications of AFC are separation control and circulation control of an airfoil. Separation control is the ability to delay flow separation by ensuring near-wall fluid has sufficient momentum to continue its motion, thus increasing the

airfoil's effective stall angle of attack (AoA) [77]. This is commonly achieved either through leading- or trailing-edge suction/blowing or both [77]. Circulation control, on the other hand, is a method of generating lift using the Coanda effect that utilizes tangential jet blowing on, typically, the trailing edge of an airfoil [77], which is often made circular/rounded (such as in Figure 2.3) to promote such an effect [79, 80]. Both of these techniques typically involve the usage of non-zero mass-flux actuators, such as steady jets, pulsed jets, and sweeping jets (oscillatory blowing).

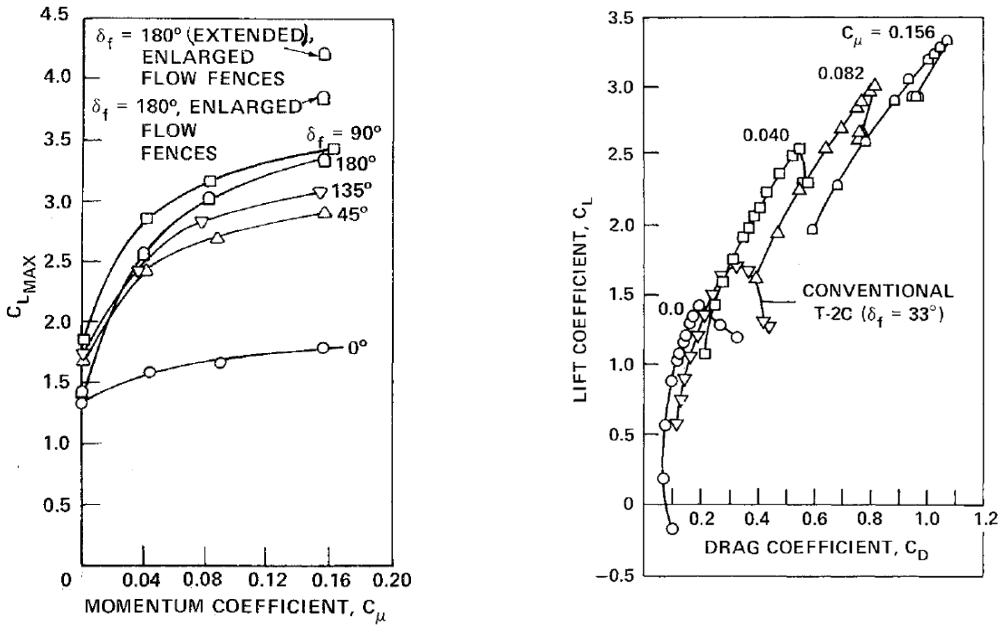

(a) Max C_L vs momentum coefficient at various flap deflection angles [79]

(b) Drag polars for 180° flap deflection [79]

Figure 2.4: Lift and drag curves for the circulation control airfoil of Figure 2.3. From [79]; reprinted by permission of the American Institute of Aeronautics and Astronautics, Inc.

Looking more closely at circulation control, an array of research suggests it is theoretically possible to modify the lift coefficient on an airfoil by as much as a factor of 2 [79, 80, 81, 82, 83, 84]. In [79], wind tunnel investigations on both two-dimensional and three-dimension airfoils employing tangential upper-surface blowing over a cylindrically-shaped trailing edge (Figure 2.3) demonstrated lift coefficient gains of double to triple that of a conventional flapped airfoil, as shown in Figure 2.4a. In [81] and [82], flight testing with a full-scale A-6A demonstrator aircraft

developed in collaboration with Grumman verified the results of [79], where it was found that doubling of the lift coefficient compared to a conventional flapped A-6A at the same AoA was possible. The works of Loth and Boasson [80] and McLachlan [84] further reinforce these potential lift coefficient gains. For example, in [84], the author carried out an experimental study of a circulation control airfoil with leading- and trailing-edge blowing, furthering experimental findings from [85]. A two-dimensional, uncambered airfoil was placed in a water tunnel and all combinations of steady leading- and trailing-edge blowing were measured at a fixed AoA of 0°. It was found that trailing-edge blowing was particularly effective at controlling the section lift coefficient, where it was possible to reach values of nearly 3 before lift unsteadiness occurred.

Considering the potential benefits such AFC techniques could have for an aircraft configuration such as a tiltwing, where the wings are pushed to their aerodynamic limits during a maneuver like transition, these concepts will be employed in this thesis when investigating optimal transition trajectories to provide insight into the theoretical aerodynamic requirements of a tiltwing during transition.

2.7 Uncertainty Modeling and Robustness Analysis

The mathematical model of a system is never perfect. Whether due to intentional modeling assumptions or a lack of understanding of the system's underlying behavior (to some degree), there will always remain unmodeled dynamics that are present in any real system. In the context of control system design, it is therefore necessary that a controller be sufficiently robust to deal with such modeling uncertainties (i.e. ensure the closed-loop system remains stable). Note that the bold notation for vectors and matrices is not used in this section.

2.7.1 Representing Uncertainty

Model uncertainty can be grouped into broadly two classes: parametric (real-valued) uncertainty and dynamic (frequency-dependent) uncertainty [86]. Parametric uncertainty describes a system model in which its structure (e.g. transfer function) is known but certain (real) parameters within the model are not. In contrast, dynamic uncertainty describes a system that contains unknown dynamics (typically at high frequencies), and this type of uncertainty is therefore usually complex-valued. In either case, uncertainty is commonly represented in a normalized form by some uncertain Δ, where $|\Delta| \leq 1$ for parametric (real) uncertainty and $\|\Delta\|_\infty \leq 1$ for dynamic (frequency-dependent) uncertainty.

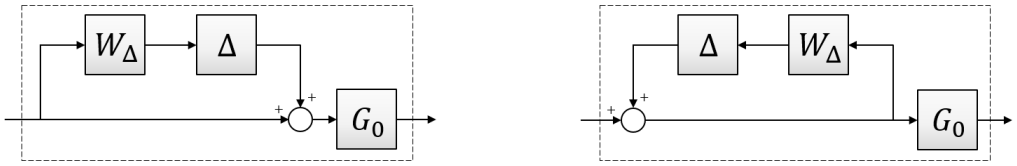

(a) Multiplicative uncertainty

(b) Inverse multiplicative uncertainty

Figure 2.5: Two types of uncertainty representations for an uncertain system G (represented at the system input)

There are many ways in which uncertainty can be modeled in a system. Two examples are given in Figure 2.5, which shows multiplicative uncertainty (Figure 2.5a) and inverse multiplicative uncertainty (Figure 2.5b)[6]. Multiplicative (input) uncertainty results in an uncertain system transfer function G of the form $G = G_0(I + W_\Delta\Delta)$, where G_0 represents the nominal system and W_Δ is a weighting transfer function that captures the magnitude of the (possibly frequency-dependent) uncertainty. As can be seen by the form of G, this type of uncertainty representation is convenient for modeling gain uncertainty in the system. In contrast, inverse multiplicative (input) uncertainty results in a system of the form $G = G_0(I - W_\Delta\Delta)^{-1}$ and is therefore

[6]Note that both of these examples show the uncertainty represented at the system input. Alternatively, uncertainty may be represented at the system output. This does not make a difference for SISO systems (scalar transfer functions) but does matter for MIMO systems.

useful for modeling pole uncertainty. Note that the negative sign in this expression does not matter since Δ can be positive or negative.

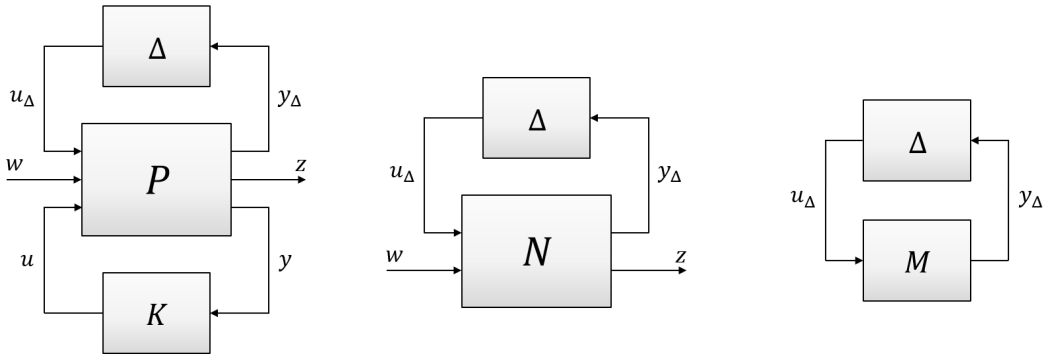

(a) General control configuration (b) $N\Delta$-structure (c) $M\Delta$-structure

Figure 2.6: General block diagram configurations with Δ uncertainty

There exist three standard block diagram arrangements that are useful for general controller synthesis and robust stability analysis. These are shown in Figure 2.6. In the general control configuration of Figure 2.6a, the generalized (open-loop) plant P is in feedback with a controller K and uncertainty Δ. P is a block transfer matrix from $[u_\Delta \quad w \quad u]^T$ to $[y_\Delta \quad z \quad y]^T$, where u and y are the plant inputs and outputs, respectively, w and z represent any exogenous inputs/outputs to the system (such as reference inputs or performance quantities to be minimized), and u_Δ and y_Δ are perturbations due to the uncertainty Δ. By grouping the perturbations and exogenous inputs/outputs, P is commonly partitioned as:

$$\begin{bmatrix} \begin{bmatrix} y_\Delta \\ z \end{bmatrix} \\ y \end{bmatrix} = \begin{bmatrix} P_{11} & P_{12} \\ P_{21} & P_{22} \end{bmatrix} \begin{bmatrix} \begin{bmatrix} u_\Delta \\ w \end{bmatrix} \\ u \end{bmatrix} \tag{2.50}$$

By combining P and K in closed-loop, we end up with the $N\Delta$-structure of Figure 2.6b, where N is related to P and K through a lower linear fractional transformation (LFT), that is:

$$N = F_l(P, K) := P_{11} + P_{12}K(I - P_{22}K)^{-1}P_{21} \tag{2.51}$$

In a similar manner to P, N is commonly written in a partitioned form as:

$$\begin{bmatrix} y_\Delta \\ z \end{bmatrix} = \begin{bmatrix} N_{11} & N_{12} \\ N_{21} & N_{22} \end{bmatrix} \begin{bmatrix} u_\Delta \\ w \end{bmatrix} \tag{2.52}$$

The $N\Delta$-structure is useful for analyzing the robust performance of a system (i.e. ensuring stability *and* performance specifications are met in light of given model uncertainty) and is thus common in the context of controller synthesis. However, if only interested in analyzing robust stability, we will instead use the $M\Delta$-structure of Figure 2.6c, where $M = N_{11}$.

Lastly, the uncertainty Δ can be defined/represented as either *structured* or *unstructured*. Structured uncertainty means that the perturbations are restricted to certain directions/terms in the system and is most generally represented as a block-diagonal matrix, i.e. $\Delta = \text{diag}\{\Delta_i\}$. Furthermore, the blocks in Δ may be repeated and/or may be real-valued, meaning there is additional structure. In contrast, unstructured uncertainty means that perturbations are defined/exist in all possible directions and is thus represented as a "full" complex perturbation matrix Δ. The latter can be useful for obtaining a simple uncertainty model [86], but it can also come at the cost of increased conservativeness since the uncertainty present in physical systems generally tends to exhibit some kind of structure.

2.7.2 Robust Stability for Unstructured Uncertainty

Using the $M\Delta$-structure of Figure 2.6c, we can derive a robust stability condition for a system in the presence of an uncertainty Δ. In this section, we consider the case in which Δ is unstructured.

If we first consider a SISO system, the Nyquist stability condition requires that

the loop transfer function $M\Delta$ not encircle the -1 point for all Δ, that is:

$$|M(j\omega)\Delta(j\omega)| < 1, \quad \forall\omega \tag{2.53}$$

However, noting that $|M\Delta| \leq |M| \cdot |\Delta|$ and, by definition, $\|\Delta\|_\infty \leq 1 \Rightarrow |\Delta(j\omega)| \leq 1, \forall\omega$ for scalar $\Delta(s)$, it follows that robust stability is therefore guaranteed for:

$$|M(j\omega)| < 1, \quad \forall\omega \tag{2.54}$$

which is the familiar result of the small-gain theorem [86]. Note that this condition assumes that $\Delta(s)$ and $M(s)$ are stable transfer functions. The former is an assumption that is always made when defining a dynamic uncertainty $\Delta(s)$, while the latter implies nominal stability of the closed-loop system.

As an example, we will use (2.54) to derive a robust stability condition for an uncertain SISO system $G(s)$ in the presence of a multiplicative uncertainty. See Figure 2.7 below, which shows $G(s)$ in feedback with a controller K. This diagram can be used to derive the $M\Delta$-structure for this system by finding the transfer function from u_Δ to y_Δ, giving:

$$\frac{y_\Delta(s)}{u_\Delta(s)} = M(s) = \frac{W_\Delta(s)K(s)G_0(s)}{1 - K(s)G_0(s)} \tag{2.55}$$

Hence, robust stability is guaranteed for:

$$|W_\Delta(j\omega)T_0(j\omega)| < 1, \quad \forall\omega \tag{2.56}$$

where $T_0(s)$ is the nominal closed-loop transfer function.

In consideration of MIMO systems, (2.54) can be expressed more generally as [86]:

$$\|M\|_\infty < 1 \qquad \Leftrightarrow \qquad \bar{\sigma}(M(j\omega)) < 1, \quad \forall\omega \tag{2.57}$$

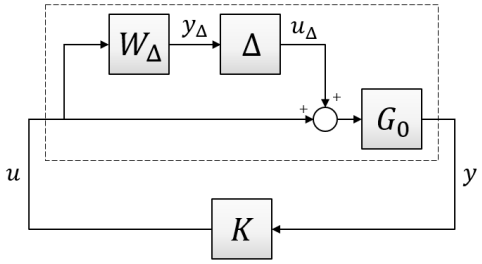

Figure 2.7: Uncertain system G represented by a multiplicative uncertainty W_Δ in feedback with a controller K. This diagram can be used to derive the $M\Delta$-structure for this system.

noting that the H_∞ norm of a general transfer matrix $M(s)$ is the supremum of its maximum singular value $\bar{\sigma}$.

2.7.3 The Structured Singular Value

The robust stability condition given in (2.57) provides an exact condition (i.e. necessary and sufficient) for systems with unstructured uncertainty (i.e. "full" perturbations). However, for many systems, the uncertainty Δ exhibits some sort of structure, which means that (2.57) is often conservative. We will now consider the *structured singular value*, which provides a more general and less conservative approach to proving robust stability by taking into account the structure of Δ.

The structured singular value (SSV), often denoted μ, is a function that provides a generalization of the maximum singular value $\bar{\sigma}$ and the spectral radius ρ [86]. The SSV for a complex matrix M is defined as:

$$\mu(M) := \frac{1}{\min\{k_m | \det(I - k_m M\Delta) = 0\}} \tag{2.58}$$

for structured Δ, $\|\Delta\|_\infty \leq 1$. If no such structured Δ exists, then $\mu(M) = 0$. In words, we are looking for the smallest scaling factor k_m such that a structured, normalized perturbation Δ causes the closed-loop system to go unstable[7]. Hence, a smaller μ is

[7]The condition $\det(I - k_m M\Delta) = 0$ comes from the generalized Nyquist theorem applied to the $M\Delta$-structure [86].

considered "good", while a larger μ is considered "bad", with $\mu = 1$ indicating that the full range of Δ is just large enough to destabilize the system. Note that computing $\mu(M)$ itself is generally difficult, so instead various analytical and numerical methods are often employed to find (tight) upper and lower bounds on its value.

To get a sense for some general (non-tight) upper and lower bounds on μ as well as its relationship to $\bar{\sigma}$ and ρ, the following turns out to be true [86]: 1) $\mu(M) = \rho(M)$ for $\Delta = \delta I$, where δ is a complex scalar; and 2) $\mu(M) = \bar{\sigma}(M)$ when Δ is a full (complex) matrix. It therefore follows that μ is bounded by ρ and $\bar{\sigma}$ for complex perturbations, that is:

$$\rho(M) \leq \mu(M) \leq \bar{\sigma}(M) \tag{2.59}$$

This is because Δ in the first case has the fewest degrees of freedom (less conservative), while Δ in the second case has the most degrees of freedom (more conservative).

The significance of the definition of the SSV is that it fundamentally depends on the structure of Δ itself. It can therefore be shown [86] that a necessary and sufficient condition for robust stability for a general system arranged in the $M\Delta$-structure is:

$$\mu(M(j\omega)) < 1, \quad \forall\omega \tag{2.60}$$

This can be thought of as a generalization of (2.57). Analyzing the robust stability of a system using the SSV is often referred to as μ-analysis.

2.7.4 H_∞ and μ-synthesis Control

Two control design approaches that utilize the robustness analysis techniques presented in the previous three subsections are H_∞ optimal control and μ-synthesis. H_∞ optimal control is the general problem of finding a feedback controller K that minimizes the H_∞ norm of $N = F_l(P, K)$, where N comes from the $N\Delta$-structure for the general control configuration (see Figure 2.6). In contrast, μ-synthesis involves

finding a feedback controller K that minimizes $\mu(N)$. There is currently no direct method to synthesize a μ-optimal controller [86], but a common approach that works for complex perturbations is known as DK-iteration, which involves solving a series of scaled H_∞ control problems. The μ-synthesis approach tends to be less conservative than H_∞ control but also tends to produce higher-order controllers, whereas the order of an H_∞ controller will always match that of the generalized plant $P(s)$. Contrasting other control approaches such as LQR, both of these control techniques take system robustness and modeling uncertainty explicitly into account during the design process. Algorithms for carrying out H_∞ and μ-synthesis design are available in commercial software packages such as MATLAB's Robust Control Toolbox. These make use of either Riccati or linear matrix inequality (LMI) based methods for solving the H_∞ optimal control problem. LMI-based methods have largely superseded the Riccati formulae due to their efficiency and scalability, a result of transforming the problem into a convex optimization.

Chapter 3

Tiltwing Aircraft Modeling and Dynamics

3.1 Introduction

This chapter presents the unique tiltwing aircraft configuration used in this thesis. An overview of the physical aircraft, which is a 3D-printed scaled demonstrator model of the Rolls-Royce eVTOL concept, is given in Section 3.2. Section 3.3 then presents the complete modeling of this aircraft's forces and moments to build a 6-DoF dynamical model, followed by a brief discussion of some practical implementation details. Validation of this model's aerodynamics is then carried out in Section 3.5, where stability derivatives computed through numerical linearization are compared to predicted values from three programs (AVL, XFLR5, and OpenVSP) that make use of the vortex-lattice and panel methods for aerodynamic calculations. Lastly, an analysis of the aircraft's dynamics is given in Section 3.6.

3.2 Overview of the Aircraft

The aircraft chosen for this research is a $1/6^{\text{th}}$ scale model of a unique tiltwing design—the Rolls-Royce eVTOL concept [87]—equipped with four propellers on the main wing and two on the tail, as shown in Figures 3.1 and 3.2 below, the original full-scale concept having been presented at Farnborough Airshow in 2018 (see Figure 1.3b). The purpose of the scaled model is to serve as a research platform for the theoretical control requirements of the full-scale concept, which has not yet flown. The aircraft's propellers and engine nacelles are fixed to the lifting surfaces, and both the main wing and tail rotate to orient the propellers vertically for hover mode and horizontally for cruise. This design places the aircraft's CG very close to the main wing, rendering its dynamic performance in cruise to be similar to that of a conventional fixed-wing aircraft [32]. Additionally, in consideration of the full-scale design, this allows for the wing motors to be optimized for a high thrust condition during hover flight, while the tail motors are optimized for cruise flight. Contrasting typical scale models commonly seen in research (e.g. see [88], [89], and [90]), this demonstrator aircraft is both geometrically and dynamically scaled. Specifically, Froude number was used to match gravitational and inertial force ratios[1], allowing for scaling of design parameters such as power, wing loading, and disk loading. This makes its dynamics and control-design challenges comparable to that of its full-scale counterpart, such as relatively low thrust margins and comparatively slow dynamics[2].

A note on low thrust margins: While off-the-shelf drones often have thrust-to-weight ratios of 2 or 3 : 1 (or higher) [91, 92], this demonstrator aircraft sees a thrust-to-weight ratio of around just 1.5 : 1. Such a low thrust margin means little

[1]Note that Reynolds number is therefore not matched between the scaled model and the full-scale aircraft.

[2]Specifically, this dynamic scaling allows for translating velocities and time scales between the scaled model and the full-scale aircraft to characterize the dynamics and test controllers. However, note that since Reynolds number is not matched, flow conditions will be fundamentally different.

control headroom, particularly in hover and transition, which fundamentally limits the achievable performance of the aircraft in terms of maneuverability and its ability to reject wind gust disturbances. ■

(a) Tiltwing indoors with harness attached

(b) Tiltwing outside ready for testing

Figure 3.1: Tiltwing Aircraft

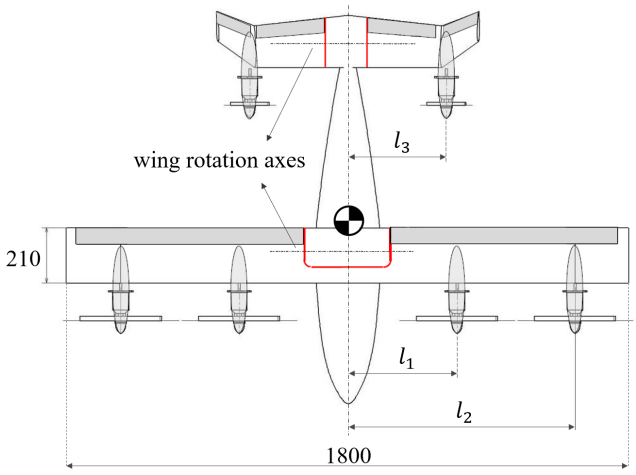

(a) Top View

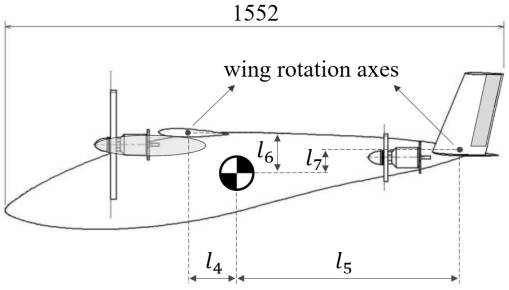

(b) Side View

Figure 3.2: Scaled Model Schematic (dimensions in mm)

Classical control surfaces on this aircraft include flaperons (combined ailerons and flaps), an elevator, and two rudders (one on each vertical tail). For this scaled demonstrator aircraft, each control surface is driven by a combination of servos, while each propeller is driven by an individual brushless motor. Additionally, two linear actuators (LACs) are used to drive the main wing and tail around a hinge to control

their rotation angles. In total, there are 12 control inputs (six props, four control surface deflection angles, and two wing angles). This scaled model is a fully 3D-printed prototype with a carbon fiber-reinforced wing structure that weighs 7.45 kg and has a wingspan of 1.8 m. It has a design cruise speed of 45 mph (20 $\mathrm{m\,s^{-1}}$), while the full-scale aircraft has a design cruise speed of around 150 mph (67 $\mathrm{m\,s^{-1}}$). When equipped with a fully-electric powertrain, the full-scale tiltwing is designed to carry a payload of 450 kg, including 4 passengers plus a pilot, and has a maximum take-off weight (MTOW) of 2600 kg. A summary of both the full-scale ("FS") and scaled model's ("SM") parameters and specifications is given in Tables 3.1–3.3 below. Note that the wing and tail prop disk areas are identical for the scaled model, contrasting the full-scale aircraft, simply due to hardware availability.

Table 3.1: Tiltwing Aircraft Specifications

Parameter	**Symbol**	**Value (SM)**	**Value (FS)**	**Units**
Mass	m	7.45	2600	kg
Total payload	–	n/a	450	kg
Number of passengers	–	n/a	4 + pilot	–
Wing area	S	0.285	15	$\mathrm{m^2}$
Wing prop disk area	$A_{prop,W}$	0.086	5.31	$\mathrm{m^2}$
Tail prop disk area	$A_{prop,T}$	0.086	3.14	$\mathrm{m^2}$
Number of wing props	n_W	4	4	–
Number of tail props	n_T	2	2	–
Max total thrust	T_{max}	122	56,000	N
Wing and tail angle range	$[\xi_{min}, \xi_{max}]$	$[0, 90]$	$[0, 90]$	deg
Cruise speed	V_{cruise}	45	150	mph

Table 3.2: Aircraft Prop Moment Arms (in m)

Moment Arm	**Value (SM)**	**Value (FS)**
l_1	0.353	2.17
l_2	0.752	4.65
l_3	0.382	2.34
l_4	0.158	0.60
l_5	0.658	4.60
l_6	0.20	0.99
l_7	0.130	0.54

Table 3.3: Moments of Inertia (in $\mathrm{kg\,m^2}$)

Inertia	**Value (SM)**	**Value (FS)**
I_{xx}	0.820	8914
I_{yy}	0.763	8294
I_{zz}	1.335	14512
I_{xz}	0.037	402

3.2.1 Assumed Mission Profile

Figure 3.3 illustrates the assumed mission profile for this aircraft that we will be designing for in this thesis. The aircraft begins with vertical take-off into hover. In this mode, the wings are kept fixed in a vertical orientation, and the primary control inputs are the six propulsors producing differential thrust demands to maneuver the aircraft. Additionally, the ailerons are employed to produce yawing moment demands by leveraging the prop wake over the (vertically-oriented) main wings. The elevator and rudders are not used in this flight mode. Next, the aircraft transitions from hover to cruise flight (forward transition). During this maneuver, all control inputs are employed, noting that wing angle demands are based on pre-determined reference trajectories. In cruise mode, the wings are fixed horizontally and the control inputs simplify greatly, consisting of classical control surfaces and a synchronized prop speed demand for all six propulsors[3]. After the cruise segment is completed, the aircraft transitions back to hover flight (backward transition), once again employing all control inputs during the maneuver. Finally, the aircraft hovers (same operation setup as before) until landing vertically, thus completing the mission.

[3]For the full-scale aircraft, the wing motors are instead shut off during cruise and only the tail motors operate for efficiency. This was not mirrored on the scaled demonstrator model due to hardware constraints.

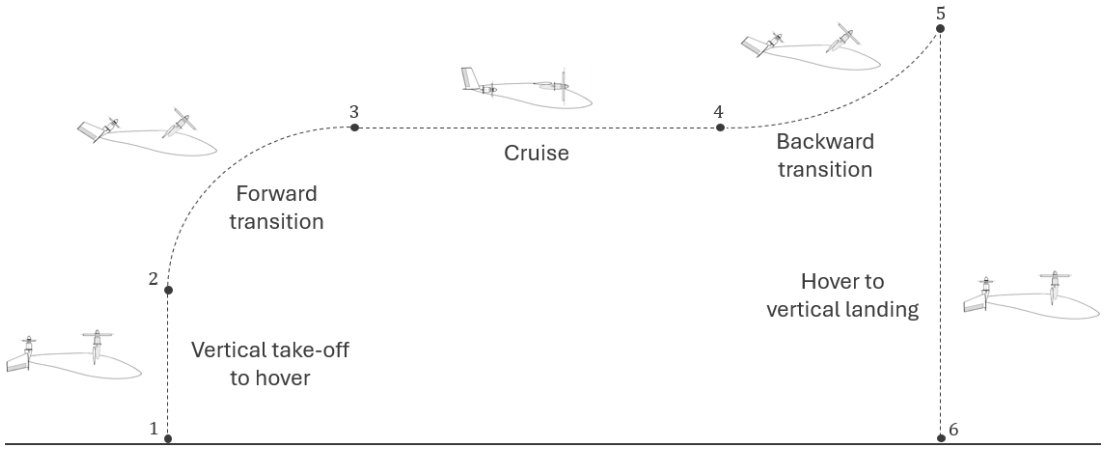

Figure 3.3: Assumed Mission Profile

3.3 Mathematical Modeling

In Section 2.2, the 6-DoF equations for a rigid body were derived for a generic aircraft. In this section, the net external forces and moments ($\mathbf{F}_{net}$ and $\mathbf{M}_{net}$) appearing in (2.6) will now be defined for this specific tiltwing configuration and its effector model. These forces and moments include: gravity ($\mathbf{F}_G$), thrust forces ($\mathbf{F}_T$), thrust moments ($\mathbf{M}_T$), reactions torques ($\mathbf{M}_Q$), gyroscopic moments ($\mathbf{M}_g$), and aerodynamic forces and moments ($\mathbf{F}_A$ and $\mathbf{M}_A$, respectively). Hence,

$$\begin{aligned} \mathbf{F}_{net} &= \mathbf{F}_G + \mathbf{F}_T + \mathbf{F}_A \\ \mathbf{M}_{net} &= \mathbf{M}_T + \mathbf{M}_Q + \mathbf{M}_g + \mathbf{M}_A \end{aligned} \tag{3.1}$$

Note that all coordinate frames used are as described in Section 2.2.1.

3.3.1 Gravity

Gravity is modeled in the inertial (Earth-fixed) frame as the vector $[0 \quad 0 \quad mg]^T$ which must then be transformed into the aircraft's body axes. This is accomplished by applying $\mathbf{T}_{be}$ from (2.4) and results in the following force vector in the aircraft's

body frame:

$$\mathbf{F}_G = \begin{bmatrix} -mg(s_\theta) & mg(s_\phi c_\theta) & mg(c_\phi c_\theta) \end{bmatrix}^T \tag{3.2}$$

where g is the acceleration due to gravity.

3.3.2 Propeller Forces and Moments

3.3.2.1 Thrust Forces

Due to the tilting wing and tail, thrust forces can be present in both the x- and z-directions of the aircraft's body frame:

$$\mathbf{F}_T = \begin{bmatrix} \sum T_{i,W} \cos(\xi_W) + \sum T_{i,T} \cos(\xi_T) \\ 0 \\ -\sum T_{i,W} \sin(\xi_W) - \sum T_{i,T} \sin(\xi_T) \end{bmatrix} \tag{3.3}$$

where T_i denotes the thrust of a single prop, subscripts W and T denote the wing and tail, respectively, and the angle ξ is the wing angle measured from the horizontal body axis as shown in Figure 3.4b.

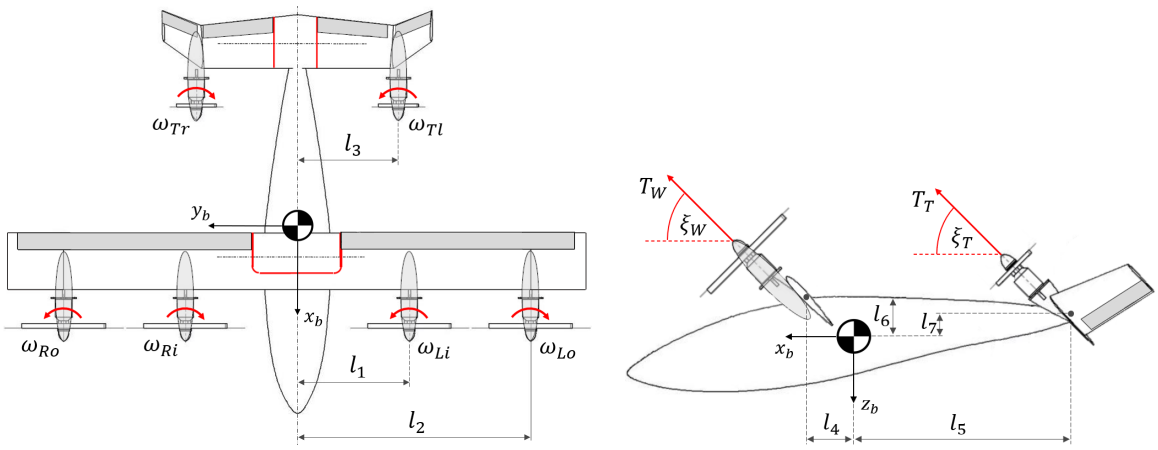

(a) Lateral arms and prop directions

(b) Longitudinal arms and wing angles

Figure 3.4: Tiltwing Moment Arms, Prop Directions, and Wing Angles

3.3.2.2 Thrust Moments

Non-zero moment arms l_1 to l_7 produce moments about the aircraft's CG during thrust generation (see Figure 3.4):

$$\mathbf{M}_T = \begin{bmatrix} (l_1(T_{Li} - T_{Ri}) + l_2(T_{Lo} - T_{Ro}))\sin(\xi_W) + l_3(T_{Tl} - T_{Tr})\sin(\xi_T) \\ \sum T_{i,W}\left(l_4\sin(\xi_W) - l_6\cos(\xi_W)\right) - \sum T_{i,T}\left(l_5\sin(\xi_T) + l_7\cos(\xi_T)\right) \\ (l_1(T_{Li} - T_{Ri}) + l_2(T_{Lo} - T_{Ro}))\cos(\xi_W) + l_3(T_{Tl} - T_{Tr})\cos(\xi_T) \end{bmatrix} \tag{3.4}$$

where subscript Li indicates the left-inner wing prop, Ro indicates the right-outer wing prop, Tl indicates the left tail prop, etc.

3.3.2.3 Reaction Torques

Reaction torques are a result of the aerodynamic loading on a rotating prop, which induces a torque of equal magnitude on the aircraft in the opposite direction:

$$\mathbf{M}_Q = \begin{bmatrix} (Q_{Lo} - Q_{Li} + Q_{Ri} - Q_{Ro})\cos(\xi_W) + (Q_{Tl} - Q_{Tr})\cos(\xi_T) \\ 0 \\ (-Q_{Lo} + Q_{Li} - Q_{Ri} + Q_{Ro})\sin(\xi_W) + (-Q_{Tl} + Q_{Tr})\sin(\xi_T) \end{bmatrix} \tag{3.5}$$

where Q denotes the motor torque required to drive the propeller at a given speed. Note that the signs chosen are based upon the rotational directions of the propellers.

3.3.2.4 Prop Thrust and Torque Modeling

In (3.3)–(3.5), the prop thrust and torque are generically defined by the variables T and Q, respectively. However, the input request to the motors' electronic speed controllers (ESCs) will be a pulse-width modulation (PWM) signal based on a demanded motor speed ω. To compute the prop thrust and torque from ω, quadratic correlations are used based on correlation coefficients C_i of this tiltwing's propellers

determined from wind tunnel testing [93]:

$$\begin{aligned} C_T &= C_{T_0} + C_{T_1} J + C_{T_2} J^2 \\ C_P &= C_{P_0} + C_{P_1} J + C_{P_2} J^2 \end{aligned} \tag{3.6}$$

where C_T and C_P are the prop thrust and power coefficients, respectively, and J is the advance ratio given by:

$$J = \frac{\|\mathbf{V}_{\infty_\perp}\|}{n D_{prop}} \tag{3.7}$$

where $\|\mathbf{V}_{\infty_\perp}\|$ is the magnitude of the freestream airspeed vector perpendicular to the prop disk plane, D_{prop} is the prop diameter, and n is the prop speed in rev s^{-1} such that $n = \omega/2\pi$. The thrust and power coefficients plotted against advance ratio are given in Figure 3.5. The prop thrust and torque are then computed from the following standard correlations derived from momentum theory [33, 34]:

$$\begin{aligned} T &= C_T \rho n^2 D_{prop}^4 \\ Q &= C_Q \rho n^2 D_{prop}^5 \end{aligned} \tag{3.8}$$

where $C_Q = C_P/2\pi$ and ρ is the air density.

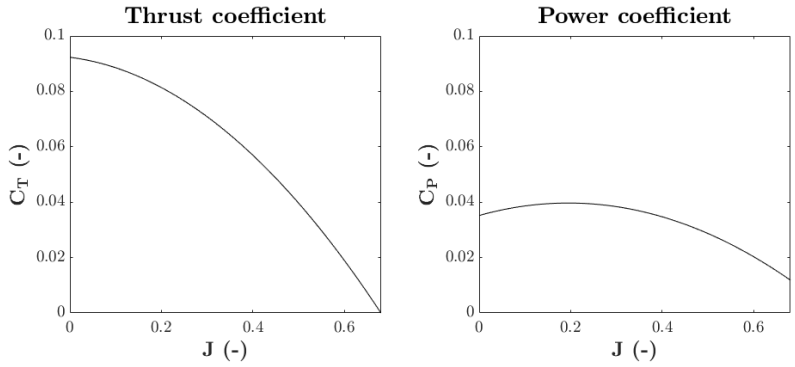

Figure 3.5: Propeller Coefficients vs Advance Ratio

3.3.2.5 Gyroscopic Moments

Gyroscopic moments arise due to a motor/prop's angular momentum vector $\mathbf{H}_{prop,i}$ changing direction due to the aircraft's body rotation rate $\mathbf{\Omega}_b$ and the rotation rate of the wing $\dot{\xi}_j$ that the motor/prop is attached to. This is modeled as:

$$\mathbf{M}_{g,i} = \mathbf{H}_{prop,i} \times \left(\mathbf{\Omega}_b + [0 \quad \dot{\xi}_j \quad 0]^T \right) \tag{3.9}$$

where,

$$\mathbf{H}_{prop,i} = \begin{bmatrix} \omega_i I_{prop} \cos(\xi_j) & 0 & -\omega_i I_{prop} \sin(\xi_j) \end{bmatrix}^T \tag{3.10}$$

and where I_{prop} is the combined rotational inertia of the motor and prop. Carrying this out for all six propulsors on this tiltwing aircraft gives the total gyroscopic moment:

$$\mathbf{M}_g = \begin{bmatrix} \overline{\omega}_W I_{prop}(q + \dot{\xi}_W) \sin(\xi_W) + \overline{\omega}_T I_{prop}(q + \dot{\xi}_T) \sin(\xi_T) \\ -\overline{\omega}_W I_{prop} \left(p \sin(\xi_W) + r \cos(\xi_W) \right) - \overline{\omega}_T I_{prop} \left(p \sin(\xi_T) + r \cos(\xi_T) \right) \\ \overline{\omega}_W I_{prop}(q + \dot{\xi}_W) \cos(\xi_W) + \overline{\omega}_T I_{prop}(q + \dot{\xi}_T) \cos(\xi_T) \end{bmatrix} \tag{3.11}$$

where $\overline{\omega}_W = (\omega_{Li} - \omega_{Lo} - \omega_{Ri} + \omega_{Ro})$ and $\overline{\omega}_T = (\omega_{Tl} - \omega_{Tr})$.

3.3.3 Aerodynamic Forces and Moments

The aerodynamics of this tiltwing aircraft is arguably the most complex part of this model. Due to this configuration's rotating wings combined with having a large proportion of its wing area within the propeller slipstream, it is necessary to capture local aerodynamic effects at various sections along the aircraft (arising from propeller-wing interactions). This means that, in addition to different flow field conditions, locally-computed forces (e.g. lift and drag) will contribute additional aerodynamic moments about the aircraft's CG due to the moment arm between the CG and a

modeled section's local aerodynamic center[4] (AC). Hence, each aerodynamic section is modeled separately, and their individual contributions are then added together to determine the total aerodynamic forces and moments ($\mathbf{F}_A$ and $\mathbf{M}_A$) acting on the aircraft. A schematic of the various individual aero sections being modeled for this tiltwing aircraft is shown in Figure 3.6a.

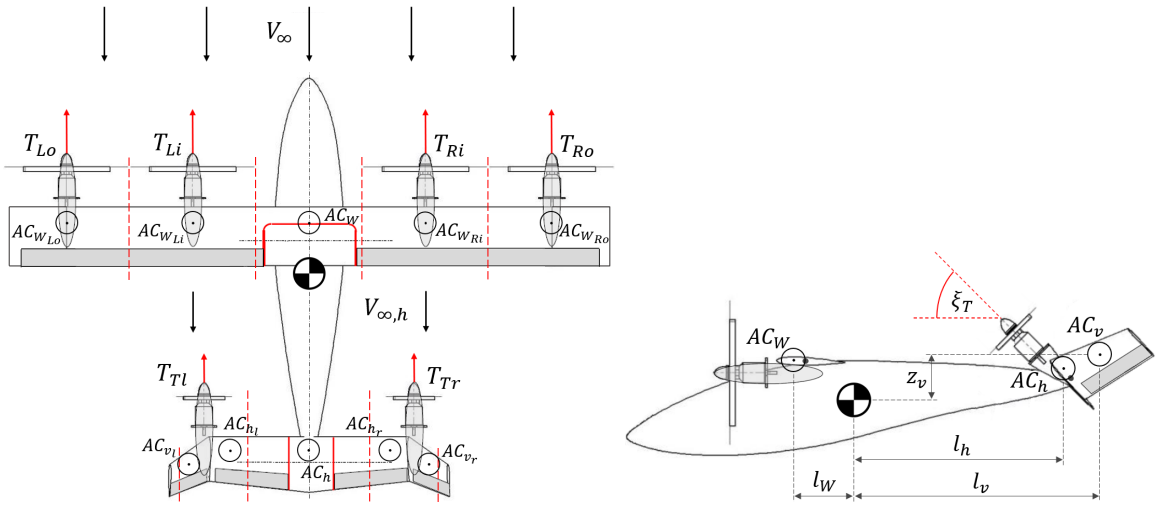

(a) Wetted/unwetted aero sections

(b) Aero moment arms

Figure 3.6: Tiltwing Aerodynamic Modeling

3.3.3.1 Aerodynamic Coefficients and Stability Derivatives

The aircraft's longitudinal aerodynamics are modeled using non-dimensional coefficients which give the aerodynamic lift, drag, and pitching moment (L, D, M) acting on the aircraft in the aerodynamic stability frame of reference [32]:

$$C_L = \frac{L}{\bar{q}S} \quad C_D = \frac{D}{\bar{q}S} \quad C_m = \frac{M}{\bar{q}S\bar{c}} \tag{3.12}$$

where $\bar{q} = \frac{1}{2}\rho\|\mathbf{V}\|^2$ is the dynamic pressure, S is the wing planform area, and $\bar{c}$ is the wing mean chord. Similarly, the aircraft's lateral aerodynamics are modeled using non-dimensional coefficients for the aerodynamic side force, rolling moment,

[4]It is assumed that the location of each airfoil's aerodynamic center is approximately the same as its center of pressure.

and yawing moment (Y, L, N) in the stability frame [32]:

$$C_Y = \frac{Y}{\bar{q}S} \quad C_l = \frac{L}{\bar{q}Sb} \quad C_n = \frac{N}{\bar{q}Sb} \tag{3.13}$$

where b is the wingspan. Note that L is used here to denote both the lift force and the rolling moment.

To derive expressions for these aerodynamic coefficients given in (3.12) and (3.13), the following effects were assumed negligible for this tiltwing aircraft:

- Incidence rates $\dot{\alpha}$ and $\dot{\beta}$, which serve to model time-changing vortex fields [32]. Ignoring these follows from the assumption that this aircraft will not be doing high-performance maneuvering
- Coupling between the longitudinal and lateral aerodynamics. Ignoring this is a good approximation for aircraft with xz-plane symmetry [35]. It is exact for the lateral dynamics and approximate to first-order terms for the longitudinal dynamics
- Derivatives of order higher than first-order. This is a common assumption that is valid for most aircraft [32, 35]

Following these assumptions, the aerodynamic coefficients for this aircraft can be described as functions of the following states and inputs:

$$\begin{aligned} C_L &= f(\alpha, q, \delta_f, \delta_e) \quad C_Y = f(\beta, p, r, \delta_a, \delta_r) \\ C_D &= f(\alpha, q, \delta_f, \delta_e) \quad C_l = f(\beta, p, r, \delta_a, \delta_r) \\ C_m &= f(\alpha, q, \delta_f, \delta_e) \quad C_n = f(\beta, p, r, \delta_a, \delta_r) \end{aligned} \tag{3.14}$$

and expanded with a first-order Taylor series using the notion of stability derivatives

[32, 35], for example:

$$C_L = C_{L_0} + C_{L_\alpha}\alpha + C_{L_q}\hat{q} + C_{L_{\delta_f}}\delta_f + C_{L_{\delta_e}}\delta_e \tag{3.15}$$

where C_{L_0} is the steady-state zero-incidence zero-input lift coefficient, $\hat{q} = q\bar{c}/2u$ is dimensionless pitch rate, and the stability derivatives are defined as:

$$C_{L_\alpha} = \frac{\partial C_L}{\partial \alpha} \qquad C_{L_q} = \frac{\partial C_L}{\partial \hat{q}} \qquad C_{L_{\delta_f}} = \frac{\partial C_L}{\partial \delta_f} \qquad C_{L_{\delta_e}} = \frac{\partial C_L}{\partial \delta_e} \tag{3.16}$$

In the above equations, δ_e and δ_r are the elevator and rudder deflections, respectively, while δ_f and δ_a denote flap and aileron deflections, respectively. The latter two variables are an abstraction from the physical model, which implements two flaperons. The real flaperon inputs δ_{fa} are determined from a superposition of a collective flap input and a differential aileron input:

$$\delta_{fa} = \delta_f \pm \delta_a \tag{3.17}$$

The longitudinal coefficients for this aircraft's main wing and tail are computed using two-dimensional lookup tables in α and δ and were determined from wind tunnel testing[5] [93], while the fuselage and lateral-direction coefficients are computed using standard analytical and empirical expressions for stability derivatives given in [32] and [35]. Note that the wind tunnel data only covers the pre-stall regime, as AoA constraints during transition will be imposed later via the reference trajectories tracked by the controller. For this aircraft, the main wing uses a NACA 63(2)-615 airfoil, while the tail wing uses a NACA 0009 airfoil (for both the horizontal and vertical stabilizers). Plots of the wing and tail longitudinal aero coefficients vs AoA at zero control surface deflection angle are shown in Figure 3.7, while plots of the

[5]This data applies to both the scaled model and the full-scale aircraft.

wing and tail lift coefficients for various control surface deflection angles are shown in Figure 3.8. Table 3.4 summarizes the lateral-direction stability derivative expressions used.

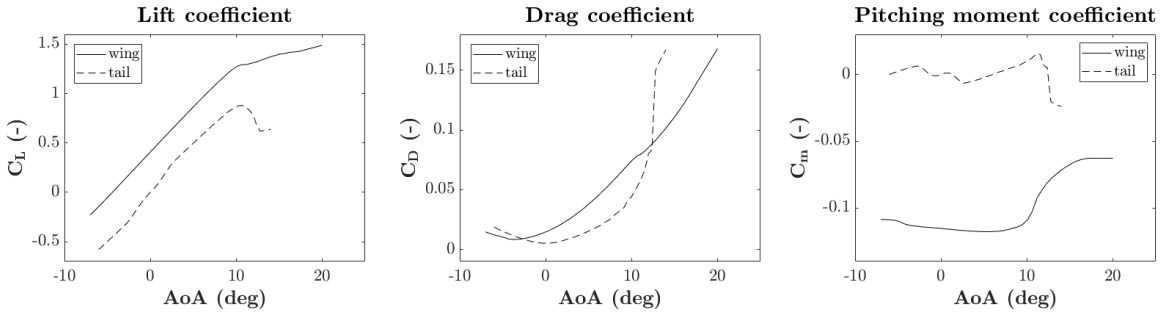

Figure 3.7: Longitudinal Aero Coefficients vs Angle of Attack ($\delta = 0$)

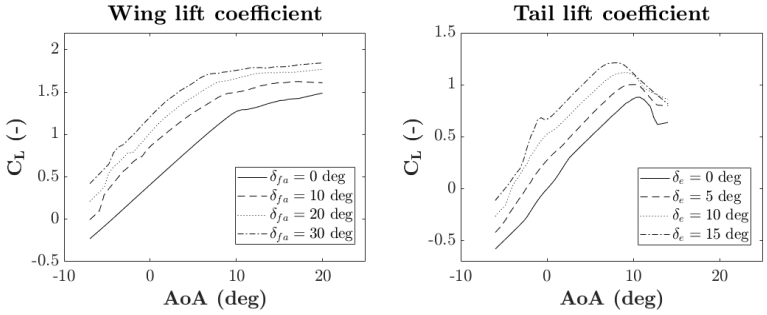

Figure 3.8: Wing and Tail C_L vs Angle of Attack (range of δ)

While lookup tables are used to determine the longitudinal wing and tail coefficients vs α and δ, pitch rate effects (e.g. C_{L_q}) must still be captured. However, instead of modeling C_{L_q} empirically (as is being done for the analogous lateral-direction terms), this effect can be modeled using the lookup table data by simply adjusting the effective AoA entering the lookup tables by the change in local incidence ($\Delta\alpha$) due to the pitch rotation of the aircraft [32]:

$$\Delta\alpha_{q,i} = \arctan\left(\frac{q l_i}{u_i}\right) \tag{3.18}$$

where subscript i denotes the ith surface (e.g. the wing or horizontal tail), l_i is the longitudinal moment arm between AC_i and the aircraft's CG (see Figure 3.6b), and

u_i is the forward airspeed component on that surface. This expression is derived by forming the velocity triangle involving the slew rate of surface i due to the aircraft's rotation rate and the surface's forward airspeed [32]. Note that for conventional fixed-wing aircraft, it is common to approximate these $\Delta\alpha_q$ as small angles (thus omitting the arctan function) [32]; however, this assumes that u_i is much larger than the slew rate of the surface, an assumption that does not necessarily hold at certain points within the flight envelope of a tiltwing aircraft.

For the analogous lateral-direction terms (e.g. C_{Y_p} and C_{Y_r}) that are being modeled empirically, dimensionless roll and yaw rates ($\hat{p}$ and $\hat{r}$, respectively) are required as part of the definition/derivation of these stability derivatives. These come from standard expressions in [32] and are defined as[6]:

$$\hat{p} = \frac{pb}{2u_i} \qquad \hat{r} = \frac{rb}{2u_i} \tag{3.19}$$

Downwash is an important effect being modeled, which serves to reduce the freestream AoA on the tail wing (α_h) due to lift production at the main wing. This is given by the following empirical correlation [32]:

$$\begin{aligned} \alpha_h &= \alpha - \epsilon \\ \epsilon &= \epsilon_0 + \frac{d\epsilon}{d\alpha}\alpha \end{aligned} \tag{3.20}$$

where ϵ is the downwash angle, α in this context represents the freestream incidence angle on the main wing, and ϵ_0 and $d\epsilon/d\alpha$ are constants. However, different from a conventional fixed-wing aircraft, it was chosen to taper off the effect of downwash with increasing α due to the high incidence angles seen by the rotating wings of this

[6]Note that for this model, a typo was discovered at the end of the project where the horizontal component of the vertical tail moment arm was used instead of $b/2$ in the computation of $\hat{r}$. Fortunately, the numerical difference between these values for this aircraft is small, and comparison of the resulting stability derivative values reveals differences on the order of just 10^{-4}. Hence, the impact of this error has been deemed negligible.

tiltwing aircraft. Specifically, it was chosen to apply a multiplier m to the downwash angle ϵ that linearly decays from 1 to 0 as α varies from 6° to 10°.

Fuselage contribution to the longitudinal aerodynamics primarily involves additional skin and parasite drag as well as a positive pitching moment during non-zero incidence. The latter is modeled as the following stability derivative:

$$C_{m_{\alpha,f}} = \frac{2\nu_{hor}}{S\bar{c}} \tag{3.21}$$

where the subscript f indicates fuselage contribution and ν_{hor} is the equivalent width volume of the fuselage [32].

Table 3.4: Lateral-Direction Stability Derivatives

–	β	p	r	δ_r
C_Y	$-C_{L_{\alpha,v}}\left(1-\frac{d\sigma}{d\beta}\right)\eta_v\frac{S_v}{S}$	$-2C_{L_{\alpha,v}}\frac{\eta_v S_v}{bS}z_{v_s}$	$2C_{L_{\alpha,v}}\frac{\eta_v S_v}{bS}l_{v_s}$	(LUT)
C_l	$C_{Y_\beta}\frac{z_{v_s}}{b}$	$-\frac{C_{L_{\alpha,W}}}{8}+C_{Y_p}\frac{z_{v_s}}{b}$	$\frac{C_{L_0}}{4}+C_{Y_r}\frac{z_{v_s}}{b}$	$C_{Y_{\delta_r}}\frac{z_{v_s}}{b}$
C_n	$-\frac{2\nu_{vert}}{Sb}-C_{Y_\beta}\frac{l_{v_s}}{b}$	$-\frac{C_L}{8}-C_{Y_p}\frac{l_{v_s}}{b}$	$-\frac{C_{D_0}}{4}-C_{Y_r}\frac{l_{v_s}}{b}$	$-C_{Y_{\delta_r}}\frac{l_{v_s}}{b}$

For the lateral-direction stability derivatives summarized in Table 3.4, subscript v indicates a quantity related to the vertical tail (which many of these expressions depend on), $\eta_v = \bar{q}_v/\bar{q}$ is a scaling factor to relate the dynamic pressure on the vertical tail to freestream values, $d\sigma/d\beta$ is the sidewash derivative analogous to downwash but for sideslip angle, and ν_{vert} is the equivalent height volume of the fuselage [32]. These expressions have been organized in the table such that each row-column pair indicates the corresponding stability derivative, e.g. the top-left element gives the expression for C_{Y_β}, etc. All of these stability derivatives have been scaled relative main wing reference quantities (e.g. S and b) and were derived such that they represent total-aircraft values; the latter means that, while individual aero sections are still modeled to capture differing flow conditions, no additional moment contributions (e.g. due

to side force) need to be accounted for in the case of these lateral terms. Since the vertical tail uses the same NACA airfoil as the horizontal stabilizer, $C_{L_{\alpha,v}}$ is the same lift-curve slope as given in Figure 3.7 (the "tail" curve). Similarly, $C_{Y_{\delta_r}}$ was determined directly from lookup table data ("LUT"), as the rudder geometry matches that of the elevator. Note that δ_a has not been explicitly shown in Table 3.4. This is because the aileron contribution to lateral-direction aero coefficients is automatically accounted for by the longitudinal lookup tables combined with the local aero section modeling; hence, flaperon deflections computed on the main wing already generate rolling and yawing moments that would be equivalently determined from more conventional aileron effectiveness modeling. Lastly, z_{v_s} and l_{v_s} indicate the projection of the vertical tail moment arms z_v and l_v (see Figure 3.6b) onto the local stability axes and are therefore functions of the local AoA. This is typical for conventional fixed-wing aircraft [32] but, different from conventional fixed-wing aircraft, z_v and l_v for this tiltwing are not constant due to this aircraft's tilting tail. Hence, these quantities depend on the tail hinge geometry and are thus additionally a function of the tail wing angle ξ_T. Note that AC_W and AC_h have been assumed to be very close to the rotation axes of the wings, so l_W and l_h are reasonably assumed to be constant.

3.3.3.2 Propeller-Wing Interactions

To model the effects of each propeller's wake on its respective section of the aircraft, it is common practice to model the prop-induced velocity using momentum theory [94, 95, 96, 24, 97, 98]. The prop-induced velocity serves to modify the local incidence velocity on a particular section of the aircraft. Hence, once the prop-induced velocity is calculated, a modified local incidence velocity $\mathbf{V}_{wet}$ (referred to here as a "wetted" velocity) is calculated for each propeller and used to determine local dynamic pressures and angles of incidence. For this aircraft, approximately 80% of the

wing and tail is wetted by propeller wake, while the other 20% remains unwetted (i.e. freestream conditions). The fuselage is assumed to be entirely unwetted, while the flaperons are assumed to be 100% wetted due to their placement along the main wing. (Refer again to Figure 3.6a.)

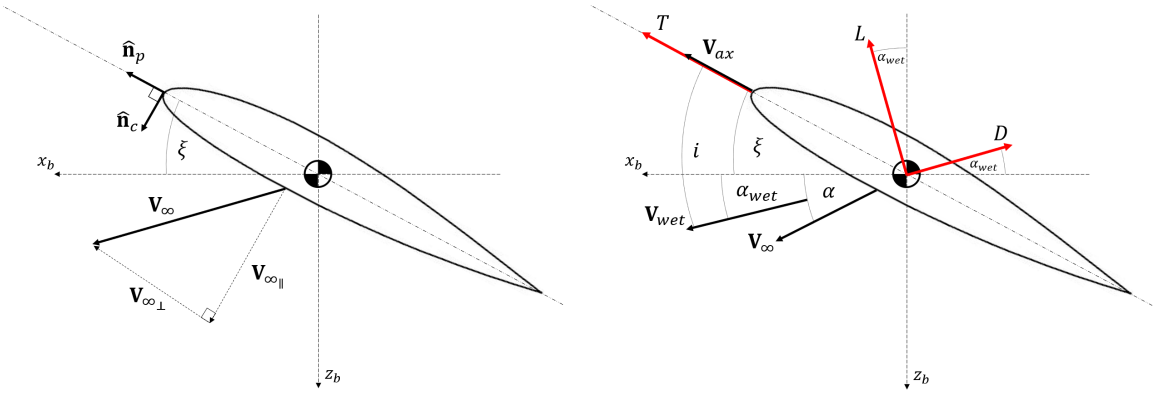

(a) Airspeed vector prop-plane decomposition

(b) Velocities and Angles of Incidence

Figure 3.9: Airspeed Vector Diagrams

Figure 3.9a shows unit vectors perpendicular to the prop disk plane ($\hat{\mathbf{n}}_p$) and the wing chord ($\hat{\mathbf{n}}_c$). These can be written mathematically in the aircraft's body frame as:

$$\begin{aligned} \hat{\mathbf{n}}_p &= \begin{bmatrix} \cos(\xi) & 0 & -\sin(\xi) \end{bmatrix}^T \\ \hat{\mathbf{n}}_c &= \begin{bmatrix} \sin(\xi) & 0 & \cos(\xi) \end{bmatrix}^T \end{aligned} \tag{3.22}$$

The freestream velocity vector $\mathbf{V}_\infty$ can be decomposed into a component perpendicular to the prop disk plane $\mathbf{V}_{\infty_\perp} = (\mathbf{V}_\infty \cdot \hat{\mathbf{n}}_p) \cdot \hat{\mathbf{n}}_p$ and a component parallel to the prop disk plane $\mathbf{V}_{\infty_\parallel} = (\mathbf{V}_\infty \cdot \hat{\mathbf{n}}_c) \cdot \hat{\mathbf{n}}_c + \mathbf{V}_{\infty_y}$, where $\mathbf{V}_{\infty_y}$ contains only the lateral (second) component of the freestream velocity vector. The modified local incidence velocity vector $\mathbf{V}_{wet}$ is simply given by the sum of the freestream velocity component parallel to the prop disk plane and the resulting axial velocity $\mathbf{V}_{ax}$ due to the combined axial inflow ($\mathbf{V}_{\infty_\perp}$) and prop-induced velocity:

$$\mathbf{V}_{wet} = \mathbf{V}_{\infty_\parallel} + \mathbf{V}_{ax} \tag{3.23}$$

In this work, an idealized model of the prop's slipstream is assumed that represents the prop-induced velocity far downstream from the prop. From momentum theory [33, 34], the resulting axial velocity vector is thus computed from:

$$\mathbf{V}_{ax} = \sqrt{\frac{2T}{\rho A_{prop}} + \|\mathbf{V}_{\infty_\perp}\|^2} \cdot \hat{\mathbf{n}}_p \tag{3.24}$$

where A_{prop} is the prop disk area. It can be seen by inspection that if there is no propeller thrust, then $\mathbf{V}_{wet} = \mathbf{V}_\infty$ as expected.

The freestream AoA α is the angle formed between the freestream velocity $\mathbf{V}_\infty$ and the horizontal body axis [35]. Due to the propeller-wing interactions, there are induced AoAs locally on each section of the wing and tail surfaces, generically denoted α_{wet}, which are analogously measured between the wetted velocity vectors and the body x-axis (see Figure 3.9b). Due to the wing and tail rotation angles, the actual incidence on a generic surface section is given by:

$$i = \alpha_{wet} + \xi \tag{3.25}$$

With each surface section's $\mathbf{V}_{wet}$ and i calculated, their individual aerodynamic forces $\mathbf{F}_A^{stab} = [D \quad Y \quad L]^T$ and moments $\mathbf{M}_A^{stab} = [L \quad M \quad N]^T$ in the stability frame can be computed by rearranging (3.12) and (3.13).

3.3.3.3 Stability-to-Body Frame Transformation

To complete the aerodynamics, the computed aerodynamic force and moment vectors must be transformed from the stability frame in which they were derived to the aircraft's body frame for use in the 6-DoF equations using a standard rotation matrix. Typically, the freestream AoA α is the appropriate rotation angle to choose. However, due to wetted aero effects, local stability frames will be oriented relative their local wetted velocity vector $\mathbf{V}_{wet}$ as opposed to the freestream velocity (see Figure 3.9b).

Hence, for the aerodynamic forces and moments computed on these sections of the aircraft, the necessary angle of rotation instead becomes the local α_{wet}, resulting in the following counter-clockwise rotation matrix:

$$\mathbf{R}_{bs} = \begin{bmatrix} \cos(\alpha_{wet}) & 0 & -\sin(\alpha_{wet}) \\ 0 & 1 & 0 \\ \sin(\alpha_{wet}) & 0 & \cos(\alpha_{wet}) \end{bmatrix} \tag{3.26}$$

3.4 Model Implementation

This 6-DoF nonlinear model was implemented in both MATLAB and Simulink. While the flight conditions and dynamics of this tiltwing aircraft vary widely between hover, cruise, and transition, no explicit distinction was made in the modeling approach of Section 3.3 between different flight modes, such as hover (thrust-borne flight) or cruise (wing-borne flight); hence, this mathematical model represents the tiltwing's dynamics continuously at all points in the flight envelope, and differing "flight modes" can be inferred simply based on the current value of the wing angles ξ_W and ξ_T on which these dynamics depend. However, in consideration of the aerodynamics presented in Section 3.3.3, many quantities (such as AoA) are not explicitly defined outside of the context of conventional fixed-wing flight, such as during pure hover (i.e. zero airspeed). Hence, in consideration of, primarily, hover flight, it was necessary when implementing this model to account for such edge cases unique to this tiltwing aircraft. For example, in the case of an undefined AoA when $u = w = 0$, it was chosen to assign an arbitrary value for α (such as $\alpha = 0$) when $\|\mathbf{V}\|$ becomes sufficiently small, which is a valid assignment in this context since the correspondingly small dynamic pressure would make any aerodynamic contributions negligible at these conditions.

3.5 Model Validation

Validation of this model's aerodynamics was achieved by comparing numerical results in MATLAB/Simulink with those predicted by three programs (AVL[7], XFLR5[8], and OpenVSP[9]) that make use of potential flow-based computational fluid dynamics (CFD) methods to carry out aerodynamic and flight-dynamic analysis of rigid aircraft of arbitrary configuration. Specifically, by linearizing the complete 6-DoF model of this tiltwing aircraft at its cruise operating point condition, its corresponding (total-aircraft) stability derivatives were compared to those predicted using these CFD methods for this tiltwing's given surface geometries. While this does not validate the model's aerodynamics in all flight phases (e.g. hover and transition), this approach gives good indication that the aerodynamic model (the most complex part of this 6-DoF model) has been set up properly. AVL and XFLR5 make use of the vortex-lattice method (VLM), which is a useful technique for obtaining initial estimates of aerodynamic stability derivatives and dynamic modes. However, VLM also has certain known limitations inherent in its basic assumptions (potential flow and thin-airfoil theory), such as its inability to accurately deal with general fuselage geometries. In particular, it has been shown [99] that the lateral stability derivatives computed using VLM may demonstrate poor agreement with those determined from empirical models (such as those used in Section 3.3.3.1), especially at flight conditions with higher angles of attack. In contrast, OpenVSP uses an alternative approach known as the panel method, which differs in its abilities by being capable of modeling thick airfoils and non-lifting geometries (e.g. a fuselage) but still makes certain limiting assumptions that can impact modeling accuracy [100]. Nevertheless, both VLM and panel method can be useful tools for approximate model validation, and hence, the goal of these comparisons is to verify a rough order-of-magnitude cor-

[7]https://web.mit.edu/drela/Public/web/avl/. Accessed February 2025.
[8]https://www.xflr5.com/xflr5.htm. Accessed March 2025.
[9]https://openvsp.org/. Accessed April 2025.

relation between this 6-DoF model's aerodynamics and values predicted using these alternative computational techniques.

To determine this aircraft's stability derivatives from a numerical linearization, it is necessary to develop analytical expressions for the linearized state-spaces of this aircraft's longitudinal and lateral dynamics. Consider the state variable vector $\mathbf{x}_{lon} = [\theta \quad u \quad w \quad q]^T$. The longitudinal dynamics can be linearized about the point $(\mathbf{x}_0, \mathbf{u}_0)$ and written as $\dot{\mathbf{x}}_{lon} = \mathbf{A}_{lon}\mathbf{x}_{lon} + \mathbf{B}_{lon}\mathbf{u}$, where:

$$\mathbf{A}_{lon} = \begin{bmatrix} 0 & 0 & 0 & 1 \\ -gc_{\theta_0} & X_u & X_w & 0 \\ -gs_{\theta_0} & Z_u & Z_w & u_0 + Z_q \\ 0 & M_u & M_w & M_q \end{bmatrix} \tag{3.27}$$

and where $X_u := (1/m)\partial X/\partial u$, $M_w := (1/I_{yy})\partial M/\partial w$, etc., are known as the *dimensional* stability derivatives. Similarly, the lateral dynamics with state vector $\mathbf{x}_{lat} = [\phi \quad v \quad p \quad r]^T$ can be linearized, giving:

$$\mathbf{A}_{lat} = \begin{bmatrix} 0 & 0 & 1 & 0 \\ gc_{\theta_0} & Y_v & Y_p & Y_r - u_0 \\ 0 & (\frac{L_v + i_x N_v}{1 - i_x i_z}) & (\frac{L_p + i_x N_p}{1 - i_x i_z}) & (\frac{L_r + i_x N_r}{1 - i_x i_z}) \\ 0 & (\frac{N_v + i_z L_v}{1 - i_x i_z}) & (\frac{N_p + i_z L_p}{1 - i_x i_z}) & (\frac{N_r + i_z L_r}{1 - i_x i_z}) \end{bmatrix} \tag{3.28}$$

where $i_x = I_{xz_s}/I_{xx_s}$, $i_z = I_{xz_s}/I_{zz_s}$, and where the subscript s here indicates the moments of inertia about the aircraft's stability frame axes (instead of body axes). In (3.27) and (3.28) above, X, Y, and Z represent the net external forces ($F_{x_{net}}$, $F_{y_{net}}$, $F_{z_{net}}$) but excluding gravity, and L, M, and N represent the net external moments ($M_{x_{net}}$, $M_{y_{net}}$, $M_{z_{net}}$). This slight difference in notation is used in this subsection for consistency with literature on this topic [32, 35].

The dimensional stability derivatives found in $\mathbf{A}_{lon}$ and $\mathbf{A}_{lat}$ can be connected to

the more familiar *non-dimensional* stability derivatives introduced in Section 3.3.3.1 (e.g. C_{L_α}) by the relationships given in Tables 3.5 and 3.6 [32, 35], where the subscript 0 indicates values at the linearization point $(\mathbf{x}_0, \mathbf{u}_0)$. Note that I_{yy} is equivalent to I_{yy_s} since the stability axes and body axes share the same y-axis. Hence, to determine the (total-aircraft) non-dimensional stability derivatives of this 6-DoF model, numerical linearization was carried out in MATLAB using `linearize()` based on a steady-state cruise operating point determined using `findop()`. The expressions in Tables 3.5 and 3.6 were then used to compute these non-dimensional stability derivatives based on the numerical values for the elements of $\mathbf{A}_{lon}$ and $\mathbf{A}_{lat}$ determined from the linearization.

Table 3.5: Longitudinal non-dimensional stability derivatives in terms of dimensional quantities

–	α	q
Z	$C_{L_\alpha} = -\frac{mu_0 Z_w}{\bar{q}_0 S} - C_D(\alpha_0)$	$C_{L_q} = -\frac{2mu_0 Z_q}{\bar{q}_0 S\bar{c}}$
M	$C_{m_\alpha} = \frac{I_{yy} u_0 M_w}{\bar{q}_0 S\bar{c}}$	$C_{m_q} = \frac{2I_{yy} u_0 M_q}{\bar{q}_0 S\bar{c}^2}$

Table 3.6: Lateral non-dimensional stability derivatives in terms of dimensional quantities

–	β	p	r
Y	$C_{Y_\beta} = \frac{mu_0 Y_v}{\bar{q}_0 S}$	$C_{Y_p} = \frac{2mu_0 Y_p}{\bar{q}_0 Sb}$	$C_{Y_r} = \frac{2mu_0 Y_r}{\bar{q}_0 Sb}$
L	$C_{l_\beta} = \frac{I_{xx_s} u_0 L_v}{\bar{q}_0 Sb}$	$C_{l_p} = \frac{2I_{xx_s} u_0 L_p}{\bar{q}_0 Sb^2}$	$C_{l_r} = \frac{2I_{xx_s} u_0 L_r}{\bar{q}_0 Sb^2}$
N	$C_{n_\beta} = \frac{I_{zz_s} u_0 N_v}{\bar{q}_0 Sb}$	$C_{n_p} = \frac{2I_{zz_s} u_0 N_p}{\bar{q}_0 Sb^2}$	$C_{n_r} = \frac{2I_{zz_s} u_0 N_r}{\bar{q}_0 Sb^2}$

Figures 3.10 and 3.11 show the modeled aircraft geometries in each of the three programs. For both AVL and XFLR5, the fuselage geometry was omitted, as this is more effective since slender-body assumptions do not apply for this tiltwing's fuselage. An identical vortex lattice was specified for both programs: 26 chordwise and 52 spanwise vortices for the main wing, 15 chordwise and 30 spanwise vortices for the horizontal tail, and 15 chordwise and 15 spanwise vortices for each vertical tail.

All vortices were distributed along these lifting surfaces using a cosine distribution. In contrast, for OpenVSP, the approximate fuselage geometry was included since panel method can more effectively deal with such geometries. Additionally, instead of specifying a vortex lattice, the (triangulated) mesh geometry for panel method is generated automatically. For this model, it produced a mesh with $1,690$ nodes, $3,368$ loops, and $5,052$ edges.

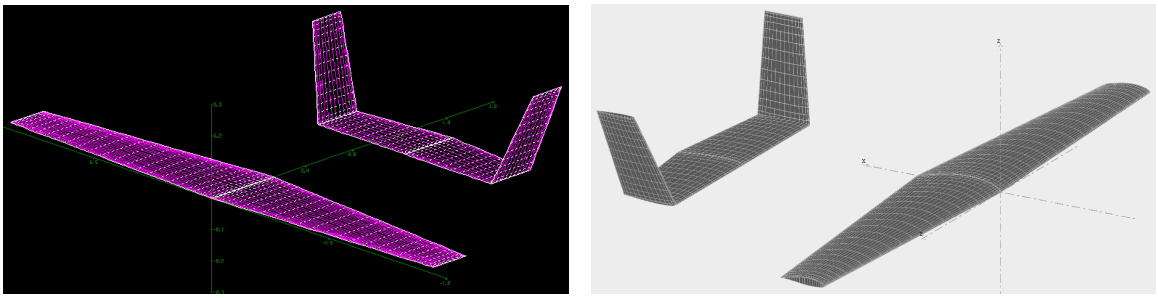

(a) In AVL (b) In XFLR5

Figure 3.10: Tiltwing aircraft wing and tail geometries set up for VLM

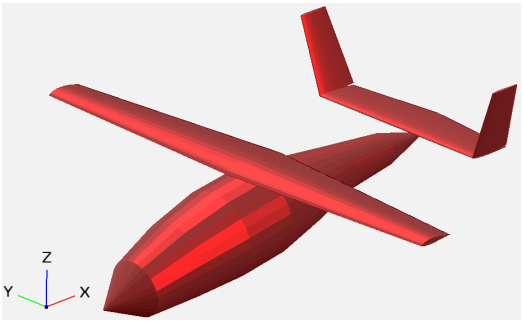

Figure 3.11: Tiltwing aircraft wing, tail, and fuselage geometries in OpenVSP set up for panel method

Tables 3.7 and 3.8 summarize the results of the non-dimensional stability derivative values found for this 6-DoF model alongside the values computed in AVL, XFLR5, and OpenVSP. The operating point selected for this analysis corresponds to a trimmed constant-altitude cruise condition with an airspeed of $u_0 = 20.6\ \mathrm{m\,s^{-1}}$. It can be seen that the longitudinal stability derivatives are in good agreement across all four sources, while there is much more variation in the predicted lateral stability derivative values, including between the three CFD programs themselves (which is interesting to note and is indicative of the expected uncertainty of these values). However, these lateral

stability derivatives remain within a rough order-of-magnitude (with the exception of C_{Y_p} and C_{n_p}) and, without wind tunnel data (as is available for the longitudinal terms), more uncertainty is already expected in the lateral coefficients that the flight controller will need to be robust enough to deal with. Additionally, the greatest complexities to be solved with this tiltwing aircraft lie within the transition maneuver itself, which is symmetric laterally and thus relies primarily on the longitudinal terms in this model. Hence, for the purposes of this research, it was chosen to rely on the empirical correlations from Section 3.3.3.1 for the lateral-direction terms, and the completed 6-DoF model's aerodynamic model was deemed satisfactory.

Table 3.7: Comparison of longitudinal stability derivative values

Derivative	**6-DoF Model**	**AVL**	**XFLR5**	**OpenVSP**
C_{L_α}	6.64	7.49	7.75	8.91
C_{m_α}	−1.10	−1.28	−1.77	−1.05
C_{L_q}	7.47	9.95	11.76	6.97
C_{m_q}	−49.56	−47.08	−52.86	−66.62

Table 3.8: Comparison of lateral stability derivative values

Derivative	**6-DoF Model**	**AVL**	**XFLR5**	**OpenVSP**
C_{Y_β}	−1.51	−0.74	−0.70	−1.17
C_{l_β}	−0.04	−0.20	−0.16	−0.32
C_{n_β}	0.74	0.27	0.26	0.12
C_{Y_p}	3.73	−0.07	−0.23	−0.25
C_{l_p}	−0.64	−0.70	−0.73	−0.89
C_{n_p}	−0.13	0.003	−0.04	−0.06
C_{Y_r}	1.67	0.60	0.60	−0.12
C_{l_r}	0.15	0.38	0.36	0.33
C_{n_r}	−0.75	−0.25	−0.23	−0.37

3.6 Analysis of the Dynamics

This section looks at the poles of this tiltwing aircraft at a range of operating points to build an understanding of the dynamic behavior of this configuration. We start with an analysis of the cruise dynamics, followed by hover mode. We then explore how these dynamics vary during transition by looking at the loci of the poles while considering perturbations in the aircraft's CG.

3.6.1 Cruise

To analyze the cruise dynamics, the model was linearized at the same operating point as in Section 3.5 ($u_0 = 20.6\ \mathrm{m\,s^{-1}}$) as well as two additional operating points for comparison ($u_0 = 19.5\ \mathrm{m\,s^{-1}}$ and $u_0 = 21.5\ \mathrm{m\,s^{-1}}$), which were chosen as small perturbations around the nominal cruise operating point to understand the variation of the aircraft's dynamics at expected cruise speeds. Using these operating points, the eigenvalues of $\mathbf{A}_{lon}$ and $\mathbf{A}_{lat}$ were computed. Figure 3.12 shows maps of all longitudinal and lateral poles, while Table 3.9 summarizes the pole characteristics at the nominal operating point ($u_0 = 20.6\ \mathrm{m\,s^{-1}}$). For this tiltwing aircraft in cruise, standard dynamic modes for conventional fixed-wing aircraft [32, 35] are exhibited. Comparing the different operating points, it can be seen that the poles move further left (i.e. quicken/become more stable) with increasing airspeed; this includes the (unstable) spiral mode, in which case its time constant increases (i.e. it becomes less unstable) as it moves left toward the origin.

Table 3.9: Dynamic modes in cruise (nominal, $u_0 = 20.6\ \mathrm{m\,s^{-1}}$)

Mode	**Frequency** ($\mathrm{rad\,s^{-1}}$)	**Damping Ratio**	**Time Constant** (s)
Short period	6.39	0.66	0.24
Phugoid	1.18	0.33	2.59
Rolling	4.58	1	0.22
Dutch roll	8.66	0.24	0.49
Spiral	0.11	-1	8.8

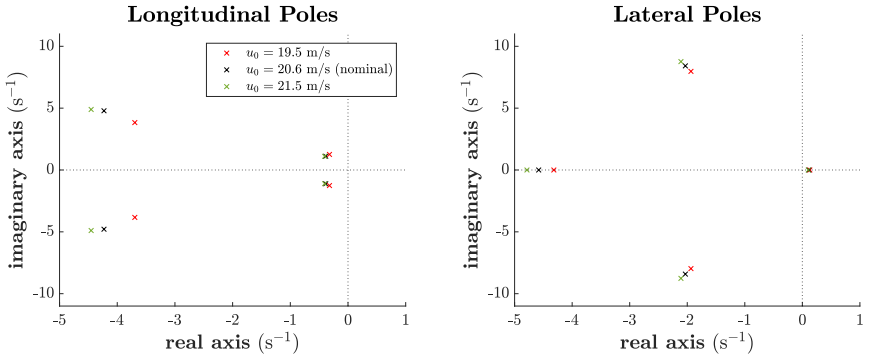

Figure 3.12: Cruise mode pole locations (three operating points)

3.6.2 Hover

For hover mode, the poles were computed similarly but at a steady-state pure hover operating point with the wings at 77° (their equilibrium hover angle[10]). Note that a small forward airspeed of 0.01 $\mathrm{m\,s^{-1}}$ was applied to avoid edge-case discontinuities during the linearization process. The hover mode poles are plotted in Figure 3.13 (along with the nominal cruise poles for comparison) and their characteristics summarized in Table 3.10. The dynamic modes in hover flight cannot be described by the same terminology as for conventional fixed-wing cruise, but it is interesting to compare the differences in the pole locations. Specifically, the slower longitudinal oscillatory mode of cruise (phugoid) is now an unstable, very slow and minimally-

[10]This is due to lift generation from the prop-induced velocities; a wing angle of 90° in hover results in the aircraft slowly drifting backward.

damped mode in hover (Mode 3), while the remaining poles have split into two real modes (Modes 1 and 2). Similar longitudinal pole locations can be found for other tiltwing aircraft, such as in [101] and [102]. For the lateral dynamics, the distribution of the poles is not too dissimilar to the cruise case, but the locations have changed slightly. The oscillatory mode (cruise's "dutch roll") has moved to the left, making it quicker (Mode 5), while the two real modes have moved to the right (Modes 4 and 6), quickening the unstable real pole and slowing the stable one. This somewhat contrasts the lateral pole locations found in [101] and [102], where the oscillatory mode (Mode 5) is instead found to be unstable.

Table 3.10: Dynamic modes in hover

Mode	**Frequency** (rad s^{-1})	**Damping Ratio**	**Time Constant** (s)
1, 2	4.99, 0.08	1, 1	0.2, 13
3	2.65	−0.004	103
4	0.62	1	1.6
5	6.46	0.53	0.29
6	0.95	−1	1.05

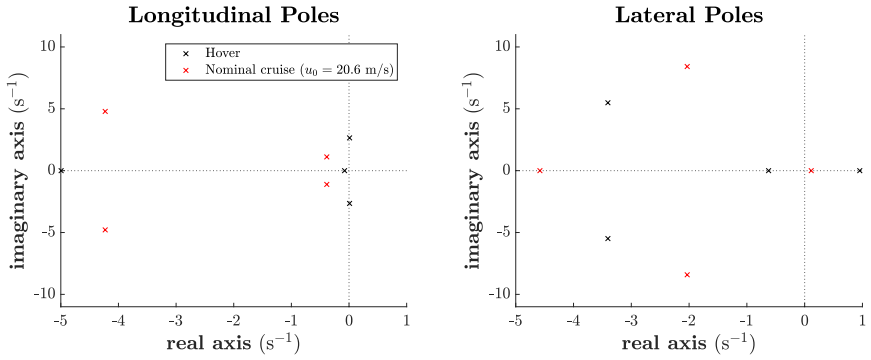

Figure 3.13: Hover mode pole locations (with nominal cruise for comparison)

3.6.3 Transition

For transition, it was of interest to understand how this aircraft's dynamics vary during the conversion between hover and cruise flight. To investigate this, the model

was linearized at a plethora of points along the maneuver (shown in Figure 3.14)[11] and the poles plotted to form their "loci" throughout transition. Additionally, to further understand these loci, the procedure was repeated for differing CG values corresponding to a ±2.6% change in the CG position (from nominal) relative the length of the aircraft, which represents the variation in CG position within the upper and lower design loading limits of the full-scale aircraft. The results are shown in Figure 3.15, where the color shift from orange to blue indicates moving from hover to cruise and is correlated with the colored points in Figure 3.14. Note that, different from the hover and cruise results, these pole locations do not represent equilibrium operating points since the system was linearized along an accelerating maneuver.

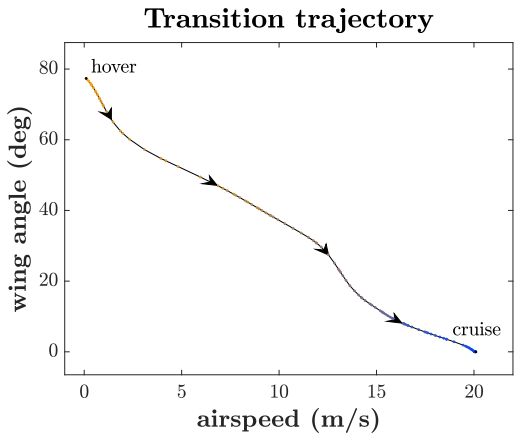

Figure 3.14: Transition trajectory used for linear analysis

The most notable observation is the general spread of these pole locations, exemplifying the large shift in dynamic behavior throughout transition. For example, for the nominal CG position, we see time constants ranging from as quick as 0.05 s to as slow as 39.7 s as well as damping ratios reaching as small as 0.25. Furthermore, there is a notable difference in the spread of these poles across the varying CG positions. Specifically, the loci of the real poles are more condensed along the real axis for more forward-shifted CG positions, while in contrast, the complex poles have real parts

[11]This trajectory comes from computing a feasible transition trajectory for this aircraft using the method that will be discussed in Chapter 4.

that become more spread out. This indicates a strong correlation between the variation in open-loop dynamics throughout transition and the physical loading limits of this aircraft, and this characteristic will be shown to be significant when considering feasible transition trajectories in Chapter 4. To see more detail near the origin, Figure 3.16 shows the same plots but zoomed-in to the region indicated by the red rectangles in Figure 3.15. It can be observed that minimum damping ratios get even smaller for a forward-shifted CG position (indicated by the angle of the complex poles relative the imaginary axis), while the real parts of the complex poles once again become more spread out, just as was observed for the poles in Figure 3.15. In contrast, a further-aft CG sees higher minimum damping ratios, and the wider spread of the real roots results in faster unstable dynamics.

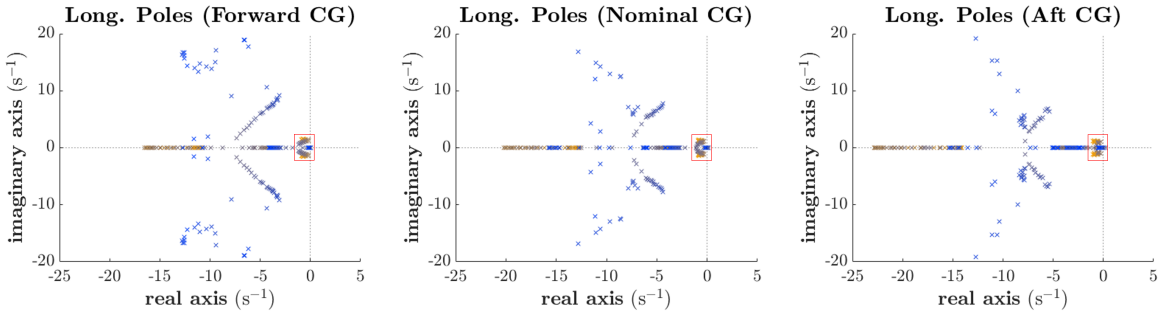

Figure 3.15: All longitudinal pole locations throughout transition with ±2.6% CG position variations. The color shift from orange to blue indicates moving from hover to cruise.

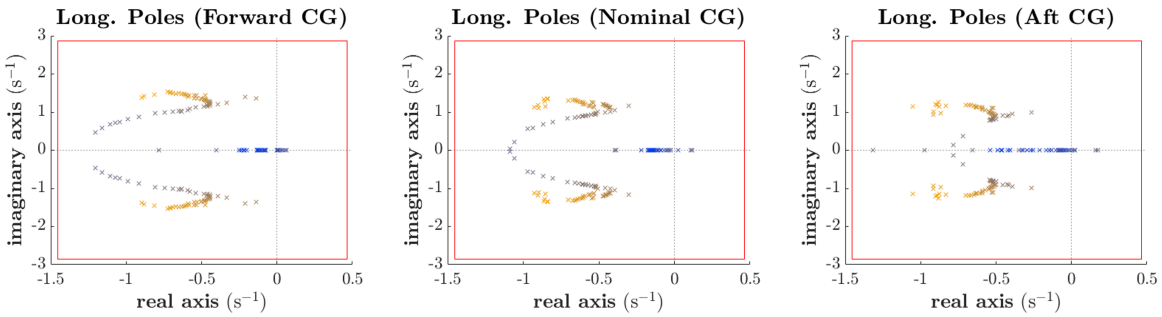

Figure 3.16: Longitudinal pole locations throughout transition – zoomed near origin. The color shift from orange to blue indicates moving from hover to cruise.

Chapter 4

Transition Trajectory Optimization

4.1 Introduction

The purpose of this chapter is to investigate optimal transition trajectories for the tiltwing aircraft of interest in this thesis and assess relative merits of employing advanced AFC concepts. As discussed in Section 1.1, the transition maneuver for tiltwing aircraft is particularly challenging due to the shifting wing angle while the aircraft rapidly changes airspeed [24], which often pushes the local AoA of the wings to their aerodynamic stall regime and is thus problematic for aircraft controllability, performance, and safety. While this may be somewhat eased in forward (hover-to-cruise) transition due to the high thrust condition to accelerate the aircraft (which helps reduce the local AoA), backward (cruise-to-hover) transition sees a conversely low thrust condition, which can greatly constrain what is achievable control wise. As an example, the Canadair CL-84 tiltwing had pilot-noticeable buffeting at rates of descent above 700 ft/min in transition, which was found to be due to low normal velocity damping [16] and made it difficult for pilots to track the glide slope. Additionally, it was later found that similar flight conditions gave rise to non-minimum phase zeros at low speeds [103], providing further explanation for such pilot difficulty;

it is well known that the presence of non-minimum phase zeros places fundamental limits on achievable closed-loop performance [104] and hence what is achievable with stability augmentation systems. For the CL-84, these effects made limiting the rate of descent a necessity. Similarly, as will also be shown in this chapter, backward transition is often constrained to increase in altitude throughout the maneuver to avoid wing aerodynamic stall, but this constraint is not ideal for air taxi operations, as it would increase separation distances and thus limit aircraft density in the operational airspace. Further, smooth maneuvers and minimal altitude variations may also be preferred for passenger comfort.

More recently, Chauhan and Martins [24] explored the take-off-to-cruise transition trajectory for a tandem tiltwing VTOL aircraft based on the Airbus A^3 Vahana. Using gradient-based optimization with and without stall constraints, it was demonstrated that minimum-energy trajectories involve flying near the stall AoA, though the energy penalty for avoiding stall was practically negligible. This extensive and thorough study used a 2-DoF model with the forward transition maneuver starting from take-off, and so constant-altitude constraints at height conditions were not investigated. Pradeep and Wei [105] also investigated the transition maneuver of the Airbus A^3 Vahana using a fixed final-time multiphase optimization. The general cruise-to-landing maneuver for this tandem tiltwing configuration was explored for various initial cruise speeds where it was found that distributing the maximum possible delay to the cruise phase before beginning the transition phase allowed for minimum energy consumption. However, their model assumed that in the transition phase, the rotation of the wings from cruise (horizontal) to hover (vertical) occurs in negligible time, and so the specific details of the aircraft's performance in backward transition were not explored. Yeh et al. [106] investigated the ability to predict optimal tiltwing take-off trajectories using a machine learning-based inverse mapping. They used the same 2-DoF open-source model provided by Chauhan and Martins

[24], and it was found that optimal trajectories could be predicted with 99% relative accuracy using long short-term memory (LSTM) prediction. However, the focus of this study was primarily about predicting known results using a trained machine learning model, though this framework made significant reductions to computation time. Other authors that made significant reductions in computation time were Doff-Sotta et al. [74, 75], where successive convexification methods were used to achieve fast convergence of optimal solutions by representing the nonlinear tiltwing dynamics with DC functions. Only recently has the tiltwing transition trajectory been explored using a 3-DoF [107] and 6-DoF [108] model. In [107], dynamic investigations were carried out to understand the operational limits of, similarly, the Airbus A^3 Vahana during various types of transition maneuvers, while in [108] failure modes were investigated for the same configuration for the backward transition maneuver. Other papers exist which explore trajectory optimization for the transition maneuver but for tail-sitter aircraft configurations [95, 94, 109, 96]. For example Maqsood and Go [95] study a tail-sitter with a variable-incidence wing (i.e. the propeller is fixed to the nose of the aircraft while the wing can rotate). It was found that incorporating the variable-incidence wing as a control variable allowed for better performance of the aircraft throughout transition and avoided wing aerodynamic stall. Kubo and Suzuki [94] found that their tail-sitter aircraft could not avoid climbing to prevent aerodynamic stall during the backward transition maneuver without using high-lift devices such as flaps or slats. Such high-lift devices were not necessary for forward transition.

Considering the particular difficulty of the transition maneuver for tiltwing aircraft—the balancing act between thrust and lift generation combined with pushing the aircraft to its aerodynamic limits—a natural question that arises is whether clever aerodynamic effects could be utilized to benefit such an aircraft configuration. One idea is the usage of fluid injection AFC to manipulate the lift and drag coefficients of the

wing of the aircraft. For example, it was discussed in Section 2.6 how circulation control has been shown to have the potential to increase the lift coefficient of an airfoil by as much as a factor of 2 [79, 82, 80, 84, 81, 83]. Of course, such a system would add certain weight and energy penalties in a practical application [81, 110]. For example, Jabbal et al. [110] found that separation control via the use of pulsed air jet actuators added a total systems architecture mass of 150 kg for leading-edge blowing and 30 kg for trailing-edge blowing when sized for an Airbus A320. These numbers are equivalent to 38% the A320 slat mass and 4% the flap mass, respectively. Additionally, power off-take from the integrated-drive generators was as high as 21% of total power for leading-edge blowing and 4% for trailing-edge blowing. Nevertheless, apart from previous work by the author [28], a theoretical investigation of the potential effects and benefits of applying AFC to a tiltwing aircraft in the context of the transition trajectory has not presently been done.

This chapter explores trajectory optimization of the transition maneuver for the unique tiltwing aircraft of interest in this thesis and studies the theoretical benefits of augmenting such an aircraft with fluid injection AFC. In consideration of air taxi operations (air traffic, passenger comfort, limited energy storage), the goal of this optimization is to determine minimum-energy trajectories that satisfy stall and constant-altitude constraints with and without AFC, if possible. Additionally, finding constant-altitude solutions would serve to expand the theoretically achievable flight envelope for such a tiltwing aircraft (if otherwise constrained to climbing trajectories). Hence, in this chapter, we wish to address the following research questions:

1. What does the minimum-height-change transition trajectory look like with or without AFC?

2. What is the benefit of AFC in minimizing height change during transition?

3. What is the sensitivity of the CG position on the feasibility of transition tra-

jectories?

It is assumed that AFC could be implemented by modifying the lift and drag coefficients on the main wing of the aircraft, which provides insight into the theoretical aerodynamic requirements of a tiltwing during transition. Additionally, variations in the placement of the aircraft's CG are investigated to understand the physical loading limits of this aircraft during transition. In a previous study [28], a similar transition trajectory optimization was investigated with the scaled version of this aircraft but with a 2-DoF model. The present work extends this model to 3-DoFs for the full-scale aircraft (to include pitch dynamics), and both forward and backward transition are investigated using this 3-DoF model to address the gap in the current research. Section 4.2 briefly reviews the tiltwing's model and dynamics (now in the context of 3-DoFs) and introduces the simplified model of the AFC to modify the lift and drag coefficients, Section 4.3 gives an overview of the optimization problem setup and how it was implemented, and Section 4.4 presents the results.

4.2 Simplified Model and Dynamics

To provide a wider context for this work, it was chosen to carry out these investigations using the full-scale aircraft. Due to the dynamic scaling between the demonstrator aircraft and its full-scale counterpart, the same model given in Chapter 3 still applies but with the full-scale aircraft's parameters and specifications used instead (see Tables 3.1–3.3). However, incorporating the complete 6-DoF dynamical model of Chapter 3 is both computationally costly and unnecessary for the purposes of this optimization due to the lateral symmetry of the transition maneuver, so it was instead chosen to use a reduced 3-DoF model of this aircraft's longitudinal dynamics. Hence, this section simply outlines the key differences between the complete 6-DoF model of Chapter 3 and this 3-DoF simplification as well as introduces the AFC model. Figure 4.1 below

shows the dimensions of the full-scale aircraft for reference. Note that, as variations in the aircraft's CG will be investigated in this chapter, the prop moment arms given in Table 3.2 represent their values at the nominal CG position. Also note that bold notation for vectors will be omitted in this chapter.

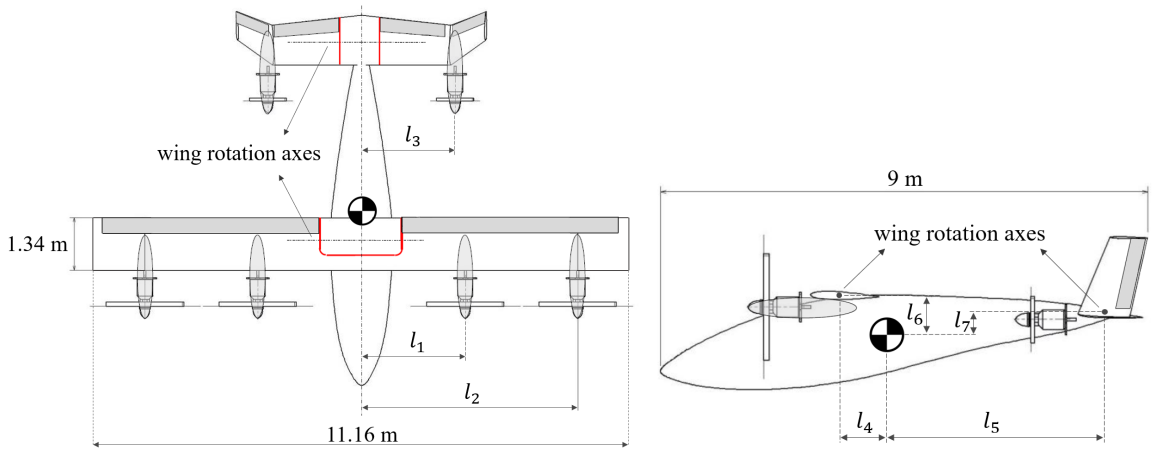

(a) Top View (b) Side View

Figure 4.1: Full-Scale Aircraft Schematic

4.2.1 Three-Degree-of-Freedom Rigid Body Equations

A V-γ coordinate system was chosen for this reduced 3-DoF model, where V is the airspeed and γ is the flight path angle. Applying Newton's Second Law to a 3-DoF rigid body gives the following state equations in this coordinate system [30]:

$$\begin{aligned}
\dot{V} &= \frac{F_{x_{net}}}{m}\cos(\theta-\gamma) + \frac{F_{z_{net}}}{m}\sin(\theta-\gamma) \\
\dot{\gamma} &= \frac{1}{V}\left(\frac{F_{x_{net}}}{m}\sin(\theta-\gamma) - \frac{F_{z_{net}}}{m}\cos(\theta-\gamma)\right) \\
\dot{\theta} &= q \\
\dot{q} &= \frac{M_{y_{net}}}{I_{yy}}
\end{aligned} \tag{4.1}$$

where the variables have their usual meanings as given in Chapters 2 and 3. Due to the 3-DoF simplification, the net external forces and moments acting on this tiltwing

aircraft now simply include gravity (F_G), thrust forces (F_T), thrust moments (M_T), and aerodynamic forces and moments (F_A and M_A, respectively).

4.2.2 Thrust Forces and Moments

Considering the symmetry involved in this simplified 3-DoF model, it was chosen to lump the thrusts together into single quantities; hence, T_W and T_T represent the total combined wing and tail thrusts, respectively. With this notation, the thrust forces can be written as:

$$\begin{aligned} F_{T_x} &= T_W \cos\xi_W + T_T \cos\xi_T \\ F_{T_z} &= -T_W \sin\xi_W - T_T \sin\xi_T \end{aligned} \tag{4.2}$$

Similarly, the pitching moment due to thrust is written as:

$$M_{T_y} = T_W \left(l_4 \sin\xi_W - l_6 \cos\xi_W\right) - T_T \left(l_5 \sin\xi_T + l_7 \cos\xi_T\right) \tag{4.3}$$

4.2.3 Aerodynamics

Only the aerodynamic lift, drag, and pitching moment are relevant in this 3-DoF longitudinal model. Furthermore, the only control surface that will be used is the elevator. Hence, the longitudinal aero coefficients for this aircraft can be simplified to be functions of the following states and inputs:

$$C_L = f(\alpha, q, \delta_e) \quad C_D = f(\alpha, q, \delta_e) \quad C_m = f(\alpha, q, \delta_e) \tag{4.4}$$

Propeller-wing interactions are computed as described in Section 3.3.3.2. However, due to the lumped thrust quantities (T_W and T_T), it is necessary for accurate prop wake modeling to separate these into the individual prop contributions when

computing the resulting axial velocity vector:

$$V_{ax} = \sqrt{\frac{2T_i}{\rho A_{prop}} + V_{\infty_\perp}^2} \cdot \hat{n}_p \tag{4.5}$$

where T_i represents the thrust generated from an individual prop, i.e. $T_i = T_W/n_W$ when modeling the slipstream velocity for the main wing and $T_i = T_T/n_T$ when modeling the slipstream velocity for the tail, where n indicates the number of props.

4.2.4 Active Flow Control Model

A simple model for the fluid injection AFC is being used that serves to modify the lift coefficient of the main wing of the aircraft [79, 82, 80, 84, 81, 83, 111, 112]. As discussed in Section 2.6, wind and water tunnel testing with leading- and trailing-edge fluid injection AFC [79, 80, 84, 83] as well as flight testing with trailing-edge fluid injection AFC [81, 83] suggest it is theoretically possible to modify the lift coefficient on an airfoil by as much as a factor of 2. Hence, it was chosen to model the fluid injection AFC by applying a factor η to the wing's C_L such that:

$$C_{L_{total}} = (1 + \eta)C_L \tag{4.6}$$

Since changes in C_L will cause an additional induced drag, C_D is modified as well using a drag polar term involving this factor η [113]. It was also desired to be able to directly modify the wing's drag coefficient itself without changing its C_L. This is modeled by simply adding an additional term μ to its original drag coefficient (physically, this could represent a deployable flap or other high-drag device) [32]. Because this AFC model modifies the local lift and drag coefficients of the main wing of the aircraft, pitching moment effects are automatically captured by this model due to the moment arm between the wing's center of pressure and the aircraft's CG. Direct effects on C_m due to AFC could also be modeled, particularly for the stall regime where pitching

moments may become highly unsteady. However, as will be discussed in Section 4.3, explicit stall constraints were implemented in this optimization to avoid the stall regime altogether. Additionally, it is assumed that the trajectories generated in this work could be used as reference inputs for an outer-loop airspeed flight controller; hence, in such a setup, a high bandwidth closed-loop control on pitch would exist that would be able to deal with any uncertainty on C_m in this model.

4.3 Optimization Methodology and Implementation

Although there has been a recent success in convex approximations of VTOL trajectory optimization for specific cases [72], in the most generic sense it is a difficult nonlinear non-convex programming problem that requires a full nonlinear optimization. While alternative techniques such as successive convexification could be employed [73, 74, 75] which allows for faster convergence, it was chosen to utilize general-purpose solvers in this work to avoid unnecessary complexity in the optimization problem setup. Hence, the trajectory optimization software chosen for this research was GPOPS-II [70], a general-purpose MATLAB software for solving nonlinear optimal control problems. GPOPS-II employs variable-order adaptive LGR quadrature orthogonal collocation to approximate the optimal control problem as a sparse NLP, which is then solved using the NLP solver IPOPT (refer to Section 2.5).

4.3.1 Optimization Problem Setup

The desired objective of this optimization is to minimize energy use throughout the transition maneuver subject to the constraints of the tiltwing's dynamics ($\dot{x}$), states (x), and inputs (u), and any path and boundary constraints (C_P and C_B, respec-

tively). This was formulated mathematically as:

$$\begin{aligned}
\min_{t,u(t)} \quad & J\left(t, x(t), u(t)\right) = \int_{t_0}^{t_F} w_1 \left(\frac{T_W}{T_{W,max}}\right)^2 + w_2 \left(\frac{T_T}{T_{T,max}}\right)^2 + w_3 \left(\frac{q}{q_{max}}\right)^2 dt \\
& \text{s.t.} \quad \dot{x} = f(x,u) \\
& \qquad x_{min} \leq x \leq x_{max} \\
& \qquad u_{min} \leq u \leq u_{max} \\
& \qquad C_{P,min} \leq C_P\left(t, x(t), u(t)\right) \leq C_{P,max} \\
& \qquad C_{B,min} \leq C_B\left(t_0, t_F, x(t_0), x(t_F)\right) \leq C_{B,max}
\end{aligned} \tag{4.7}$$

where $x = [V \quad \gamma \quad \theta \quad q]^T$, $u = [T_W \quad T_T \quad \xi_W \quad \xi_T \quad \delta_e \quad \eta \quad \mu]^T$, $f(x,u)$ is a vector function containing the nonlinear dynamics of the tiltwing, and w_i are cost weightings to be chosen. The specific cost function chosen is to reflect the fact that we wish to additionally minimize nominal thrust consumption while serving as a proxy for energy. This is to ensure maximum headroom for flight control action during transition that might be required to compensate for disturbances such as wind gusts. Additionally, a penalty on the pitch rate q was added to the cost function to further help with convergence in this 3-DoF model. Physically, this can be interpreted as a moment penalty and helped ensure smoother, less aggressive pitch maneuvers were achieved. In terms of the implication this might have on potential handling qualities of the aircraft, these are considerations for an inner-loop pitch control design, and in this work it is assumed that the trajectories generated could be used as reference inputs for an outer-loop airspeed flight controller; hence, if we were to allow the optimizer to create very aggressive pitch rate trajectories, this would instead limit control headroom in the case of disturbances such as wind gusts. The cost weightings for the wing and tail thrusts were chosen to be $w_1 = (1 + \frac{l_4}{l_5})^{-1}$ and $w_2 = 1 - w_1$, respectively, and are based on the relative weight loading on the wing and tail during hover as a function of the CG placement. The pitch weighting was nominally chosen to be $w_3 = 1$ and

was adjusted as necessary to help with convergence.

State and input constraints are summarized in Tables 4.1 and 4.2. Constant-altitude constraints were implemented by adjusting the constraints on the flight path angle γ, while AFC was implemented by adjusting the constraint bounds set for the inputs η and μ. Note that input dynamics and/or explicit rate constraints were not considered in this work. The original intention was to put explicit rate constraints on the tilt angles (the particularly slow inputs), but this unfortunately led to infeasible solutions with GPOPS-II. So instead, solutions considered were those that achieved maximum tilt angle rates of 13 deg/s for the main wing and 26 deg/s for the tail, which were judged to be feasible for the real platform. For the elevator, it is assumed that it can respond fast enough to forego the need for explicit rate constraints due to its assumed use for much faster inner-loop pitch control dynamics. Similarly, for both wing and tail rotors, it is assumed that the electric motors in the aircraft are sized to act for use as flight control actuators for inner-loop stability purposes and therefore respond sufficiently fast relative the speed of the aircraft's outer-loop airspeed dynamics, thereby also not requiring explicit rate constraints in the context of this optimization. The path constraints chosen for both forward and backward transition were upper and lower limits on α_{wet} of $\pm 15^\circ$ for the main wing and $\pm 8^\circ$ for the tail wing (to avoid stall) and a fixed starting point for the wing angles ξ_W and ξ_T. The starting wing angles for forward transition were set to 75° (the hover equilibrium point[1]) and 1° for backward transition (numerically this worked better than setting the wing angles to 0°). Note that this constraint on the initial wing angles was implemented as a path constraint instead of a boundary constraint because boundary constraints in GPOPS-II can only be functions of t or x (see (4.7) above). These constraints are summarized in Table 4.3. Lastly, Table 4.4 summarizes all initial and final conditions set along with their respective bounds (Δ values). For all runs, the

[1]Note that the hover equilibrium wing angle for the full-scale aircraft differs from the scaled model (77°) due to differing tail prop sizes; refer to Table 3.1.

final time t_F was set to 20 ± 19 s, allowing for a time horizon range of 1–39 s. This was chosen to leave the time horizon effectively unconstrained, and most trajectories tended to span approximately 15–20 s. For all forward transition results, the initial velocity was set to exactly 0.1 $\mathrm{m\,s^{-1}}$ and the final velocity was set to 67 ± 5 $\mathrm{m\,s^{-1}}$ (the aircraft's nominal cruise speed). The remaining states were set to zero, though final Δ values in pitch and pitch rate of 0.1 rad and 0.1 $\mathrm{rad\,s^{-1}}$, respectively, were chosen. These final Δ values for airspeed, pitch, and pitch rate were set to help with convergence in this two-point boundary value problem. For all backward transition results, the initial and final states are broadly the opposite. The exception is the Δ flight path angle, which was relaxed due to constant-altitude being infeasible in this direction, and the Δ pitch rate, which was set to zero to avoid a numerical spike at the last time step in the mesh. Additionally, optimization was eased by using analytical correlations for the aerodynamic coefficients of the tail wing derived from the wind tunnel data; however, validation of the nominal trajectories was done by re-running the optimization with the original lookup tables while using the solution trajectory as the starting point for the optimization. Note that all mesh and collocation setup parameters within GPOPS-II itself were left at their default values as given in the GPOPS-II reference manual [114]. Specifically, the default setup employs an *hp*-adaptive mesh refinement method, as discussed in Section 2.5 and described in [71], that divides the mesh interval into sub-intervals until a relative error estimate meets a specified tolerance. The default mesh tolerance used in this optimization was 10^{-3}.

Table 4.1: State constraints

State	**Symbol**	**Bounds**	**Units**
Airspeed	V	$[0, 75]$	$\mathrm{m\,s^{-1}}$
Flight path angle	γ	$[-\pi/2, \pi/2]$[a]	rad
Pitch angle	θ	$[-0.1, 0.1]$	rad
Pitch rate	q	$[-0.1, 0.1]$	$\mathrm{rad\,s^{-1}}$

[a] In the case of constant altitude, γ bounds were set to zero.

Table 4.2: Input constraints

Input	**Symbol**	**Bounds**	**Units**
Wing thrust	T_W	$[0, 50]$	kN
Tail thrust	T_T	$[0, 6]$	kN
Wing angle	ξ_W	$[0, \pi/2]$	rad
Tail angle	ξ_T	$[0, \pi/2]$	rad
Elevator	δ_e	$[-0.3, 0.3]$	rad
AFC lift factor	η	$[-0.5, 0.5]$[a]	–
AFC drag coefficient	μ	$[0, 1]$[b]	–

[a,b] In the case of no AFC, η and μ bounds were set to zero.

Table 4.3: Path constraints

Constraint	**Symbol**	**Bounds**	**Units**
Wing wetted AoA	$\alpha_{wet,W}$	$[-15, 15]$	deg
Tail wetted AoA	$\alpha_{wet,T}$	$[-8, 8]$	deg
Initial wing angles (forward)	ξ_0	75	deg
Initial wing angles (backward)	ξ_0	1	deg

Table 4.4: Boundary conditions

Condition	**Symbol**	**Value**	**Δ**	**Units**
Initial time	t_0	0	–	s
Final time	t_F	20	± 19	s
Initial states (forward)	x_0	$[0.1, 0, 0, 0]$	–	mixed[a]
Final states (forward)	x_F	$[67, 0, 0, 0]$	$\pm [5, 0, 0.1, 0.1]$	mixed
Initial states (backward)	x_0	$[67, 0, 0, 0]$	–	mixed
Final states (backward)	x_F	$[0.1, 0, 0, 0]$	$\pm [5, \pi/2, 0.1, 0]$	mixed

[a] The units for x are the same as for those in Table 4.1.

4.3.2 Computing a Guess Trajectory

Providing a sufficiently accurate guess trajectory as a starting point for the optimization was very important in achieving convergence in this study. In theory, computing a guess for the state trajectories is straightforward, as the state equations $\dot{x} = f(x, u)$ can be solved numerically for a given initial condition. However, to do so requires supplying trajectories for the inputs, which are also unknown. Hence, a similar procedure was taken as in Panish and Bacic [28], in which an exponential trajectory was

assumed for the wing angle, and the thrust was found at each time step that satisfied a constant-altitude constraint. However, due to the increase in complexity from 2-DoF to 3-DoF, it was necessary here to formulate this procedure more generally as the following NLP:

$$\begin{aligned} \min_{u} \quad & M_{y_{net}}^2 \\ \text{s.t.} \quad & u_{min} \leq u \leq u_{max} \\ & C \leq 0 \\ & C_{eq} = 0 \end{aligned} \tag{4.8}$$

where the net pitching moment $M_{y_{net}}$ consists of thrust and aerodynamic moments, and C and C_{eq} represent nonlinear inequality and equality constraints, respectively. With exponential trajectories already assumed for the wing angles ξ_W and ξ_T, (4.8) seeks to solve for the remaining inputs $u = [T_W \quad T_T \quad \delta_e]^T$ (no AFC in the guess). When solving this NLP for forward transition, the Earth-frame net vertical forces were chosen to be constrained to zero (i.e. to ensure constant-altitude):

$$C_{eq} = F_{x_{net}} \sin(\theta) - F_{z_{net}} \cos(\theta) \tag{4.9}$$

while for backward transition, constant-altitude was infeasible and instead an inequality constraint of:

$$C = \dot{V} \tag{4.10}$$

was set to ensure the aircraft was slowing down at all points.

To implement this, a fixed-step 4th-order Runge-Kutta solver with 100 time steps was used. This was chosen instead of variable step solvers, as the latter frequently stalled due to edge-case discontinuities and lookup tables present in this model. The assumed exponential trajectories for ξ_W and ξ_T took the following form:

$$\xi(t) = Ae^{b(t/t_F)} + B \tag{4.11}$$

where different values for A, B, and b were chosen to distinguish between a forward or backward guess trajectory solution. A fixed final-time t_F and initial conditions x_0 were specified, and (4.8) was then solved at each time step using MATLAB's `fmincon()` to determine the inputs. Lastly, the state equations for $x = [V \quad \gamma \quad \theta \quad q]^T$ (given in (4.1)) were formulated and numerically integrated.

4.4 Results

This section presents and discusses the main results from the trajectory optimization performed in this study. Results are presented for both forward and backward transition, each looking at CG variations with and without AFC. Three CG positions were investigated: nominal, forward, and aft. Forward and aft CG positions represent a 2.6% change in CG position relative the length of the aircraft, which corresponds to ± 0.23 m from the nominal point. This range of positions represents the change in CG position within the upper and lower design loading limits of the aircraft (i.e. nominal 225 kg payload ± 225 kg). All optimization constraints set are as summarized in Tables 4.1–4.4 in Section 4.3 unless otherwise specified.

4.4.1 Forward Transition

4.4.1.1 Without AFC

Forward transition results without AFC are presented in Figures 4.2 and 4.3 for each CG position. In general, convergence was straightforward and constant-altitude trajectories could be achieved without issue. It can be seen that changes in CG position minimally affect the optimal trajectory with a few exceptions. Most notable are the differences in main wing AoA, where a more forward CG position sees a higher peak AoA, pushing it closer to its 15° stall limit at max payload ("Forward CG"). Conversely, the tail wing sees a smaller variation in AoAs with different CG positions.

These correlations can be explained by the corresponding thrust demands. A further-forward CG position demands a lower total wing thrust on average, resulting in a lower prop-induced velocity and therefore a higher wetted AoA (refer to (3.23) and (3.24)). The opposite is true for the tail, though with thrusts an order of magnitude smaller than the wing, the resulting wetted AoA differences are much smaller. Another exception is the elevator demands. It can be seen that the elevator begins to saturate at the beginning of the trajectory for more forward-shifted CG positions. In fact, a more negative elevator deflection is commanded overall in this case despite higher tail thrusts on average, suggesting that both the tail props and elevator are trying to compensate for a positive pitching moment resulting from the forward-shifted CG position. This would also explain why a lower wing thrust is demanded for more forward-shifted CG positions despite there being a smaller moment arm between the wing props and the aircraft's CG. It is interesting to note the V-shaped response of the total wing thrust in these results. This is likely due to a combination of the pitching moment balance and a means to lower the overall energy cost where possible. It can be seen that there is an inverse correlation between the magnitudes of T_W and T_T depending on CG position in this region, and this dip in total thrust occurs past the peak of the wing AoA curve, suggesting there is freedom at this point to reduce the total thrust without violating stall constraints (again recall that higher thrust can help reduce local wetted AoA). Of additional note is the shape of the tail wing angle curve relative the main wing angle. Specifically, it was found that it was necessary for the tail wing to rotate more quickly to avoid aerodynamic stall as the aircraft accelerates forward. This is because the much lower thrust produced by the tail props (due to having a larger moment arm to work with) results in lower prop-induced velocities and therefore higher wetted AoAs (as mentioned previously); hence, it is necessary to get the tail angle horizontal as quickly as possible before the freestream airspeed becomes too large. This was a pattern first discovered when

computing the guess trajectory (Section 4.3.2). Lastly, it can be seen that there is a drop in thrust at the end of the maneuver. In fact, this pattern can be observed for all runs to some degree. We believe this is likely a result of the energy-based cost function: If the aircraft has enough momentum to reach the specified final conditions for the problem, then this would allow the thrusts to drop right at the end of the maneuver to marginally improve the total cost. This is further substantiated by the fact that this pattern is present in all runs; despite differences in the problem setup (initial/final conditions, state and input constraints) to produce different trajectories, the cost function always remains the same.

4.4.1.2 With AFC

In comparison, forward transition results with AFC are presented in Figures 4.4 and 4.5. Constant-altitude forward trajectories were already achievable without AFC, but here it can be observed that the use of AFC in this context serves to maximize the performance of the aircraft and decrease its total transition time. Specifically, it can be observed that the main wing AoAs now peak at the 15° stall limit for all CG positions, made possible by temporarily decreasing the lift coefficient through η. Hence, it would appear that pushing the aircraft close to its stall AoA serves to minimize energy throughout transition. These findings are consistent with those found in [24]. A similar pattern can be observed for the tail wing, with the exception of the "Forward CG" position. However, the lower wetted AoA in this case can be partially attributed to the higher thrust demand from the tail props, analogous to the non-AFC case. Also analogous to the non-AFC case are the elevator demands, where the nominal and forward-shifted CG positions see the elevator saturating at the start of the maneuver. What is different, however, is the elevator demand for the aft-shifted CG position. Here, more negative values are demanded compared to the non-AFC case. This can be attributed to the higher wetted tail AoA (nearly double

compared to the non-AFC case), suggesting that the elevator is using its control authority to adjust the tail wing's lift coefficient to push the tail AoA up to its stall limit, similar to η for the main wing. This also further explains why the "Forward CG" tail AoA does not reach its 8° limit—the elevator demand is most saturated in this case. In terms of overall transition time, it is found to be marginally quicker than the non-AFC case at approximately 14 s. Additionally, there is a smaller variation in total transition time between each CG position. In fact, smaller variations in the trajectory for each CG position can be observed in general compared to the non-AFC case, suggesting that using AFC has the additional benefit of increasing the aircraft's robustness to CG variations in terms of its expected performance during transition. Note that μ bounds were set to zero for these runs instead of $0 \leq \mu \leq 1$; it was found otherwise that μ demanded small intermittent spikes on the order of 10^{-4}, which were numerically insignificant and thus omitted.

4.4.2 Backward Transition

4.4.2.1 Without AFC

Backward transition results without AFC are presented in Figures 4.6 and 4.7 for each CG position. Here, it was not possible to achieve convergence with constant-altitude constraints, unlike for forward transition. Tightening the bounds on the flight path angle γ was attempted but either didn't work or resulted in numerically sensitive trajectories, so γ was left unconstrained. This necessary climbing-trajectory behavior was consistent with findings from a previous investigation [28]. Additionally, the pitch weighting w_3 was adjusted to 0.1 for all runs to help produce numerically smoother trajectories. In comparing these backward climb-transition trajectories, it was found that CG placement had an effect on the total change in altitude. Specifically, it can be observed that the further aft the CG position, the smaller the total altitude change, with a maximum difference of 35.1 m and an average total altitude increase

of 259.7 m. Additionally, a further-aft CG position results in a marginally quicker transition time, with a range of 12–13.2 s across these three runs. Interestingly, it can be observed that the tail AoAs reach their peak values from the middle to end of the maneuver and in correlation with the pitch angle curve leveling out. This suggests that a slowly increasing pitch angle is preferred up to the limit in which stall constraints are reached. In parallel, it can be observed that the tail thrusts and elevator ramp up in demand to produce a more negative pitching moment on the aircraft. However, saturation of the elevator and higher tail thrust demands occur for further-aft CG positions, likely due to the smaller tail prop moment arm and thus smaller pitch control effectiveness of these actuators. Lastly, while the tail wing reaches its stall limit in all runs, the main wing AoAs never peak above just over 8°, and there is minimal variation in their curves for each CG position. As will be further evidenced in the backward transition results with AFC, it is believed that this behavior (the tail AoAs saturating but the wing AoAs not) indicates that the primary factor affecting the feasibility of constant-altitude backward transition is the AoA constraints in the optimization problem formulation. Secondarily, the specific aircraft configuration itself plays a role in this result, both because the tail props' larger moment arm results in a lower thrust condition, which in turn means a lower contribution to reducing the wetted AoAs on the tail wing, and due to the low stall limit of the tail airfoil itself (as will be seen next).

4.4.2.2 With AFC

In comparison, backward transition results with AFC are presented in Figures 4.8 and 4.9. Similarly to the non-AFC case, the pitch weighting w_3 was set to 0.1. Most significantly, the use of AFC did not make it possible to achieve constant-altitude transition, and hence γ constraints were relaxed similarly to the non-AFC case. This result contradicts the conclusion found in [28], where AFC made constant-altitude

backward transition possible. However, the key difference here is the extension to a 3-DoF model and the addition of modeling the tail wing. Specifically, it can be observed that the tail AoAs reach their stall limit for nearly the entirety of the maneuver, but in parallel, the wing AoAs appear to "saturate" at roughly 9° despite having a stall limit of 15°. Thus, what can be observed is that the aircraft is being pushed to the limits of its performance through the use of AFC (as for the forward case), but its primary limiting factor has become the 8° stall angle on the tail. To confirm this, an alternative backward transition trajectory solution was produced (see Figures 4.10 and 4.11) using a different airfoil for the tail wing (NACA 63(2)-615, same as the main wing) with a higher stall angle of attack (15°). It was found in this case that constant-altitude backward transition was possible with the use of AFC and an aft-shifted CG. In addition, while the tail AoA similarly saturates for the majority of this constant-altitude trajectory, the main wing AoA now varies freely within its bounds. These results suggest that a redesign of the tail wing of this unique tiltwing aircraft (specifically, a higher stall limit) could better leverage the benefits of using AFC in this application. Nevertheless, this aircraft still sees benefits from using AFC in its current design. Specifically, while constant-altitude cannot be achieved, the use of AFC serves to provide a 24.2% reduction in the average total altitude change compared with the non-AFC case. In consideration of operating in urban environments and potentially dense airspace, this could still justify the complexity of including an AFC system in this aircraft. Additionally, the use of AFC narrows the general variation in performance of the aircraft for each CG position (just as for forward transition), similarly suggesting that the use of AFC increases the aircraft's robustness to CG position changes in terms of its expected transition performance. This is particularly noticeable in the airspeed, altitude, and wing angle curves of Figure 4.8. An interesting feature to note is the temporary dip in tail AoAs (Figure 4.9). The cause of this can be traced back to the moment the elevator saturates, requiring the

tail thrusts to ramp up to maintain pitching moment authority. The increase in tail thrust temporarily decreases the tail AoAs due to the increased prop-induced velocities, but as the pitch angles continue to increase, the tail AoAs eventually increase again. Similarly to the non-AFC case, when the tail AoAs reach their stall limit (once again), the only option is to level out the change in pitch angle. Another interesting feature to note is how both the wing and tail thrusts are held almost constant for the first half of this trajectory. We believe this observation can be explained by the saturation of the tail AoA. Because the tail airfoil is sitting right at its stall limit, a further reduction in thrust would violate stall constraints (due to thrust's effect on local wetted AoA). And any increase in thrust (to reduce the tail AoA) would only reduce deceleration of the aircraft and increase total energy consumption. Lastly, it can be observed that the generation of additional drag through μ is demanded at the end of the trajectory (interestingly correlated with the pitch angle curve leveling out). However, this can be simply attributed to a means of reducing the total energy cost by more quickly slowing down the aircraft at the end of the transition maneuver. Specifically, compared with a solution with μ bounds set to zero in the nominal case, the use of μ here reduces the transition time by just under 1 s and results in a cost decrease of 13.7%. Interestingly, in the case that the aircraft's performance is not limited by the tail wing's stall limit (Figures 4.10 and 4.11), the minimum-energy trajectory prefers a high μ value initially to slow down the aircraft combined with a transition time nearly triple the length in order to achieve constant-altitude backward transition.

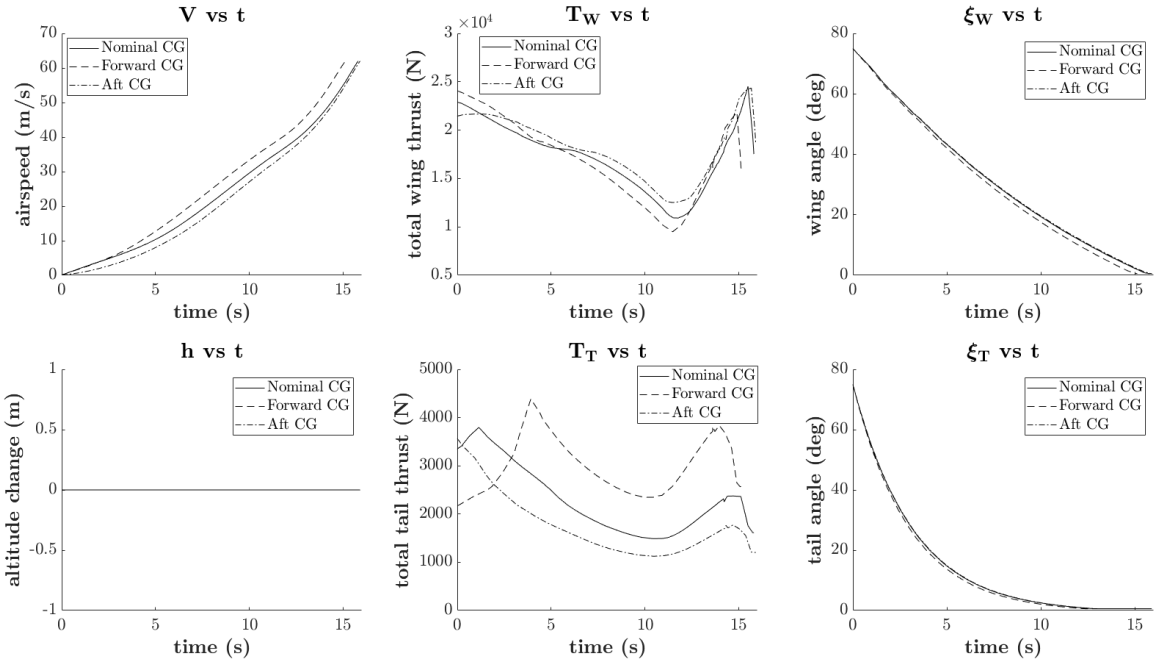

Figure 4.2: Forward trajectories without AFC – airspeed, altitude, thrusts, and wing angles. CG variation of ±2.6% the aircraft's length

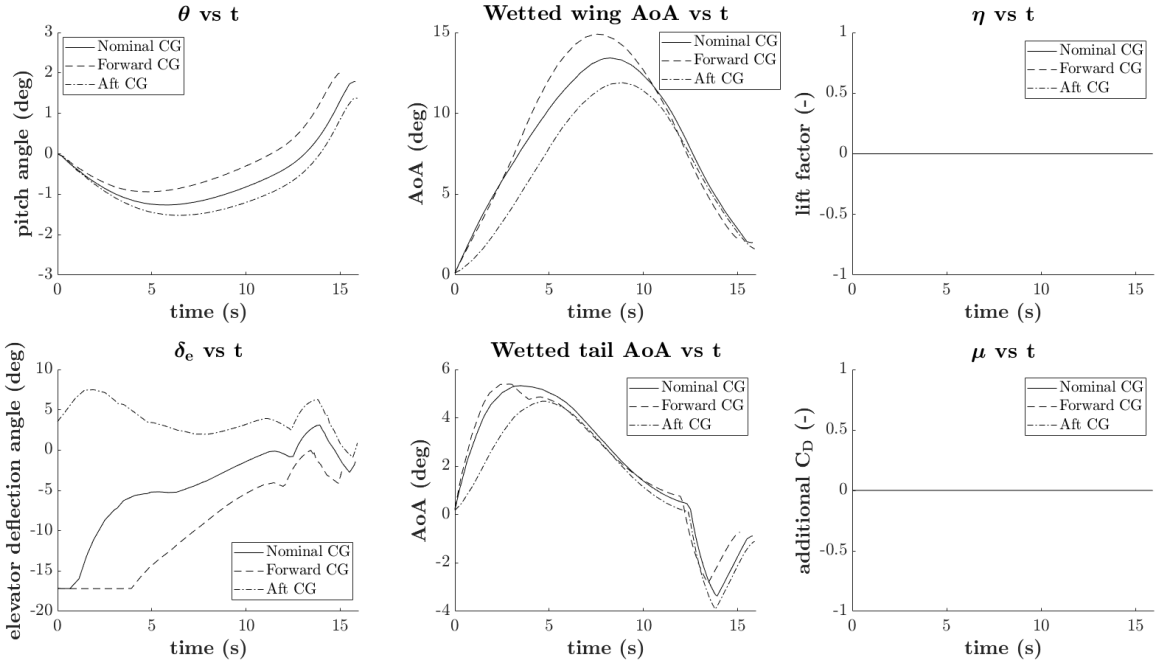

Figure 4.3: Forward trajectories without AFC – pitch angle, elevator deflection, wetted AoAs, and AFC demands. CG variation of ±2.6% the aircraft's length

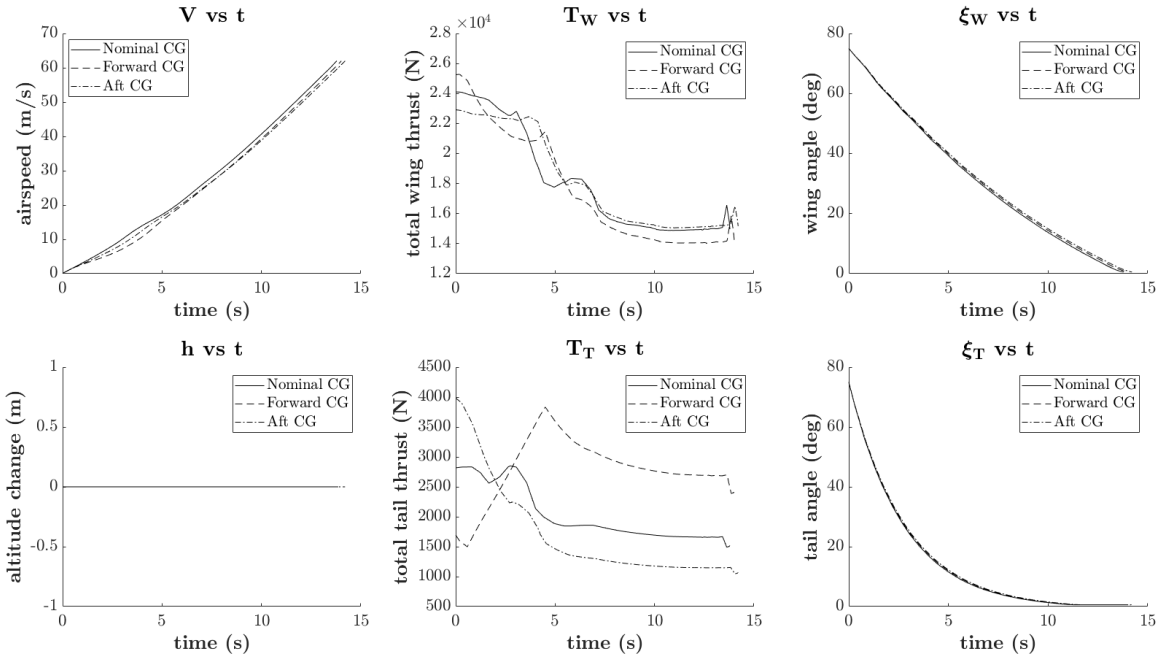

Figure 4.4: Forward trajectories with AFC – airspeed, altitude, thrusts, and wing angles. CG variation of ±2.6% the aircraft's length

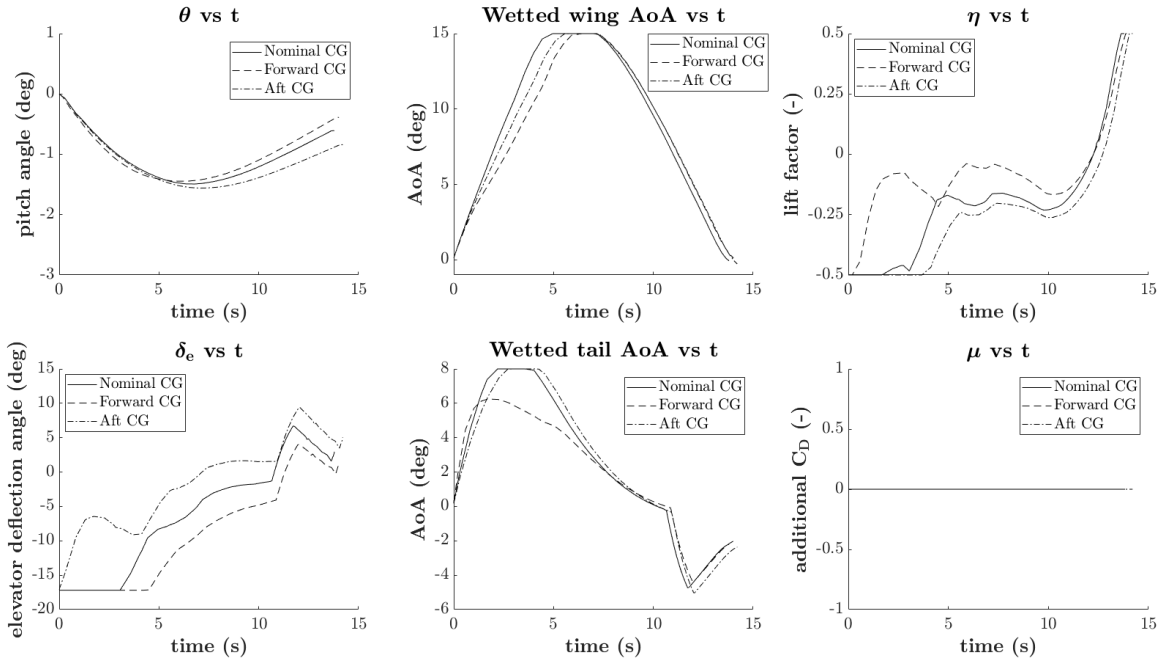

Figure 4.5: Forward trajectories with AFC – pitch angle, elevator deflection, wetted AoAs, and AFC demands. CG variation of ±2.6% the aircraft's length

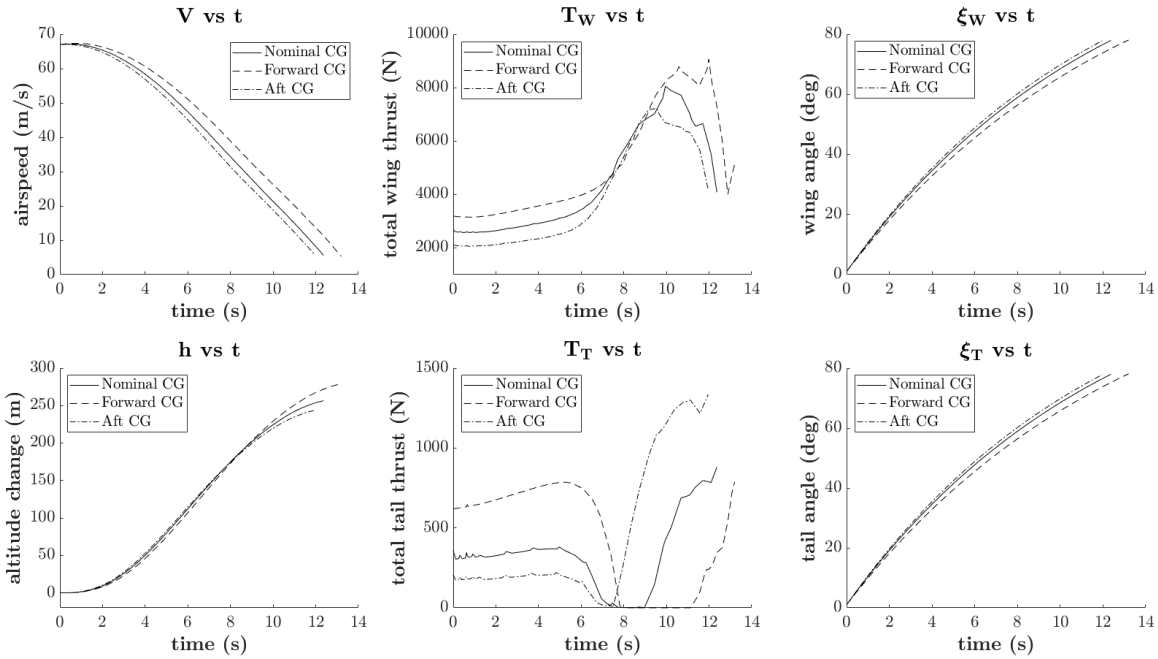

Figure 4.6: Backward trajectories without AFC – airspeed, altitude, thrusts, and wing angles. CG variation of ±2.6% the aircraft's length

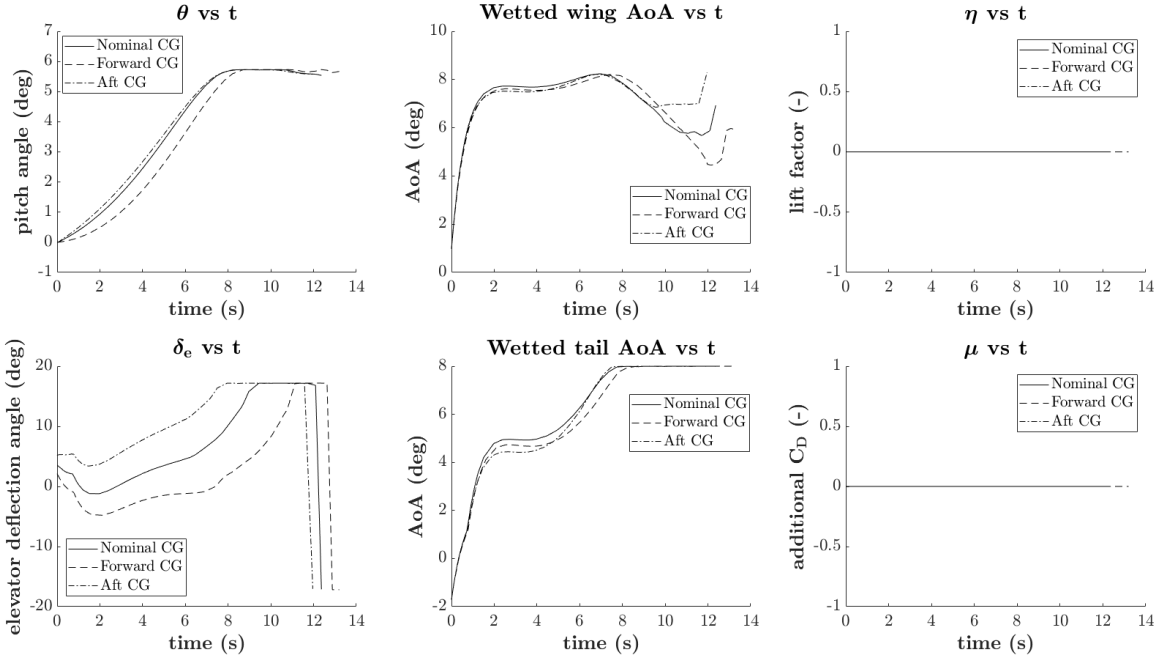

Figure 4.7: Backward trajectories without AFC – pitch angle, elevator deflection, wetted AoAs, and AFC demands. CG variation of ±2.6% the aircraft's length

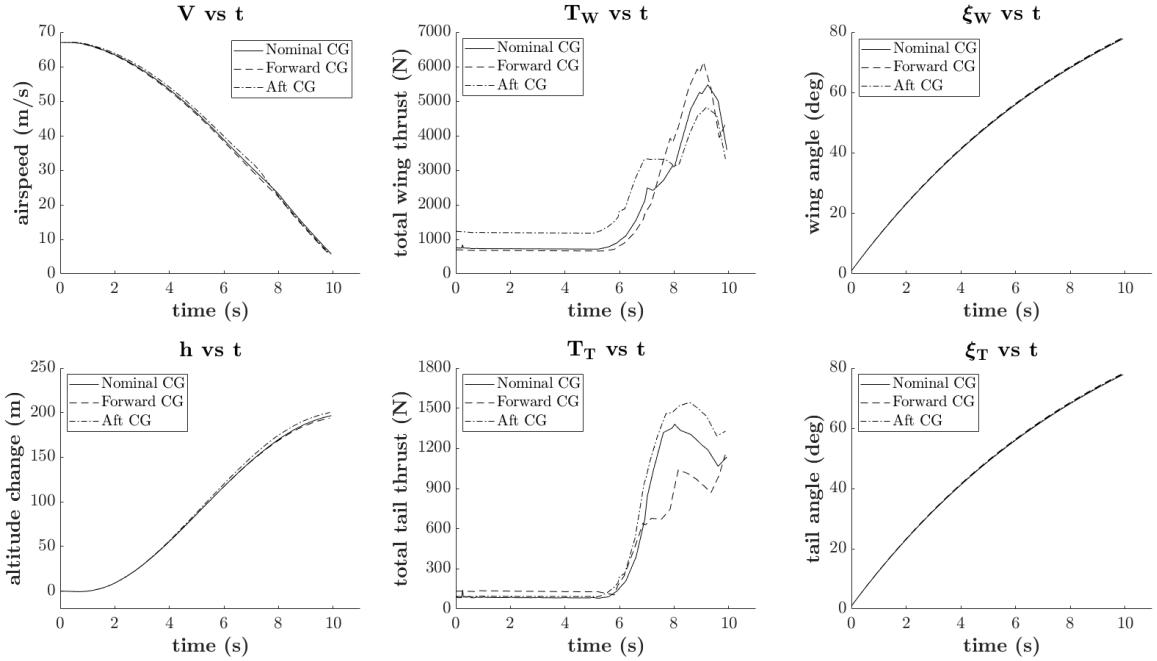

Figure 4.8: Backward trajectories with AFC – airspeed, altitude, thrusts, and wing angles. CG variation of ±2.6% the aircraft's length

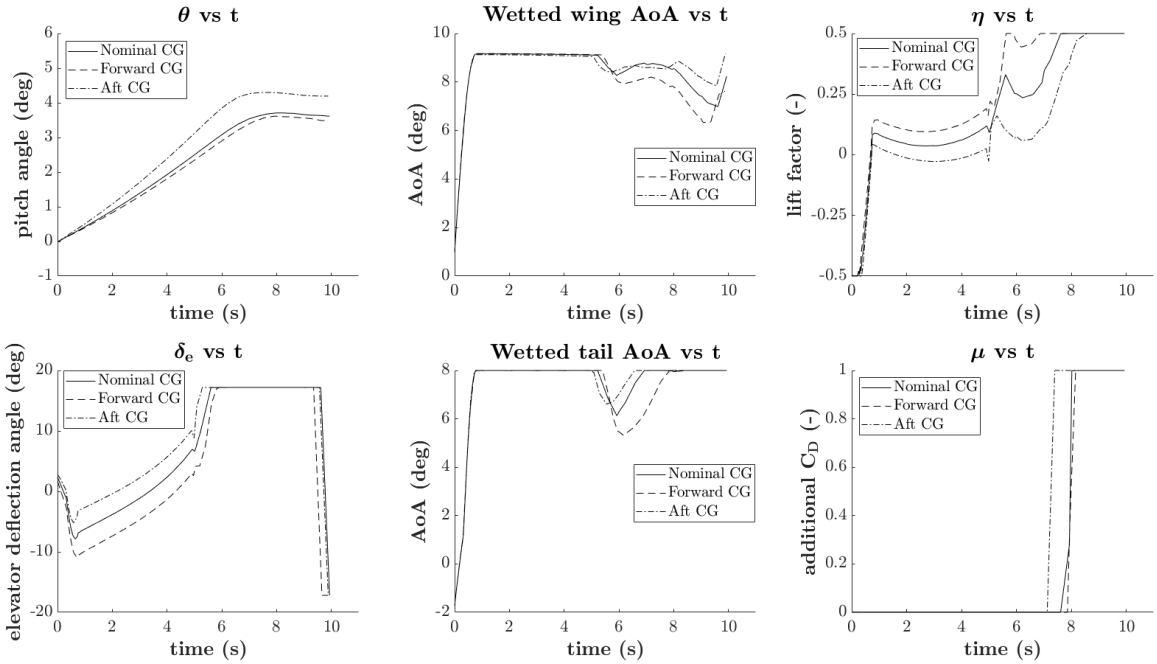

Figure 4.9: Backward trajectories with AFC – pitch angle, elevator deflection, wetted AoAs, and AFC demands. CG variation of ±2.6% the aircraft's length

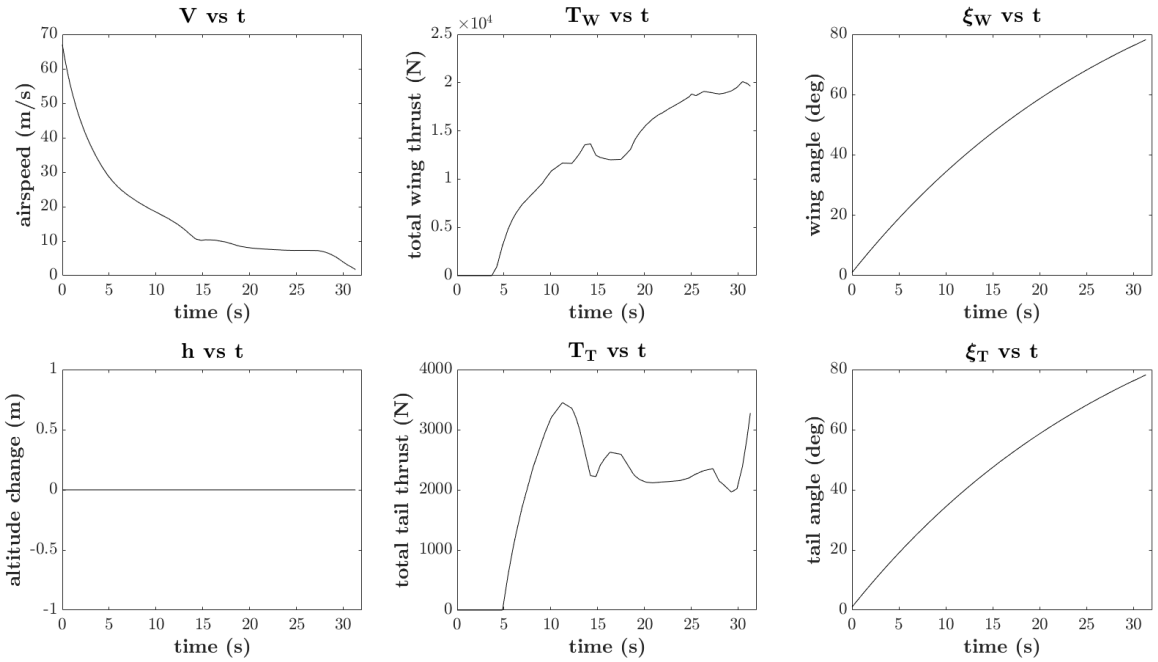

Figure 4.10: Backward trajectory with NACA 63(2)-615 tail airfoil, AFC, and aft-shifted CG – airspeed, altitude, thrusts, and wing angles

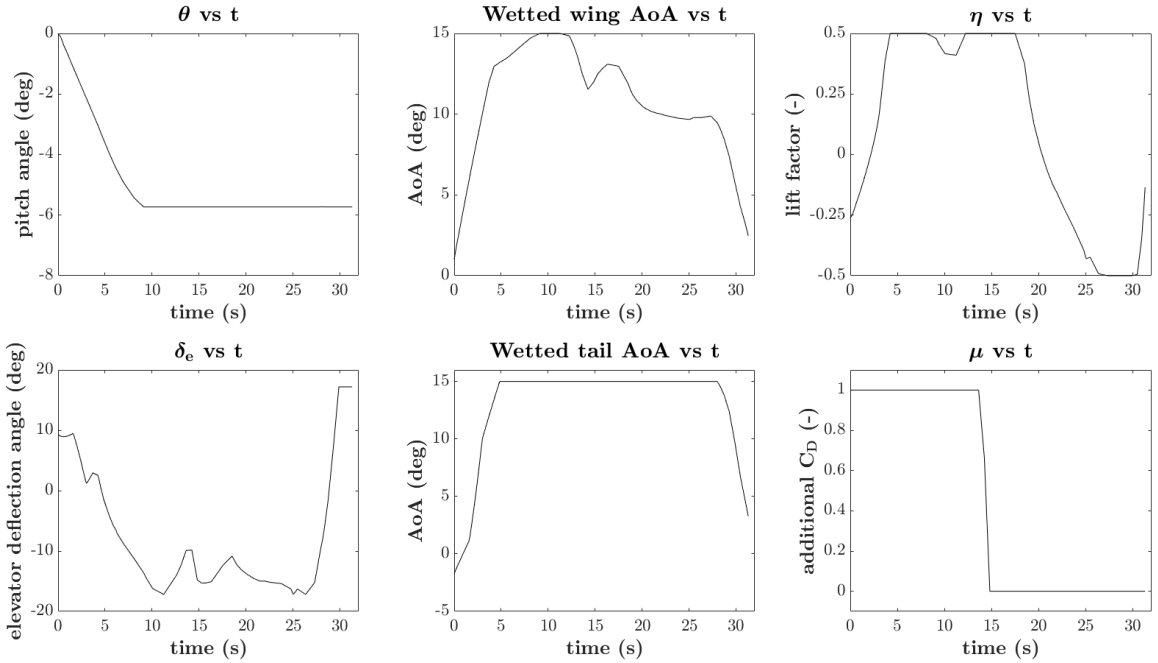

Figure 4.11: Backward trajectory with NACA 63(2)-615 tail airfoil, AFC, and aft-shifted CG – pitch angle, elevator deflection, wetted AoAs, and AFC demands

4.5 Conclusion

This chapter explored transition trajectory optimization for this unique tiltwing design and studied the theoretical benefits of augmenting such an aircraft with fluid injection AFC. The goal of this optimization was to find minimum-energy trajectories that satisfied constant-altitude and stall constraints with and without AFC, if possible, and to present 3-DoF solutions and both forward and backward transition trajectories to address the gap in the current research. AFC was implemented by modifying the lift and drag coefficients on the main wing of the aircraft, providing insight into the theoretical aerodynamic requirements of a tiltwing during transition, and, additionally, variations in the placement of the aircraft's CG were investigated to understand the physical loading limits of the aircraft during this maneuver.

It was found that constant-altitude transition could be achieved for forward transition without the need for AFC. However, the use of AFC served to maximize the performance of the aircraft, decrease the total transition time, and increase the robustness of the aircraft to CG variations in terms of expected performance during transition. The backward transition maneuver gained these same benefits of using AFC, but it was not possible to achieve constant-altitude transition, even with the use of AFC. Specifically, the tail airfoil's stall constraint was found to be a limiting factor in backward transition performance, and the use of AFC was not sufficient to overcome this design limit. A comparative constant-altitude backward transition solution using a different airfoil with a higher stall limit suggests that a redesign of the tail wing of this unique tiltwing aircraft could better leverage the benefits of using AFC in this application. Nevertheless, the use of AFC in the aircraft's current design allowed for a 24.2% reduction in the average total altitude change during the backward transition maneuver.

In the future, it would be of interest to investigate acceleration constraints and/or include an acceleration penalty in the cost function in consideration of passenger

comfort in an air taxi application. It would also be of interest to implement input dynamics and/or explicit rate constraints, particularly for the tilt angle actuation. Additionally, improved mathematical modeling of the AFC would allow for better comparison of the optimization results to the physical control limits of fluid injection AFC on the real tiltwing aircraft. Lastly, this work focused solely on the transition maneuver itself without explicit considerations of vertical take-off and landing or approaching, for example, steady-state cruise conditions. Future work could investigate a multiphase optimization approach in which separate flight phases could be defined for hover, transition, and cruise, and fortunately the GPOPS-II framework used in this work would allow for direct expansion to this in the future. Such a problem setup would allow for smoother connection between hover and cruise mode operating point conditions as well as the ability to model, for example, the energy cost associated with vertical take-off and landing before/after transition.

Chapter 5

Inner-Loop Design and Testing

5.1 Introduction

The flight control of a tiltwing aircraft is a particularly challenging problem due to the significant variation of flight dynamics as the aircraft transitions from hover to cruise flight. This necessitates an attitude controller that is operable over a wide range of flight conditions. A standard method of solving this is to use a gain scheduling controller, where the nonlinear aircraft dynamics are linearized at many points within the full flight envelope and individual controller tunings are found for each operating point. Plenty of examples of this control strategy applied to tiltwing aircraft for the attitude control design exist [115, 116, 117, 118, 119, 120, 121, 122, 123]. The advantage of such a method is the ability to employ classical linear control design techniques at each linearization point. The drawback, however, is the onerous task on the designer of tuning gains at many operating points and the effect of gain switching on the system. An alternative approach is to implement nonlinear control laws such as feedback linearization, which removes the need for gain scheduling altogether (at least within the context of the feedback linearization controller). Hence, among the numerous studies of tiltwing flight control design and testing within the

last two decades [124], full-envelope attitude control design through the application of dynamic inversion-based control laws (i.e. feedback linearization or NDI) has become increasingly more popular in recent years [88, 89, 125, 90, 126, 97, 127, 128, 129]. However, of these examples, only three have investigated flight testing [88, 89, 90].

Cetinsoy et al. [88] consider a quad tiltwing design and pair their NDI attitude controller with linear PID control laws. Due to the decoupling of the underlying dynamics through the NDI transformation, three individual SISO control loops are implemented with virtual controls to regulate the roll, pitch, and yaw of the aircraft. Masuda and Uchiyama [89] also consider a quad tiltwing design but instead choose to investigate NDI paired with H_∞ control with a focus on robustness to wind disturbances. The same controller was then modified a year later in [125], employing μ-synthesis instead of H_∞ to improve the transient response due to reduced conservativeness in the robust design. Rohr et al. [90] take a different approach by pairing their NDI with a nonlinear control law, namely a quaternion-based attitude controller that maps attitude errors to desired body-frame angular accelerations using a cascaded P-PID-structure. This was then applied to a scaled hobbyist model of the Canadair CL-84 tiltwing. Lastly, Autenrieb et al. [126] take an adaptive control approach by investigating the use of an adaptive neural network to help mitigate inversion errors, while Milz et al. [97, 127, 128, 129] investigate INDI as well as sensory NDI[1] applied to the Airbus A^3 Vahana.

The goal of this thesis is to design a full-envelope flight controller for the unique tiltwing aircraft configuration of present interest that is operable in all modes of flight (hover, cruise, and transition) through the application of feedback linearization. Additionally, it is of interest to understand how a model-based inversion design (such as NDI, contrasting INDI) performs with a configuration such as a tiltwing, where

[1] A variant of INDI that similarly utilizes sensor-based estimation but without requiring incremental control commands. Note that both INDI and sensory NDI therefore suffer in performance when sensors are noisy.

there is much greater uncertainty in the aerodynamic modeling compared to, for example, a conventional fixed-wing aircraft. To this end, this chapter presents the first step in this endeavor: the inner-loop attitude controller design. NDI is chosen for the attitude rate control of the aircraft, while off-line weighted pseudo-inverse control allocation is implemented to enable versatility in the control design and which provides fault-tolerant capabilities due to the over-actuated nature of this system. The performance of this inner-loop design is then predicted through simulation and assessed through hover flight testing, and the controller's robustness is explored in the context of motor failure. Section 5.2 begins with a general overview of the complete full-envelope flight controller that will be developed over the next three chapters (Chapters 5, 6, and 7). Section 5.3 then discusses the design of the inner-loop attitude controller. Simulation results are then presented in Section 5.4 followed by flight testing in Section 5.5.

5.2 Full-Envelope Flight Controller Overview

Figure 5.1 below shows a schematic of the general full-envelope flight controller architecture used in this research. Motivated by the goal of developing a full-envelope design for this tiltwing aircraft through the application of feedback linearization, the general flight controller structure chosen employs an inner-/outer-loop cascade design under assumptions of time-scale separation involving an inner-loop attitude controller and an outer-loop velocity controller that are combined with control allocation. The inner-loop attitude controller continuously regulates the rotational axes of the aircraft to ensure stability at all points within the flight envelope by outputting generalized moment demands ($\mathbf{M}$). The outer-loop velocity controller tracks motion commands demanded from the pilot ($\mathbf{V}_{ref}$) and outputs generalized force demands ($\mathbf{F}$) in addition to supplying attitude reference inputs ($\mathbf{A}_{ref}$) to the inner-loop. Hence, both

of these controllers work in a virtual input space (generalized forces and moments) and interface with control allocation to enable versatility in the design and help ensure validity over the complete flight envelope. This setup additionally provides fault-tolerant capabilities due to the over-actuated nature of this system (i.e. more effectors than degrees of freedom). Finally, a control allocation scheme converts the demanded virtual controls (**F** and **M**) to real input demands (**u**) while taking into account actuator effectiveness throughout the flight envelope.

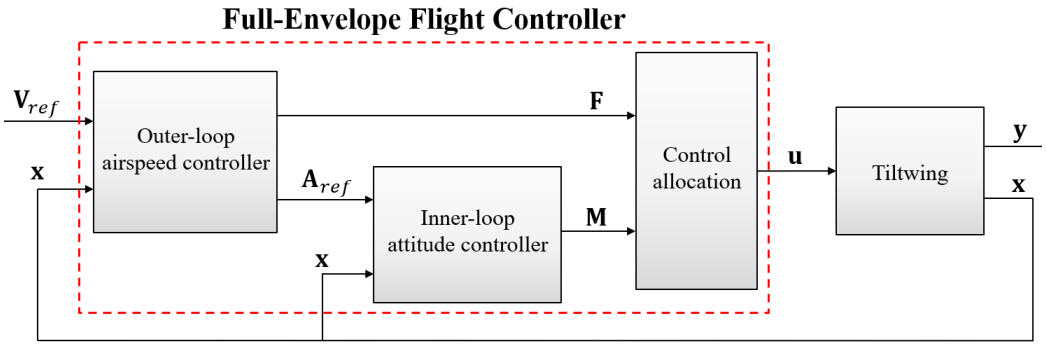

Figure 5.1: General full-envelope flight controller architecture

5.3 Inner-Loop Design

5.3.1 Controller Architecture

Following from the general full-envelope flight controller architecture of Figure 5.1, a schematic is given in Figure 5.2 that illustrates the specifics of the inner-loop flight controller design. Firstly, the "Inner-loop attitude controller" subsystem of Figure 5.1 now consists of two subsystems in Figure 5.2: an NDI attitude rate controller cascaded with an Euler angle tracking controller. Roll and pitch angle references (ϕ_{ref}, θ_{ref}) along with a yaw rate reference ($\dot{\psi}_{ref}$) from the pilot are converted within the Euler angle tracking subsystem to body-frame rotation rate references ($\mathbf{\Omega}_{b,ref}$) for the NDI controller. The NDI controller, which regulates the attitude rates of the aircraft across the entire flight envelope, outputs its virtual moment demands (**M**)

to the control allocation. In parallel, for the purposes of testing and evaluation in hover mode, an additional reference input from the pilot (F_z) demands a vertical force request in the analogous virtual force space (i.e. $\mathbf{F} = F_z$). The control allocation then maps these collective virtual controls to the tiltwing's real actuator inputs $\mathbf{u} = [\omega_{Li} \quad \omega_{Lo} \quad \omega_{Ri} \quad \omega_{Ro} \quad \omega_{Tl} \quad \omega_{Tr} \quad \delta_a \quad \delta_e \quad \delta_r \quad \delta_f \quad \xi_W \quad \xi_T]^T$. Note that, in the context of the control allocation, the flap and wing angle inputs (δ_f, ξ_W, and ξ_T) are treated as system parameters and are thus not allocated based on the virtual controls. Furthermore, for hover-mode flight, the elevator and rudder inputs (δ_e and δ_r) are not used, while the ailerons (δ_a) are used to contribute to, primarily, yaw control. As was briefly mentioned in Section 5.2, because this NDI attitude rate controller does not directly interact with the real inputs $\mathbf{u}$ and because this tiltwing is an over-actuated system, it is possible to perfectly allocate the virtual controls even in light of actuator failures, such as a motor power failure. Lastly, to deal with actuator saturation, excess virtual moments $\delta\overline{\mathbf{M}}$ (i.e. those which cannot be achieved) are fed back from the control allocation to the NDI attitude rate controller as part of an integral anti-windup strategy.

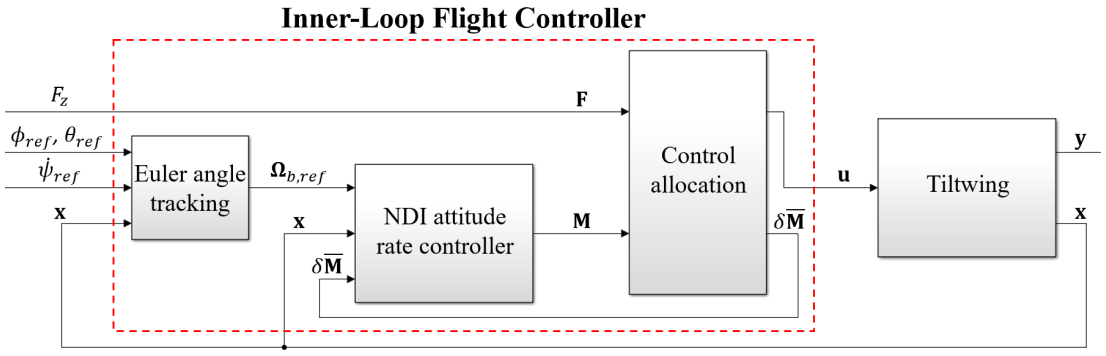

Figure 5.2: Inner-loop flight controller composed of inner-loop attitude controller subsystem (NDI attitude rate control + Euler-angle tracking) and control allocation.

5.3.2 NDI Attitude Rate Controller

To apply NDI control to the attitude rates of this aircraft, we first write the attitude rate dynamics in the input-affine form $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{M}$:

$$\begin{bmatrix} \dot{p} \\ \dot{q} \\ \dot{r} \end{bmatrix} = \begin{bmatrix} \dot{p}_{passive} \\ \dot{q}_{passive} \\ \dot{r}_{passive} \end{bmatrix} + \begin{bmatrix} \frac{-I_{zz}}{I_{xz}^2 - I_{xx}I_{zz}} & 0 & \frac{-I_{xz}}{I_{xz}^2 - I_{xx}I_{zz}} \\ 0 & \frac{1}{I_{yy}} & 0 \\ \frac{-I_{xz}}{I_{xz}^2 - I_{xx}I_{zz}} & 0 & \frac{-I_{xx}}{I_{xz}^2 - I_{xx}I_{zz}} \end{bmatrix} \begin{bmatrix} M_x \\ M_y \\ M_z \end{bmatrix} \tag{5.1}$$

$$\mathbf{y} = \begin{bmatrix} p & q & r \end{bmatrix}^T$$

where $\mathbf{M} = [M_x \quad M_y \quad M_z]^T$ represents generalized moment inputs that depend on the real inputs $\mathbf{u}$ (i.e. the effector model) and where $[\dot{p}_{passive} \quad \dot{q}_{passive} \quad \dot{r}_{passive}]^T$ are the passive attitude rate derivatives $\dot{\mathbf{\Omega}}_{b,passive}$ (i.e. the composition of $\dot{\mathbf{\Omega}}_b$ that only depends on the current state of the system, not on the real inputs $\mathbf{u}$) such that:

$$\dot{\mathbf{\Omega}}_{b,passive} = \mathbf{J}^{-1}\left(-\mathbf{\Omega}_b \times (\mathbf{J}\mathbf{\Omega}_b) + \mathbf{M}_{A_{passive}}\right) \tag{5.2}$$

$\mathbf{M}_{A_{passive}}$ above indicates the passive aerodynamic moments (i.e. all aerodynamic moments *not* due to the aircraft's control surfaces nor induced by prop wake), which are the only external moments that do not depend explicitly on the input. These moments can be expressed using total-aircraft moment coefficient quantities[2] as:

$$\mathbf{M}_{A_{passive}} = \mathbf{R}_{bs}(\alpha) \cdot 0.2\bar{q}S \begin{bmatrix} bC_l & \bar{c}C_m & bC_n \end{bmatrix}^T \tag{5.3}$$

where α is the freestream AoA, $\bar{q}$ is the freestream dynamic pressure, the multiplier of 0.2 is due to the 20% unwetted split designated for this model, and $\mathbf{R}_{bs}(\alpha)$ converts these aerodynamic moments from the stability frame to the aircraft's body frame.

[2]Note that the complete/expanded expression for (5.3) sees individual term contributions from the main wing, tail, and fuselage as per the 6-DoF model described in Chapter 3.

By equating a new input $\mathbf{v}$ to $\dot{\mathbf{x}}$, (5.1) can be simply rearranged as:

$$\mathbf{M} = \mathbf{g}^{-1}(\mathbf{x})\left(\mathbf{v} - \mathbf{f}(\mathbf{x})\right) \tag{5.4}$$

noting that $\mathbf{g}(\mathbf{x})$ is simply the inverse of the aircraft's symmetric inertia tensor $\mathbf{J}$ and is therefore always invertible. As discussed in Section 2.3.2, $\mathbf{v}$ is a new input vector generated by a linear controller acting on a (decoupled) linear system whose state vector is $\mathbf{z} = \mathbf{x}$. However, to ensure zero steady-state error, this linear system was augmented with integrators such that $\mathbf{z} = \begin{bmatrix} p & q & r & \int p & \int q & \int r \end{bmatrix}^T$, rendering the linear system to be:

$$\begin{aligned} \dot{\mathbf{z}} &= \begin{bmatrix} \mathbf{0}_{3\times3} & \mathbf{0}_{3\times3} \\ \mathbf{I}_{3\times3} & \mathbf{0}_{3\times3} \end{bmatrix} \mathbf{z} + \begin{bmatrix} \mathbf{I}_{3\times3} \\ \mathbf{0}_{3\times3} \end{bmatrix} \mathbf{v} \\ \mathbf{y} &= \begin{bmatrix} \mathbf{I}_{3\times3} & \mathbf{0}_{3\times3} \end{bmatrix} \mathbf{z} \end{aligned} \tag{5.5}$$

Under the assumption of full-state feedback availability, the LQR controller design methodology was utilized to implement a control law of the form $\mathbf{v} = \mathbf{K}_{LQR}(\mathbf{z}_{ref} - \mathbf{z})$. LQR has desirable inherent robustness guarantees [130], but it is important to note that in practice the rotation rates are obtained by an extended Kalman filter (EKF) embedded in the on-board inertial measurement unit (IMU) that fuses the sensor information from real gyroscope, accelerometer, and magnetometer measurements. While no robustness guarantees can therefore be formally given in this particular setup [131] without carrying out loop transfer recovery procedures [132], the closed-loop bandwidth of our flight controller design is much lower than the loop rate of the IMU + EKF; hence, effective full-state feedback is assumed in this design. In consideration of the effect of model error and measurement error on this approach, the robustness of this design to aerodynamic modeling uncertainty will be explored in Chapter 7. Additionally, the controller's ability to handle external disturbances,

such as wind gusts, will depend on the system's achievable bandwidth.

Following the state feedback control loop on $\mathbf{z}$, the resulting reference-to-output closed-loop transfer matrix has the following (decoupled) form:

$$\mathbf{T}_{pqr}(s) = \begin{bmatrix} \frac{K_P^p s + K_I^p}{s^2 + K_P^p s + K_I^p} & 0 & 0 \\ 0 & \frac{K_P^q s + K_I^q}{s^2 + K_P^q s + K_I^q} & 0 \\ 0 & 0 & \frac{K_P^r s + K_I^r}{s^2 + K_P^r s + K_I^r} \end{bmatrix} \tag{5.6}$$

Note the presence of zeros defined by the LQR controller's proportional and integral gains in the form of $K_P^i s + K_I^i$, which could potentially saturate demands to step responses. To mitigate this, a reference pre-filter of the form:

$$PF_i(s) = \frac{K_I^i}{K_P^i s + K_I^i} \tag{5.7}$$

is therefore utilized for each channel (p, q, and r) to loop shape the response by canceling the unwanted zeros (which are known exactly due to being functions of the controller gains and not model parameters).

To visualize this setup, a block diagram of this NDI attitude rate controller in closed-loop with the tiltwing aircraft and control allocation is shown in Figure 5.3, which demonstrates how the NDI control law results in an equivalent linear system (the "NDI Linearization Loop"). This is analogous to Figure 2.2 from Section 2.3.

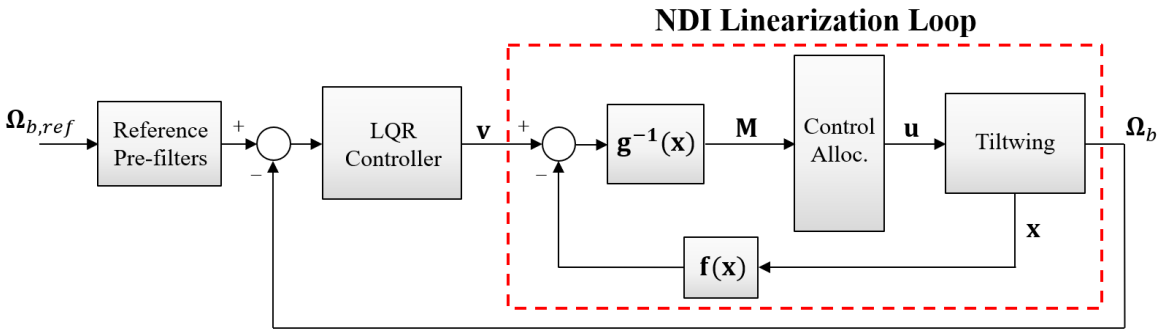

Figure 5.3: NDI attitude rate controller in closed-loop with the tiltwing aircraft and control allocation.

5.3.3 Euler Angle Tracking

Applying feedback linearization to the attitude rates of an aircraft is mathematically convenient as the relative degree for each state is simply 1, allowing for direct inversion of the dynamics due to the decoupling matrix $\mathbf{E}(\mathbf{x})$ simplifying to $\mathbf{E}(\mathbf{x}) = \mathbf{g}(\mathbf{x})$ (i.e. NDI). However, this means that body-frame attitude rate references must be supplied to this NDI loop, while it is desired to control Earth-frame Euler angles directly.

The solution is to first simply cascade an additional feedback loop outside the NDI controller to generate Earth-frame attitude rate references for given Euler angle commands. It was found that a simple proportional control was sufficient for this task. For this aircraft, it is desired to control the roll (ϕ) and pitch (θ) angles, while a direct (Earth-frame) yaw rate command ($\dot{\psi}_{ref}$) from the pilot is preferred for regulating aircraft heading. Hence, proportional controllers are only required for ϕ and θ:

$$\dot{\phi}_{ref} = K_P^{\phi}(\phi_{ref} - \phi) \qquad \dot{\theta}_{ref} = K_P^{\theta}(\theta_{ref} - \theta) \tag{5.8}$$

Next, it is necessary to convert these Euler angle rate demands to body-frame rotation rate demands. To do this, $\dot{\phi}_{ref}$ and $\dot{\theta}_{ref}$ are combined with the $\dot{\psi}_{ref}$ command from the pilot, and these three Euler angle rate quantities are transformed to body-frame attitude rate reference inputs for the NDI controller by applying the Earth-to-body frame rotation matrix $\mathbf{R}_{be}$ from (2.4) and repeated here:

$$\mathbf{R}_{be} = \begin{bmatrix} 1 & 0 & -s_\theta \\ 0 & c_\phi & s_\phi c_\theta \\ 0 & -s_\phi & c_\phi c_\theta \end{bmatrix} \tag{5.9}$$

The complete mapping from the Euler angle commands to the body-frame attitude rate references is represented by the "Euler angle tracking" subsystem in Figure 5.2. Note that this setup of combining these Euler angle proportional controllers with

the Earth-to-body frame transformation fully accounts for the nonlinear relationship between Euler angles and body-frame attitude rates. Hence, similar to the NDI controller of Section 5.3.2, this Euler angle tracking setup is defined over the complete flight envelope.

5.3.4 Control Allocation

The purpose of the control allocation is to map the commanded virtual controls $\boldsymbol{\tau}_c = [F_z \quad M_x \quad M_y \quad M_z]^T$ (i.e. a vector containing $\mathbf{F}$ and $\mathbf{M}$) to the real actuator inputs $\mathbf{u}$. It was discussed in Section 2.4 that there are many ways to approach this problem. For this design, it was chosen to linearize the nonlinear effector model $\mathbf{f}_B(\mathbf{x}, \mathbf{u})$ and utilize an off-line weighted pseudo-inverse method due to its ease of implementation, computational efficiency, pre-determined closed-form solution, and ability to easily deal with actuator failures/faults [43].

To begin, we assign $\boldsymbol{\tau}_c = \mathbf{f}_B(\mathbf{x}, \mathbf{u})$—that is, we equate the commanded virtual controls to the nonlinear effector model—and apply a first-order Taylor expansion:

$$\boldsymbol{\tau}_c = \mathbf{f}_B(\mathbf{x}, \mathbf{u}) \approx \mathbf{f}_B(\mathbf{x}_0, \mathbf{u}_0) + \left.\frac{\partial \mathbf{f}_B}{\partial \mathbf{x}}\right|_{(\mathbf{x}_0, \mathbf{u}_0)} (\mathbf{x} - \mathbf{x}_0) + \left.\frac{\partial \mathbf{f}_B}{\partial \mathbf{u}}\right|_{(\mathbf{x}_0, \mathbf{u}_0)} (\mathbf{u} - \mathbf{u}_0) \tag{5.10}$$

where the subscript 0 indicates values at a nominal operating point. Within the context of the control allocation, the system's states are constant. Hence, $(\mathbf{x} - \mathbf{x}_0) = 0$ and (5.10) can be rearranged as:

$$\delta\boldsymbol{\tau}_c = \left.\frac{\partial \mathbf{f}_B}{\partial \mathbf{u}}\right|_{(\mathbf{x}_0, \mathbf{u}_0)} \delta\mathbf{u} = \mathbf{B}(\mathbf{x}_0, \mathbf{u}_0)\delta\mathbf{u} \tag{5.11}$$

where $\delta\boldsymbol{\tau}_c = \boldsymbol{\tau}_c - \mathbf{f}_B(\mathbf{x}_0, \mathbf{u}_0)$, $\delta\mathbf{u} = \mathbf{u} - \mathbf{u}_0$, and $\mathbf{B}$ is known as the control effectiveness matrix. For this tiltwing aircraft, $\dim(\delta\mathbf{u}) > \dim(\delta\boldsymbol{\tau}_c)$ and so $\mathbf{B}(\mathbf{x}_0, \mathbf{u}_0)$ is not strictly invertible. However, the generalized inverse QP from (2.45) can be applied here to

find a solution for $\delta\mathbf{u}$:

$$\begin{aligned} \min_{\delta\mathbf{u}} \quad & \frac{1}{2}(\delta\mathbf{u} - \delta\mathbf{u}_p)^T\mathbf{W}(\delta\mathbf{u} - \delta\mathbf{u}_p) \\ \text{s.t.} \quad & \delta\boldsymbol{\tau}_c = \mathbf{B}(\mathbf{x}_0, \mathbf{u}_0)\delta\mathbf{u} \end{aligned} \tag{5.12}$$

where $\mathbf{W}$ is a weighting matrix to be defined and $\delta\mathbf{u}_p$ is taken to be zero. Hence, without input constraints, this generalized inverse QP has the following explicit solution:

$$\begin{aligned} \delta\mathbf{u} &= \mathbf{C}\delta\boldsymbol{\tau}_c \\ \mathbf{C} &= \mathbf{W}^{-1}\mathbf{B}^T(\mathbf{B}\mathbf{W}^{-1}\mathbf{B}^T)^{-1} \end{aligned} \tag{5.13}$$

where $\mathbf{C}$ is the generalized inverse of $\mathbf{B}$. Note that in the case in which the weighting matrix $\mathbf{W} = \mathbf{I}$ (and since $\delta\mathbf{u}_p = 0$), the solution will be written as $\delta\mathbf{u} = \mathbf{B}^+\delta\boldsymbol{\tau}_c$ to denote the Moore-Penrose pseudo-inverse.

The weighting matrix $\mathbf{W}$ was taken to be a diagonal matrix whose elements were $1/u_{i,max}^2$, where u_i denotes the ith element of the tiltwing's input vector $\mathbf{u}$, to normalize the inputs relative their maximum values for a balanced allocation. These weightings can be adjusted as desired, and hence to address actuator failure, it is possible to "cancel" actuators from the weighted pseudo-inverse solution by simply setting their relative weighting in $\mathbf{W}$ to be very large such that their values in $\mathbf{W}^{-1}$ approach zero. The commanded virtual controls are then redistributed to the remaining actuators/elements in $\delta\mathbf{u}$.

Lastly, this solution was derived assuming no input constraints to enable a closed-form solution. To deal with actuator saturation, actuator bounds and rate limits must be explicitly applied after the pseudo-inverse, and a suitable integral anti-windup strategy is required for the inner-loop controller's integrators.

5.3.4.1 Integral Anti-Windup Strategy

A conventional integral anti-windup strategy is to feed the difference between commanded system inputs and saturated inputs back into a controller's integrators [133]. This is challenging to implement for a controller that outputs commands in a virtual moment space. One option is to determine the attainable moment set (AMS) of the aircraft [48, 51]. However, this is an onerous task for an over-actuated system such as this tiltwing, and the AMS would change if any actuators failed[3]. An alternative [55] is to apply saturations to real actuator inputs and invert the pseudo-inverse solution within the control allocation to determine the excess virtual moments (i.e. those moments which cannot be achieved by the real actuators). Once determined, these can be fed back into a controller's integrators as normal. Furthermore, the computation of excess virtual moments within the control allocation can be done on-line and remains valid during actuator failure. This was the integral anti-windup strategy chosen for this NDI attitude rate controller.

To compute these excess virtual moments within the control allocation, we consider (5.13). In the case that input constraints are exceeded, saturated inputs ($\delta\mathbf{u}_{sat}$) are subtracted from unsaturated inputs ($\delta\mathbf{u}$) to give:

$$\delta\mathbf{u} - \delta\mathbf{u}_{sat} = \mathbf{C}\left(\delta\boldsymbol{\tau}_c - \delta\boldsymbol{\tau}_{c,sat}\right) \tag{5.14}$$

where $\delta\boldsymbol{\tau}_{c,sat}$ denotes the corresponding saturated virtual controls. By expanding this out using definitions from (5.11), the offset terms $\mathbf{u}_0$ and $\mathbf{f}_B(\mathbf{x}_0, \mathbf{u}_0)$ cancel, leaving us with:

$$\delta\overline{\mathbf{u}} = \mathbf{C}\delta\overline{\boldsymbol{\tau}}_c \tag{5.15}$$

where $\delta\overline{\mathbf{u}} = \mathbf{u} - \overline{\mathbf{u}}$ and $\delta\overline{\boldsymbol{\tau}}_c = \boldsymbol{\tau}_c - \overline{\boldsymbol{\tau}}_c$, and where the overline denotes constrained

[3]Additionally, this tiltwing aircraft sees a coupling between feasible moment generation and force (thrust) generation, furthering the difficulty of finding its AMS.

values. $\mathbf{C}$ can be inverted using a Moore-Penrose pseudo-inverse to give the excess virtual controls vector:

$$\delta\overline{\boldsymbol{\tau}}_c = \mathbf{C}^{+}\delta\overline{\mathbf{u}} \tag{5.16}$$

Lastly, recalling that $\boldsymbol{\tau}_c$ is composed of $\mathbf{F}$ and $\mathbf{M}$, the excess virtual moments $\delta\overline{\mathbf{M}}$ can be extracted as a subvector from $\delta\overline{\boldsymbol{\tau}}_c$. Figure 5.4 below shows a block diagram layout of this anti-windup strategy. Note that when feeding $\delta\overline{\mathbf{M}}$ back into the NDI controller's integrators, it is passed through the NDI transformation by multiplying by the inverse inertia tensor $\mathbf{J}^{-1}$ (i.e. the $\mathbf{g}(\mathbf{x})$ term from (5.1)).

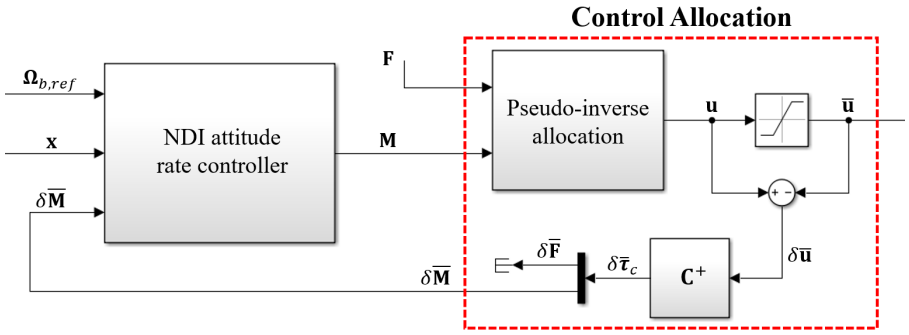

Figure 5.4: Integral Anti-Windup Strategy

To determine $\overline{\mathbf{u}}$, actuator constraints cannot simply be applied as would be expected for SISO systems. For MIMO systems, it is necessary to preserve the directionality of the unconstrained input vector $\mathbf{u}$ [49]. This means that re-scaling just one actuator demand to satisfy input constraints requires re-scaling all actuator demands by the same factor. An example of this is visualized in Figure 5.5, where it can be seen that input u_2 is saturating but not input u_1. However, if only u_2 is constrained and u_1 is not also scaled accordingly, then the original direction of the vector $\mathbf{u}$ is not preserved (represented by the gray dashed vector).

Mathematically, implementing this in general looks like multiplying $\mathbf{u}$ by the in-

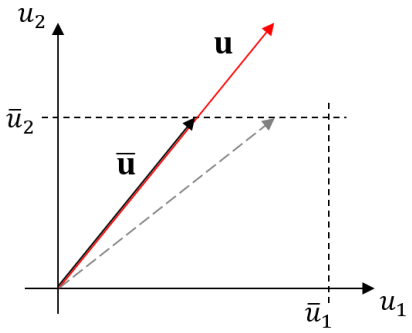

Figure 5.5: Constrained MIMO input vector scaling example

verse of the maximum constraint violation factor:

$$\bar{\mathbf{u}} = \mathbf{u} \cdot \frac{1}{\max\left\{\frac{u_i}{\bar{u}_i}\right\}} \tag{5.17}$$

Note that in practice, $\delta\mathbf{u}$ is scaled instead of $\mathbf{u}$ since $\delta\mathbf{u}$ is the actual control command output from the pseudo-inverse algorithm at each time step. Additionally, edge cases must be considered to avoid divide-by-zero errors.

5.3.4.2 Vertical Take-Off Schedule

While the NDI attitude controller remains valid over the whole flight domain, this control allocation uses a local linear approximation to enable real-time computation and as such only remains valid at a single operating point (and a small region around that point). Hence, it is necessary to linearize the control effector model at a range of operating points across the nominal flight envelope and smoothly interpolate between these points based on a scheduled variable[4]. This is perhaps expected for transition, where the tiltwing aircraft's dynamics vary significantly throughout the maneuver, but it also turned out to be necessary for vertical take-off and landing. Hence, in

[4]Note that, while this may seem to contradict the general approach of choosing a feedback linearization-based design so that a single controller is operable over the complete flight domain, the difference here is that these linearizations need only pre-computed once for a given control effector model (and this process can be easily automated). In contrast, a gain scheduled controller that linearizes a system over a range of points would need re-tuned at each operating point if the system model ever changes, a much more laborious task.

consideration of hover mode, the nonlinear control effector model was linearized at a range of points from zero thrust up to the steady-state hover thrust (i.e. throughout vertical take-off), and the scheduled variable was chosen to be the total thrust request (specifically, the virtual vertical force request F_z) since this tiltwing aircraft's control effectiveness scales with prop speeds in hover.

As a first approximation, these operating points were computed using a simple thrust and moment balance involving the prop inputs while ignoring additional aerodynamic effects, and this was found to be sufficient for stabilizing the aircraft during take-off. Considering Figure 5.6, a vertical force and pitching moment balance in a NED coordinate frame gives:

$$\begin{aligned} F_z &= -T_W \sin\xi_W - T_T \sin\xi_T \\ M_y &= T_W \left(l_4 \sin\xi_W - l_6 \cos\xi_W\right) - T_T \left(l_5 \sin\xi_T + l_7 \cos\xi_T\right) \end{aligned} \tag{5.18}$$

where T_W and T_T are the total combined wing prop and tail prop thrusts, respectively. To achieve a pitching moment balance, M_y is set to zero and F_z is chosen based on the number of linearization points desired. This square system can then be inverted to determine T_W and T_T, noting that the wing angles are fixed during hover. Once these total combined thrusts are determined for a particular F_z request, the nominal inputs vector $\mathbf{u}_0$ can be determined and used in the complete nonlinear effector model to determine the nominal control effectiveness matrix $\mathbf{B}(\mathbf{x}_0, \mathbf{u}_0)$ and virtual control offset $\mathbf{f}_B(\mathbf{x}_0, \mathbf{u}_0)$, all previously introduced in (5.10) and (5.11).

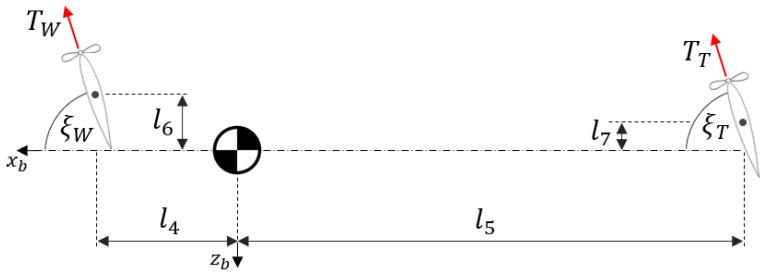

Figure 5.6: Simplified thrust and moment balance during take-off

While this first approximation was sufficient (which provides some indication of the system's robustness to aerodynamic uncertainty in hover mode), it was later chosen to improve upon the accuracy of these operating points by solving this force and moment balance using the complete nonlinear effector model; that is, by finding solutions in $\mathbf{u}$ to the following equation:

$$\mathbf{f}_B\left(\mathbf{x}_0, \mathbf{u}\right) - \begin{bmatrix} F_z & 0 & 0 & 0 \end{bmatrix}^T = \mathbf{0} \tag{5.19}$$

which was then solved at each operating point using MATLAB's `fsolve()`.

To implement this vertical take-off schedule, the nominal control effectiveness matrices, virtual control offset vectors, and nominal inputs were placed into lookup tables and accessed based on the F_z request being demanded. To balance accuracy with memory overhead on the microcontroller board used to deploy the flight controller, 20 operating points were computed and stored, and smooth interpolation was achieved using a linear point-slope algorithm within the lookup tables.

5.4 Simulations

5.4.1 Controller Tuning

For the NDI attitude rate controller, the LQR performance indices were chosen to be:

$$Q = \begin{bmatrix} 0.1 \cdot \mathbf{I}_{3\times 3} & \mathbf{0}_{3\times 3} \\ \mathbf{0}_{3\times 3} & 10 \cdot \mathbf{I}_{3\times 3} \end{bmatrix} \quad R = 0.007 \cdot \begin{bmatrix} I_{xx} & 0 & 0 \\ 0 & I_{yy} & 0 \\ 0 & 0 & I_{zz} \end{bmatrix} \tag{5.20}$$

where I_{xx}, I_{yy}, and I_{zz} were selected to scale the bandwidth of the controller relative the inertia of each rotational axis. This controller tuning sets the roll and pitch rate bandwidths to approximately 7.2 $\mathrm{rad\,s^{-1}}$ and the yaw rate bandwidth to 6.25

rad s^{-1} with all three attitude rates achieving an infinite gain margin and a minimum phase margin of 68°. Note that, considering this is the inner-most loop of a cascade controller design (where time-scale separation assumptions will become important), this controller tuning is slower than would be desired. However, during initial hover flight testing with the aircraft, a control-structural resonance mode was discovered when running the controller with a higher bandwidth. This effect is shown in Figure 5.7, where oscillations in the yaw response can be seen that were caused by structural vibrations. These vibrations were observed in both wings, although primarily in the main wing, and would increase in amplitude with increasing thrust. It was specifically found that the aileron inputs used to control the yaw of the aircraft were exciting this structural resonance mode, which was around 4 Hz. To fix this, it was necessary to reduce the NDI controller bandwidth to its current tuning (around 1.1 Hz) to avoid exciting this resonant frequency. This structural resonance, inherent in the design of the scaled demonstrator aircraft[5], thus serves as a fundamental performance limitation for the full-envelope flight controller design in this thesis. In consideration of the vehicle design itself, this was of course a rapid prototype vehicle with a 3D-printed airframe and off-the-shelf hobbyist parts; a more carefully-designed model would have such a structural resonance much outside the bandwidth of the inner-loop. For example, a stiffening of the wing structure through a redesign of the scaled model would theoretically increase this resonant frequency, thus enabling a higher NDI controller bandwidth, but this was not feasible within the time constraints of the project.

With the NDI control loop tuning finalized, the gains chosen for the Euler angle tracking controllers (roll and pitch) were $K_P^\phi = 1.21$ and $K_P^\theta = 1.39$, which set the roll and pitch angle closed-loop bandwidths to approximately 1 rad s^{-1} to maximize the speed of the response while also providing a level of time-scale separation with

[5]And presumably present for the full-scale aircraft as well, though likely at a different resonant frequency.

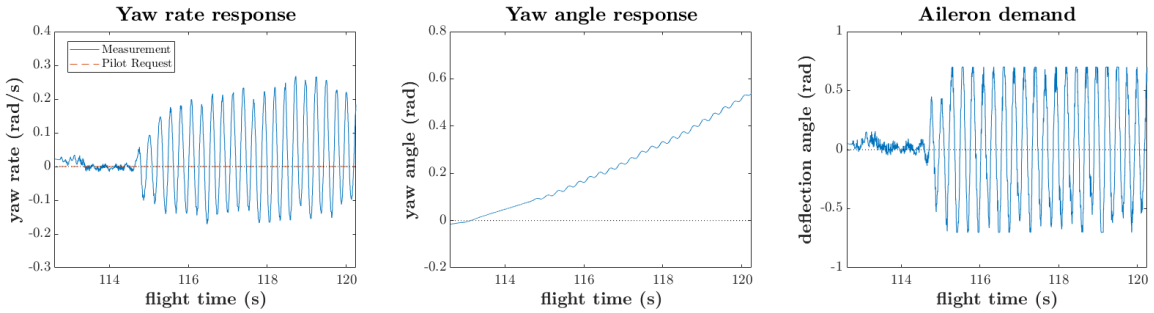

Figure 5.7: Yaw rate, yaw angle, and aileron demand during structural vibrations excited by the NDI attitude rate controller with sufficiently-high bandwidth

the NDI loop for robustness. This tuning resulted in gain and phase margins of approximately 17 dB and 71°.

5.4.2 Results

Figure 5.8 below shows attitude angle step responses of 0.1 rad for roll and pitch and a yaw rate step response of 0.1 rad s^{-1} during pure hover. For each step response, the corresponding attitude rate response (or angle, in the case of the yaw rate step) is shown in the axes just beneath it, while corresponding prop and aileron demands are shown separately in Figure 5.9. Table 5.1 below summarizes the performance metrics of each step response. Of particular note is the longer settling time and higher overshoot in roll compared with pitch, though pitch sees a slightly longer rise time. Also note that the yaw rate response is not directly comparable to the roll and pitch angle steps due to no explicit yaw angle tracking loop. In general, performance was found to be satisfactory, though response speed was lower than desired to avoid exciting the 4 Hz structural resonance mode of the airframe.

To demonstrate the control allocation's ability to compensate for actuator failure and the controller's effectiveness at dealing with such a disturbance, a simulation was run in which the left-inner wing motor "fails" at $t = 5$ s. Figures 5.10 and 5.11 below show the output responses of the attitude angles and rates during this failure as well as the prop and aileron demands required to compensate for the temporarily

imbalanced system. It can be seen that the maximum roll deviation is within 0.1 rad, and the other two attitude angle responses are even less at just 0.048 rad for pitch and 0.071 rad for yaw. For the prop speeds, it can be seen that each prop must change speed to compensate for both the thrust and moment imbalance. The greatest demand is on the left-outer wing prop which requires a 71% increase in thrust generation to compensate.

Table 5.1: Performance metrics of the simulated step responses given in Figure 5.8

Step Response	**rise time**	**2% settling time**	**overshoot**
roll	1.5 s	12.5 s	19%
pitch	2.5 s	4.5 s	0%
yaw rate	0.7 s	2 s	5%

Table 5.2: Performance metrics of the attitude responses from tethered flight data given in Figure 5.14. See Section 5.5.2 for details.

Step Response	**rise time**	**2% settling time**	**overshoot**
roll	1.6 s	–	–
pitch	2.1 s	–	–
yaw rate	0.9 s	1.9 s	26%

Table 5.3: Performance metrics of the attitude responses from untethered flight data given in Figure 5.20. See Section 5.5.3 for details.

Step Response	**rise time**	**2% settling time**	**overshoot**
roll	1.4 s	1.9 s	0%
pitch	1.7 s	2.1 s	0%
yaw rate	0.8 s	–	–

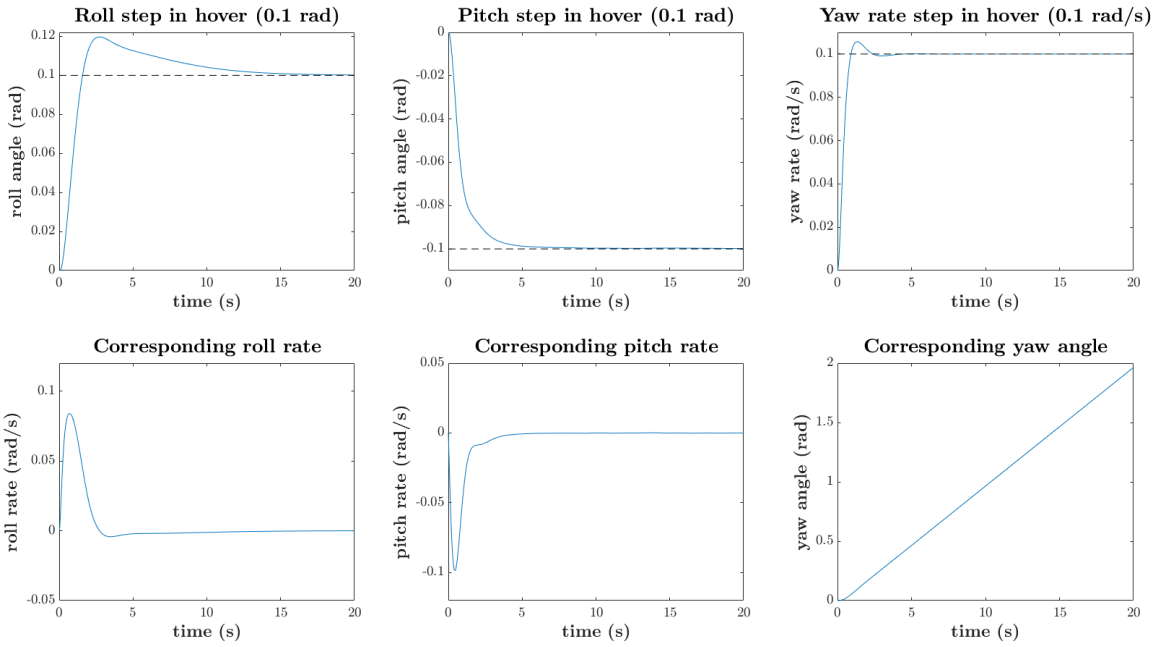

Figure 5.8: Roll and pitch step responses of 0.1 rad and yaw rate step response of 0.1 rad s^{-1} in hover simulations. Horizontal dashed line indicates step input value.

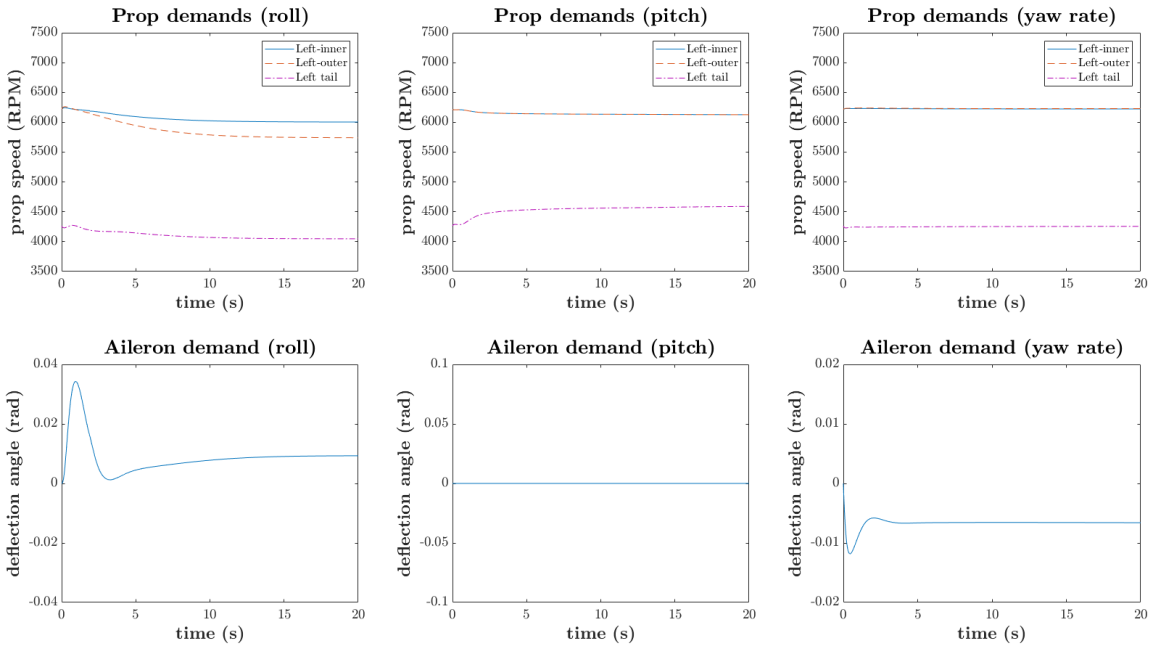

Figure 5.9: Prop and aileron demands corresponding to each attitude step response from Figure 5.8

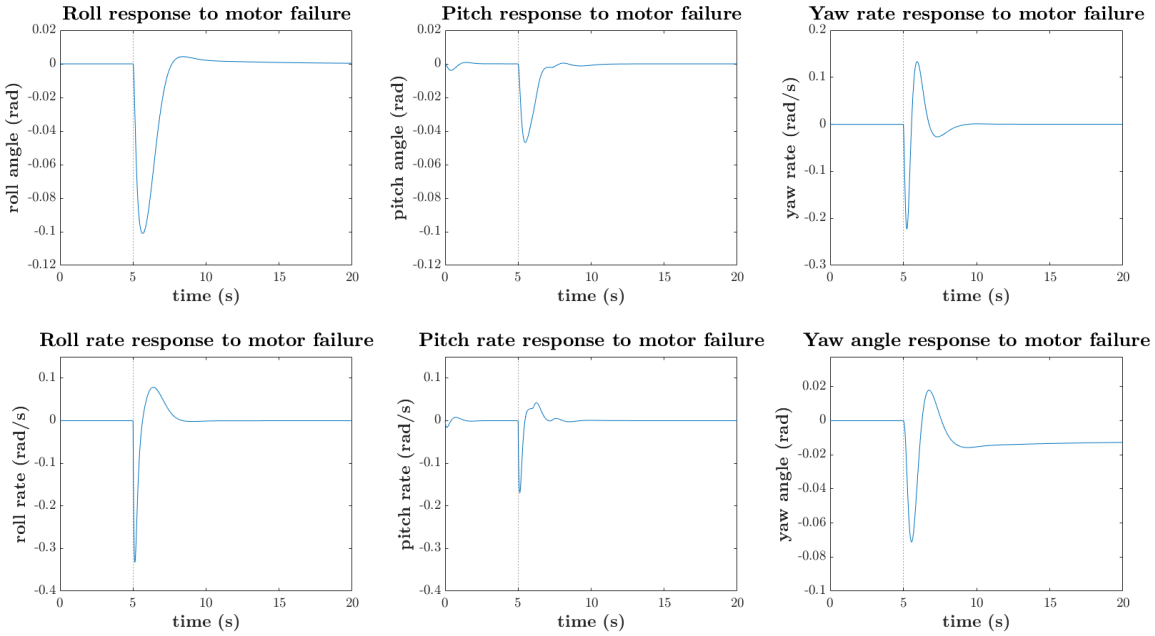

Figure 5.10: Attitude responses to failure of left-inner wing prop at $t = 5$ s in hover simulations. Vertical dotted line indicates moment of failure.

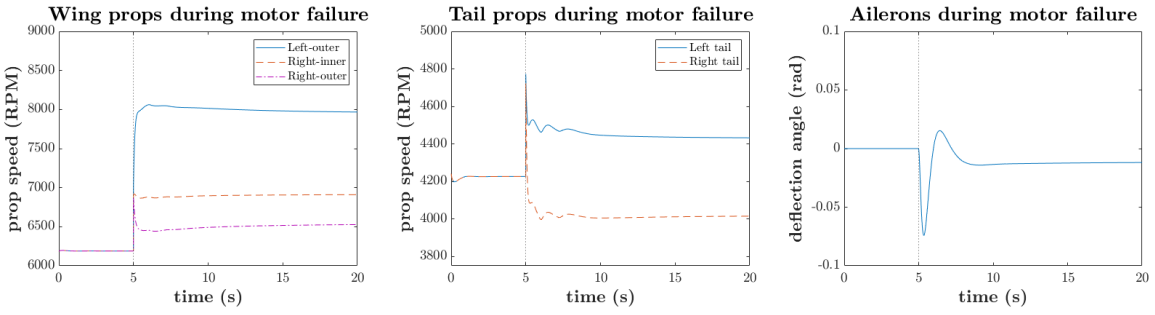

Figure 5.11: Prop and aileron demands during failure of left-inner wing motor at $t = 5$ s in hover simulations. Vertical dotted line indicates moment of failure.

5.5 Flight Testing

5.5.1 Experimental Setup

Flight testing of the scaled demonstrator model was conducted both indoors (tethered) and outdoors (untethered). For indoor tethered testing, experiments took place in a large warehouse space where a harness was secured around the aircraft's wing section and fuselage and a thick climbing rope was used as a tether (see Figure 5.12a). This tether rope was looped through a hook in the ceiling and held by an operator on the ground (see Figure 5.12b). To avoid interference from the tether, the length of the rope was continuously adjusted to ensure it remained slack during flight while also not drooping too far and contacting one of the propellers. In practice, this was not always perfect (and these specific instances will be addressed), but overall it was found to be sufficient for testing the validity of this control design and to ensure that all systems were working as expected before testing untethered outdoors. The designated testing area itself was approximately 6.75×6.75 m wide, which equates to roughly 3.75 aircraft wing span lengths per dimension.

(a) Scaled tiltwing demonstrator with harness

(b) Tether with ceiling mount

Figure 5.12: Indoor (tethered) testing setup

For outdoor untethered flight, testing was conducted in a large field West of Cotswold Airport in Kemble with permission from the local British Model Flying Association (BMFA) club that operates there[6]. In obvious contrast to indoor testing, there was sufficient room for maneuverability and no rope interference, which helped for validating test results from indoor flight. Testing took place on a clear day with relatively calm winds. Figure 5.13 shows pictures of the aircraft during outdoor testing.

(a) Tiltwing hovering during flight test

(b) Tiltwing outside ready for untethered testing

Figure 5.13: Outdoor (untethered) testing setup

The aircraft itself was powered by three lithium-polymer (LiPo) batteries: One for the two left wing motors, one for the two right wing motors, and one for the tail motors plus the remaining electronics (servos, LACs, and the on-board flight controller hardware). For the on-board hardware, a Pixhawk 4 flight controller board was chosen for this system due to its integrated IMU with a pre-built EKF and easy pairing with MATLAB's support packages for rapid deployment of custom Simulink models. For tethered testing, while everything was powered and running locally, the Pixhawk 4 board was additionally hard-wired via a 10 m USB cable for real-time data logging and communication with a laptop computer. In contrast, for untethered testing, there was no real-time communication with the laptop, and data logging was recorded locally on an SD card. During flight (both tethered and untethered), the

[6]The Cirencester & District Model Aircraft Club (CDMAC).

aircraft was operated by a pilot using a remote control transmitter that communicated with the Pixhawk 4 via a Spektrum radio receiver. To control the aircraft in hover flight, pilot demands included a total thrust request (specifically, the virtual vertical force F_z), roll and pitch angle requests, and a yaw rate request. Apart from the data logging and laptop communication, there were no other software/hardware differences between the tethered and untethered flight testing setups.

5.5.2 Tethered Flight

Figures 5.14 and 5.15 below show attitude responses and corresponding prop and aileron demands from an example tethered hover flight test. This test began with the aircraft suspended in air by the tether and then quickly reached the full hover flight condition after startup. Roll, pitch, and yaw rate measurements show that references demanded by the pilot were tracked well. Specifically, roll inputs saw a rise time of about 1.6 s, while pitch inputs saw a rise time of approximately 2.1 s. Both of these are in good agreement with simulation results (see Table 5.1). Yaw rate inputs were also in good agreement with simulation results (0.9 s rise time and 1.9 s settling time) with the exception of its overshoot, which was around 26%. These results are summarized in Table 5.2. Note that settling times and overshoots for roll and pitch responses were difficult to ascertain from tethered flight data due to the limited size of the indoor testing area; non-zero attitude angles could not be continuously held and additional input corrections from the pilot were required to prevent the aircraft from drifting after completing step inputs. Nevertheless, pilot feedback described the system as responsive and was able to comfortably maintain controlled flight within the boundaries of the indoor testing area.

Figures 5.16 and 5.17 show attitude responses and prop and aileron demands during an example take-off. In these plots, the first vertical dotted line indicates when the main gear first began to lift-off from the ground, while the second vertical

dotted line indicates when the aircraft reached full hover. It can be seen that stability is maintained throughout the take-off maneuver and that the vertical take-off schedule works. Attitude deviations are kept within 0.03 rad for roll and yaw while pitch is kept within 0.1 rad. Note that this is an example of where influence from the tether harness can be observed in the results due to coordinating the length of the tether rope while the aircraft lifted off the ground. Oscillations in pitch due to pendulum-like motion can be observed with a frequency of about 1.25 Hz.

Lastly, Figures 5.18 and 5.19 show an example run of a motor failure during flight. In this case, the left-inner wing motor was "failed" during mid-hover, just as for the simulated case (here at flight time $t = 184.3$ s). It can be seen that the maximum deviations in roll and pitch immediately after failure were within just 0.067 and 0.065 rad, respectively, which is in good agreement with the simulation results. It must be noted, however, that this is another example where the influence of the tether came into play, since the temporary loss of thrust caused the aircraft to fall briefly and be partially supported by the rope. However, it is believed that the roll response is similar to what would be observed in free flight since the geometry of the harness doesn't hold any significant rolling moment. This, however, is not true for the pitching moment; hence, the pitch angle response to motor failure is likely to be different in a theoretical free-flight test case.

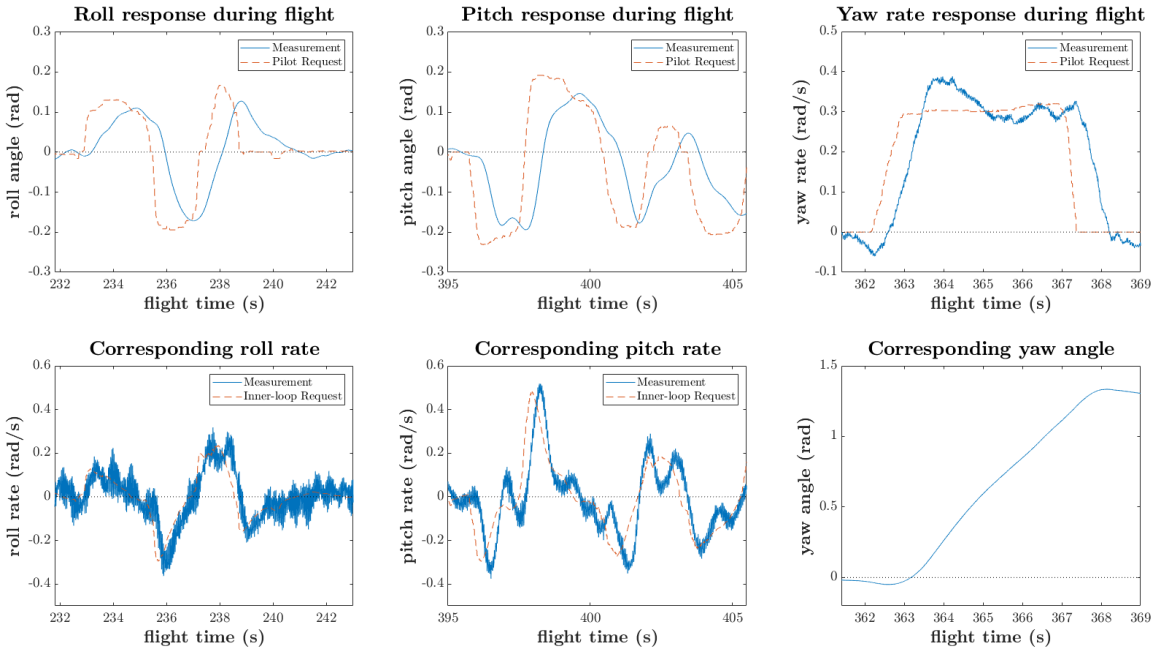

Figure 5.14: Roll, pitch, and yaw rate responses during example flight

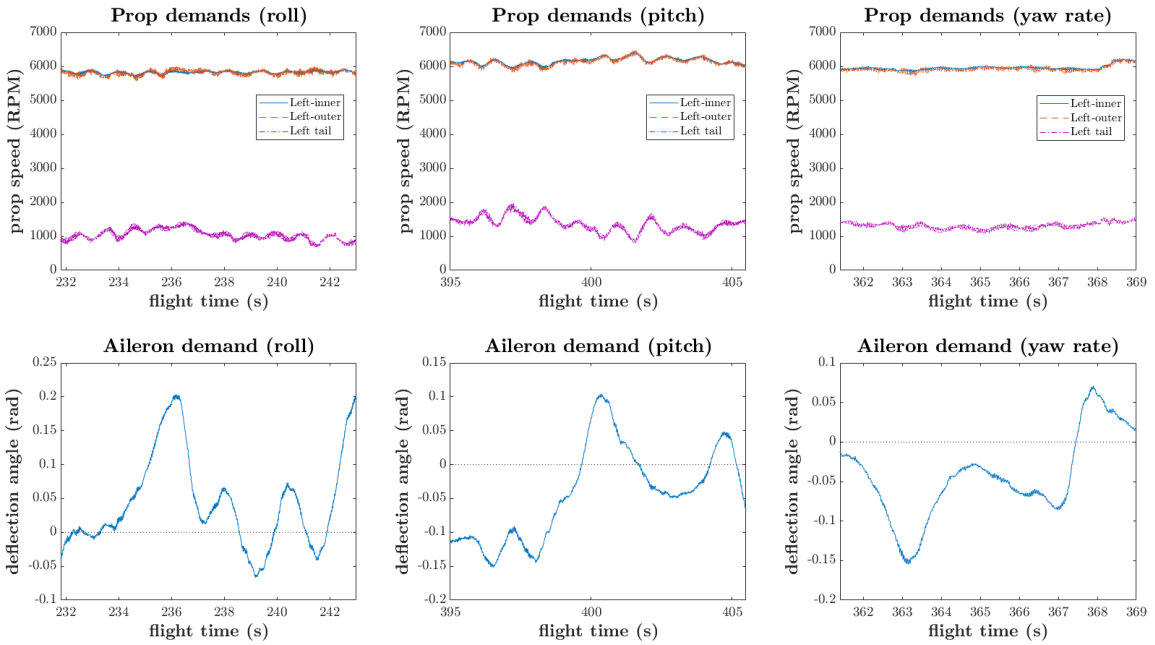

Figure 5.15: Prop and aileron demands corresponding to each attitude response from Figure 5.14

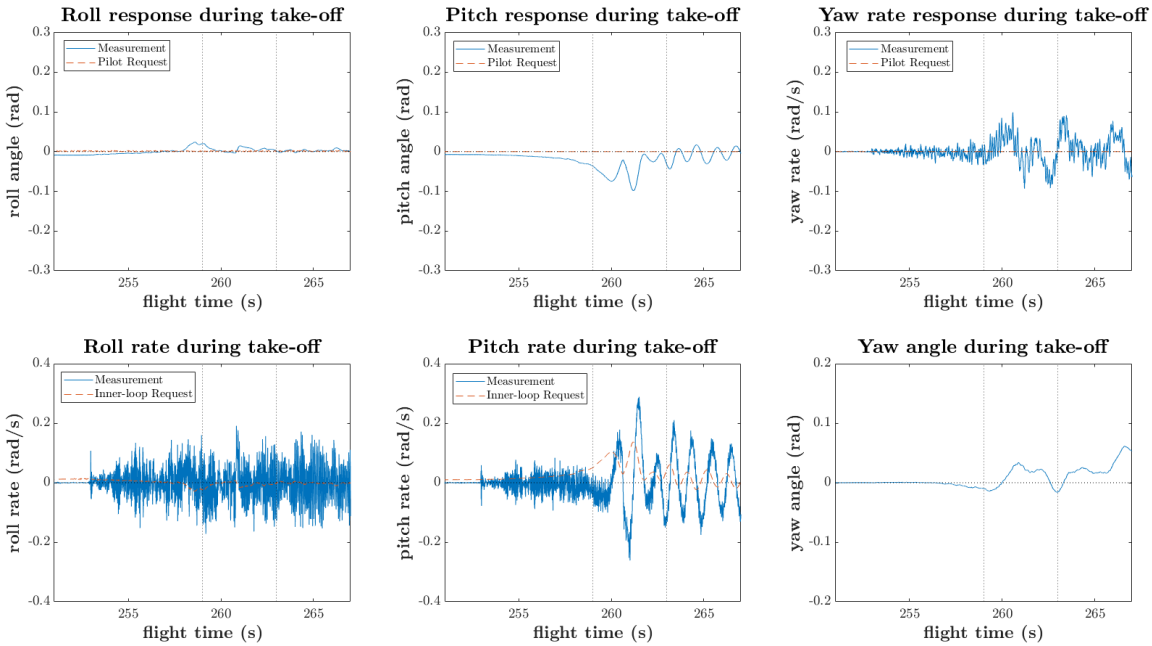

Figure 5.16: Attitude responses during example take-off. The first vertical dotted line indicates first moment of lift-off while the second vertical dotted line indicates full hover.

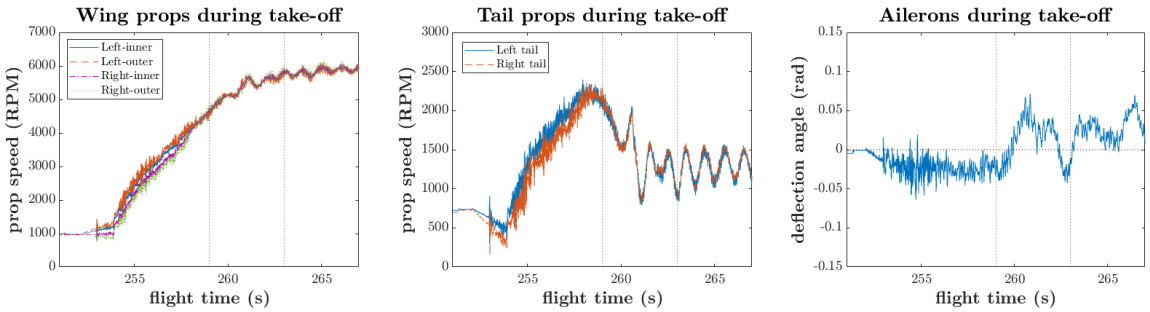

Figure 5.17: Prop and aileron demands during example take-off. The first vertical dotted line indicates first moment of lift-off while the second vertical dotted line indicates full hover.

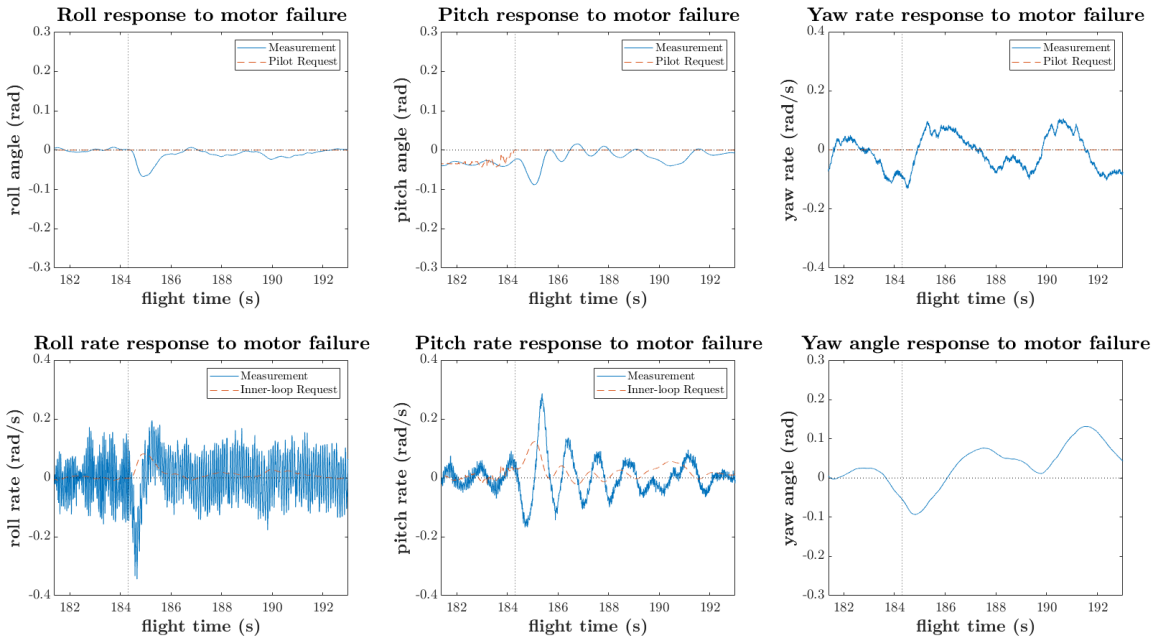

Figure 5.18: Attitude responses to failure of left-inner wing prop at flight time $t = 184.3$ s. Vertical dotted line indicates moment of failure.

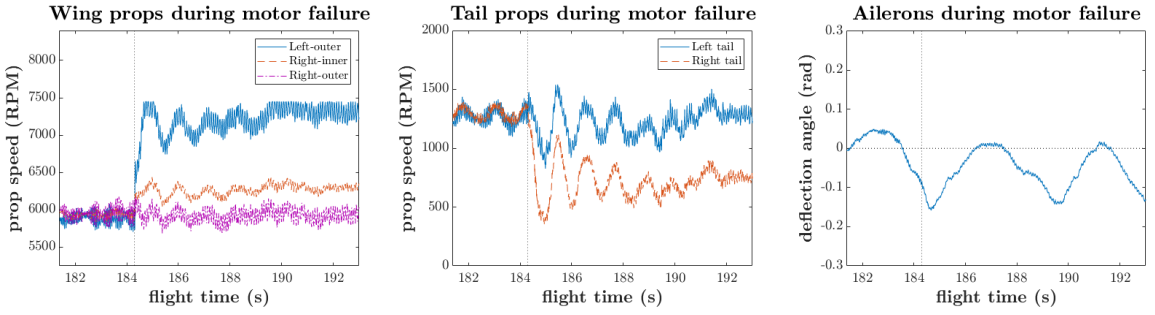

Figure 5.19: Prop and aileron demands during failure of left-inner wing prop at flight time $t = 184.3$ s. Vertical dotted line indicates moment of failure.

5.5.3 Untethered Flight

Figures 5.20 and 5.21 below show example roll, pitch, and yaw rate maneuvering during untethered flight. These are plotted alongside their respective rate responses (or angle, in the case of yaw control). This testing took place in the mid-afternoon on a clear day with mild winds around 5 knots. The most immediate feature in the data to note is the large pitch angle demands from the pilot, which allude to the significant effect that the wind had on the aircraft during flight. Despite these relatively mild wind conditions, large pitch angles were required to maintain ground position during hover. This was largely attributed to the vertically-oriented wings of this tiltwing configuration. Firstly, the pre-chosen wing angle was initially set at 90°, which could have been trimmed lower to help compensate. And secondly, the main wing in particular ultimately (and perhaps unsurprisingly) was found to behave like a large sail in the wind, creating a sort of "barn door" effect. As a result of these large pitch angles, combined with the aircraft's low thrust margin, the aircraft was frequently pushed close to its limits thrust-wise. Despite all this, the aircraft remained stable even in light of some occasional strong wind gusts, indicating good robustness to such disturbances, and the test was deemed a success.

Performance-wise, it can be seen that pilot references were tracked well overall, and the untethered response characteristics (Table 5.3) are generally consistent with the results from both simulations and tethered flight (Tables 5.1 and 5.2). However, there are a few notable differences to address. Firstly, the pitch response rise time was found to be faster (1.7 s) than both simulation (2.5 s) and tethered flight (2.1 s). With the pitch response in tethered flight already being marginally faster than simulation predictions, it is believed that the actual closed-loop bandwidth for the physical aircraft may be higher than what is predicted by the model. However, following this, the faster pitch response in untethered flight vs tethered flight further indicates that the tether rope was potentially influencing the pitch response more

than originally suspected (causing the performance to be marginally slower when tethered in comparison to free flight[7]). Secondly, the settling times for both roll and pitch were found to be much quicker than simulation predictions, with the largest discrepancy occurring in the roll response (a 2% settling time of 1.9 s vs 12.5 s in simulation), which also appears to have no overshoot unlike the roll response predicted in simulation (19%). However, it was ultimately difficult to ascertain these particular metrics (settling time and overshoot) in an exact sense, as the reference inputs from the pilot were not always constant and wind disturbances were present. For the same reasons, it was not possible to determine a settling time and overshoot for the yaw response in the data that was collected. Hence, the settling time and overshoot values in Table 5.3 should be interpreted as approximate.

To demonstrate the free-flight take-off response, figures 5.22 and 5.23 show attitude responses and prop and aileron demands during take-off. In these plots, the first vertical dotted line indicates when the main gear first began to lift-off from the ground, while the second vertical dotted line indicates when the aircraft reached full hover. Here, the effect of the wind on the aircraft's yaw orientation due to its vertically-oriented wings can be spotted as the weight of the aircraft was lifted from the ground, causing the aircraft to begin rotating in alignment with the wind. Furthermore, the pitch response is much smoother and there are no more pitching oscillations present during take-off as there were in the example take-off maneuver from tethered flight testing due to no rope interactions.

[7]Note that this was not noticeable from a piloting perspective; pilot feedback indicated the aircraft felt equally responsive in both cases.

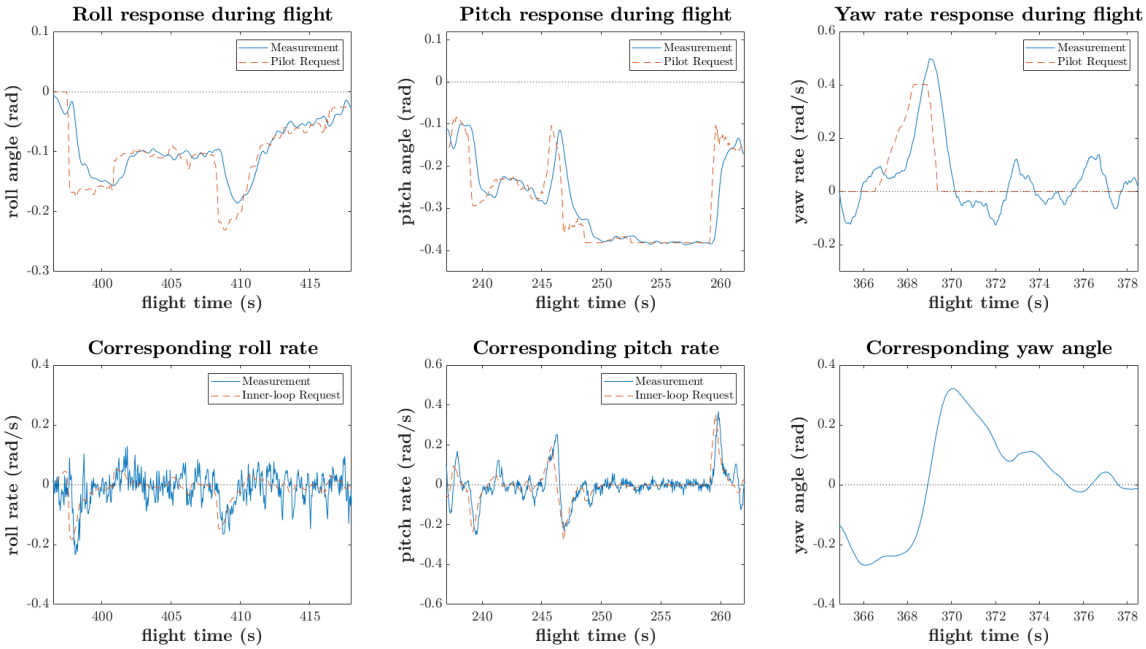

Figure 5.20: Example roll, pitch, and yaw rate maneuvering during hover flight

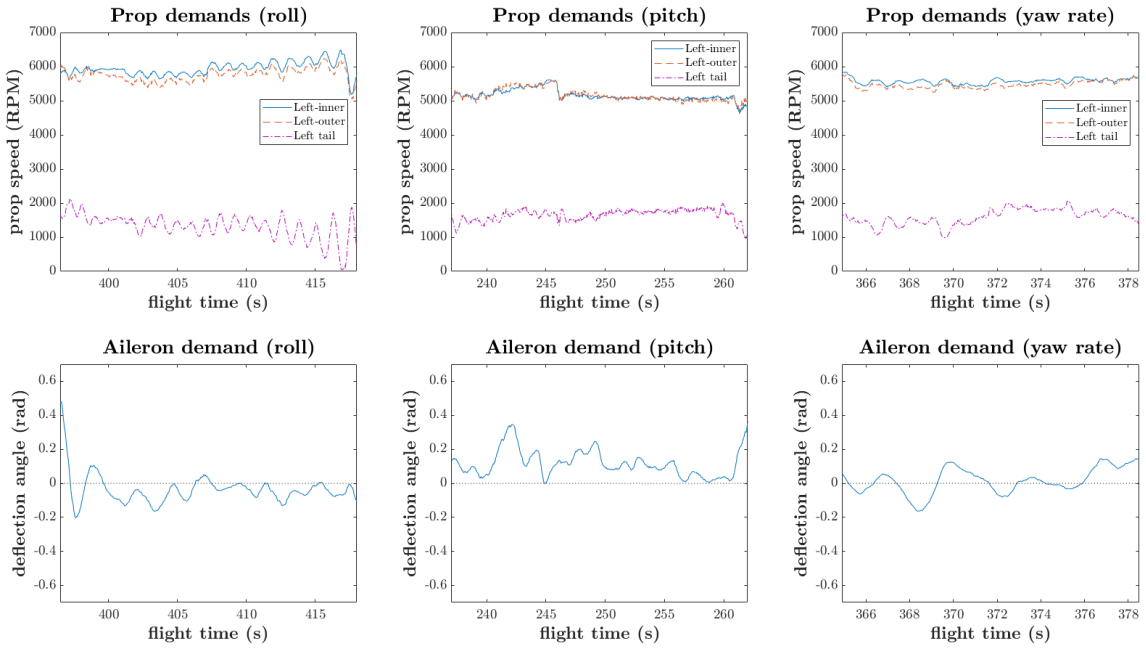

Figure 5.21: Prop and aileron demands corresponding to each attitude maneuver from Figure 5.20

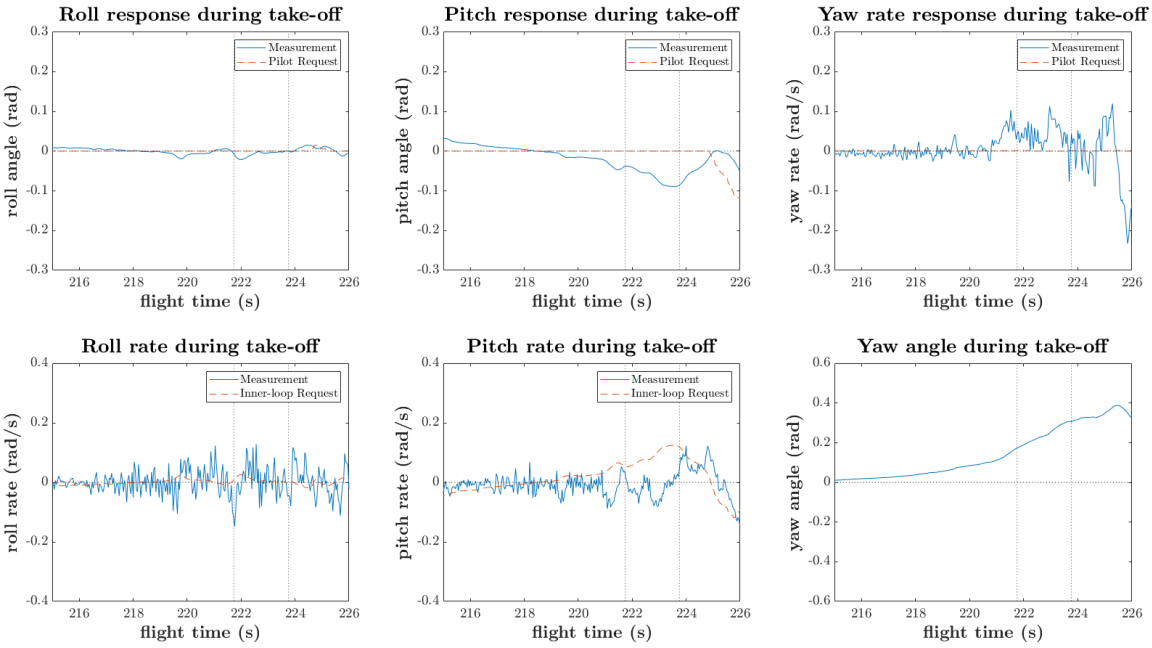

Figure 5.22: Attitude responses during take-off. The first vertical dotted line indicates first moment of lift-off while the second vertical dotted line indicates full hover.

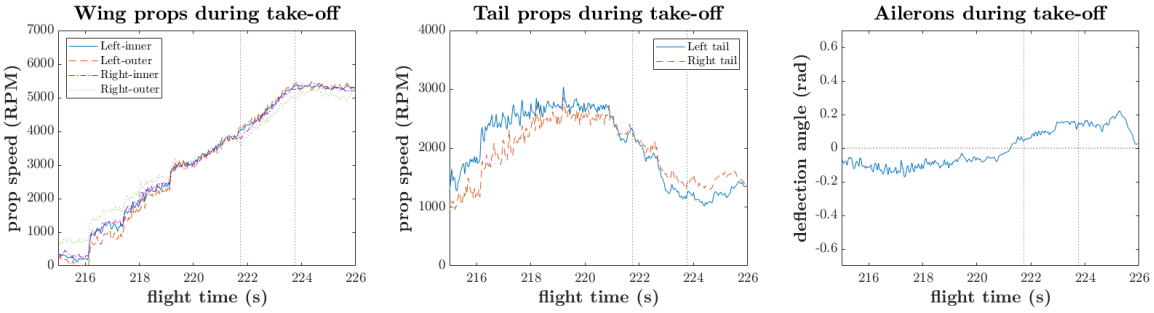

Figure 5.23: Prop and aileron demands during take-off. The first vertical dotted line indicates first moment of lift-off while the second vertical dotted line indicates full hover.

5.6 Conclusion

This chapter presented the inner-loop attitude controller design for the unique tiltwing configuration of interest in this thesis, marking the first step in the development of a full-envelope flight controller for this aircraft. NDI was paired with LQR and combined with off-line weighted pseudo-inverse control allocation to create a full-envelope attitude controller capable of fault-tolerance due to the over-actuated nature of this system. Controller performance was found to be satisfactory, though inner-loop bandwidth was limited to avoid exciting a structural resonance mode of the 3D-printed airframe. Flight testing results were generally in good agreement with those predicted in simulations, including in response to motor failure, the latter which demonstrated that full controllability of the aircraft could be retained in the presence of such a disturbance. Furthermore, good robustness to wind gust disturbances was demonstrated in untethered flight. Thus far, the simulations and flight testing results presented have only covered controller performance in hover flight. In Chapter 7, the aircraft's full envelope (hover, cruise, and transition) will be demonstrated in simulations, while the outer-loop velocity controller will be tested in additional hover flights.

Chapter 6

Generalized Full-Envelope Outer-Loop Control Strategies

6.1 Introduction

This chapter tackles the general problem of finding a full-envelope solution for the outer-loop velocity control design. Emphasis is placed on designing for a tiltwing, the focus of this thesis, but the methods presented in this chapter are generalizable to a wide range of transition VTOL configurations. As was demonstrated in Chapter 5, full-envelope solutions are commonly achieved for the inner-loop attitude control using inversion-based methods such as feedback linearization. This is made possible due to the form of the attitude rate dynamics, which are affine in the moment inputs ($\mathbf{M}$) and full-state linearizable such that the nonlinear equations can be directly inverted. In contrast, the outer-loop velocity dynamics do not admit the same convenient mathematical form. Specifically, common transition VTOL configurations, such as tiltwings, tiltrotors, and tiltducts, typically exhibit a lack of (independent) force actuation (i.e. thrust) in both the vertical and horizontal directions, resulting in dynamics that are not generally input-affine. This is because full controllability of the

(Earth-frame) velocity states only remains possible by coupling the thrust demand T with an angle at which it will act, such as the thrust angle ξ or the pitch angle θ (if considering just the longitudinal velocities). This results in velocity dynamics that see non-linearly coupled inputs typically in the form $T\cos(\xi)$, etc. In addition, the nonlinear dependence on thrust due to prop-wing aerodynamic interactions (commonly exhibited by these configurations) adds further complexity by making it such that no closed-form solution for the inputs exists. These concepts are illustrated in the following example.

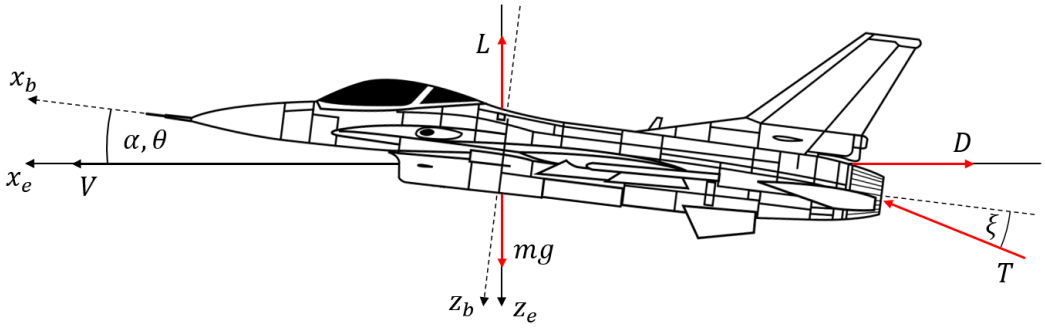

Figure 6.1: 2-dimensional thrust-vectored aircraft [134]

Example 6.1. 2-D thrust-vectored aircraft.

Consider the longitudinal plane of a generic thrust-vectored aircraft as shown in Figure 6.1. In this 2-D thrust-vectored setup, a thrust input T is commanded at thrust angle input ξ, producing forces in both the body x- and z-directions:

$$F_x = T\cos(\xi) \qquad F_z = -T\sin(\xi) \tag{6.1}$$

While both inputs are non-linearly coupled and are therefore not affine in the Earth-frame velocity dynamics, this system still admits a closed-form solution. For given force demands $F_{x,ref}$ and $F_{z,ref}$, the inputs can be solved to be:

$$T = \sqrt{F_{x,ref}^2 + F_{z,ref}^2} \qquad \xi = -\arctan\left(\frac{F_{z,ref}}{F_{x,ref}}\right) \tag{6.2}$$

However, it is assumed that the aircraft in this example sees no coupling between the thrust T and/or the thrust angle ξ and its aerodynamics. Unfortunately, this is not the case in general for transition VTOL configurations (consider the tiltwing aircraft model presented in Chapter 3), where the actual dependence on T and ξ (and also, importantly, θ) in this problem is highly nonlinear, making a strictly closed-form solution not possible. ■

Reduction to a SISO system could theoretically circumvent these issues by decoupling T from ξ and θ in the design, resulting in input-affine velocity dynamics that can therefore be feedback linearized. However, while feedback linearization has the benefit of linearizing a system across a whole domain (instead of a single point), limitations can still exist in terms of the validity of the decoupling matrix and physical range of actuation based on the physics of the real system. In the case of thrust- to wing-borne transition VTOL aircraft, where independent thrust actuation in both the vertical and horizontal directions cannot be assumed in general, decoupling the problem and applying a feedback linearization to control one of the Earth-frame velocities using solely T requires either switching between relative degree 1 and 2 systems or results in a decoupling matrix that becomes singular at certain points in the flight envelope, limiting generality of the feedback linearization approach. This concept is demonstrated in the following example.

Example 6.2. Changing relative degrees of the outer-loop velocity dynamics.

Consider once again the thrust-vectored aircraft of Figure 6.1. In addition to no coupling between the thrust T and the aircraft's aerodynamics (i.e. unwetted conditions), we make the following assumptions:

- The (Earth-frame) vertical velocity V_z will be much smaller than the horizontal velocity V_x. Therefore, the flight path angle γ is approximately zero
- The pitch angle θ, and therefore the angle of attack α, is approximately zero.

Hence, aircraft C_L and C_D are treated as constant

Consider choosing $\mathbf{u} = T$ only, and the thrust angle ξ is treated as a system parameter instead. This choice of input and the above assumptions allow the Earth-frame (NED) velocity dynamics to be expressed in an input-affine form as:

$$\begin{bmatrix} \dot{V}_x \\ \dot{V}_z \end{bmatrix} = \begin{bmatrix} -\frac{1}{m}D \\ g - \frac{1}{m}L \end{bmatrix} + \begin{bmatrix} \frac{1}{m}\cos(\xi) \\ -\frac{1}{m}\sin(\xi) \end{bmatrix} T \tag{6.3}$$

where,

$$D = \frac{1}{2}\rho\left(V_x^2 + V_z^2\right) SC_D \qquad L = \frac{1}{2}\rho\left(V_x^2 + V_z^2\right) SC_L \tag{6.4}$$

This (now) single-input system requires selecting a single output. Consider choosing the system output to be $\mathbf{y} = V_z$ to help regulate the aircraft's altitude. We explore two possible cases for feedback linearization of this system, $\xi > 0$ and $\xi = 0$, that mathematically represent examples of hover and cruise mode, respectively.

In the case that $\xi > 0$, it is readily seen that the system relative degree is $\rho = 1$, since $\dot{\mathbf{y}} = \dot{V}_z$ directly depends on $\mathbf{u} = T$. Following the feedback linearization procedure thus reveals:

$$\mathbf{E}(\mathbf{x}) = -\frac{1}{m}\sin(\xi) \tag{6.5}$$

and the feedback-linearized system from new input $\mathbf{v}$ to output $\mathbf{y}$ is just a single integrator (i.e. $1/s$). Importantly, note that $\mathbf{E}(\mathbf{x})$ becomes singular when $\xi = 0$. However, more generally, there is a loss of control authority as ξ gets small (i.e. closer to cruise mode).

Now consider the system dynamics when $\xi = 0$. We see that the $\mathbf{g}(\mathbf{x})$ term reduces to $\mathbf{g}(\mathbf{x}) = [1/m \quad 0]^T$. Hence, T no longer directly affects $\dot{V}_z$, and differentiation of the output $\mathbf{y}$ reveals that the system relative degree has become $\rho = 2$. Applying

feedback linearization to the system under these conditions gives:

$$\mathbf{E}(\mathbf{x}) = -\frac{1}{m^2}\rho V_x S C_L \tag{6.6}$$

and the feedback-linearized system from new input $\mathbf{v}$ to output $\mathbf{y}$ is now a double integrator (i.e. $1/s^2$).

Note that, while this example has considered $\mathbf{y} = V_z$, a similar result is found for $\mathbf{y} = V_x$. ■

Beyond such mathematical constraints, there are additional complexities with the outer-loop velocity control that come down to a question of how we might want these aircraft to operate (conventional hover/cruise vs continuously thrust-vectored) and the chosen design philosophy for the flight controller (a gain-scheduled/switched-mode controller vs a full-envelope design without switching). We remark that there exists a fundamental duality between the operation of these aircraft in hover and cruise modes. In hover mode (with a fixed thrust angle ξ), the thrust T primarily controls the vertical velocity V_z, while control of the horizontal velocity V_x is coupled with the pitch angle θ, analogous to a quadcopter drone. Contrast this with cruise mode, where the behavior is essentially the opposite: T now controls V_x, while V_z is broadly controlled by θ (e.g. via the elevator). This swap in control influence between the inputs and outputs suggests the need for separate (i.e. mode-specific) flight controllers in order to maintain conventional hover and cruise operation (i.e. fixed ξ) in each respective flight mode. This is further substantiated by the example just seen, in which the system relative degree changes in conjunction with $\mathbf{E}(\mathbf{x})$ becoming singular as the aircraft transitions from hover to cruise, a result of the fundamental dynamics of the problem; to achieve a full-envelope design would require implementing a switching strategy to accommodate the two different decoupling terms and their associated feedback-linearized systems, or alternatively, designing different (i.e. mode-specific)

flight controllers in order to switch between the output variable relevant to that flight mode ($\mathbf{y} = V_z$ for hover or $\mathbf{y} = V_x$ for cruise).

Hence, to achieve a full-envelope design without switching (with a lack of independent thrust actuation), a common approach in literature is to draw the analogy (either directly or indirectly) to a thrust-vectoring problem, such as that of Example 6.1, where the direction of the aircraft's thrust is controlled relative the aircraft's body axes (i.e. using ξ as a control variable). However, while seemingly intuitive, this approach has potential drawbacks. Firstly, using ξ as a control variable is likely not practical, as it assumes that the thrust angle input has a sufficiently-high bandwidth to maintain controllability of the aircraft's velocity dynamics. On a full-scale platform, this would place potentially large force requirements on the tilt actuation system, which is unlikely to be practically realizable when considering large wing (or nacelle) inertias and the need to robustly reject wind gust disturbances. Secondly, there remains the question of how we want these aircraft to operate. While this continuously thrust-vectored approach mathematically enables a full-envelope design without switching, it creates potentially undesirable behavior in flight, namely continuous adjustment of the thrust angle in all flight modes. Using a tiltwing as an example, this would mean constant (non-zero) adjustment of the wing angle to maintain controllability, which is particularly disadvantageous in, for example, cruise mode where conventional fixed-wing operation is likely desired.

The problems and challenges presented thus far render full-envelope outer-loop design a complex task without either assuming independent thrust actuation in both the vertical and horizontal directions, assuming the thrust angle input has a sufficiently-high bandwidth to be used as a flight control actuator, and/or implementing a gain schedule or other switching behavior, and current literature on the topic reflects this. The numerous examples of full-envelope outer-loop control design applied to transition VTOL aircraft in literature can be grouped into

those that use either a gain schedule or other switching behavior [88, 115, 116, 117, 118, 119, 120, 126, 135, 136, 123, 137] or take a "unified" control approach [90, 97, 122, 138, 139, 140, 141, 142, 143, 144, 127, 128, 129]. In the latter case, feedback linearization is commonly employed to achieve a full-envelope design with a single controller. Lombaerts et al. [141] investigate a unified control law comparing both NDI and INDI applied to a conceptual Urban Air Mobility (UAM) eVTOL configuration equipped with separate lift and cruise rotors. A full-envelope solution is considered by choosing to linearize both the attitude and velocity states of the aircraft simultaneously. However, due to the configuration choice, no limitation in terms of a single thrust direction is considered in the context of thrust- to wing-borne vehicles. Yayla et al. [142] take an adaptive approach to feedback linearization applied to a fixed-wing tri-rotor, a configuration equipped with two wing props that rotate forward in addition to a tail prop that rotates laterally, as shown in Figure 6.2. In a similar manner, attitude states are linearized simultaneously with the altitude

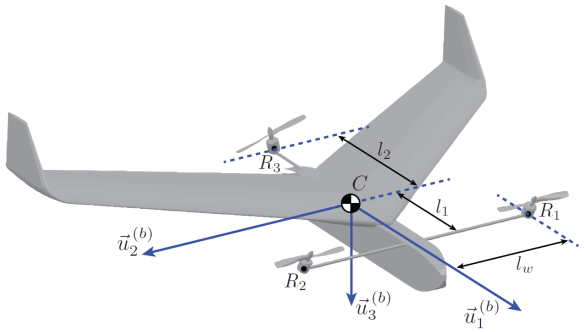

Figure 6.2: Fixed-wing tri-rotor configuration from Yayla et al. [142] (©Metehan Yayla; reprinted with permission.)

dynamics. However, their configuration's tail prop provides continuous thrust actuation in the vertical direction, and their choice of virtual inputs leverages this fact when at low wing-rotor angles to maintain invertibility of their control allocation matrix. Raab et al. [143] present a general unified control strategy for transition VTOL aircraft configurations, also by means of linearizing the attitude and velocity states simultaneously using INDI. However, while their generalized approach considers pos-

sible internal dynamics to the system and the use of virtual controls to resolve them, it does not address the inherent lack of controllability of both vertical and horizontal airspeed when thrust actuation is only available in a single direction. Following this, their given example assumes such continuous vertical thrust actuation by choice of a tiltrotor design that is equipped with fixed, vertically-oriented props. The most convincing unified control approach is presented by Milz et al. [128]. In this paper, which builds upon previous works [97, 127], sensory NDI is applied to the attitude and velocity states of a model of the Airbus A^3 Vahana and successful forward transition is achieved with this thrust- to wing-borne tiltwing configuration. However, the wing tilt angle itself is used as a control variable (i.e. the thrust angle) to enable full controllability of both the vertical and horizontal velocities with a single thrust direction, which has potentially certain disadvantages in terms of achievable bandwidth (particularly on a full-scale aircraft), as has been discussed. Furthermore, the results with this thrust-vectored approach display erratic wing angle commands during transition and non-zero wing angles in cruise. Only recently have the potential bandwidth limitations been improved upon in [129], where the pitch angle is incorporated with the tilt angle demand to utilize the faster bandwidth of the attitude dynamics and increase the overall thrust angle response. Results are improved, though no results demonstrating feasible backward transition are presented.

This chapter looks at finding ways around the fundamental challenges described in this introduction:

1. Lack of independent thrust actuation in both longitudinal directions
2. Nonlinear coupling between the thrust, thrust and pitch angles, and the aircraft's aerodynamics
3. Decoupling the problem results in changing relative degrees of the airspeed dynamics

4. Employing the thrust angle as a control variable has potentially certain disadvantages in terms of achievable bandwidth and desired flight behavior

and seeks to find a full-envelope outer-loop flight controller that is valid in all flight modes for general thrust- to wing-borne transition VTOL aircraft. We will begin by setting up the general mathematical form of the problem. We then present three novel approaches ("strategies") to tackling this problem, made possible by choosing virtual inputs and outputs that are defined within a rotating reference frame. The content of this chapter presents the theory of these novel approaches; in Chapter 7, these strategies will be applied to the tiltwing configuration of interest in this thesis to be compared and evaluated.

6.2 General Problem Setup

Consider the longitudinal NED Earth-frame velocities of an aircraft assuming lateral states are zero (refer to (2.6)):

$$\begin{aligned} V_x &= u\cos(\theta) + w\sin(\theta) \\ V_z &= -u\sin(\theta) + w\cos(\theta) \end{aligned} \tag{6.7}$$

Note that these velocities (V_x and V_z) represent Earth-frame quantities but in a reference frame always aligned with the heading of the aircraft. Additionally, the pitch angle θ (as well as the thrust angle ξ) is treated as a system parameter in the present context. Applying Newton's Second Law to a pointmass in this frame (see Figure 6.3) gives the following nonlinear state equations:

$$\begin{aligned} \dot{V}_x &= \frac{1}{m}F_{x_{net}}\cos(\theta) + \frac{1}{m}F_{z_{net}}\sin(\theta) \\ \dot{V}_z &= g - \frac{1}{m}F_{x_{net}}\sin(\theta) + \frac{1}{m}F_{z_{net}}\cos(\theta) \end{aligned} \tag{6.8}$$

where $F_{x_{net}}$ and $F_{z_{net}}$ are the net external horizontal and vertical forces, respectively, acting in the body-frame but excluding gravity. In consideration of possible wet-

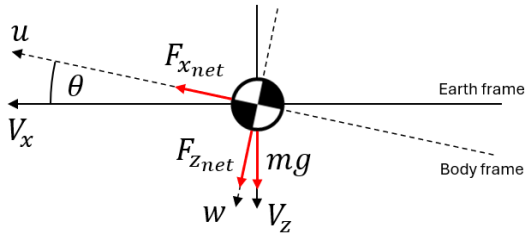

Figure 6.3: Generic 2-DoF pointmass free-body diagram

ted/unwetted aerodynamics (or, more generally, any active/passive forces), we decompose these net forces as $F_{x_{net}} = F_{x,p} + F_x$, etc., where $F_{x,p}$ denotes any passive forces (i.e. not input-dependent) and F_x denotes a generalized force input (analogous to $\mathbf{M}$ for the inner-loop design). With this notation, (6.8) can be therefore expanded and rearranged into the following nonlinear state-space in an input-affine form:

$$\begin{bmatrix} \dot{V}_x \\ \dot{V}_z \end{bmatrix} = \begin{bmatrix} \frac{1}{m} F_{x,p} \cos(\theta) + \frac{1}{m} F_{z,p} \sin(\theta) \\ g - \frac{1}{m} F_{x,p} \sin(\theta) + \frac{1}{m} F_{z,p} \cos(\theta) \end{bmatrix} + \begin{bmatrix} \frac{1}{m} \cos(\theta) & \frac{1}{m} \sin(\theta) \\ -\frac{1}{m} \sin(\theta) & \frac{1}{m} \cos(\theta) \end{bmatrix} \begin{bmatrix} F_x \\ F_z \end{bmatrix} \tag{6.9}$$

Let us now consider applying feedback linearization control to this system. If the system output is selected to be $\mathbf{y} = [V_x \quad V_z]^T$, then the system relative degree matches the order of the system and a full-state linearization is possible through direct inversion of (6.9). The resulting control law becomes:

$$\begin{bmatrix} F_x \\ F_z \end{bmatrix} = \begin{bmatrix} \frac{1}{m} \cos(\theta) & -\frac{1}{m} \sin(\theta) \\ \frac{1}{m} \sin(\theta) & \frac{1}{m} \cos(\theta) \end{bmatrix} \left(\mathbf{v} - \begin{bmatrix} \frac{1}{m} F_{x,p} \cos(\theta) + \frac{1}{m} F_{z,p} \sin(\theta) \\ g - \frac{1}{m} F_{x,p} \sin(\theta) + \frac{1}{m} F_{z,p} \cos(\theta) \end{bmatrix} \right) \tag{6.10}$$

where $\mathbf{v}$ is a new input vector generated from a linear controller acting on the feedback-linearized system, and F_x and F_z can be mapped to real actuator inputs through a chosen control allocation scheme. In this particular instance, the decoupling matrix $\mathbf{E}(\mathbf{x}) = \mathbf{g}(\mathbf{x})$ only depends on θ and is always invertible, suggesting

that this control law is theoretically valid for any pitch angle over the complete flight envelope. However, there is an issue. While general thrust- to wing-borne transition VTOL aircraft can produce forces (i.e. thrust) in both the x- and z-directions of the body frame, thus necessitating accounting for both F_x and F_z contribution, the existence of both virtual force inputs additionally suggests that these forces (thrust) can be produced in both directions independently, which is not the case in general. This limitation will arise once dealing with the allocation of F_x and F_z, at which point the effector model will be overconstrained. Foregoing using the thrust angle ξ as a control variable, this leaves us in an interesting predicament that we will look to solve through the careful selection of different virtual inputs and outputs.

6.3 Selecting a Generalized Force Input and Compound Velocity Output

The solution proposed in this research is to select a new virtual input F_c, which represents a generalized force along the direction of the aircraft's thrust vector, as shown in Figure 6.4a[1]. Because thrust- to wing-borne transition VTOL aircraft are most generally characterized by a thrust vector that rotates depending on flight mode, this suggests that F_c is defined within a rotating reference frame and is a function of the thrust angle ξ (i.e. the wing angle, in the case of a tiltwing; refer to Figures 3.4b and 3.9b to visualize this). Hence, body-frame generalized forces can be expressed in terms of this new virtual input as:

$$F_x = F_c \cos(\xi) \qquad F_z = -F_c \sin(\xi) \tag{6.11}$$

[1]Once again, note that a generalized force is considered here instead of solely thrust to account for possibly input-dependent aerodynamics, such as propeller-wing interactions, as is common for these transition VTOL configurations.

Unfortunately, in the most general sense these are not complete expressions, as "active" forces can appear in the perpendicular direction too (such as lift generated by prop-induced velocities, which is a function of thrust). This is illustrated by the following example.

Example 6.3. Deriving F_c for a simplified 2-DoF tiltwing.

Consider a 2-DoF pointmass model of a tiltwing that models the Earth-frame velocity dynamics as in (6.8). For simplicity, the model will consist of a single aerodynamic section representing the main wing of the aircraft, and there will only be a single thrust source T (consider Figure 3.9b as an example). In contrast to the other examples thus far, the aerodynamics are assumed to be fully wetted. For such a model, the tiltwing's net external forces consist solely of thrust and aerodynamic forces, giving the following state equations:

$$\begin{aligned} \dot{V}_x &= \frac{1}{m}T\cos(\xi+\theta) - \frac{1}{m}D_{wet}\cos(\alpha_{wet}-\theta) + \frac{1}{m}L_{wet}\sin(\alpha_{wet}-\theta) \\ \dot{V}_z &= g - \frac{1}{m}T\sin(\xi+\theta) - \frac{1}{m}D_{wet}\sin(\alpha_{wet}-\theta) - \frac{1}{m}L_{wet}\cos(\alpha_{wet}-\theta) \end{aligned} \tag{6.12}$$

where,

$$D_{wet} = \frac{1}{2}\rho V_{wet}^2 S C_D \qquad L_{wet} = \frac{1}{2}\rho V_{wet}^2 S C_L \tag{6.13}$$

and,

$$V_{wet} = \sqrt{\frac{2T}{\rho A_{prop}} + V_x^2 + V_z^2} \tag{6.14}$$

Hence, the subscript *wet* emphasizes wetted quantities, i.e. terms that are a function of the thrust T due to prop-wing interactions.

Deriving F_c for this system gives:

$$F_c = T - D_{wet}\cos(\alpha_{wet}+\xi) + L_{wet}\sin(\alpha_{wet}+\xi) \tag{6.15}$$

where it can be seen that wetted aero terms (which depend on T) are present in

addition to T itself. What about in the perpendicular direction then? Deriving this for this system reveals perpendicular forces of $D_{wet}\sin(\alpha_{wet}+\xi)+L_{wet}\cos(\alpha_{wet}+\xi)$. Hence, T itself has disappeared, but its influence remains via the wetted aero terms that are still present in this direction. ■

Therefore, an additional term must be added to the general setup for completeness, here we will call $F_{c_\perp}$, which represents any input-dependent forces in a direction perpendicular to the direction of thrust. This makes the complete expressions for the body-frame generalized forces:

$$F_x = F_c\cos(\xi) - F_{c_\perp}\sin(\xi) \qquad F_z = -F_c\sin(\xi) - F_{c_\perp}\cos(\xi) \tag{6.16}$$

It is important to note that $F_{c_\perp}$ is implicitly a function of F_c (i.e. $F_{c_\perp} = f(F_c)$) and is therefore not an independent input in the most general case, since input-dependent aerodynamics are a function of thrust as was just seen in the above example. Substituting the expressions for F_x and F_z from (6.16) into the state-space given in (6.9) gives:

$$\begin{bmatrix}\dot{V}_x\\ \dot{V}_z\end{bmatrix} = \begin{bmatrix}\frac{1}{m}F_{x,p}\cos(\theta)+\frac{1}{m}F_{z,p}\sin(\theta)\\ g-\frac{1}{m}F_{x,p}\sin(\theta)+\frac{1}{m}F_{z,p}\cos(\theta)\end{bmatrix} + \begin{bmatrix}\frac{1}{m}\cos(\xi+\theta)\\ -\frac{1}{m}\sin(\xi+\theta)\end{bmatrix}F_c + \begin{bmatrix}-\frac{1}{m}\sin(\xi+\theta)\\ -\frac{1}{m}\cos(\xi+\theta)\end{bmatrix}F_{c_\perp} \tag{6.17}$$

where trigonometric identities for the sum and difference of two angles have been used. With this selection of a generalized force input, direct thrust actuation is always available while still accounting for the effect of both body x- and z-direction generalized forces in the model.

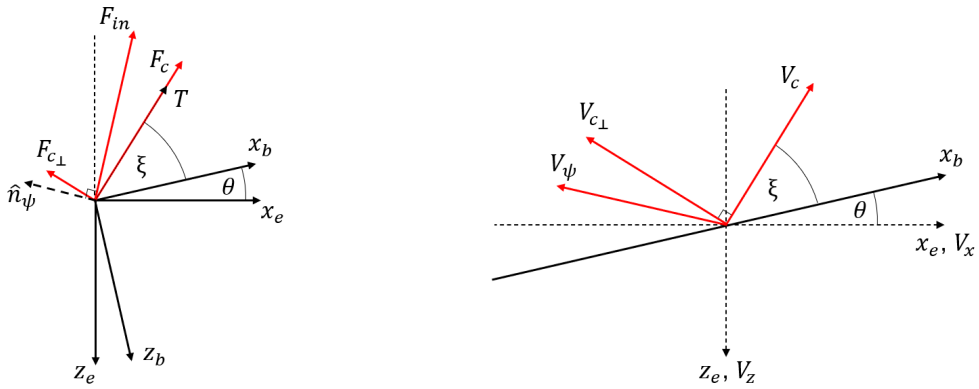

(a) Force Diagram (b) Velocity Diagram

Figure 6.4: Force and Velocity Diagrams

This solves the problem of direct thrust actuation, but we now must consider our output selection. Clearly, the now single-input[2] 2-state system is not full-state linearizable, but it is still input-output linearizable for a judicious choice of output. An intuitive choice might be to choose V_z alone, for example, as we may be primarily concerned with maintaining the aircraft's altitude[3], but doing so results in the issue of changing relative degrees as was seen in Example 6.2. Therefore, we are interested in finding a virtual output that bypasses this mathematical limitation. In a similar manner to choosing F_c, the solution proposed in this research is to choose a compound velocity output V_c defined as:

$$V_c = V_x \cos(\xi + \theta) - V_z \sin(\xi + \theta) \tag{6.18}$$

This compound velocity output represents a velocity pointing in the direction of the thrust vector in an analogous sense to the generalized force input F_c (see Figure 6.4b). As will be seen in the next section, this selection of input and output results in a decoupling matrix that is constant and invertible over the whole flight domain.

[2]Again, recall that $F_{c_\perp}$ is not an independent input due to being implicitly a function of F_c in the most general case, so F_c is treated as the sole input to the system in (6.17).

[3]Controlling the aircraft's horizontal airspeed could be a secondary objective in this example.

6.4 Strategy 1 (simplified case)

The nonlinear dependence of $F_{c_\perp}$ on F_c means that the system described in (6.17) is not strictly affine in F_c. It turns out that this will not limit our ability to apply feedback linearization to this system. However, for the purpose of introducing this controller concept, we will first simplify the analysis by initially assuming an idealized case in which $F_{c_\perp} = 0$. We will then later reintroduce $F_{c_\perp}$ and explore how it changes our derivation.

6.4.1 Controller Definition

Assuming the simplified case of $F_{c_\perp} = 0$, which corresponds to no input-dependent forces appearing in the perpendicular direction[4], (6.17) and (6.18) compose an input-affine system in the form:

$$\begin{aligned} \dot{\mathbf{x}} &= \mathbf{f}(\mathbf{x}) + \mathbf{g}(\mathbf{x})\mathbf{u} \\ \mathbf{y} &= \mathbf{h}(\mathbf{x}) \end{aligned} \tag{6.19}$$

where $\mathbf{u} = F_c$ and $\mathbf{y} = V_c$. Differentiating the output $\mathbf{y}$ reveals that this system has relative degree 1, leading to the feedback linearization terms $\mathbf{E}(\mathbf{x})$ and $\mathbf{b}(\mathbf{x})$ being defined as:

$$\begin{aligned} \mathbf{E}(\mathbf{x}) &= L_{\mathbf{g}}\mathbf{h}(\mathbf{x}) = 1/m \\ \mathbf{b}(\mathbf{x}) &= L_{\mathbf{f}}\mathbf{h}(\mathbf{x}) = \dot{V}_{x,p}\cos(\xi+\theta) - \dot{V}_{z,p}\sin(\xi+\theta) \end{aligned} \tag{6.20}$$

where $\dot{V}_{x,p}$ represents the first element of $\mathbf{f}(\mathbf{x})$ (i.e. a "passive" quantity), etc. Of note is that the resulting decoupling matrix $\mathbf{E}(\mathbf{x})$ (a scalar in this case) is constant and only depends on the aircraft's mass, and it is therefore invertible over the complete flight domain. With these terms derived, the following standard feedback linearization

[4]An example of this in practice could be the case in which it is assumed there are no wetted aerodynamics, such as for the aircraft model given in Example 6.2.

control law can be formulated:

$$\mathbf{u} = \mathbf{E}^{-1}(\mathbf{x})(\mathbf{v} - \mathbf{b}(\mathbf{x})) \tag{6.21}$$

where $\mathbf{v}$ is the new input computed from a linear control law acting on the following linear system:

$$\begin{aligned} \dot{\mu} &= \mathbf{v} \\ \mathbf{y} &= \mu \end{aligned} \tag{6.22}$$

where $\mu = V_c$ represents the system's external dynamics.

6.4.2 Internal Dynamics

The relative degree of this system is 1, but the order of the system is 2. Hence, this transformation only represents an input-output linearization, and there remains dynamics internal to the system that are unobservable by this transformation. Recall from Section 2.3.3 that a diffeomorphism $\boldsymbol{\phi}(\mathbf{x}) = [\mu \quad \psi]^T$ can be formed in normal coordinates μ and ψ, where μ represents the external dynamics of the system and ψ represents the internal dynamics. For this system, $\mu = V_c$, and ψ must be found such that the coordinate transformation $\boldsymbol{\phi}(\mathbf{x})$ is invertible. Due to the set of the columns of $\mathbf{g}(\mathbf{x})$ trivially satisfying involutivity for this SISO system (i.e. $\mathbf{g}(\mathbf{x})$ is a column vector), it is possible to choose ψ such that $L_{\mathbf{g}}\psi = 0$, which gives the following PDE:

$$\frac{\partial \psi}{\partial V_x}\cos(\xi + \theta) - \frac{\partial \psi}{\partial V_z}\sin(\xi + \theta) = 0 \tag{6.23}$$

A solution to this equation can be found to be:

$$\psi = -V_x \sin(\xi + \theta) - V_z \cos(\xi + \theta) \tag{6.24}$$

which represents a velocity perpendicular to V_c. Hence, we will denote this velocity $\psi = V_{c_\perp}$ (see Figure 6.4b). The Jacobian matrix of this diffeomorphism $\boldsymbol{\phi}(\mathbf{x})$ is then:

$$\frac{\partial \boldsymbol{\phi}}{\partial \mathbf{x}} = \begin{bmatrix} \cos(\xi+\theta) & -\sin(\xi+\theta) \\ -\sin(\xi+\theta) & -\cos(\xi+\theta) \end{bmatrix} \tag{6.25}$$

which is non-singular for any ξ or θ. Hence, this state transformation is a global diffeomorphism and the inverse transformation $\boldsymbol{\phi}^{-1}(\mu, \psi)$ is given by:

$$\begin{aligned} V_x &= \mu\cos(\xi+\theta) - \psi\sin(\xi+\theta) \\ V_z &= -\mu\sin(\xi+\theta) - \psi\cos(\xi+\theta) \end{aligned} \tag{6.26}$$

And the normal form equations for this system are therefore:

$$\begin{aligned} \dot{\mu} &= -g\sin(\xi+\theta) + \frac{1}{m}F_{x,p}\cos(\xi) - \frac{1}{m}F_{z,p}\sin(\xi) + \frac{1}{m}F_c \\ \dot{\psi} &= -g\cos(\xi+\theta) - \frac{1}{m}F_{x,p}\sin(\xi) - \frac{1}{m}F_{z,p}\cos(\xi) \end{aligned} \tag{6.27}$$

As with any internal dynamics, it would be necessary to ensure that they remain stable throughout the complete flight envelope for a particular aircraft. (6.27) above can be used to verify the (local) stability of ψ by analyzing the zero-dynamics (i.e. when $\mu = \dot{\mu} = 0$; refer to Section 2.3.3). However, stability alone does not guarantee that they will converge to the desired reference trajectory, as we will see next. Hence, we will alternatively look for a way to control these dynamics.

6.4.3 Controlling the Internal Dynamics through Pitch Reference Commands

Consider a reference input demand $V_{c_{ref}}$ to this controller:

$$V_{c_{ref}} = V_{x_{ref}}\cos(\xi+\theta) - V_{z_{ref}}\sin(\xi+\theta) \tag{6.28}$$

where $V_{x_{ref}}$ and $V_{z_{ref}}$ are the actual outer-loop reference input demands from the pilot. Because of control of V_c through F_c (via (6.21)), V_c will be driven to $V_{c_{ref}}$ in steady-state (provided sufficient integral action to ensure zero steady-state error). However, due to the existence of the internal dynamics ψ in the inverse transformation of (6.26), V_c being driven to $V_{c_{ref}}$ alone does not guarantee that V_x and V_z will be driven to $V_{x_{ref}}$ and $V_{z_{ref}}$, respectively, as both $\mu = V_c$ and $\psi = V_{c_\perp}$ must be driven to their respective references. However, the latter will instead evolve according to (6.27). Hence, to ensure complete tracking of the commanded Earth-frame velocities, something must drive ψ to ψ_{ref} (i.e. ψ evaluated at $V_{x_{ref}}$ and $V_{z_{ref}}$) in addition to F_c driving V_c to $V_{c_{ref}}$.

The solution proposed in this research is to leverage the use of pitch control to regulate the system's internal dynamics. The ability for θ to regulate ψ is analogous to a conventional quadcopter's longitudinal velocity control being coupled with pitch angle (as described earlier). What is unique in this control strategy, however, is since $\psi = V_{c_\perp}$ is a function of thrust angle (and thus points in a different direction depending on flight mode), θ will automatically vary its control influence between V_x in hover mode and V_z in cruise, thus generalizing the conventional hover/cruise operation described at the beginning of this chapter.

Multiple strategies can be found to regulate ψ using θ. For this research, it has already been assumed that a sufficiently-fast inner-loop attitude controller exists, allowing for θ to be controlled through desired θ_{ref} values. Hence, for this first controller definition, we will apply the same strategy originally introduced in [26], where θ_{ref} is determined by finding the necessary change in pitch angle that bridges the gap between current body-frame velocities (u and w) and desired Earth-frame

velocity references ($V_{x_{ref}}$ and $V_{z_{ref}}$):

$$\begin{aligned} V_{x_{ref}} &= u\cos(\theta_{ref}) + w\sin(\theta_{ref}) \\ V_{z_{ref}} &= -u\sin(\theta_{ref}) + w\cos(\theta_{ref}) \end{aligned} \tag{6.29}$$

where this relationship is obtained from (6.7). (6.29) above can then be rearranged to solve for θ_{ref} as[5]:

$$\theta_{ref} = \arctan\left(\frac{-uV_{z_{ref}} + wV_{x_{ref}}}{uV_{x_{ref}} + wV_{z_{ref}}}\right) \tag{6.30}$$

Note that this control strategy effectively serves as a means of course correction by regulating the internal dynamics in ψ-space and should therefore be regarded as an outer-loop relative the speed of the dynamics of the feedback linearization controller proposed earlier. To ensure sufficient time-scale separation between these two loops (as well as with the inner-loop attitude controller)[6], the bandwidth of θ_{ref} is set using a first-order lag $1/(1+sT)$, where T is a time constant to be chosen appropriately.

While simple and effective, this strategy has the downside of (typically) requiring different first-order time constants depending on flight mode/flight condition and has been found to be sensitive to this choice. In Section 6.6, we will present an alternative, model-based design of $V_{c_\perp}$ using a pseudo-control allocation scheme. However, with this control strategy now presented, we will first look at the consequences of reintroducing the $F_{c_\perp}$ term.

[5]Note that this mapping in (6.30) is singular when both u and w are zero (i.e. pure hover). However, issues can be avoided if θ_{ref} is simply set to its equilibrium hover value (e.g. zero) when the aircraft's total velocity is zero.

[6]Note that, while this may seem to imply a bandwidth difference of two orders of magnitude between the θ_{ref}-loop and the inner-loop attitude controller, in practice F_c is demanded in parallel with the inner-loop's virtual moment commands due to the inherent force and moment coupling exhibited by these transition VTOL configurations, so the F_c-loop is not bandwidth limited by the tuning of the attitude loop and the bandwidth of the θ_{ref}-loop can be made faster.

6.5 Strategy 1 (general case)

Considering the more general case with $F_{c_\perp}$ no longer assumed to be zero, the coordinate transformation $\boldsymbol{\phi}(\mathbf{x})$ from Section 6.4.2 is still a valid diffeomorphism, and the V_c-$V_{c_\perp}$ dynamics with $F_{c_\perp}$ now included look like:

$$\begin{aligned}\dot{V}_c &= -g\sin(\xi+\theta) + \frac{1}{m}F_{x,p}\cos(\xi) - \frac{1}{m}F_{z,p}\sin(\xi) + \frac{1}{m}F_c \\ \dot{V}_{c_\perp} &= -g\cos(\xi+\theta) - \frac{1}{m}F_{x,p}\sin(\xi) - \frac{1}{m}F_{z,p}\cos(\xi) + \frac{1}{m}F_{c_\perp}\end{aligned} \tag{6.31}$$

where $F_{c_\perp}$ now appears in the $\dot{V}_{c_\perp}$ equation and $\dot{V}_c$ remains unchanged.

The first remark is that since $\dot{V}_c$ remains unchanged, it turns out that the original controller definition from Section 6.4.1 remains valid. This is because the non-affine portion of the input F_c in the original system (via $F_{c_\perp}$) is eliminated in the $\dot{V}_c$ equation when transforming the system to the V_c-$V_{c_\perp}$ frame, and it can be seen that the exact same feedback linearization terms $\mathbf{E}(\mathbf{x})$ and $\mathbf{b}(\mathbf{x})$ from (6.20) appear through direct inversion of the $\dot{V}_c$ equation.

However, the second remark is that the current choice of $\psi = V_{c_\perp}$ no longer satisfies $L_{\mathbf{g}}\psi = 0$, since the dynamics of ψ now depend on the input $\mathbf{u} = F_c$ through the $F_{c_\perp}$ term. Of course, $\psi = V_{c_\perp}$ is still a valid expression for the internal dynamics, since it completes the coordinate transformation $\boldsymbol{\phi}(\mathbf{x})$, but a question naturally arises: What expression for ψ would therefore satisfy $L_{\mathbf{g}}\psi = 0$? Unfortunately, finding a closed-form/analytical expression for this in a general case is not possible. However, it is possible to describe its direction in the longitudinal plane to at least gain intuition into what this would be. First, we know that the "net" external generalized force input vector, here denoted F_{in}, will necessarily be the vector sum of F_c and $F_{c_\perp}$, and we can therefore describe its magnitude as $\sqrt{F_c^2 + F_{c_\perp}^2}$ and its angle (relative F_c) as $\arctan(F_{c_\perp}/F_c)$. This is visualized in Figure 6.4a. Then, in order to satisfy $L_{\mathbf{g}}\psi = 0$, it is necessarily the case that ψ points in a direction perpendicular to F_{in}. We will

denote this direction to be $\hat{n}_\psi$ with corresponding internal dynamics (velocity) state V_ψ (see Figures 6.4a and 6.4b). Hence, by definition, no input forces appear in the direction of V_ψ and therefore the input $\mathbf{u}$ will not be present in the (theoretical) $\dot{V}_\psi$ equation. In any case, the original selection of $\psi = V_{c_\perp}$ still adequately describes the internal dynamics of the system and is sufficient for deriving the system's zero-dynamics for the purpose of stability analysis.

Lastly, to control $V_{c_\perp}$, Section 6.4.3 presented a kinematics-based mapping for θ_{ref}. However, the $\dot{V}_{c_\perp}$ equation has changed and now includes $F_{c_\perp}$. Fortunately, it turns out that this mapping still works for driving $V_{c_\perp}$ to its respective reference, as the kinematic relationship in which the mapping was based remains unchanged[7]. That said, we will now explore a more formal, model-based design of $V_{c_\perp}$.

6.6 Strategy 2: Formal Design of $V_{c_\perp}$ Using a Pseudo-Control Allocation Scheme

Section 6.4.3 presented a kinematics-based mapping for regulating the internal dynamics that is simple and effective but has certain downsides, such as being sensitive to the choice of time constant in the necessary first-order lag. In this section, we present an alternative, model-based design for controlling $V_{c_\perp}$ through θ_{ref} commands by making use of a pseudo-control allocation scheme.

6.6.1 Motivation and Problem Setup

The problem with formally designing a (full-envelope) controller for $V_{c_\perp}$ through θ_{ref} commands is the highly nonlinear dependence of $\dot{V}_{c_\perp}$ on θ. As has been discussed, if it is desired to implement a feedback linearization design to achieve this, it is (gener-

[7]It is necessary, however, that the first-order time constant be chosen carefully to avoid additional interaction between these loops now that there is a stronger coupling between V_c and $V_{c_\perp}$ through the presence of $F_{c_\perp} = f(F_c)$.

ally speaking) necessary that the inputs be affine, which is of course not the case for $\dot{V}_{c_\perp}$ and θ. One way around this is to manipulate the way in which the input enters the system through a technique called dynamic extension [40, 145, 146], in which a first-order state on the input is introduced to create a new affine input through integration. Unfortunately, when applied to the $V_{c_\perp}$-dynamics, this technique frequently forms additional internal dynamics that are unstable, rendering this approach unfit. Another full-envelope approach could be to do a Lyapunov-based design for $V_{c_\perp}$ to find an expression for θ that drives the error dynamics to zero at all points in the flight envelope. However, due to the complexity of the outer-loop dynamics, this quickly becomes analytically intractable for this system. On-line numerical approaches, such as a nonlinear root solver, could be tried instead, but these often turn out to be too inefficient (in a general case) for a real-time application[8].

In this work, we present an alternative approach to enable a full-envelope design through the use of a pseudo-control allocation scheme. Where a typical control allocation seeks to map virtual inputs to real actuator inputs, we instead consider an intermediate mapping of virtual inputs to a new/different set of virtual inputs (hence, a "pseudo-control" allocation). Such a mapping from virtual to real inputs in a standard control allocation is made possible by the fact that the mapping itself contains no dynamics and is thus simply determined from the aircraft's effector model (e.g. as was seen in Section 5.3.4). In contrast, in consideration of the $V_{c_\perp}$-dynamics, there are no real[9] inputs available in the $V_{c_\perp}$-direction to be theoretically allocated. However, we propose treating the pitch angle θ like a "real" input (in the context of the outer-loop design) and collecting up all terms involving θ in the $V_{c_\perp}$-dynamics

[8]Recent work involving successive convexification receding horizon control [75] could be a viable alternative, but this method was not explored in this project.

[9]At least no reasonable ones; flaps, for example, would contribute to $F_{c_\perp}$ and therefore affect $\dot{V}_{c_\perp}$, but they do not provide sufficient control authority and cannot be used under all flight conditions.

(see (6.31)) into a single quantity to form a new virtual input $\overline{F}_{c_\perp}$, resulting in:

$$\dot{V}_{c_\perp} = \frac{1}{m}\overline{F}_{c_\perp} \tag{6.32}$$

where,

$$\overline{F}_{c_\perp} = -mg\cos(\xi+\theta) - F_{x,p}\sin(\xi) - F_{z,p}\cos(\xi) + F_{c_\perp} \tag{6.33}$$

This is made possible by the inner-/outer-loop time-scale separation assumptions, which render $\theta \approx \theta_{ref}$. Hence, no dynamics therefore exist between θ and $\overline{F}_{c_\perp}$, and a direct mapping between these values can be found. Note that $F_{x,p}$, etc., have been assumed to be a function of θ due to primarily involving aerodynamic terms and, through the same time-scale separation assumptions, are thus affected by changes in θ since $\Delta\theta \approx \Delta\alpha$ within a given time step. By virtue of this new virtual input selection, the system in (6.32) has been rendered linear (as all terms involve θ)[10], and a linear control law for $\mathbf{u} = \overline{F}_{c_\perp}$ can be implemented to control $\mathbf{y} = V_{c_\perp}$ that will remain valid across the whole flight domain. What remains then is defining and inverting the mapping from $\overline{F}_{c_\perp}$ to θ using a chosen control allocation method.

To get a sense for how this approach fits in to the rest of the control design and the original/standard control allocation, Figure 6.5 presents a diagram of the flow of variable assignments from the outer-loop $V_{c_\perp}$ control to the aircraft's real input demands $\mathbf{u}$.

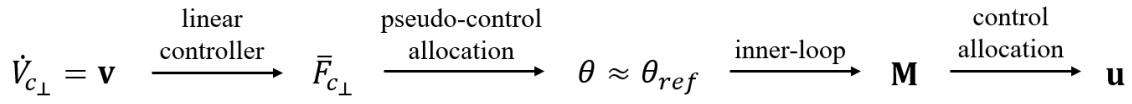

Figure 6.5: Pseudo-control allocation scheme diagram showing flow of variable assignments for $V_{c_\perp}$ control using θ

[10]Note that if some terms remain that do not involve θ, this procedure still in the least renders the system input-affine.

6.6.2 Allocation Method

For this pseudo-control allocation scheme, it was chosen to numerically form the mapping (a "pseudo-effector" model) from θ to $\overline{F}_{c_\perp}$ and linearize this mapping across the expected flight envelope. This was accomplished by applying the following first-order Taylor expansion:

$$\overline{F}_{c_\perp} \approx \overline{F}_{c_\perp}(\mathbf{x}_0, \mathbf{u}_0) + \left.\frac{\partial \overline{F}_{c_\perp}}{\partial \theta}\right|_{(\mathbf{x}_0, \mathbf{u}_0)} (\theta - \theta_0) = \overline{F}_{c_{\perp,0}} + B_\theta(\theta - \theta_0) \tag{6.34}$$

where $\mathbf{x}_0$ and $\mathbf{u}_0$ indicate the nominal system states and inputs, respectively, and θ_0 indicates the nominal value of θ at the same operating point, etc. Therefore, with θ_{ref} assumed to be equal to θ through time-scale separation assumptions, this (SISO) pseudo-control allocation scheme mapping can be completed and assigned to θ_{ref} through direct inversion:

$$\theta_{ref} = \frac{1}{B_\theta}(\overline{F}_{c_\perp} - \overline{F}_{c_{\perp,0}}) + \theta_0 \tag{6.35}$$

to enter the inner-loop attitude controller as a pitch reference command. To implement this, B_θ, $\overline{F}_{c_{\perp,0}}$, and θ_0 are pre-computed and scheduled along the expected flight trajectory of the aircraft.

6.6.3 Discussion

The advantage of this approach is that it is model-based, and the pseudo-control allocation constants (B_θ, $\overline{F}_{c_{\perp,0}}$, and θ_0) simply need to be pre-computed based on the "pseudo-effector" model. Specifically, if the aircraft model ever changes, the (linear) controller itself does not need re-tuned—the pseudo-control allocation constants would simply need re-computed once. This contrasts the kinematics-based mapping for θ_{ref} of Section 6.4.3, which is (generally) quite sensitive to the first-order time

constant selection, and the time constant typically needs to be re-tuned if the system model ever changes. The obvious trade-off with the proposed pseudo-control allocation method, though, is increased complexity in the initial setup.

It is thus important to note that, while this approach is model-based and thus foregoes the need for manual tuning of a first-order time constant, this control strategy still effectively serves as a means of course correction by regulating the internal dynamics in ψ-space. Therefore, similar to the kinematics-based mapping for θ_{ref} of Section 6.4.3, this control strategy should be regarded as an outer-loop relative the speed of the dynamics of the feedback linearization controller for V_c and, hence, should be tuned accordingly (to ensure sufficient time-scale separation). A primary reason for this requirement is due to the lack of accounting for any cross-coupling between F_c and θ in the V_c-$V_{c_\perp}$ dynamics in, what has been thus far, a decoupled approach. Hence, we will next consider a coupled MIMO design.

6.7 Strategy 3: Coupled MIMO Approach

Up to this point, we've considered the novel approach of enabling a feedback linearization design of the outer-loop velocity control for general transition VTOL aircraft by choosing virtual inputs and outputs in a rotating reference frame that results in a globally-valid transformation of the velocity states. However, we have thus far only considered a decoupled approach to the problem, where V_c and $V_{c_\perp}$ are individually controlled by their respective inputs (F_c and $\theta \approx \theta_{ref}$), and any cross-coupling has been ignored (e.g. how θ demands affect the V_c-dynamics). This inherently limits the potential performance of this design, such as in terms of achievable bandwidth and tuning versatility, due to the coupling/interaction between these two loops. Hence, a coupled MIMO design should be considered, but, due to the unique problems faced by these transition VTOL configurations (as discussed in this chapter's introduction),

such a design is rendered analytically intractable in a general case, requiring numerical methods.

6.7.1 MIMO Pseudo-Control Allocation Motivation

The most direct approach to this problem would be to consider an on-line nonlinear root solver. For example, consider the V_c-$V_{c_\perp}$ dynamics from (6.31) and rewrite them as:

$$\begin{aligned} \dot{V}_c &= f_1(F_c, \theta) \\ \dot{V}_{c_\perp} &= f_2(F_c, \theta) \end{aligned} \tag{6.36}$$

where f_1 and f_2 are functions representing the nonlinear dynamics of V_c and $V_{c_\perp}$, respectively, with inputs F_c and θ. While feedback linearization, for example, generally requires the inputs to be affine, this is ultimately for the purpose of making the nonlinear dynamics analytically invertible in a general case. However, this inversion process is mathematically equivalent to simply finding the solutions (i.e. roots) to a set of nonlinear equations, and so there is no reason why this couldn't be done for a general nonlinear system through a root solving algorithm[11]—in fact, it can be shown [145, 146] that such a numerical inversion of a set of general nonlinear equations results in the same linear system in Brunovsky form (i.e. from $\mathbf{v}$ to $\mathbf{y}$) that would be expected after applying an equivalent feedback linearization transformation. Hence, a linear MIMO control law could be designed for V_c and $V_{c_\perp}$ (in a new input $\mathbf{v}$, as has been done for feedback linearization), and the inputs F_c and θ could then be found by solving (6.36) using a nonlinear root solving algorithm at each time step.

Unfortunately, while this concept works in theory (and simulation), the general complexity of the outer-loop velocity dynamics typically renders this method impractical in a general case due to its inability to run in real time[12]. While exceptions to

[11] Provided a solution exists, which could be determined/confirmed by utilizing the implicit function theorem [147].

[12] This would of course depend on the available computing resources, but these could be limited

this could be found depending on a specific aircraft configuration and level of complexity/fidelity of its model, we are instead interested in finding a general approach that is applicable to a wide range of transition VTOL configurations. Hence, in this work, we propose utilizing the pseudo-control allocation concept introduced in Section 6.6 (combined with using the virtual inputs F_c and θ in the V_c-$V_{c_\perp}$ frame) in a MIMO context to make this system solvable in real time.

6.7.2 Full-Envelope MIMO Design

To begin, we once again take the V_c-$V_{c_\perp}$ dynamics from (6.31), but we now rewrite them with new virtual inputs $\overline{F}_c$ and $\overline{F}_{c_\perp}$ by collecting up all terms containing the original virtual inputs F_c and θ:

$$\begin{aligned} \dot{V}_c &= \frac{1}{m}\overline{F}_c \\ \dot{V}_{c_\perp} &= \frac{1}{m}\overline{F}_{c_\perp} \end{aligned} \tag{6.37}$$

where,

$$\begin{aligned} \overline{F}_c &= -mg\sin(\xi+\theta) + F_{x,p}\cos(\xi) - F_{z,p}\sin(\xi) + F_c \\ \overline{F}_{c_\perp} &= -mg\cos(\xi+\theta) - F_{x,p}\sin(\xi) - F_{z,p}\cos(\xi) + F_{c_\perp} \end{aligned} \tag{6.38}$$

in an embedded application relying on an on-board microcontroller.

Hence, $\overline{F}_c$ and $\overline{F}_{c_\perp}$ can be expanded in the inputs F_c and θ using a first-order Taylor series:

$$\begin{aligned}\overline{F}_c &\approx \overline{F}_c(\mathbf{x}_0, \mathbf{u}_0) + \left.\frac{\partial \overline{F}_c}{\partial F_c}\right|_{(\mathbf{x}_0,\mathbf{u}_0)} (F_c - F_{c_0}) + \left.\frac{\partial \overline{F}_c}{\partial \theta}\right|_{(\mathbf{x}_0,\mathbf{u}_0)} (\theta - \theta_0) \\ &= \overline{F}_{c_0} + B_{11}\delta F_c + B_{12}\delta\theta \\ \overline{F}_{c_\perp} &\approx \overline{F}_{c_\perp}(\mathbf{x}_0, \mathbf{u}_0) + \left.\frac{\partial \overline{F}_{c_\perp}}{\partial F_c}\right|_{(\mathbf{x}_0,\mathbf{u}_0)} (F_c - F_{c_0}) + \left.\frac{\partial \overline{F}_{c_\perp}}{\partial \theta}\right|_{(\mathbf{x}_0,\mathbf{u}_0)} (\theta - \theta_0) \\ &= \overline{F}_{c_{\perp,0}} + B_{21}\delta F_c + B_{22}\delta\theta\end{aligned} \tag{6.39}$$

This can be organized into the following matrix system:

$$\begin{bmatrix} \overline{F}_c \\ \overline{F}_{c_\perp} \end{bmatrix} = \begin{bmatrix} \overline{F}_{c_0} \\ \overline{F}_{c_{\perp,0}} \end{bmatrix} + \begin{bmatrix} B_{11} & B_{12} \\ B_{21} & B_{22} \end{bmatrix} \begin{bmatrix} \delta F_c \\ \delta\theta \end{bmatrix} \tag{6.40}$$

which can be written more compactly as:

$$\overline{\mathbf{F}} = \overline{\mathbf{F}}_0 + \mathbf{B}_{\overline{\mathbf{F}}}\delta\mathbf{u} \tag{6.41}$$

Hence, analogous to Section 6.6.2 (and with θ_{ref} assumed to be equal to θ through time-scale separation assumptions), this (MIMO) pseudo-control allocation scheme mapping can be completed and assigned to $\mathbf{u} = [F_c \quad \theta_{ref}]^T$ through direct inversion:

$$\mathbf{u} = \mathbf{B}_{\overline{\mathbf{F}}}^{-1}\left(\overline{\mathbf{F}} - \overline{\mathbf{F}}_0\right) + \mathbf{u}_0 \tag{6.42}$$

Note that the element B_{11} is simply unity by definition (see (6.38)), and B_{22} is equivalent to B_θ from Section 6.6.2. Hence, this MIMO setup has introduced two new terms (B_{12} and B_{21}) that account for the cross-coupling between F_c and θ in the V_c-$V_{c_\perp}$ dynamics. To implement this, the linearization quantities $\mathbf{B}_{\overline{\mathbf{F}}}$, $\overline{\mathbf{F}}_0$, and $\mathbf{u}_0$ are

pre-computed and scheduled along the expected flight trajectory of the aircraft, just as before.

6.7.3 Discussion

At this point, it is interesting to observe that mathematically this pseudo-control allocation method (whether MIMO or SISO) could be equivalently set up and implemented as a direct inversion control law (i.e. NDI), in which case it can be viewed as equivalent to a feedback linearization of the system $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{u})$ that has been made input-affine through linearization with respect to the inputs $\mathbf{u}$ while remaining a nonlinear function of the system state $\mathbf{x}$ and then evaluated at the operating point $\mathbf{x} = \mathbf{x}_0$. To illustrate this, first consider the linearization of $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{u})$ with respect to only $\mathbf{u}$:

$$\dot{\mathbf{x}} \approx \mathbf{f}(\mathbf{x}, \mathbf{u}_0) + \left.\frac{\partial \mathbf{f}}{\partial \mathbf{u}}\right|_{(\mathbf{x}, \mathbf{u}_0)} \delta\mathbf{u} \tag{6.43}$$

We could more familiarly express this as:

$$\dot{\mathbf{x}} = \mathbf{f}_0(\mathbf{x}) + \mathbf{g}_0(\mathbf{x})\delta\mathbf{u} \tag{6.44}$$

which thus represents the nonlinear system dynamics in an input-affine form. Hence, $\mathbf{f}_0$ and $\mathbf{g}_0$ represent the input-affine nonlinear system terms $\mathbf{f}$ and $\mathbf{g}$ for a given $\mathbf{u}_0$, and these terms are thus explicitly (nonlinear) functions of $\mathbf{x}$. Therefore, when evaluated at the current system state (i.e. operating point) $\mathbf{x} = \mathbf{x}_0$, these dynamics become:

$$\dot{\mathbf{x}} = \mathbf{f}_0(\mathbf{x}_0) + \mathbf{g}_0(\mathbf{x}_0)\delta\mathbf{u} \tag{6.45}$$

With this in mind, consider the (MIMO) pseudo-control allocation scheme of Section 6.7.2. By substituting (6.40) into (6.37), the V_c-$V_{c_\perp}$ dynamics can be written in the

same form as (6.45):

$$\begin{bmatrix} \dot{V}_c \\ \dot{V}_{c_\perp} \end{bmatrix} = \begin{bmatrix} \frac{1}{m}\overline{F}_{c_0} \\ \frac{1}{m}\overline{F}_{c_{\perp,0}} \end{bmatrix} + \begin{bmatrix} \frac{1}{m}B_{11} & \frac{1}{m}B_{12} \\ \frac{1}{m}B_{21} & \frac{1}{m}B_{22} \end{bmatrix} \begin{bmatrix} \delta F_c \\ \delta\theta \end{bmatrix} \tag{6.46}$$

wherefore,

$$\mathbf{f}_0(\mathbf{x}_0) = \frac{1}{m}\overline{\mathbf{F}}_0 \qquad \mathbf{g}_0(\mathbf{x}_0) = \frac{1}{m}\mathbf{B}_{\overline{\mathbf{F}}} \tag{6.47}$$

This is to say that, while the terms $\overline{\mathbf{F}}_0$ and $\mathbf{B}_{\overline{\mathbf{F}}}$ are pre-computed and scheduled, this pseudo-control allocation scheme represents more than just a collection of linearized terms, as these terms ultimately represent a parameterization of the nonlinear system dynamics with respect to the system state (i.e. $\mathbf{f}_0(\mathbf{x})$ and $\mathbf{g}_0(\mathbf{x})$) that, in this particular implementation, are then evaluated along the expected flight trajectory of the aircraft (i.e. $\mathbf{x} = \mathbf{x}_0$). This is mathematically analogous to evaluating the system dynamics at the current time step, as would be done for a typical feedback linearization, the difference being that the operating point here is pre-determined instead of based on a current measurement. Hence, we could equivalently apply feedback linearization to the input-affine (nonlinear) system of (6.46), where this system would be found to be full-state linearizable with no internal dynamics, and thus applying the standard feedback linearization control law $\mathbf{u} = \mathbf{E}^{-1}(\mathbf{x})(\mathbf{v} - \mathbf{b}(\mathbf{x}))$ with a suitable state transformation (such as $\mathbf{z} = \mathbf{x}$ in this case) results in a linear system from $\mathbf{v}$ to $\mathbf{y}$ in Brunovsky form, as would be expected.

It is therefore also worth emphasizing that this MIMO pseudo-control allocation thus shares the same benefit in terms of versatility as the SISO version presented in Section 6.6. Specifically, if the system model ever changes, the pseudo-control allocation constants ($\mathbf{B}_{\overline{\mathbf{F}}}$, $\overline{\mathbf{F}}_0$, and $\mathbf{u}_0$) would simply need re-computed once, while the (linear) controller itself would not need re-tuned. Once again, though, this comes at the cost of increased complexity in the initial setup, but this is arguably now more

justified in the MIMO case, where cross-coupling has been accounted for to enable improved performance and tuning flexibility.

Lastly, considering this system (a MIMO system with two inputs and two outputs) is full-state linearizable with no internal dynamics, it begs the question of whether it remains necessary to work in the V_c-$V_{c_\perp}$ frame (i.e. choosing V_c and $V_{c_\perp}$ as the output variables) as it was for the decoupled approaches. Indeed, provided F_c and θ have been chosen as the inputs to ensure controllability over the complete flight domain (due to the general lack of independent thrust actuation in both the vertical and horizontal directions), the use of V_c and $V_{c_\perp}$ is no longer strictly necessary, and any two (distinct) velocities could be selected as the system outputs for this MIMO design, including the original Earth-frame velocities V_x and V_z. For example, the nonlinear root solving problem suggested in (6.36) could be alternatively set up as:

$$\begin{aligned} \dot{V}_x &= g_1(F_c, \theta) \\ \dot{V}_z &= g_2(F_c, \theta) \end{aligned} \tag{6.48}$$

where g_1 and g_2 are functions representing the nonlinear dynamics of V_x and V_z, respectively, which is a reformulation of (6.17). However, it turns out that working in the V_c-$V_{c_\perp}$ frame has a notable benefit when it comes to tuning the system. The fundamental dynamics of these aircraft demonstrate that F_c (i.e. thrust) varies its (primary) control influence between the V_z-direction in hover and the V_x-direction in cruise, while θ does essentially the opposite. However, the achievable bandwidth of these two inputs is inherently different, since F_c is a product of thrust/motor speed demands and $\theta \approx \theta_{ref}$ is a result of the inner-loop attitude control design. Hence, if V_x and V_z are chosen as the output variables to be controlled (fixed Earth-frame directions), this means that the tuning of the (linear) controllers for V_x and V_z will set the bandwidth relative those (fixed) directions, while the inputs being used to regulate these velocities (i.e. F_c and θ) vary their control influence according to

the (rotating) V_c- and $V_{c_\perp}$-directions. This swap in control influence means that the maximum bandwidth of the V_x and V_z loops will be ultimately constrained by the slower of the two inputs. For example, θ will need to perform quickly enough to match the desired bandwidths of both the V_x controller tuning when in hover mode and the V_z controller tuning when in cruise mode (and vice versa for F_c). In practice, this means that adequate overall performance of the outer-loop can only be achieved if the V_x and V_z loops are tuned to broadly equal bandwidths for consistency, but in turn this means that the inherent bandwidth differences between the two inputs cannot be leveraged[13]. This unfortunately negates one of the key benefits of doing a coupled MIMO approach in the first place, which was resolving the inherent bandwidth limitations in the decoupled approaches by accounting for cross-coupling in the dynamics. Hence, if we instead choose V_c and $V_{c_\perp}$ as the output variables to be controlled, these velocities (which by definition also vary their directions between V_x and V_z) can be tuned to different bandwidths and in accordance with the achievable bandwidths of F_c and θ, respectively, therefore making it possible[14] to leverage the inputs' inherent bandwidth differences. Said another way, choosing to work in the V_c-$V_{c_\perp}$ frame allows for different loop tunings to be achieved to maximize available performance, and this will automatically reflect in each flight mode for V_x and V_z by virtue of this chosen rotating coordinate system.

6.8 Conclusion

This chapter tackled the general problem of finding a full-envelope solution for the outer-loop velocity control design of transition VTOL aircraft. To address the fundamental challenges of this endeavor, a novel approach of selecting virtual inputs and outputs defined within a rotating reference frame that is aligned with a VTOL air-

[13]Unless, ironically, a gain schedule or other switching behavior is implemented.

[14]This is ultimately because the inputs F_c and θ have been carefully selected to specifically and primarily affect a single one of these outputs (V_c and $V_{c_\perp}$, respectively).

craft's thrust vector was presented and implemented within a feedback linearization structure. This resulted in feedback-linearized systems that remained fully controllable over the complete flight-envelope without assumptions of independent thrust actuation or specific actuators used, making this control concept generalizable to a wide range of transition VTOL configurations. To implement this concept, three unique control strategies were presented. The first control strategy employed a decoupled design, where feedback linearization was applied in a SISO context, and the system's resulting internal dynamics were regulated by a kinematics-based mapping for the pitch angle commands. The second control strategy formalized this approach by implementing a model-based design for the perpendicular velocity, made possible through the use of a pseudo-control allocation scheme. Lastly, the third control strategy applied the concepts developed from the first two approaches and implemented them within a coupled MIMO setup to enable improved performance and tuning versatility. In the next chapter, the proposed strategies are applied to the tiltwing VTOL configuration of interest in this thesis, where performance will be evaluated in both simulation and flight testing.

Chapter 7

Outer-Loop Design and Completed Flight Controller

7.1 Introduction

Chapter 6 derived a general outer-loop control strategy for transition VTOL aircraft through the novel selection of thrust-aligned virtual control inputs and outputs in a rotating reference frame. This chapter now applies this method to the specific tiltwing aircraft configuration of interest in this thesis, thus completing the full-envelope flight controller design that was initiated in Chapter 5. This full-envelope flight controller is then evaluated in both simulation and flight testing and, additionally, a formal robustness analysis is carried out. Hence, Section 7.2 discusses the outer-loop design itself, now applied to this specific tiltwing aircraft. This is followed by an overview and brief discussion in Section 7.3 of (what will be) the completed full-envelope flight controller. Simulation results are then presented next in Section 7.4 followed by flight testing in Section 7.5. Lastly, the robustness of this full-envelope design is formally investigated in Section 7.6.

7.2 Outer-Loop Design

For the outer-loop design, flight testing was done using Strategy 3, while simulation results in this chapter will additionally look at Strategy 1 for comparison between these methods[1]. Hence, this section will look at the implementation of both of these designs followed by a brief (re-)discussion of the control allocation and a discussion of the lateral velocity (V_y) control for this aircraft.

7.2.1 Decoupled Design (Strategy 1)

Analogous to the inner-loop design, the only external forces (other than gravity) that do not depend explicitly on the input for this tiltwing aircraft are passive aerodynamic forces. These can be similarly expressed using total-aircraft force coefficient quantities[2] as:

$$\mathbf{F}_{A_{passive}} = \mathbf{R}_{bs}(\alpha) \cdot 0.2\bar{q}S \begin{bmatrix} -C_D & C_Y & -C_L \end{bmatrix}^T \tag{7.1}$$

where the negative signs on C_D and C_L are necessary to ensure correct stability-to-body frame conversion and, once again, the multiplier of 0.2 is due to the 20% unwetted split designated for this model. Hence, the general form system dynamics for Strategy 1 (i.e. (6.17) and (6.18)) can be applied to this tiltwing aircraft, giving:

$$\begin{bmatrix} \dot{V}_x \\ \dot{V}_z \end{bmatrix} = \begin{bmatrix} \frac{1}{m} F_{A_{x,p}} \cos(\theta) + \frac{1}{m} F_{A_{z,p}} \sin(\theta) \\ g - \frac{1}{m} F_{A_{x,p}} \sin(\theta) + \frac{1}{m} F_{A_{z,p}} \cos(\theta) \end{bmatrix} + \begin{bmatrix} \frac{1}{m} \cos(\xi_W + \theta) \\ -\frac{1}{m} \sin(\xi_W + \theta) \end{bmatrix} F_c \tag{7.2}$$

$$V_c = V_x \cos(\xi_W + \theta) - V_z \sin(\xi_W + \theta)$$

where $F_{x,p}$ and $F_{z,p}$ have been set to $F_{A_{x,p}}$ and $F_{A_{z,p}}$, respectively, to explicitly designate that the passive forces for this tiltwing aircraft are solely aerodynamic (i.e.

[1] Investigations were carried out for Strategy 2 as well, but the results have been omitted in this thesis for brevity.

[2] In a similar manner to $\mathbf{M}_{A_{passive}}$ from (5.3), the complete/expanded expression for (7.1) sees individual term contributions from the main wing, tail, and fuselage as per the 6-DoF model described in Chapter 3.

the components of $\mathbf{F}_{A_{passive}}$). Additionally, the thrust angle ξ has been replaced with ξ_W since the majority of the thrust production for this aircraft comes from the wing props. Note that for this decoupled design (Strategy 1), $F_{c_\perp}$ has been temporarily ignored for the purpose of deriving the V_c feedback linearization controller (recall from Section 6.5 that the $\dot{V}_c$ equation remains unchanged).

The expressions for the feedback linearization terms $\mathbf{E}(\mathbf{x})$ and $\mathbf{b}(\mathbf{x})$ from (6.20) will remain the same, as will the standard feedback linearization control law from (6.21) and the linear system from (6.22) representing the external dynamics for this input-output linearized system. However, similar to the inner-loop NDI controller, it was chosen to augment this linear system with an integrator to ensure zero steady-state error:

$$\begin{aligned} \dot{\boldsymbol{\mu}} &= \begin{bmatrix} 0 & 0 \\ 1 & 0 \end{bmatrix} \boldsymbol{\mu} + \begin{bmatrix} 1 \\ 0 \end{bmatrix} v \\ y &= \begin{bmatrix} 1 & 0 \end{bmatrix} \boldsymbol{\mu} \end{aligned} \tag{7.3}$$

where the external dynamics state vector is now $\boldsymbol{\mu} = [V_c \quad \int V_c]^T$. LQR was also chosen to control this linear system, with performance indices selected to be:

$$Q = \begin{bmatrix} 4.8 & 0 \\ 0 & 6.72 \end{bmatrix} \quad R = 1 \tag{7.4}$$

which sets the V_c-loop bandwidth to 2.45 $\mathrm{rad\,s^{-1}}$ with a phase margin of 76°. The speed of this loop was chosen based on pilot feedback on responsiveness and controllability of the vertical airspeed during hover flight testing. Note that since this V_c control loop (with virtual input F_c) runs in parallel with the inner-loop, there is no explicit bandwidth constraint from the attitude dynamics.

Lastly, it was necessary to select an appropriate time constant T for the first-order lag applied to θ_{ref}. This was chosen as a balance between smooth pitch de-

Table 7.1: Selected θ_{ref} lag time constant for each flight mode

Flight Mode	θ_{ref} **lag time constant**
Hover	10 s
Cruise	1 s
Forward transition	0.1 s
Backward transition	10 s

mands (sufficient time-scale separation from the F_c-loop as well as the inner-loop) and sufficiently-fast performance, and it was found that different time constants were necessary for different flight modes. Table 7.1 summarizes the time constants selected for this aircraft. Note that these values do not represent the absolute bandwidth of this loop but rather serve to modify the bandwidth of this loop relative the equivalent bandwidth of the $V_{c_\perp}$-dynamics without the θ_{ref} mapping present. To illustrate this conceptually, consider Figure 7.1 below. At a single operating point within the flight domain, the θ_{ref} mapping of (6.30) approximately behaves like a constant gain ($K_{\theta_{ref}}$), while the linearized dynamics from θ_{ref} to $V_{c_\perp}$ are approximately first-order (with time constant T_0). Hence, the overall closed-loop bandwidth of $V_{c_\perp}$ is adjusted based on the presence of the first-order lag cascaded between the θ_{ref} mapping ($K_{\theta_{ref}}$) and the open-loop velocity dynamics ($K_0/(1+sT_0)$). Numerical investigations reveal/confirm that the choice of time constants in Table 7.1 sets the closed-loop $V_{c_\perp}$ bandwidth to around 1 $\mathrm{rad\,s^{-1}}$ nominally[3].

[3]Note that, due to the nature of this design (fixed time constant value for each flight mode), the actual bandwidth of the $V_{c_\perp}$-loop varies widely depending on operating point, ranging from as low as 0.2 $\mathrm{rad\,s^{-1}}$ to as high as 1.2 $\mathrm{rad\,s^{-1}}$ during step responses. Hence, while a nominal bandwidth of 1 $\mathrm{rad\,s^{-1}}$ does not provide the level of time-scale separation with the inner-loop that would be theoretically desired, it was chosen to preserve this tuning since it worked well in simulations and flight testing would be carried out with Strategy 3.

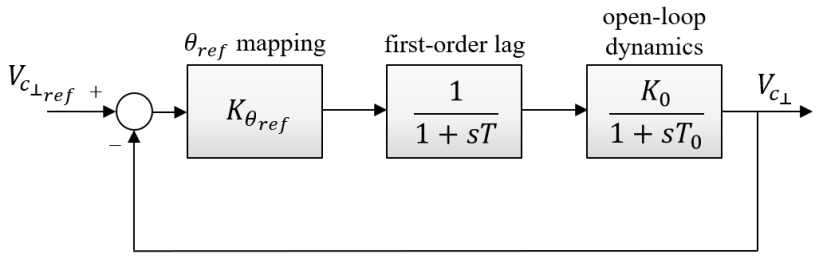

Figure 7.1: Linear approximation of internal dynamics regulation using θ_{ref} mapping (Strategy 1)

7.2.2 Coupled MIMO Design (Strategy 3)

In a similar fashion to Strategy 1, (6.38) can be applied to the dynamics of this tiltwing aircraft, giving the following pseudo-control allocation virtual inputs:

$$\begin{aligned} \overline{F}_c &= -mg\sin(\xi_W + \theta) + F_{A_{x,p}}\cos(\xi_W) - F_{A_{z,p}}\sin(\xi_W) + F_c \\ \overline{F}_{c_\perp} &= -mg\cos(\xi_W + \theta) - F_{A_{x,p}}\sin(\xi_W) - F_{A_{z,p}}\cos(\xi_W) + F_{c_\perp} \end{aligned} \tag{7.5}$$

which results in the decoupled linear system of (6.37) that can be expressed in the form $\dot{\mathbf{x}} = \mathbf{v}$. It was once again chosen to add integral action in consideration of steady-state error, resulting in:

$$\begin{aligned} \dot{\mathbf{z}} &= \begin{bmatrix} \mathbf{0}_{2\times2} & \mathbf{0}_{2\times2} \\ \mathbf{I}_{2\times2} & \mathbf{0}_{2\times2} \end{bmatrix} \mathbf{z} + \begin{bmatrix} \mathbf{I}_{2\times2} \\ \mathbf{0}_{2\times2} \end{bmatrix} \mathbf{v} \\ \mathbf{y} &= \begin{bmatrix} \mathbf{I}_{2\times2} & \mathbf{0}_{2\times2} \end{bmatrix} \mathbf{z} \end{aligned} \tag{7.6}$$

where $\mathbf{z} = \begin{bmatrix} V_c & V_{c_\perp} & \int V_c & \int V_{c_\perp} \end{bmatrix}^T$. Similarly, LQR was chosen to control this system, with performance indices selected to be:

$$Q = \begin{bmatrix} 4.8 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 6.72 & 0 \\ 0 & 0 & 0 & 10 \end{bmatrix} \quad R = \begin{bmatrix} 1 & 0 \\ 0 & 130 \end{bmatrix} \tag{7.7}$$

This design sets the V_c-loop bandwidth to 2.45 rad s^{-1} (identical tuning to Strategy 1), while the $V_{c_\perp}$-loop has a bandwidth of 0.53 rad s^{-1}. Because this is a MIMO design, any coupling between the V_c- and $V_{c_\perp}$-dynamics is now accounted for, making it no longer necessary to explicitly ensure sufficient time-scale separation between these loops. However, as $V_{c_\perp}$ is controlled by sending pitch angle references θ_{ref} to the inner-loop (while assuming $\theta \approx \theta_{ref}$), it does remain necessary to ensure sufficient time-scale separation between the $V_{c_\perp}$-loop and the inner-loop attitude controller. Nominally, this would suggest setting the $V_{c_\perp}$-loop bandwidth to less than 0.1 rad s^{-1}, but this resulted in too sluggish of performance for adequate pilot controllability in flight. Hence, the revised faster bandwidth of 0.53 rad s^{-1} was selected. While not fully time-scale separated with the inner-loop, closed-loop stability was assessed and the $V_{c_\perp}$-loop was able to achieve an adequate phase margin of 66° with this tuning[4].

7.2.3 Control Allocation (revisited)

In Section 5.3.4, the commanded virtual controls vector was chosen to include a virtual vertical force demand F_z (i.e. $\boldsymbol{\tau}_c = [F_z \quad M_x \quad M_y \quad M_z]^T$) for the purposes of testing the inner-loop controller design in hover flight. However, in consideration of the generalized full-envelope outer-loop design, this will now be revised to use the thrust-aligned virtual force F_c for operability in all flight modes:

$$\boldsymbol{\tau}_c = [F_c \quad M_x \quad M_y \quad M_z]^T \tag{7.8}$$

where F_c is constructed from the effector model body-frame forces as:

$$F_c = F_x \cos(\xi_W) - F_z \sin(\xi_W) \tag{7.9}$$

[4]Note that this compromise is ultimately a physical constraint of the particular tiltwing demonstrator model used in this thesis, where it was necessary to reduce the inner-loop bandwidth to avoid a control-structural interaction, as was discussed in Chapter 5.

7.2.3.1 Transition Trajectory Schedule

The weighted pseudo-inverse control allocation method chosen for this design is applicable in all flight modes but requires linearizing the nonlinear effector model $\mathbf{f}_B(\mathbf{x}, \mathbf{u})$ along the expected flight trajectory of the aircraft, as was seen for vertical take-off in Section 5.3.4.2. For transition, this was achieved by generating optimal transition trajectories for the scaled model using the methodology of Chapter 4. These trajectories are shown in Figures 7.2 and 7.3. Using these trajectories, a schedule based on the wing angle ξ_W was implemented in multi-dimensional lookup tables to smoothly interpolate between control effectiveness ($\mathbf{B}$) matrices corresponding to the current operating point $(\mathbf{x}_0, \mathbf{u}_0)$ of the aircraft using a linear point-slope algorithm, just as for the vertical take-off schedule.

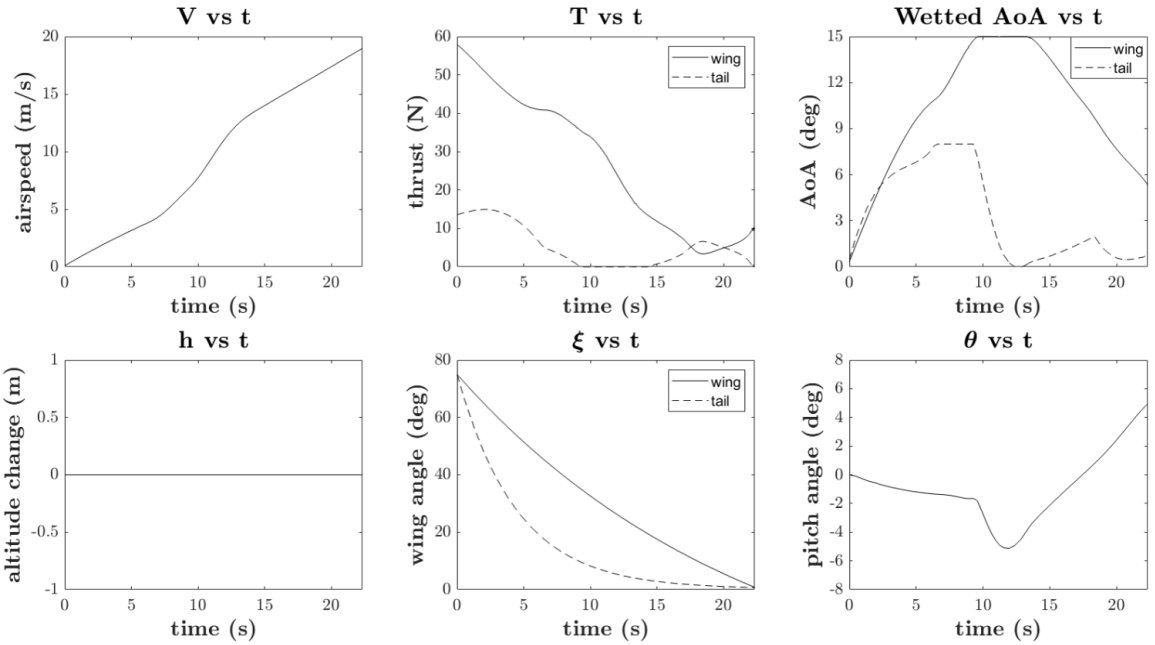

Figure 7.2: Forward transition trajectory solution for the scaled model

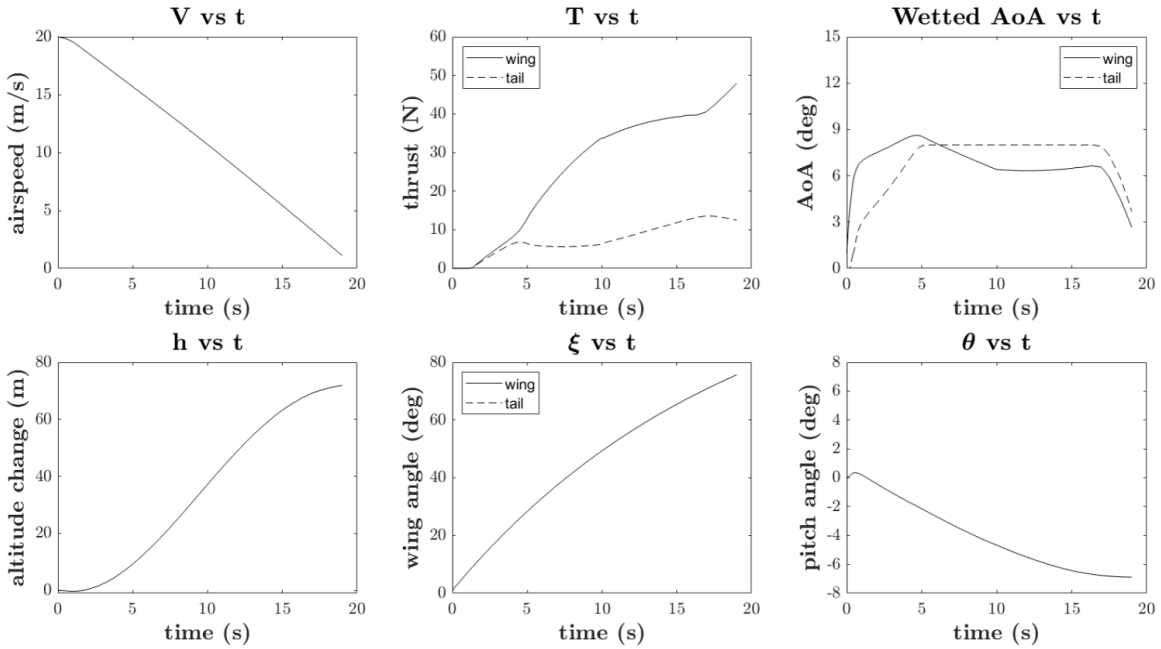

Figure 7.3: Backward transition trajectory solution for the scaled model

7.2.4 Lateral Velocity Controllers

It was chosen to split the lateral velocity (V_y) control design into separate hover and cruise mode controllers due to the fundamental difference in lateral motion operation between these flight modes. For transition, it is simply desired to maintain zero roll angle due to the lateral symmetry of the maneuver, so no explicit V_y control was implemented.

7.2.4.1 Hover Lateral Motion

In hover mode, the control of the lateral velocity V_y is coupled with the roll angle ϕ just as V_x is coupled with pitch. Hence, due to the low airspeeds in hover, a single PI controller was sufficient for this design, that is:

$$\phi_{ref} = K_P^{V_y}(V_{y_{ref}} - V_y) + K_I^{V_y} \int (V_{y_{ref}} - V_y)dt \tag{7.10}$$

For this design, the gains were selected to be $K_P^{V_y} = 0.0004$ and $K_I^{V_y} = 0.0033$, setting the closed-loop bandwidth to 0.14 $\mathrm{rad\,s^{-1}}$ and giving a phase margin of 75°. Additionally, similar to the inner-loop design, it was chosen to apply a reference pre-filter $PF_{V_y}(s) = K_I^{V_y}/(K_P^{V_y}s + K_I^{V_y})$ to cancel the resulting closed-loop zeros.

7.2.4.2 Cruise Coordinated Turn

In cruise mode, the lateral velocity V_y cannot be demanded directly. Instead, lateral motion is achieved through a coordinated turn, where the aircraft's lateral acceleration (in the body frame) and sideslip is maintained at zero. This is accomplished by rolling the aircraft to the necessary bank angle to balance the horizontal component of lift with the aircraft's centripetal acceleration in the Earth frame while maintaining constant altitude.

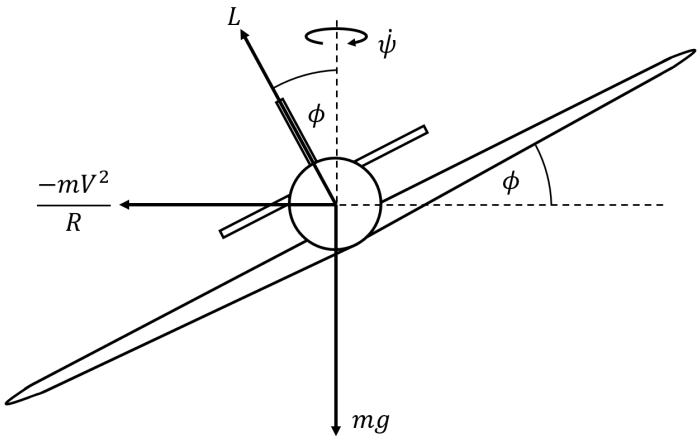

Figure 7.4: Aircraft free-body diagram during coordinated turn

The necessary roll angle can be easily derived from a simple free-body diagram of the aircraft, as shown in Figure 7.4. Noting that the aircraft's yaw rate is equal to its angular velocity about the center of turn (with turn radius R) and assuming a small pitch angle, the roll angle required to achieve a coordinated turn while maintaining

constant altitude for a given yaw rate request $\dot{\psi}_{ref}$ and longitudinal airspeed V_x is:

$$\phi_{ref} = \arctan\left(\frac{V_x \dot{\psi}_{ref}}{g}\right) \tag{7.11}$$

7.3 Completed Flight Controller Architecture

A schematic is given in Figure 7.5 of the (now) completed full-envelope flight controller. This schematic matches the layout of the general controller architecture presented in Section 5.2 (Figure 5.1) but now includes the specific signals and reference inputs for the completed design.

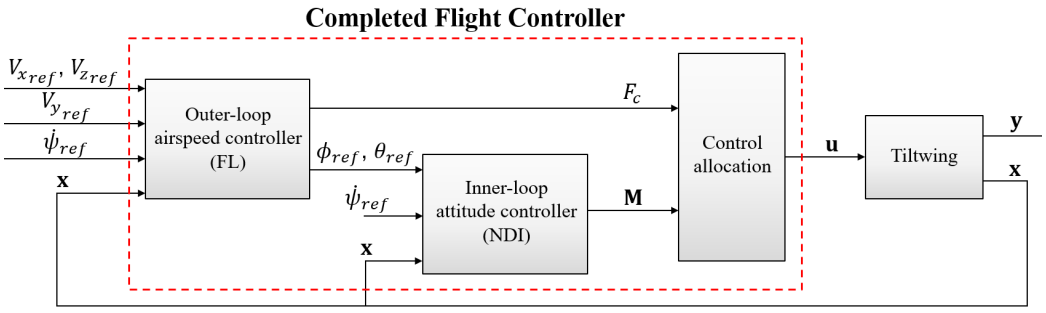

Figure 7.5: Completed full-envelope flight controller layout

The outer-loop subsystem of course contains the full-envelope V_x and V_z velocity control (either Strategy 1 or Strategy 3) along with the mode-specific controllers for regulating the lateral motion of the aircraft. Hence, the reference inputs from the pilot include $V_{x_{ref}}$, $V_{y_{ref}}$, and $V_{z_{ref}}$ along with the yaw rate request $\dot{\psi}_{ref}$. This yaw rate request also directly enters the inner-loop subsystem along with ϕ_{ref} and θ_{ref} from the outer-loop. This inner-loop subsystem of course contains the Euler angle tracking controllers for the roll and pitch angles combined with the NDI attitude rate controller for all three rotational axes. Lastly, the outer-loop outputs the virtual force F_c in parallel with the virtual moments $\mathbf{M}$ from the inner-loop to the control allocation subsystem, which converts these virtual controls to the tiltwing's real inputs $\mathbf{u}$.

7.4 Simulations

7.4.1 Strategy 1

Figures 7.6–7.9 show step responses in V_x and V_z in hover and cruise flight. 10–90% rise times (t_r) and percent overshoots (OS) are summarized in Table 7.2. Of particular note is the response speed of hover V_z and cruise V_x step responses in contrast to hover V_x and cruise V_z step responses. Specifically, the former sees rise times that are around 4 to 5 times quicker. This is because these step responses correspond to the F_c control loop, while the hover V_x and cruise V_z step responses correspond to the θ_{ref} control of the internal dynamics, which has been set to a lower bandwidth to provide an amount of time-scale separation between these loops. Additionally, there is a general trend for cruise responses to be marginally quicker than hover responses. This can be explained by the fact that hover mode operates with the wing angles set to 77° (equilibrium hover point) as opposed to 90°, so hover control of V_x and V_z sees a mixture of control from F_c and θ_{ref} (recall the definitions of V_c and $V_{c_\perp}$ from (6.18) and (6.24)). This contrasts cruise, where V_x is approximately only controlled by F_c and V_z is roughly only controlled by θ_{ref} since $\xi_W = 0$ (and θ is small). Therefore, there is less reliance on the combined effort of (and therefore, additionally, less interaction between) the F_c and θ_{ref} loops (different time-scales) to generate these velocity step responses in the cruise case. Importantly, despite the inherent coupling between V_x and V_z, corresponding velocity response variations are kept to a minimum, staying within $\pm 0.17\ \mathrm{m\,s^{-1}}$ on average.

Figures 7.10 and 7.11 show velocity, pitch angle, and wing angle plots for forward and backward transition, respectively. The optimal transition trajectory references of Section 7.2.3.1 were supplied for V_x, V_z, ξ_W, and ξ_T. In both forward and backward transition, V_x and V_z are tracked well, though there is marginally more variation in V_z. To quantify the performance, Table 7.3 shows the root-mean-square (RMS)

velocity tracking error (e_{rms}) along with its peak absolute value ($|e|_{max}$), where it can be seen that the average error is kept below 0.4 $\mathrm{m\,s^{-1}}$ with tracking in the forward direction being marginally tighter on average. To understand the behavior of this controller better, both V_c and $V_{c_\perp}$ are plotted alongside their respective references[5]. It can be seen that the F_c-loop tracks $V_{c_{ref}}$ almost perfectly—it is specifically in regards to regulating the internal dynamics $V_{c_\perp}$ that more variation can be observed, which is due to the lower bandwidth of the θ_{ref}-loop. As a slight aside, note that θ is demanded to be a negative value at the end of backward transition simply because the pre-computed reference trajectory for V_x ends at 1 $\mathrm{m\,s^{-1}}$, requiring the aircraft to pitch downward to track this forward velocity demand.

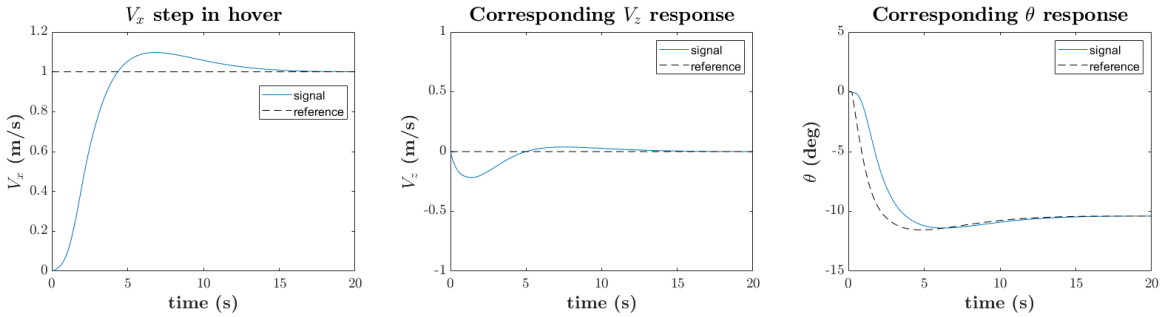

Figure 7.6: V_x step response in hover (Strategy 1)

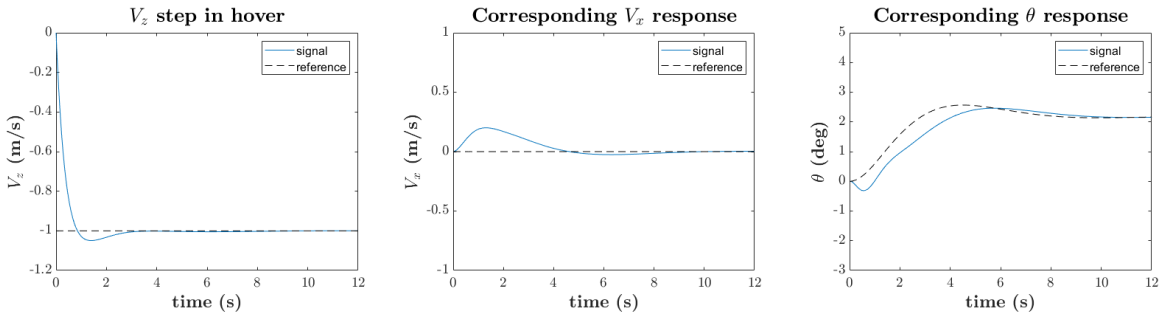

Figure 7.7: V_z step response in hover (Strategy 1)

[5]Note that, in this instance, the $V_{c_\perp}$ "reference" is more a representation of what value the internal dynamics should be converging to (in order to achieve zero steady-state error for V_x and V_z), provided sufficient regulation through θ_{ref}; recall that there is no controller that is tracking this explicit reference for Strategy 1, just the mapping for θ_{ref} combined with a first-order lag.

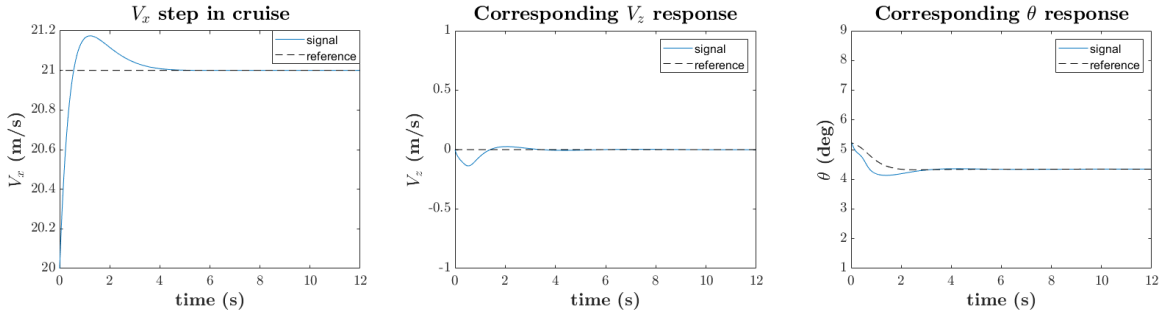

Figure 7.8: V_x step response in cruise (Strategy 1)

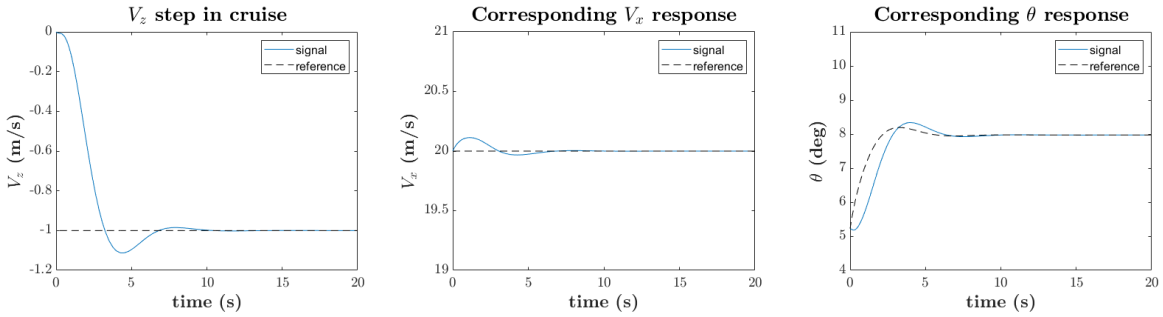

Figure 7.9: V_z step response in cruise (Strategy 1)

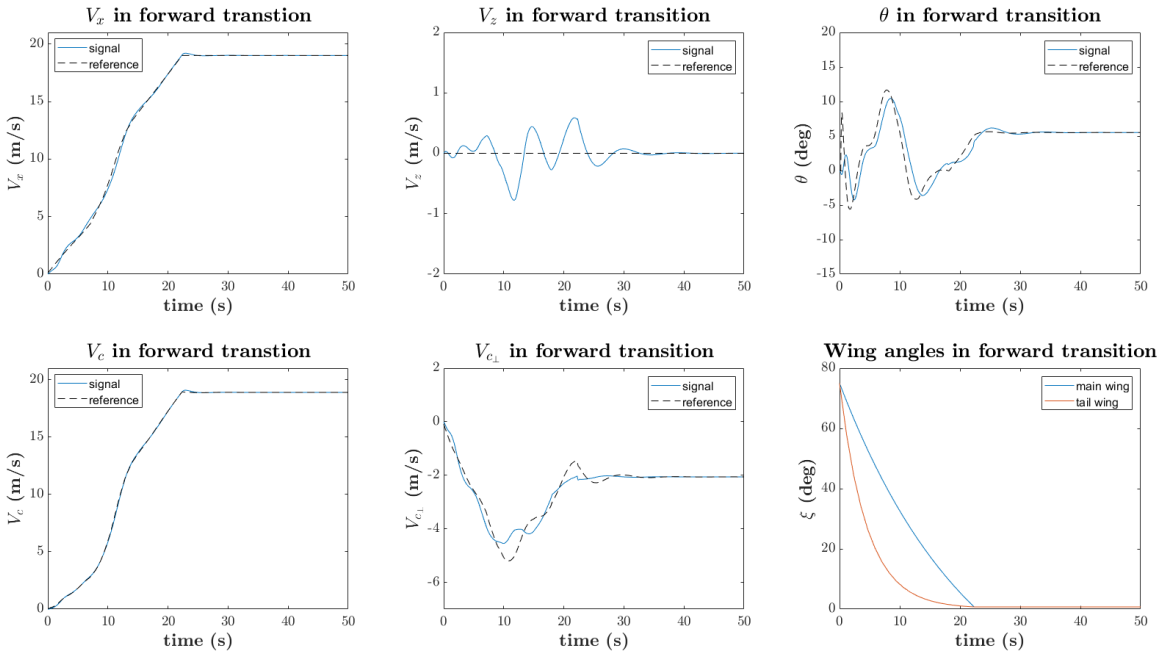

Figure 7.10: Controller performance in forward transition (Strategy 1)

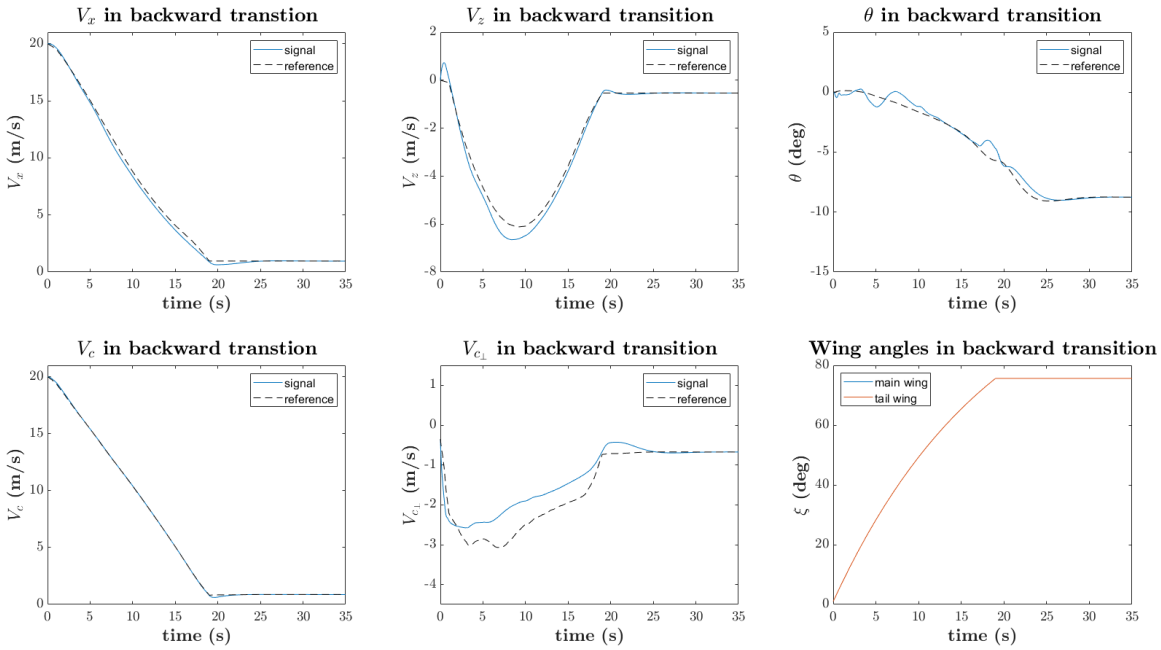

Figure 7.11: Controller performance in backward transition (Strategy 1)

Table 7.2: Step response performance characteristics (Strategy 1)

Flight Mode	**Step Variable**	t_r	OS
Hover	V_x	2.57 s	9.7%
Hover	V_z	0.56 s	4.9%
Cruise	V_x	0.42 s	17.5%
Cruise	V_z	1.92 s	11.2%

Table 7.3: Velocity tracking error during transition (Strategy 1)

Direction	**Velocity**	e_{rms}	$\lvert e \rvert_{max}$
Forward	V_x	0.26 $\mathrm{m\,s^{-1}}$	0.65 $\mathrm{m\,s^{-1}}$
Forward	V_z	0.32 $\mathrm{m\,s^{-1}}$	0.78 $\mathrm{m\,s^{-1}}$
Backward	V_x	0.37 $\mathrm{m\,s^{-1}}$	0.56 $\mathrm{m\,s^{-1}}$
Backward	V_z	0.38 $\mathrm{m\,s^{-1}}$	0.77 $\mathrm{m\,s^{-1}}$

Table 7.4: Step response performance characteristics (Strategy 3)

Flight Mode	**Step Variable**	t_r	OS
Hover	V_x	16.1 s	0%
Hover	V_z	0.51 s	7.4%
Cruise	V_x	0.47 s	20.2%
Cruise	V_z	44.1 s	0%

Table 7.5: Velocity tracking error during transition (Strategy 3)

Direction	**Velocity**	e_{rms}	$\lvert e \rvert_{max}$
Forward	V_x	$0.19\ \mathrm{m\,s^{-1}}$	$0.46\ \mathrm{m\,s^{-1}}$
Forward	V_z	$0.16\ \mathrm{m\,s^{-1}}$	$0.39\ \mathrm{m\,s^{-1}}$
Backward	V_x	$0.11\ \mathrm{m\,s^{-1}}$	$0.25\ \mathrm{m\,s^{-1}}$
Backward	V_z	$0.16\ \mathrm{m\,s^{-1}}$	$0.69\ \mathrm{m\,s^{-1}}$

7.4.2 Strategy 3

Figures 7.12–7.15 show step responses in V_x and V_z in hover and cruise flight, now for Strategy 3. Table 7.4 similarly shows 10–90% rise times (t_r) and percent overshoots (OS). Very similar comments can be made here as for Strategy 1, such as the marginally faster cruise V_x step rise time compared to the hover V_z step rise time. However, it can be observed that the hover V_x and cruise V_z step performance is noticeably slower compared to Strategy 1. While the $V_{c_\perp}$-loop was theoretically tuned to a closed-loop bandwidth of $0.53\ \mathrm{rad\,s^{-1}}$, it was revealed through numerical investigations that the actual closed-loop bandwidth for hover V_x and cruise V_z was $0.14\ \mathrm{rad\,s^{-1}}$ and $0.05\ \mathrm{rad\,s^{-1}}$, respectively. This is believed to be a result of imperfect cancellations of the outer-loop feedback linearization dynamics due to model mismatch, which is occurring in the linear control allocation when straying away from the (pre-computed) nominal operating points during these particular step responses (i.e. primarily in the $V_{c_\perp}$-direction rather than the V_c-direction[6]). This is further evidenced by the larger bandwidth discrepancy in cruise mode vs hover mode, where aerodynamic effects (and thus variations in control effectiveness) are significantly greater. While modifications could have been made, such as adjusting the outer-loop tuning or adding additional pre-computed operating points to the control allocation schedules, it was chosen to preserve the design so as to match the setup that was used during flight testing. Otherwise, performance is satisfactory, with corresponding

[6]Recall that the hover mode control allocation is scheduled with respect to F_c (V_c-direction). Additionally, while it has not been explicitly discussed, cruise mode control allocation is scheduled with respect to V_x (also V_c-direction when in this flight mode).

velocity response variations similarly kept to a minimum, staying within $\pm 0.2\ \mathrm{m\,s^{-1}}$. The exception to this is for the cruise V_x step, where the corresponding V_z response peaks at just under $0.9\ \mathrm{m\,s^{-1}}$; however, this is simply a result of the effective $0.05\ \mathrm{rad\,s^{-1}}$ $V_{c_\perp}$-loop bandwidth at this operating point.

While certain aspects of the hover and cruise step performance for Strategy 3 do not match the performance expectations set by Strategy 1, its transition performance greatly exceeds that of Strategy 1. Figures 7.16 and 7.17 show velocity, pitch angle, and wing angle plots for forward and backward transition, respectively, now for Strategy 3. The same optimal transition trajectory references of Section 7.2.3.1 were supplied just as for Strategy 1. It is immediately noticeable visually that the tracking is tighter for V_x and V_z in both forward and backward transition compared to Strategy 1, which can be more rigorously confirmed by looking at the corresponding error metrics in Table 7.5. Specifically, the RMS velocity tracking error (e_{rms}) is approximately half that of Strategy 1 on average across all cases. This is believed to be due to the inherent coupling that is accounted for in this MIMO design; in a maneuver like transition, where all longitudinal states of the aircraft are changing simultaneously, any undesired cross-coupling between the F_c (V_c) and θ_{ref} mapping ($V_{c_\perp}$) loops for the decoupled design of Strategy 1 fundamentally impacts controller performance. Another consideration is the fact that, while Strategy 1 uses different first-order lag time constants depending on flight mode, it only utilizes a single time constant value throughout the entirety of the transition maneuver (for a given direction). This therefore means that Strategy 1 (in its current implementation) does not account for changing loop bandwidths of the $V_{c_\perp}$-dynamics throughout transition, which will thus also affect performance. In contrast, the incorporation of the $V_{c_\perp}$ velocity in the inversion-based design of Strategy 3 serves to maintain a constant loop bandwidth at all operating points (inversion errors aside), making performance more consistent. This latter point also explains another interesting observation: that

the velocity tracking in the backward direction is now marginally tighter than in the forward direction (opposite Strategy 1); more generally, though, the tracking errors for Strategy 3 are simply more consistent with each other overall, with the small differences in the forward and backward directions likely being due to insignificant numerical differences (unlike Strategy 1's time constant selection). Lastly, similar conclusions to Strategy 1 regarding the controller behavior (such as the relative performance of V_c and $V_{c_\perp}$) can also be drawn for the present results involving Strategy 3.

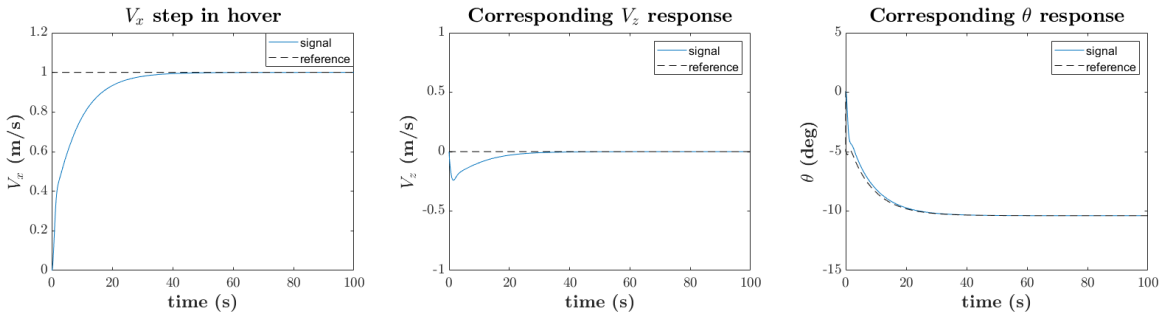

Figure 7.12: V_x step response in hover (Strategy 3)

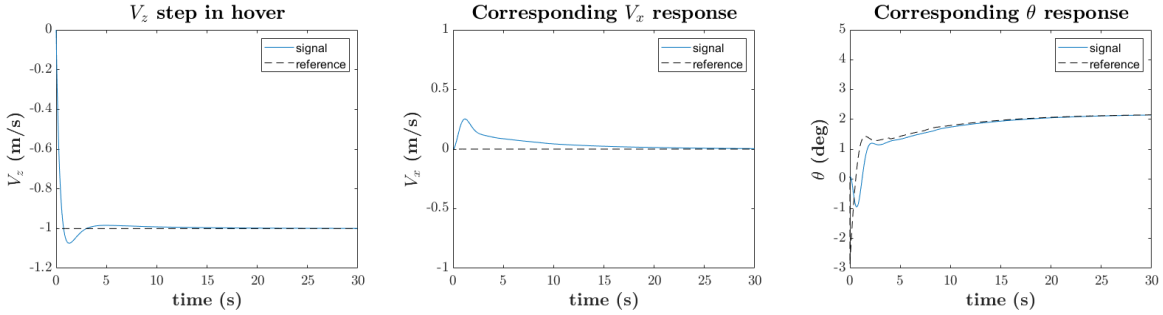

Figure 7.13: V_z step response in hover (Strategy 3)

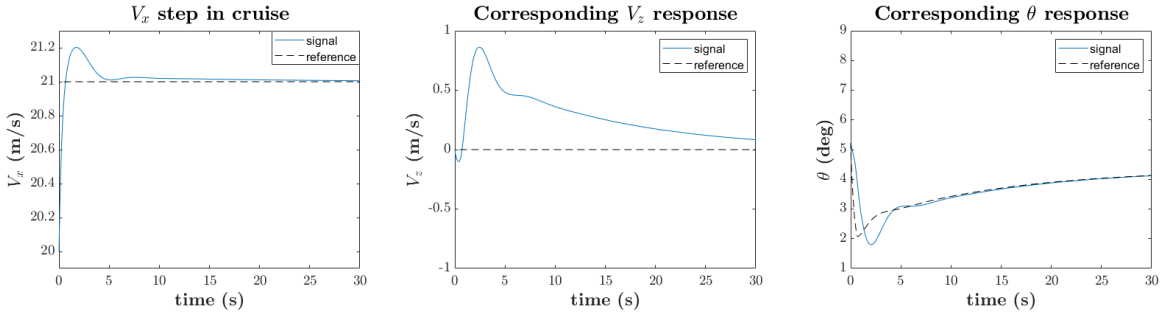

Figure 7.14: V_x step response in cruise (Strategy 3)

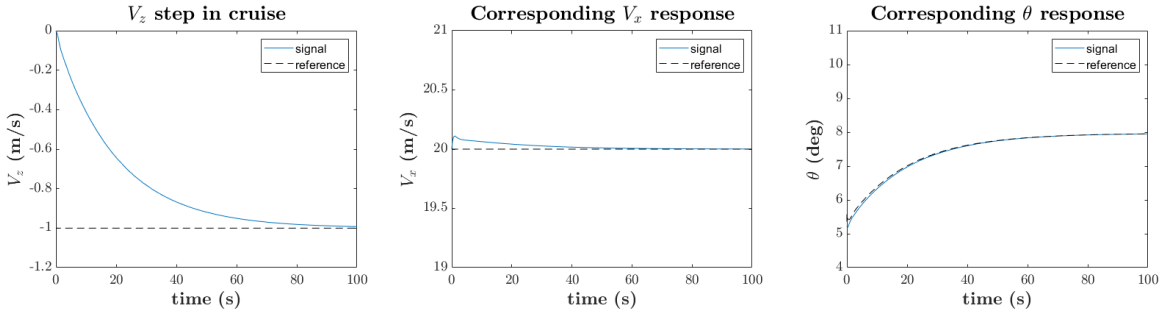

Figure 7.15: V_z step response in cruise (Strategy 3)

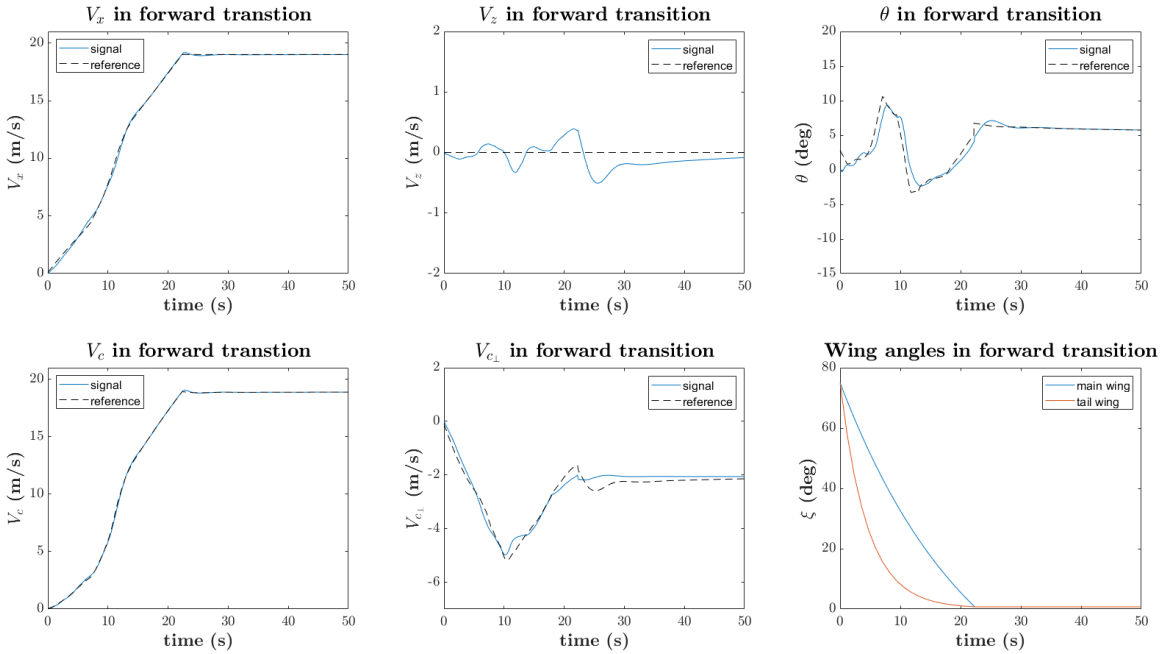

Figure 7.16: Controller performance in forward transition (Strategy 3)

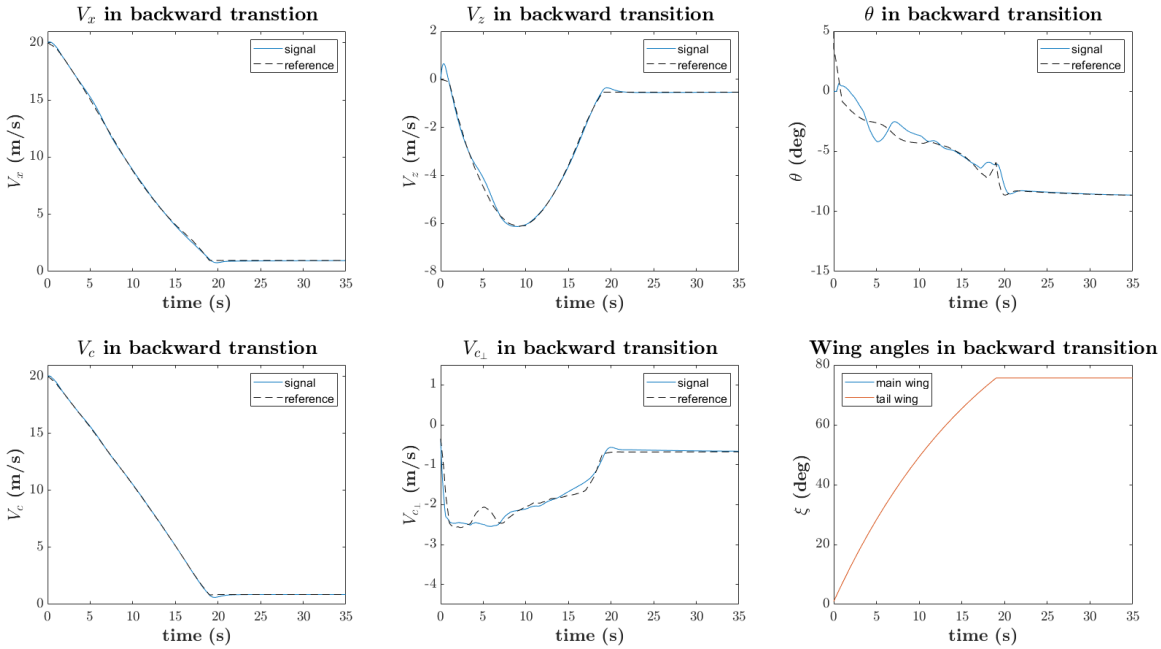

Figure 7.17: Controller performance in backward transition (Strategy 3)

7.4.3 Lateral Motion

Figures 7.18 and 7.19 present step response simulation results for the lateral motion controllers in hover and cruise, respectively. In hover, a step of 1 $\mathrm{m\,s^{-1}}$ is demanded in V_y and an inner-loop roll angle reference is therefore generated for ϕ. For cruise, a step input in yaw rate of 3 deg/s (0.0524 $\mathrm{rad\,s^{-1}}$) is demanded, which corresponds to a standard two-minute turn; a roll angle reference is then correspondingly demanded based on (7.11).

For the hover V_y step response, performance is comparable to Strategy 3's hover V_x step performance, with $t_r = 16.3$ s and $OS = 0\%$. This is perhaps expected, as the effective closed-loop bandwidths of these loops are identical at this operating point (0.14 $\mathrm{rad\,s^{-1}}$). For cruise mode, the $\dot{\psi}$ and ϕ response characteristics are consistent with the inner-loop performance characteristics discussed in Chapter 5. Both figures additionally show a plot of the Earth-frame ground position (X_e vs Y_e) to give an idea of the aircraft's motion during these maneuvers. For the hover V_y step, it can be seen that there is some variation in X_e due to regulating V_x and not position directly,

but this variation is kept minimal. For the cruise turn, it can be seen that the aircraft follows a smooth, circular trajectory corresponding to a turn radius of 382 m (which is consistent with the $V_x = 20\ \mathrm{m\,s^{-1}}$ cruise speed).

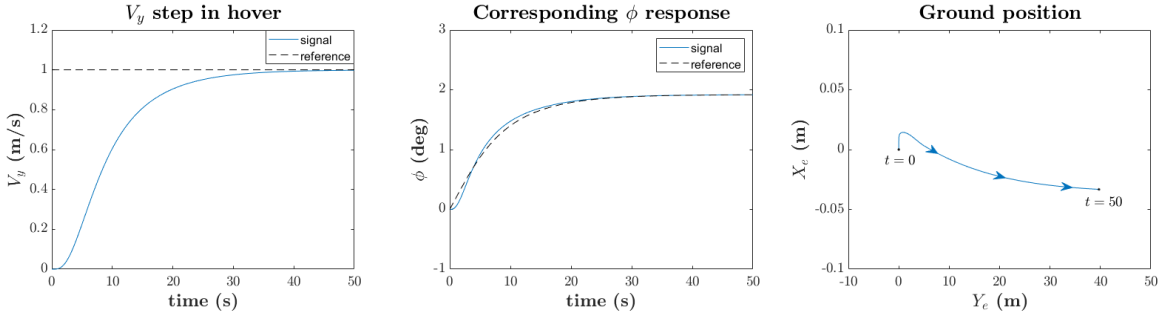

Figure 7.18: V_y step response in hover

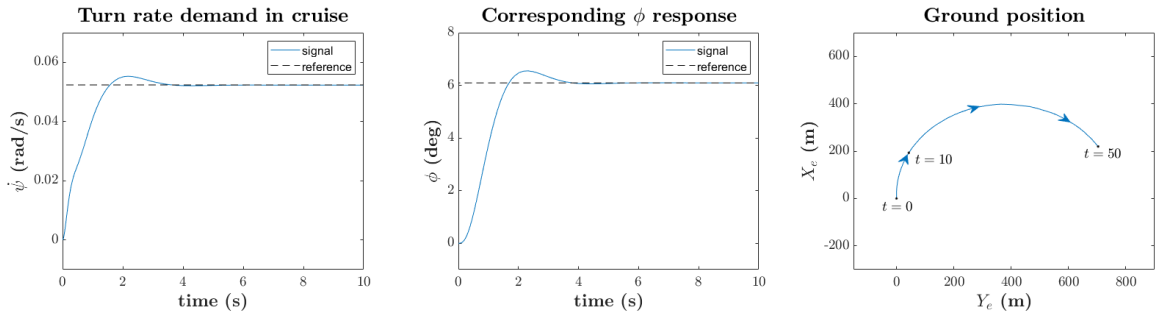

Figure 7.19: Turn rate step response in cruise ($V_x = 20\ \mathrm{m\,s^{-1}}$)

7.5 Flight Testing

7.5.1 Experimental Setup

Just as for the inner-loop, flight testing of the outer-loop controller was conducted to validate the design and evaluate its performance on the real aircraft. However, this testing was primarily conducted indoors (tethered) due to time constraints in the project. A single outdoor (untethered) flight test of the outer-loop design was conducted, but unfortunately a hardware issue during flight resulted in a crash shortly after take-off. This resulted in very little untethered closed-loop flight data being acquired, and there was not enough time at this point in the project to rebuild the

aircraft. Hence, flight testing results in this section will focus solely on the acquired indoor (tethered) flight data.

For the indoor tethered testing, the experimental setup was essentially identical to the inner-loop indoor testing setup, as described in Section 5.5.1. The only notable difference was now the need for velocity measurements. The velocity states are estimated within the Pixhawk 4's on-board EKF and primarily rely upon GPS data, which is unavailable indoors. Hence, the velocity states were estimated using other available sensor data (accelerometer, barometer, etc.), which resulted in drifty measurements that were not always accurate to the actual motion of the aircraft. However, despite this, indoor testing was still sufficient for evaluating the general performance of the outer-loop controller, including general tracking ability (even with respect to these drifty measurements) and confirming stability.

Note that for this outer-loop testing, the hover V_y controller was not implemented in order to focus on validating/evaluating the longitudinal velocity control design (the primary focus for this part of the research). Hence, the V_y velocity was simply kept open-loop, and a direct pilot demand for the roll angle control (ϕ_{ref}) was implemented just as for the inner-loop flight testing.

On the practical implementation side of things, there are a couple details worth mentioning. Firstly, to ensure a smooth start-up, the aircraft would begin each flight with only the inner-loop attitude controller in operation (hence, an identical setup to the inner-loop flight testing, where open-loop demands for vertical force/thrust as well as pitch angle were demanded from the pilot). Then, once reaching a steady hover flight condition, a switch on the remote control transmitter would be flipped to initiate the outer-loop velocity control. Hence, it was necessary to implement a bumpless transfer for the outer-loop's integrators to initialize their values based on the F_c and θ_{ref} demands at the moment of switching to closed-loop velocity control. Secondly, one of the testing procedures for the outer-loop velocity control (which will

be presented in the next section) involved rotating the aircraft's wings mid-flight to evaluate the controller's performance during rotating thrust conditions (similar to what would be experienced at the beginning of a transition flight). This was done to test the fundamental operation of this controller design, which revolves around the concept of working in a rotating reference frame, and was an opportunity to evaluate performance with a non-constant thrust direction/thrust angle. To achieve this, a switch could be flipped by the pilot (once in closed-loop velocity control) that would initiate a rotation of the wings to a secondary, pre-specified wing angle. For this testing, and in consideration of the low thrust margin of this scaled demonstrator aircraft, testing began with the wing angles set at 90°, and the secondary wing angle was set to 77°.

7.5.2 Tethered Flight

Figures 7.20 and 7.21 show example maneuvering in V_x and V_z during indoor tethered flight. Note that, due to the limited space indoors, it was not possible to hold a constant non-zero velocity demand for an extended period of time, making the precise evaluation of performance metrics challenging. However, it was still possible to evaluate general performance as well as approximate step input rise times from the acquired data. The data for the V_x maneuvering of Figure 7.20 shows good tracking of pilot demands. In fact, the performance seems more responsive than simulation predictions, and indeed the approximate 10–90% rise time for these pilot inputs was around 1.7 s, indicating a significant disparity between the simulation model and the real aircraft. The precise reason for this is inconclusive. For evaluating V_z maneuvering in hover flight, the presence of the tether rope made it challenging to get "clean" V_z step data. However, the data presented in Figure 7.21 is an example portion of a flight that demonstrates sufficient tracking of the V_z velocity[7]. Determining a precise

[7]Note that the data in Figure 7.21 was acquired while separately testing a faster $V_{c_\perp}$-loop bandwidth of 0.78 rad s^{-1}. However, that primarily affects the V_x response, and these results are therefore

rise time for the V_z response was particularly tricky for the aforementioned reasons, but a representative example can be found around $t = 195$ s. Here, the rise time is on the order of 0.5 s, which in this case is consistent with simulation predictions. Note that the seeming undershoot after this step rise (around $t = 196$ s) is due to the aircraft tugging on the rope after the descent (recall V_z is oriented positive-down in the NED frame). The effect of this can be further seen in the corresponding θ response plot, where significant pitching oscillations are initiated from this point in the data onward due to the rope being attached to the aircraft at two points along the aircraft's longitudinal axis (i.e. making it possible to create a non-zero pitching moment when pulling the rope taught). Small pitching oscillations can be seen in Figure 7.20 during the V_x maneuvering as well, but in this instance they are less significant due to very little V_z maneuvering and therefore less rope interaction.

A core concept of this controller design is its versatility to operate in any flight mode and any wing angle condition. Hence, to put this controller to the test (and get as close as possible to mimicking the start of a transition flight while tethered), a flight was carried out in which the wing angles were rotated mid-flight during hover. This is demonstrated in Figure 7.22. For this test, the wing angles began at 90° and were then commanded to rotate forward to 77° (and back). The times during which the wings were rotating (forward and backward) are indicated by the vertical dotted black lines in each plot. It can be seen that tracking performance was maintained and, importantly, the controller remained stable during these wing rotations[8]. Two things to note are the V_z velocity performance and the θ performance. Similar to the V_z maneuvering of Figure 7.21 just discussed, it was challenging to completely avoid interference from the tether rope during this maneuver. In this case, the presence of the rope resulted in somewhat favorable V_z performance, as the aircraft was partially

still representative of V_z tracking performance. Unfortunately, in all the indoor data acquired with the finalized $V_{c_\perp}$ tuning of 0.53 $\mathrm{rad\,s^{-1}}$, very little V_z maneuvering data was captured.

[8]Note that the oscillations present in the response are due to interactions with the tether rope.

held by the rope at certain moments during the maneuver (unintentionally so). Hence, in a similar manner to before, this rope interaction can be further spotted through the pendulum oscillations appearing in the corresponding θ response. Nevertheless, the aircraft remained stable and was able to return back to steady-state hovering flight. In an attempt to compare to transition simulation performance, the RMS error was additionally evaluated during the time ranges in which the wings were rotating. These metrics are summarized in Table 7.6. It can be seen that the RMS errors for V_x are quite similar to those found in transition simulations (Strategy 3); the backward V_x tracking error is marginally less favorable in this instance, but pitching oscillations (which were greater during the backward rotation) naturally affect this value due to pitch's coupling with the V_x velocity in hover flight (which can also be spotted in the V_x backward rotation response). These metrics also quantitatively reveal the favorable influence that the tether rope had on the V_z response, where the corresponding RMS errors for V_z are significantly smaller than V_x as well as those found in simulation.

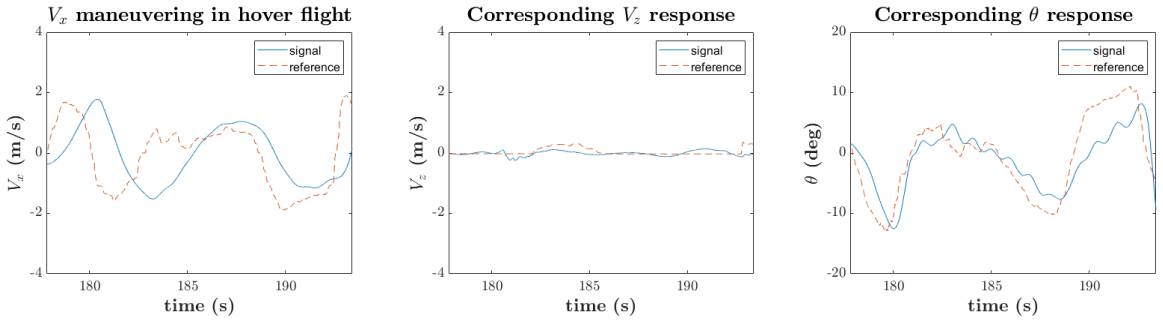

Figure 7.20: V_x maneuvering in hover flight

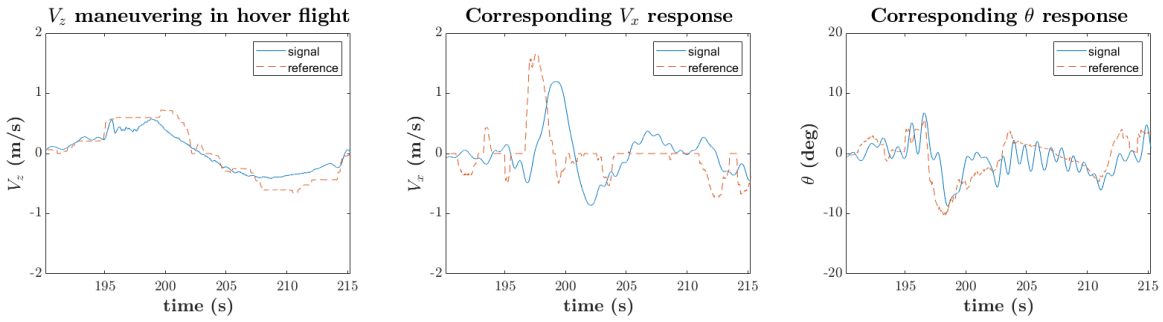

Figure 7.21: V_z maneuvering in hover flight

Table 7.6: Velocity tracking error during wing rotation tests in hover flight

Direction	**Velocity**	e_{rms}	$\|e\|_{max}$
Forward	V_x	0.23 $\mathrm{m\,s^{-1}}$	0.43 $\mathrm{m\,s^{-1}}$
Forward	V_z	0.04 $\mathrm{m\,s^{-1}}$	0.06 $\mathrm{m\,s^{-1}}$
Backward	V_x	0.25 $\mathrm{m\,s^{-1}}$	0.43 $\mathrm{m\,s^{-1}}$
Backward	V_z	0.03 $\mathrm{m\,s^{-1}}$	0.05 $\mathrm{m\,s^{-1}}$

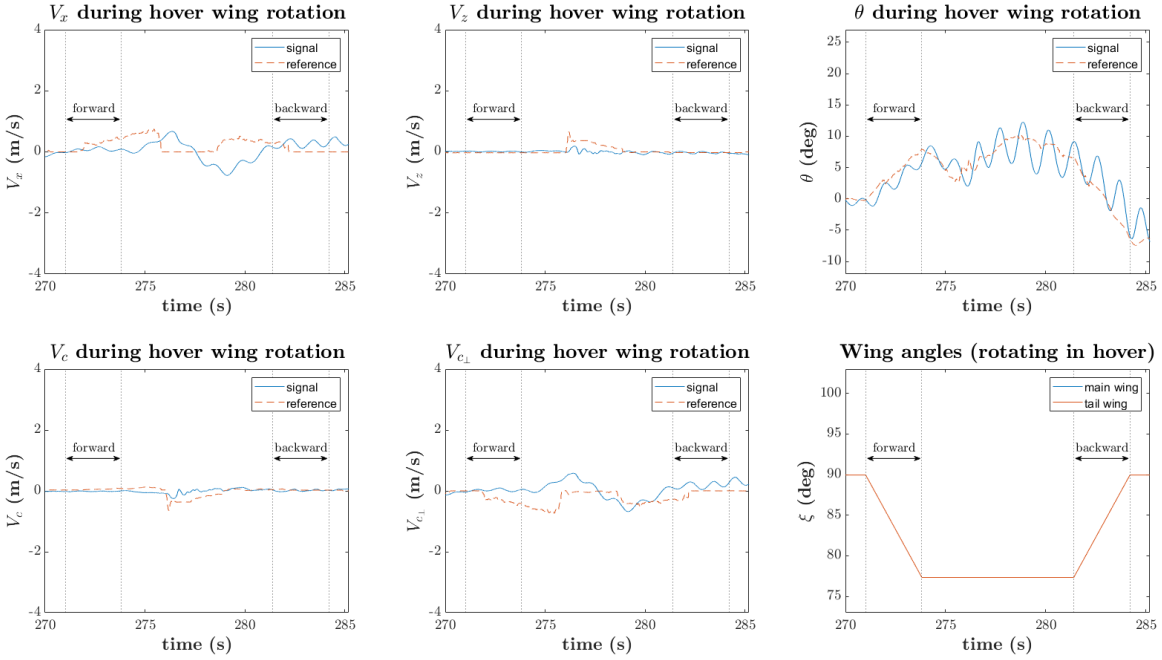

Figure 7.22: Controller performance during wing rotation tests in hover flight

7.6 Robustness Analysis

It is well known that feedback linearization is quite sensitive to disturbances and modeling uncertainties due to its inherent inversion-based design [39, 40, 42]. Ideally, the feedback linearization inversion is perfect, resulting in a linear system in Brunovsky form (a chain of integrators) [39]. However, due to uncertainty, imperfect cancellations may result in unstable closed-loop dynamics. Therefore, it is important that the linear controller acting on the inverted system (i.e. the control law for the new input $\mathbf{v}$; refer to Figure 2.2) be sufficiently robust.

In this thesis, the full-envelope flight controller design for this tiltwing aircraft employed LQR-based linear controllers outside of the feedback linearization inversions. This ensured adequate gain and phase margins for system robustness, noting the inherent robustness guarantees provided by LQR [130]. Additionally, the range of flight testing results presented in this thesis demonstrates adequate robustness under real conditions during hover flight. However, we will now formally consider the robustness of this design in the presence of (aerodynamic) modeling uncertainties to quantify/prove the uncertainty robustness bounds for this aircraft + flight controller combination.

7.6.1 Analytical Uncertainty Description for a Tiltwing

To ensure stability of the aircraft at all points within the flight envelope, it is imperative that the inner-loop attitude control be sufficiently robust in the presence of modeling uncertainties. Consider therefore the simplified longitudinal tiltwing model given in Figure 7.23, which contains two thrust inputs T_W and T_T and an elevator input δ_e. In this work, we choose to focus on the pitch rate dynamics, as that is most significant in consideration of longitudinal stability during the transition maneuver. Furthermore, we make the following assumptions in this analytical derivation:

- Fully-wetted aerodynamics
- T_W and T_T each come from a single thrust source/prop (in consideration of prop-induced velocity calculations)
- Both thrusts act at the same thrust angle ξ, and thrust/aero moment arms only exist in the horizontal direction (l_w and l_h)
- Linear lift-curve slopes are assumed with respect to both AoA and δ_e. Additionally, drag contribution on pitching moments is assumed negligible

- All aero coefficients are scaled relative the main wing reference area S, and any downwash effects on the tail wing will be ignored

Considering perturbations around a nominal trajectory, the thrusts can be written as:

$$\begin{aligned} T_W &= T_{W_0} + \delta_T \\ T_T &= T_{T_0} - \delta_T \end{aligned} \tag{7.12}$$

where T_{W_0} and T_{T_0} represent nominal values along the trajectory and δ_T represents a differential thrust input. Hence, the input vector will be defined as $\mathbf{u} = [\delta_T \quad \delta_e]^T$, and the system will have a single state $\mathbf{x} = q$. The subscript "0" will therefore be used to represent the nominal values of any states or parameters along the trajectory of the aircraft that appear in the pitch rate dynamics in addition to q (e.g. velocities and AoAs).

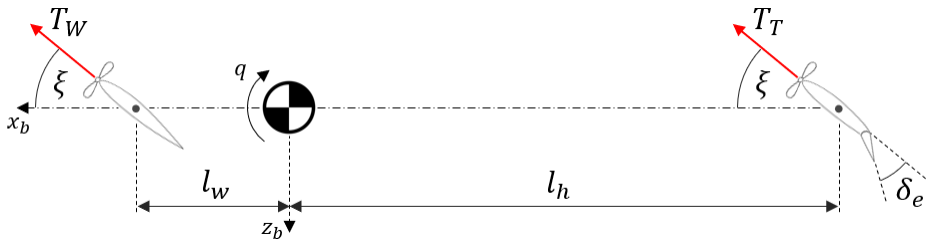

Figure 7.23: Simplified longitudinal tiltwing model for pitch rate dynamics

From the general equations of motion for an aircraft (refer to (2.6)), the pitch rate dynamics are:

$$\dot{q} = pr(I_{zz} - I_{xx})/I_{yy} + (r^2 - p^2)I_{xz}/I_{yy} + M_{y_{net}}/I_{yy} \tag{7.13}$$

For this simplified longitudinal model, p and r are assumed to be zero, and the net external pitching moment $M_{y_{net}}$ consists of both thrust (M_{T_y}) and aerodynamic (M_{A_y}) moments:

$$M_{y_{net}} = M_{T_y} + M_{A_y} \tag{7.14}$$

The thrust moments M_{T_y} can be derived to be:

$$M_{T_y} = (T_W l_w - T_T l_h)\sin(\xi) \tag{7.15}$$

Using (7.12), this can be expanded and reduced to:

$$M_{T_y} = \delta_T (l_w + l_h)\sin(\xi) \tag{7.16}$$

where it has been assumed that the nominal thrusts T_{W_0} and T_{T_0} collectively produce zero net pitching moment along the aircraft's trajectory.

The aerodynamic pitching moment M_{A_y} can be described as a function of α, q, and δ_e, and hence expanded as:

$$M_{A_y} = M^{\alpha}_{A_y} + M^{q}_{A_y} + M^{\delta_e}_{A_y} \tag{7.17}$$

where $M^{i}_{A_y}$ indicates the aerodynamic pitching moment contribution from variable i. Considering perturbations along a nominal trajectory, $\alpha = \alpha_0$ renders $M^{\alpha}_{A_y}$ a constant offset (at each operating point) in the linear input-output relationship from $\mathbf{u}$ to q. Hence, we choose to ignore the $M^{\alpha}_{A_y}$ term, but we will consider the implications of α stability in Section 7.6.2.

The $M^{q}_{A_y}$ term, which arises due to a non-zero pitch rate causing changes in effective AoA (as discussed in Section 3.3.3.1), can be thought of as a "pitch damping" term and is often lumped into a single quantity by defining the non-dimensional stability derivative C_{m_q}, for example [32]:

$$C_{m_q} = -2C_{L_{\alpha,w}}\left(\frac{l_w}{\bar{c}}\right)^2 - 2C_{L_{\alpha,h}}\left(\frac{l_h}{\bar{c}}\right)^2 \tag{7.18}$$

However, to capture the full effect for a tiltwing, we instead consider the local

lift production on each surface individually to account for possibly different wetted aero/thrust conditions on the wing and tail airfoils. The lift produced on surface i due to pitch rate q is given by:

$$L_i^q = \bar{q}_i S C_{L_{\alpha,i}} \Delta\alpha_{q,i} \tag{7.19}$$

where $\Delta\alpha_{q,i} = \arctan(ql_i/u_i)$. Therefore, the total pitching moment generated due to these lift terms is:

$$M_{A_y}^q = L_w^q \cos(\alpha_{wet,w}) l_w - L_h^q \cos(\alpha_{wet,h}) l_h \tag{7.20}$$

where $\cos(\alpha_{wet,i})$ accounts for the stability-to-body frame transformation. However, to greatly simplify the analysis, we make the assumption that $\Delta\alpha_{q,i} \approx ql_i/u_i$[9]. Therefore, given $\bar{q}_i = (1/2)\rho V_{wet,i}^2$ and $u_i/V_{wet,i} = \cos(\alpha_{wet,i})$, the above reduces to:

$$M_{A_y}^q = -\frac{1}{2}\rho S \left(V_{wet,w} C_{L_{\alpha,w}} l_w^2 + V_{wet,h} C_{L_{\alpha,h}} l_h^2 \right) q \tag{7.21}$$

Recalling that V_{wet} depends on the square root of thrust, (7.21) thus reveals that $M_{A_y}^q \propto q\sqrt{\delta_T}$.

Lastly, the $M_{A_y}^{\delta_e}$ term can be derived fairly easily using a standard $C_{L_{\delta_e}}$ stability derivative and remembering to account for the stability-to-body frame transformation:

$$M_{A_y}^{\delta_e} = -\bar{q}_h S C_{L_{\delta_e}} \delta_e \cos(\alpha_{wet,h}) l_h \tag{7.22}$$

With all the terms in $M_{y_{net}}$ fully defined (M_{T_y}, $M_{A_y}^q$, and $M_{A_y}^{\delta_e}$), we now linearize these dynamics along the expected flight trajectory of the aircraft, choosing the nom-

[9]While not the most general assumption for a tiltwing aircraft, this is arguably appropriate in this instance due to the fully-wetted aero assumption in this simplified model and will be sufficient for exemplifying the uncertain q-loop dynamics for robustness analysis purposes.

inal operating point ($\mathbf{x}_0$,$\mathbf{u}_0$) along the trajectory to be:

$$(\mathbf{x}_0, \mathbf{u}_0) = (q_0, \begin{bmatrix} \delta_{T_0} \\ \delta_{e_0} \end{bmatrix}) = (0, \begin{bmatrix} 0 \\ 0 \end{bmatrix}) \tag{7.23}$$

and the remaining trajectory-dependent quantities in $M_{y_{net}}$ (e.g. velocities and wing angles) will be treated as constant parameters in the context of these q-loop dynamics with values equal to their nominal value at each operating point. For convenience, we define the following dimensional stability derivatives:

$$M_q = \frac{1}{I_{yy}} \frac{\partial M_{y_{net}}}{\partial q}\bigg|_{(\mathbf{x}_0,\mathbf{u}_0)} \qquad M_{\delta_T} = \frac{1}{I_{yy}} \frac{\partial M_{y_{net}}}{\partial \delta_T}\bigg|_{(\mathbf{x}_0,\mathbf{u}_0)} \qquad M_{\delta_e} = \frac{1}{I_{yy}} \frac{\partial M_{y_{net}}}{\partial \delta_e}\bigg|_{(\mathbf{x}_0,\mathbf{u}_0)} \tag{7.24}$$

noting that, by definition, these represent quantities evaluated at the current operating point ($\mathbf{x}_0$,$\mathbf{u}_0$).

Because M_{T_y} and $M_{A_y}^{\delta_e}$ do not depend on the pitch rate, differentiating $M_{y_{net}}$ with respect to q and evaluating at ($\mathbf{x}_0$,$\mathbf{u}_0$) gives:

$$M_q = -\frac{\rho S}{2I_{yy}} \left(V_{wet,w_0} C_{L_{\alpha,w}} l_w^2 + V_{wet,h_0} C_{L_{\alpha,h}} l_h^2 \right) \tag{7.25}$$

In contrast, $M_{A_y}^q$ and $M_{A_y}^{\delta_e}$ do depend on δ_T (due to the wetted aerodynamics) but remain linear in q and δ_e, respectively. Therefore, because q_0 and δ_{e_0} are both zero, evaluating $\partial M_{y_{net}}/\partial \delta_T$ at ($\mathbf{x}_0$,$\mathbf{u}_0$) leaves only the contribution from M_{T_y}, hence:

$$M_{\delta_T} = \frac{1}{I_{yy}} (l_w + l_h) \sin(\xi_0) \tag{7.26}$$

Lastly, neither $M_{A_y}^q$ nor M_{T_y} depend on δ_e, therefore:

$$M_{\delta_e} = -\frac{1}{I_{yy}} \bar{q}_{h_0} S C_{L_{\delta_e}} \cos(\alpha_{wet,h_0}) l_h \tag{7.27}$$

Finally, the q-loop dynamics at each operating point along the trajectory can be written as:

$$\dot{q} = M_q q + M_{\delta_T}\delta_T + M_{\delta_e}\delta_e \tag{7.28}$$

which can be expressed more compactly as:

$$\dot{q} = M_q q + \mathbf{B}_q \mathbf{u} \tag{7.29}$$

where $\mathbf{B}_q = [M_{\delta_T} \quad M_{\delta_e}]$.

7.6.1.1 Error Modeling

In this analysis, we assume that the aerodynamic coefficients are uncertain, thereby giving $M_q = [\underline{M}_q, \overline{M}_q]$ and $M_{\delta_e} = [\underline{M}_{\delta_e}, \overline{M}_{\delta_e}]$, where $\underline{M}$ and $\overline{M}$ denote the min and max values, respectively, of the uncertain terms M_q and M_{δ_e}. Applying a simple feedback linearizing control law to (7.29) that sets the desired rate $v_q = \dot{q}$ gives:

$$\mathbf{u} = \hat{\mathbf{B}}_q^{+}(v_q - \hat{M}_q q) \tag{7.30}$$

where the symbol $\hat{}$ indicates nominal values from the controller, the superscript $+$ denotes the Moore-Penrose pseudo-inverse, and $\hat{\mathbf{B}}_q = [M_{\delta_T} \quad \hat{M}_{\delta_e}]$. Feeding this control law back into (7.29) and re-arranging gives the following uncertain/imperfect dynamics for q post-inversion:

$$\dot{q} = (M_q - B_\delta \hat{M}_q)q + B_\delta v_q \tag{7.31}$$

where,

$$B_\delta = \frac{M_{\delta_T}^2 + M_{\delta_e}\hat{M}_{\delta_e}}{M_{\delta_T}^2 + \hat{M}_{\delta_e}^2} \tag{7.32}$$

Therefore, this gives the following uncertain transfer function from v_q to q:

$$\frac{q(s)}{v_q(s)} = \frac{B_\delta}{s - (M_q - B_\delta \hat{M}_q)} \tag{7.33}$$

where it can be seen that when B_δ is nominally 1 and $\hat{M}_q = M_q$, this linear system reduces to $1/s$ as expected. This transfer function thus reveals that the effect of aerodynamic modeling uncertainty on the feedback-linearized system results in a gain uncertainty (B_δ) and a pole uncertainty ($M_q - B_\delta \hat{M}_q$). This result is consistent with similar findings from other authors [148, 149, 150]. It is interesting to observe that the gain uncertainty arises solely from uncertainty present at the system input (in this case δ_e), while the pole uncertainty is a product of both input and state uncertainty.

7.6.1.2 Robust Stability Condition

In an effort to quantify the amount of (aerodynamic) uncertainty allowable before the closed-loop system goes unstable, we can make use of the small-gain theorem result from Section 2.7.2 (refer to (2.54)) to derive a robust stability condition for this SISO system. To do this, we will first re-express (7.33) as:

$$G(s) = \frac{1 + a_1}{s + a_2} \tag{7.34}$$

where $a_1 = B_\delta - 1$ and $a_2 = B_\delta \hat{M}_q - M_q$ are uncertain parameters that quantify the (worst-case) magnitude of the modeling uncertainty. Expressing the nominal system as $G_0(s) = 1/s$, the uncertain transfer function $G(s)$ can be modeled using a (lumped) multiplicative uncertainty as:

$$G(s) = G_0(s)\left(1 + W_\Delta(s)\Delta(s)\right) \tag{7.35}$$

where,

$$W_\Delta(s) = \frac{a_1 s - a_2}{s + a_2} \tag{7.36}$$

and $\Delta(s)$ is a stable uncertain transfer function that satisfies $\|\Delta\|_\infty \leq 1$. It was shown in Section 2.7.2 that the robust stability condition for a system with multiplicative uncertainty $W_\Delta(s)$ (for unstructured Δ) is:

$$\|W_\Delta T_0\|_\infty < 1 \tag{7.37}$$

where $T_0(s)$ is the nominal closed-loop transfer function. Representing the LQR controller for the q-loop as $C(s) = K_P^q + K_I^q/s$ (i.e. effectively equivalent to a "PI" controller), this robust stability condition can thus be expanded to be:

$$\left\| \frac{a_1 K_P^q s^2 + (a_1 K_I^q - a_2 K_P^q)s - a_2 K_I^q}{s^3 + (K_P^q + a_2)s^2 + (K_I^q + a_2 K_P^q)s + a_2 K_I^q} \right\|_\infty < 1 \tag{7.38}$$

Unfortunately, finding bounds on a_1 and a_2 that satisfy this condition (i.e. the worst-case uncertainty) is analytically challenging, especially considering that these parameters depend on the system's operating point along the expected flight trajectory. Furthermore, $\Delta(s)$ is unstructured (as currently defined), which renders this condition unnecessarily conservative for this system since the uncertain parameters a_1 and a_2 are known to be real-valued and could instead be represented using a block-diagonal uncertainty $\Delta = \text{diag}\{\Delta_1, \Delta_2\}$ [86]. Hence, this motivates the use of numerical methods via the structured singular value (SSV).

7.6.2 Numerical Robustness Proof

While the analytical uncertainty description presented in Section 7.6.1 provides insight into how the presence of aerodynamic uncertainty manifests in the resulting q-loop transfer function, the fundamental complexity of these dynamics (specifically

the dimensional stability derivatives M_q and M_{δ_e} and their dependence on system operating point) makes it challenging to analytically quantify the maximum worst-case aerodynamic modeling uncertainty that the system can handle before going unstable. Hence, a better and more efficient approach is to demonstrate this numerically, and to do so, the complete nonlinear model of the tiltwing aircraft will be used.

In this work, it was chosen to take a stochastic approach via the Monte Carlo method, where each aerodynamic coefficient C_i was perturbed (in a multiplicative fashion) by some δC_i selected from a Normal distribution with zero mean and chosen standard deviation σ:

$$\delta C_i \sim \mathcal{N}(0, \sigma^2) \tag{7.39}$$

By forming the following perturbation vector:

$$\delta C_{vec} = [\delta C_D \quad \delta C_Y \quad \delta C_L \quad \delta C_l \quad \delta C_m \quad \delta C_n]^T \tag{7.40}$$

each δC_i was repeatedly sampled n_s times to produce a large uncertainty set of possible δC_{vec}. Next, linearization of the complete (6-DoF) nonlinear tiltwing model from input v_q to output q was carried out for each given δC_{vec} at a range of time steps along the aircraft's specified transition trajectory (i.e. forward or backward)[10,11]. This resulted in a set $\mathbf{G}_\Delta$ of first-order transfer functions matching the form given in (7.34) (thus confirming the analytical predictions of the previous section). Hence, using this set $\mathbf{G}_\Delta$, the expected min/max values of the uncertain (real) parameters a_1 and a_2 appearing in (7.34) were extracted/determined and used to define the uncertain closed-loop transfer function $T(s)$. Lastly, the SSV was employed to compute the robust stability margin (RSM), a normalized scaling factor that represents the

[10]Note that this approach therefore only indicates the system's robustness along the specific trajectories checked and does not consider robustness along other possible trajectories.

[11]Note also that, strictly speaking, local linearizations along a trajectory leads to a linear time-varying system, so stability with this approach cannot be guaranteed without assumptions of slow parameter variations [151, 152, 153, 154]. Therefore, validation of the results will be carried out via nonlinear simulations using the complete 6-DoF model.

allowable range of the uncertain parameters (a_1 and a_2) such that the closed-loop system remains stable. Hence, robust stability is proven when $RSM \geq 1$.

While this approach demonstrates robust stability of q for given AoAs along the trajectory (i.e. $\alpha = \alpha_0$), the stability of α itself has not been considered, which has become an internal state post-inversion of q [149]. To address this, the transfer functions in $\mathbf{G}_\Delta$ could be augmented to include the body velocity states u and w in addition to q, thus explicitly capturing the α-dynamics, but it was found that numerical convergence was difficult with so many degrees of freedom of uncertainty. Hence, it was chosen to assume α stable in the context of this μ-analysis, and to address α stability, results are then validated using the complete 6-DoF model.

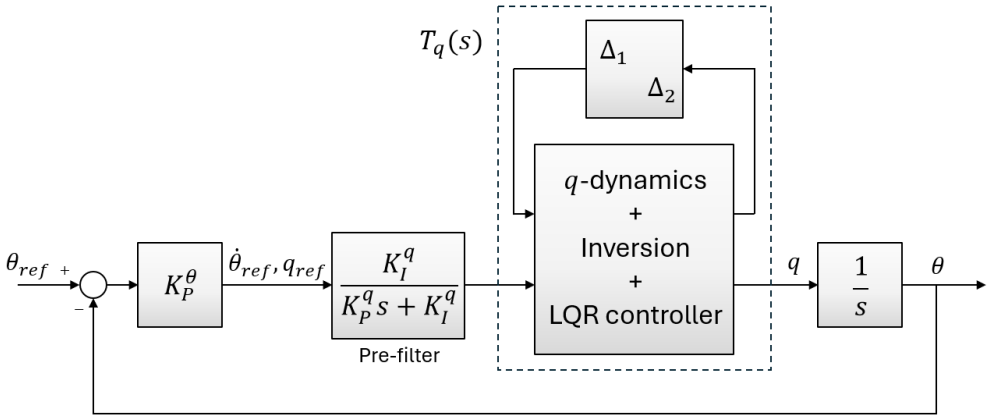

Figure 7.24: Outer-loop θ-dynamics in the presence of uncertain q-loop dynamics

For completeness of the inner-loop analysis, we additionally analyze robust stability of θ. This is accomplished by simply confirming $RSM \geq 1$ for $T_\theta(s) = \theta(s)/\theta_{ref}(s)$ in the presence of the q-loop's uncertain parameters a_1 and a_2. Denoting the q-loop uncertain closed-loop transfer function $T(s)$ as $T_q(s)$ for clarity, $T_\theta(s)$ is found by considering the block diagram in Figure 7.24, which represents $T_q(s)$ using the $N\Delta$-structure (refer to Figure 2.6b) with a diagonal uncertainty $\Delta = \text{diag}\{\Delta_1, \Delta_2\}$, giving:

$$T_\theta(s) = \frac{K_P^\theta P F_q(s) T_q(s)}{s + K_P^\theta P F_q(s) T_q(s)} \tag{7.41}$$

where $PF_q(s)$ is the q-loop reference pre-filter and, in terms of a_1 and a_2:

$$T_q(s) = \frac{(K_P^q s + K_I^q)(1 + a_1)}{s^2 + (K_P^q(1 + a_1) + a_2)\, s + K_I^q(1 + a_1)} \tag{7.42}$$

7.6.2.1 Results

Robust stability for the q-loop dynamics was proven by iteratively selecting uncertainty standard deviations σ and finding the largest bound for which $RSM \geq 1$. This was carried out for forward and backward transition separately to allow for possibly different RSM values depending on the trajectory (which is indeed what was found). For each uncertainty bound (specified by its 3σ value), each δC_i was sampled $n_s = 1000$ times to produce a corresponding δC_{vec} uncertainty set.

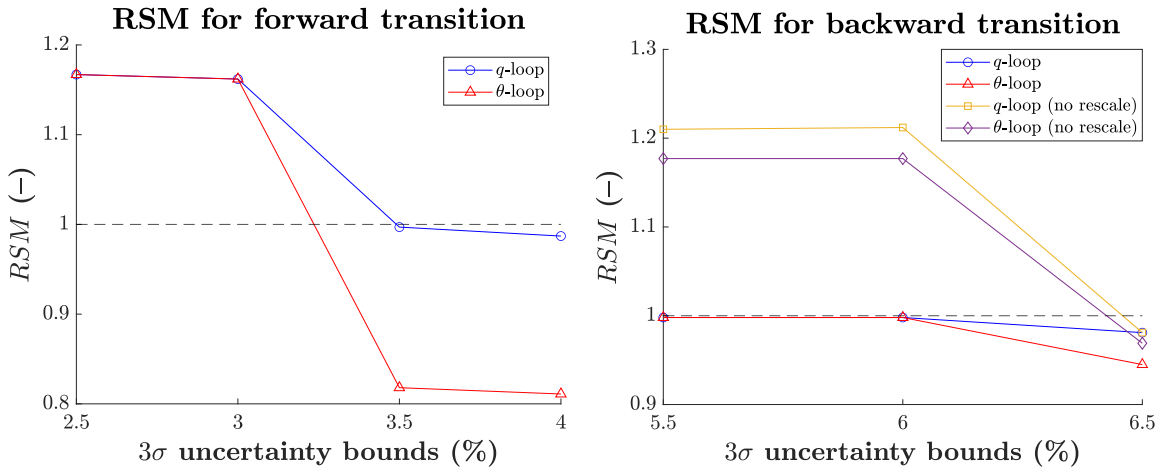

(a) Forward transition, $n_s = 10^3$ for each 3σ (b) Backward transition, $n_s = 10^3$ for each 3σ

Figure 7.25: RSM vs 3σ uncertainty bounds for tiltwing in transition

Figure 7.25 summarizes the results. For forward transition, it can be seen that q-loop robustness is guaranteed up to about $3\sigma = 3.5\%$, but the corresponding θ-loop RSM at this point drops below 1. Hence, the maximum uncertainty bound allowable for forward transition is instead $3\sigma = 3\%$. In contrast, backward transition was found to be robust up to double this value at $3\sigma = 6\%$. A possible explanation for this is that the aircraft is slowing down to re-enter thrust-borne (i.e. hover) flight

and is therefore less sensitive to variations in the aerodynamic coefficients compared to forward transition. Note that Figure 7.25b contains two sets of RSM values. Interestingly, it was found that δC_{vec} uncertainty in backward transition would map to strictly positive values for a_2 (up until $3\sigma \approx 10\%$), meaning that a_2's nominal value of zero would lie outside the computed range. Therefore, to ensure that the nominal case of $G_0(s) = 1/s$ was included, the bounds on a_2 were manually extended to include the origin. This would of course make the RSM more conservative in these cases, but it was felt that this was necessary for accuracy/completeness (and it nevertheless remained possible to demonstrate robustness). Hence, the second set of RSM values ("no rescale") represents the robustness margins computed for the original (strictly positive) set of a_2 values, plotted here for comparison. In either case, it can be seen that the maximum allowable uncertainty bound is found to be the same point ($3\sigma = 6\%$).

7.6.2.2 Validation

The fundamental basis of this proof revolves around the linearization (and thus truncation) of the system's dynamics to obtain an uncertain first-order transfer function $G(s)$. While the use of the complete (6-DoF) nonlinear tiltwing model in this numerical approach has ensured that all perturbation effects on the q-loop (and, by extension, the θ-loop) have been captured, the model reduction to a SISO system remains an approximation that ignores full 6-DoF effects (i.e. coupling between states). Additionally, it was assumed that the α-dynamics are stable within the context of this μ-analysis. Hence, we now validate the results found in this proof with the original 6-DoF model.

To do so, full 6-DoF simulations were conducted for both forward and backward transition[12] using the same δC_{vec} uncertainty sets from Section 7.6.2.1 corresponding

[12]Note that these simulations were run using the Strategy 3 outer-loop controller with the forward and backward transition trajectories from Section 7.2.3.1 supplied as reference inputs.

to the maximum 3σ uncertainty bounds found in the proof (3% for forward and 6% for backward). Hence, for each uncertainty set, $n_s = 1000$ simulations were run and the resulting state trajectories (V_x, V_z, and θ) plotted to see the general variation in the system's output response. These results are shown in Figures 7.26 and 7.27. It can be seen that not only do all state trajectories converge (i.e. remain stable), but the general spread of the trajectories remains relatively narrow, with the primary variation being seen in the θ responses due to different trimming conditions required for each (uncertain) set of aerodynamic coefficients. Hence, these simulations empirically demonstrate stability/robustness in the presence of this magnitude of δC_{vec} perturbations with the complete 6-DoF model, thus validating the robustness bounds predicted in the numerical proof.

Lastly, further investigations with these simulations revealed that the theoretical robustness bounds for this tiltwing aircraft + flight controller combination may be higher than formally predicted in the numerical proof. Figures 7.28 and 7.29 show forward and backward transition results with δC_{vec} uncertainty sets computed from $3\sigma = 9\%$ and $3\sigma = 10\%$, respectively, which were the highest uncertainties that could be applied before state trajectories began to go unstable. While a purely empirical result, it is interesting to observe that there is now a smaller disparity between the maximum allowable uncertainties for forward and backward transition compared to what was found in the numerical proof. Furthermore, the higher achievable uncertainty in these results may indicate that the numerical proof remains somewhat conservative.

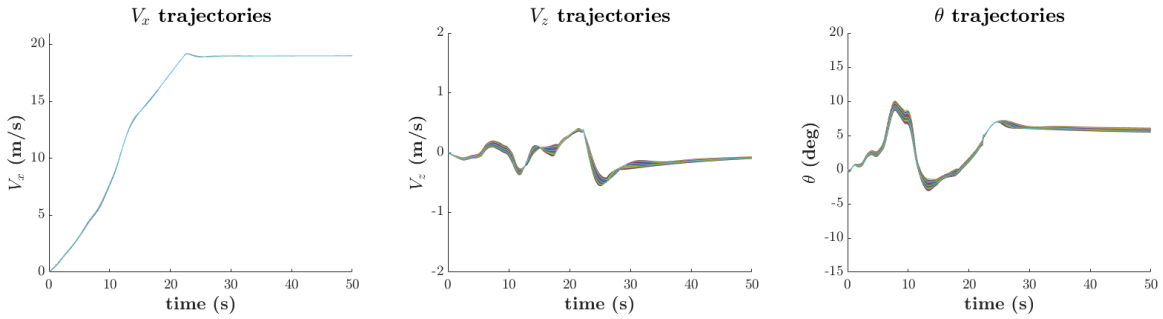

Figure 7.26: Forward transition simulations, $n_s = 10^3$ and $3\sigma = 3\%$.

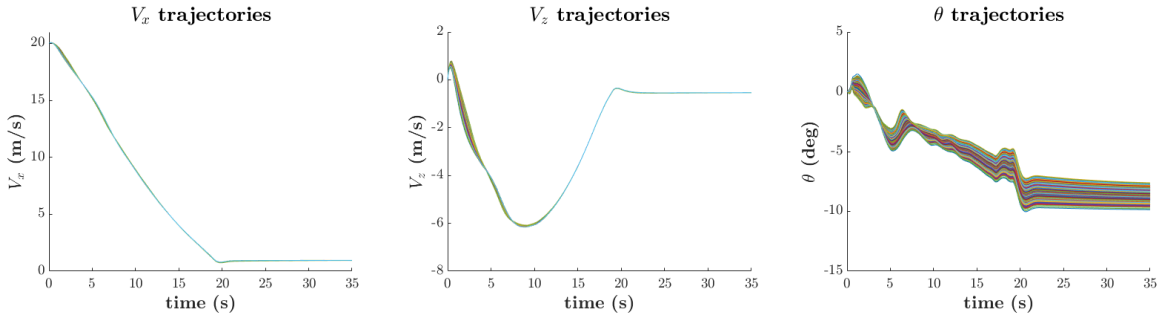

Figure 7.27: Backward transition simulations, $n_s = 10^3$ and $3\sigma = 6\%$.

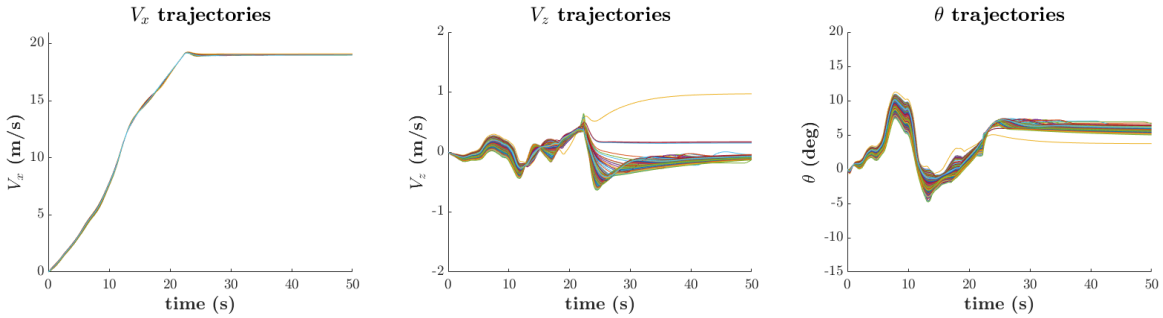

Figure 7.28: Forward transition simulations, $n_s = 10^3$ and $3\sigma = 9\%$.

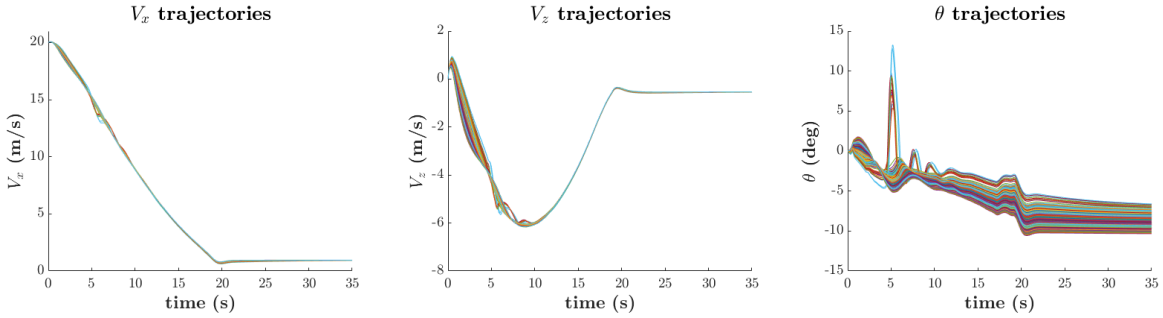

Figure 7.29: Backward transition simulations, $n_s = 10^3$ and $3\sigma = 10\%$.

7.6.3 Comparative Designs: H_∞ and μ-synthesis

As a final topic of this robustness analysis, we now consider comparative designs—H_∞ and μ-synthesis—that explicitly account for system robustness to better understand what this aircraft could theoretically achieve in terms of robust stability. Consider therefore the interconnection structure for the robust controller synthesis of feedback-linearized pitch rate given in Figure 7.30. The uncertain plant $G(s)$ of (7.34) is modeled using a multiplicative feedforward and an inverse multiplicative feedback to capture its gain and pole uncertainties, respectively, where $G_0(s) = 1/s$. This system is placed in feedback with a controller K to be synthesized, and weighting transfer functions W_i are applied to the exogenous outputs of interest to adjust the desired controller tuning. The generalized plant $P(s)$ is derived from inputs $[u_{\Delta_1} \quad u_{\Delta_2} \quad v_q]^T$ to outputs $[y_{\Delta_1} \quad y_{\Delta_2} \quad \tilde{q} \quad \tilde{\theta} \quad \tilde{v}_q \quad q]^T$, where the system input and output is $u = v_q$ and $y = q$, respectively.

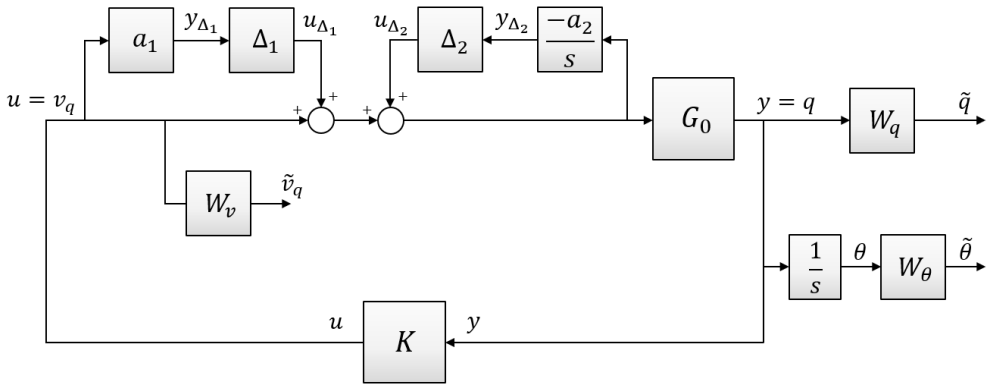

Figure 7.30: Interconnection structure for H_∞ and μ-synthesis control of feedback-linearized pitch rate.

In the previous subsection, robust stability was only able to be formally proven for this system for relatively low uncertainty bounds of $3\sigma = 3\%$ for forward transition and $3\sigma = 6\%$ for backward transition, but simulations revealed that the achievable robustness bounds may be higher. Hence, for the purposes of investigating what could be theoretically achieved with a robust controller design, we therefore consider

the forward transition $3\sigma = 9\%$ aerodynamic uncertainty set as the design point for the H_∞ and μ-synthesis controllers, which corresponds to ranges for the uncertain parameters of $a_1 = [-0.94, 1.79]$ and $a_2 = [-2.08, 6.81]$.

The weighting transfer functions W_i for controller synthesis were chosen based on the guidelines given in [155, 156], resulting in:

$$W_v = K_{dc}^v \frac{1 + (s/11)}{1 + (s/110)} \qquad W_q = \left(\frac{1}{2}\right) \frac{s + 14}{s + 0.14} \qquad W_\theta = 1000 \tag{7.43}$$

where K_{dc}^v is a design parameter that will be varied. W_v is used to adjust the controller bandwidth, where a breakpoint frequency of 11 $\mathrm{rad\,s^{-1}}$ corresponds to the 0 dB crossover frequency of the original (i.e. feedback-linearized q-loop + LQR controller) loop transfer function. W_q was constructed to shape the system sensitivity function with a maximum peak value of 2. And the value of W_θ was selected to ensure sufficient integral action[13] at a range of bandwidths.

Controllers were synthesized and the results analyzed. For best comparison with the current inner-loop LQR design, it was desired to tune these controllers to a similar bandwidth of around 7 to 8 $\mathrm{rad\,s^{-1}}$ (with the hopes of otherwise achieving improved robustness margins). This required setting $K_{dc}^v = 2000$. Unfortunately, results were not as desirable as expected (see Table 7.7 and Figure 7.31). While there was an increase in nominal phase margin (PM) between the robust designs and LQR (around 10° to 20°), the RSM values for the robust designs were comparable to what the LQR could already achieve for this range of uncertainty (around 0.83). Furthermore, the increase in nominal phase margin comes at the cost of step performance where, while there is less overshoot, the settling time has substantially increased[14].

[13]Note that, for numerical reasons, all integrators/poles at the origin within the generalized plant $P(s)$ must be approximated as $1/(s + \epsilon)$, where ϵ is a frequency that should be slower than the dynamics of the system (but not too slow) [86, 148]. For this design, $\epsilon = 11 \times 10^{-3}$ $\mathrm{rad\,s^{-1}}$ was chosen, corresponding to three decades below the 11 $\mathrm{rad\,s^{-1}}$ breakpoint frequency for W_v.

[14]Note that adjusting W_θ does not reasonably improve this behavior without additionally altering system bandwidth, and adjusting K_{dc}^v to compensate simply makes performance worse.

In the interest of finding robust controller designs that could (formally) withstand this $3\sigma = 9\%$ uncertainty set, it was found that it was possible to achieve $RSM \geq 1$ for values of $K_{dc}^{v} < 1.5$. This is demonstrated by the results for $K_{dc}^{v} = 1$ in Table 7.7 and Figure 7.32, where it can be seen that robust stability is achieved for this range of uncertainty for both designs as well as there being a further increase in phase margin. However, this tuning directly results in a significantly faster bandwidth for these controllers. In fact, a general trend was discovered when selecting each W_i that increasing the system bandwidth increases the system's robustness margins. This is an interesting discovery, since it was already found in Chapter 5 that a higher controller bandwidth is not achievable due to a structural resonance mode of the 3D-printed airframe. This therefore suggests that the achievable robustness margins for this system may be ultimately limited by the physical design of the aircraft. To further this theory, the gains for the original LQR controller were arbitrarily increased by a factor of 5, where it was found that robust stability could also be achieved for this design in the presence of this $3\sigma = 9\%$ aerodynamic uncertainty set. Specifically, this tuning for the LQR gives $RSM = 1.02$ and $PM = 85°$ while having a closed-loop system bandwidth of 47.3 $\mathrm{rad\,s^{-1}}$.

Table 7.7: Controller robustness and bandwidth metrics

Controller	K_{dc}^{v} value	PM	RSM	Closed-loop bandwidth
LQR	–	69°	0.84	7.4 $\mathrm{rad\,s^{-1}}$
H_∞	2000	80°	0.82	7.8 $\mathrm{rad\,s^{-1}}$
μ-synthesis	2000	87°	0.85	10.2 $\mathrm{rad\,s^{-1}}$
H_∞	1	103°	1.02	171[a] $\mathrm{rad\,s^{-1}}$
μ-synthesis	1	94°	1.01	51.3 $\mathrm{rad\,s^{-1}}$

[a] Note that the sensitivity function for this design flattens out for a wide range of frequencies around −3 dB, inflating the bandwidth value. Its effective bandwidth is closer to 51.3 $\mathrm{rad\,s^{-1}}$.

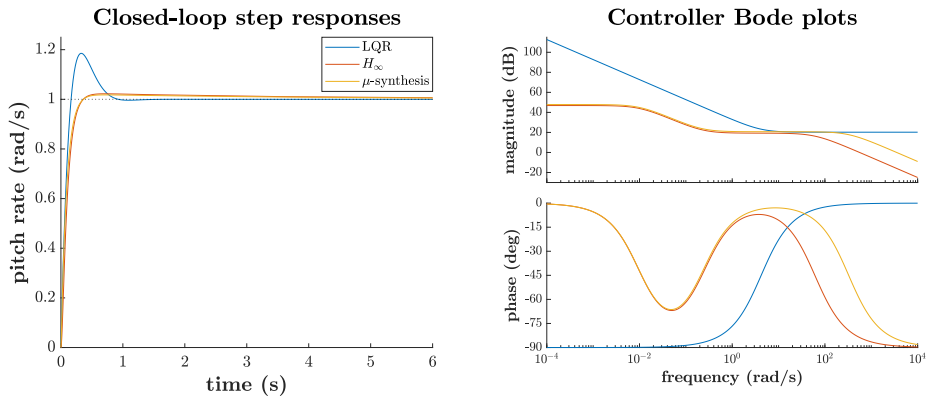

Figure 7.31: Robust controllers compared with the current LQR design, $K_{dc}^{v} = 2000$.

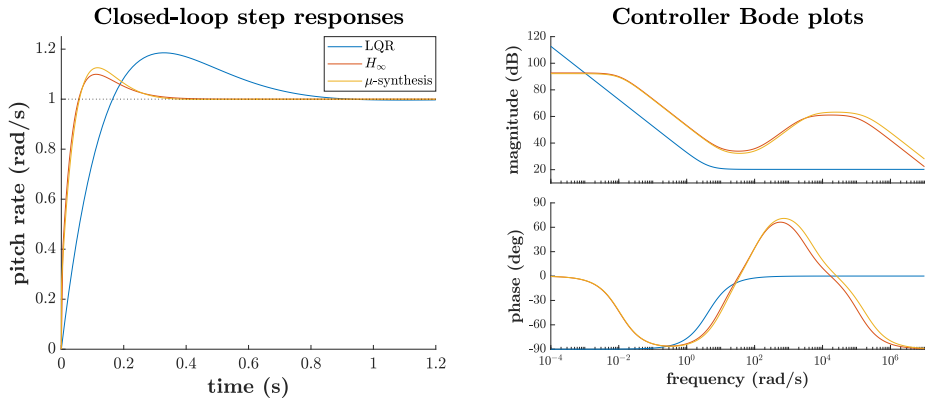

Figure 7.32: Robust controllers compared with the current LQR design, $K_{dc}^{v} = 1$.

7.7 Conclusion

This chapter concludes the design, flight testing, and analysis of the full-envelope flight controller development that was initiated in Chapter 5. The general outer-loop control strategy derived in Chapter 6 was applied to the specific tiltwing configuration of interest in this thesis and was evaluated in both simulation and flight testing, demonstrating proof-of-concept of this novel approach. Lastly, a formal robustness analysis was carried out for this system, where an analytical uncertainty model for the pitch rate dynamics of a tiltwing aircraft was derived and uncertainty bounds for this system were numerically proven. Comparative inner-loop designs employing H_∞ and μ-synthesis suggest that achievable robustness margins for this system may be ultimately constrained by the physical design of the scaled model aircraft.

Chapter 8

Summary and Conclusions

This thesis investigated full-envelope flight control design and testing of a tiltwing VTOL aircraft. The starting point was the development of a complete nonlinear 6-DoF model of the aircraft's dynamics, which captured complex aerodynamic effects such as prop-wing interactions by modeling discrete sections of the aircraft individually. The model was implemented in MATLAB and Simulink and validated by comparing its computed stability derivatives with those predicted using well-known potential flow-based CFD solvers (AVL, XFLR5, and OpenVSP). A brief analysis of the dynamics was then carried out to gain an understanding of the behavior of the aircraft in all modes of flight (hover, cruise, and transition).

Trajectory optimization was carried out to determine minimum-energy transition trajectories satisfying stall and constant-altitude constraints. Fluid injection AFC was implemented by modifying the lift and drag coefficients on the main wing of the aircraft to gain insight into the theoretical aerodynamic requirements of a tiltwing during transition. Additionally, perturbations in the aircraft's CG position were investigated to understand the physical loading limits of the aircraft during the maneuver. Results revealed that constant-altitude transition could be achieved in the forward direction but not the reverse, and the reasons for this were discussed. In addition

to providing deeper insight into the dynamics and behavior of this aircraft during transition, this methodology was later used to generate outer-loop reference inputs for the flight controller design.

A full-envelope flight controller was developed for this aircraft that employed feedback linearization in a cascade design combined with fault-tolerant control allocation. For the inner-loop, NDI was paired with LQR to regulate the attitude rates of the aircraft. A weighted pseudo-inverse control allocation scheme was then implemented to map virtual control inputs to real effector demands. Initial hover flight testing revealed a control-structural interaction that required the inner-loop bandwidth of the controller to be limited to avoid exciting this resonance mode. Despite this, controller performance was found to be satisfactory and sufficiently responsive to pilot inputs. Additionally, the controller demonstrated good robustness in response to motor failure and wind gust disturbances.

Focus then shifted to tackling the general problem of finding a full-envelope solution for the outer-loop velocity control design. To address the issues resulting from a lack of independent thrust actuation commonly exhibited by general transition VTOL configurations, a novel solution of selecting virtual control inputs and outputs in a rotating reference frame was proposed. This resulted in a globally-valid controller that is generalizable to a wide range of transition aircraft. Three unique control strategies that implement this concept were presented, each of which demonstrates a different balance of design complexity with tuning versatility and performance. Two of these three strategies were then applied to the tiltwing aircraft of this thesis, thus completing the aircraft's full-envelope flight controller design. Evaluation in simulation and hover flight demonstrated proof-of-concept of this novel outer-loop approach, and the differences in performance between the two strategies were discussed. Lastly, a formal robustness analysis revealed that the upper-bound aerodynamic modeling uncertainty for this system is relatively small, but comparative inner-loop designs em-

ploying H_∞ and μ-synthesis suggest that these bounds may be ultimately constrained by the bandwidth limitations inherent in the scaled model's physical design.

The conclusions of this thesis are:

1. Constant-altitude transition can be achieved in the forward direction but not the backward direction for this particular tiltwing configuration. While the use of AFC expands the achievable flight envelope of the aircraft and increases performance (in terms of decreased transition time, decreased altitude variation, increased robustness to CG variations), the fundamental stall constraint of the tail airfoil limits the benefits that the AFC can provide and the achievable performance of the aircraft in transition. Hence, selecting a tail airfoil with a higher stall constraint would better leverage the benefits provided by the AFC for this particular configuration.

2. The inner-loop inversion-based attitude controller proved successful in all flight modes, as demonstrated in both simulation and hover testing. It was additionally found to have good robustness in response to motor failure and wind gust disturbances. This was despite the necessity to reduce the controller's bandwidth to avoid exciting a 4 Hz structural resonance mode of the 3D-printed airframe. A further consequence of this bandwidth limitation was reduced robustness margins for this aircraft + flight controller combination, where comparative designs employing robust control methods were unable to achieve significantly improved margins without increasing the inner-loop bandwidth. This therefore suggests that a stiffening of the airframe through a redesign of the aircraft to increase the frequency of this structural resonance could improve overall controller performance and robustness by enabling a higher inner-loop bandwidth.

3. Following (1) and (2), an interesting takeaway can be drawn about the impor-

tance of the physical design of an aircraft being carried out in parallel with the control design (i.e. control co-design). This is particularly significant for transition VTOLs, as their dynamics and operation are appreciably more complex than conventional fixed-wing aircraft, and it can be therefore difficult to predict the full consequences of certain design decisions with these new and novel configurations.

4. Proof-of-concept was demonstrated for the proposed generalized outer-loop velocity control design. This was achieved through the application of two of the three strategies employing this concept, where both were evaluated in simulation and one in hover flight. The tiltwing configuration of this thesis served as an ideal platform for testing this novel approach due to its inherent lack of independent thrust actuation in both longitudinal directions and its complex aerodynamic model (a mix of wetted and unwetted aerodynamics). Between the two strategies that were evaluated, the coupled MIMO design of Strategy 3 was found to be more effective in transition due to the inherent cross-coupling that is accounted for, while the benefit of Strategy 1's decoupled approach is a simpler initial setup. That said, Strategy 3 was found to be more sensitive to variations away from the control allocation's pre-computed operating points.

Future work of interest following this thesis includes:

- Testing this full-envelope flight controller design in transition flight.
- Extending the trajectory optimization model of Chapter 4 to 6-DoFs to investigate asymmetrical failure modes of a motor/prop (generating trajectories). Following this, investigating motor/prop failure modes in transition (simulations and flight testing).
- Failure detection. While the weighted pseudo-inverse control allocation can redistribute input demands to remaining actuators (fault-tolerant), the ability

to detect a failure would be an interesting and important capability in the context of air taxis.

- Extending the generalized outer-loop control strategies to include lateral (V_y) dynamics. While this would greatly increase the complexity of the design, it would unify the lateral and longitudinal controllers, increase controller robustness to lateral disturbances (such as crosswind gusts) during transition, and potentially improve performance in cruise coordinated turns and hover lateral motion.

- Redesign of the scaled demonstrator model. It was discovered that certain physical limitations of the aircraft (e.g. tail airfoil stall limit, structural resonance of the airframe, low thrust margin) had an impact on the controller design, feasible transition trajectories, and theoretical robustness margins. It would be of interest to remove (or minimize) these effects where possible (such as through a redesign of the tail airfoil, stiffening of the wing structure, and/or sizing for larger motors) and thus reinvestigate the achievable performance of this flight controller design.

Bibliography

[1] Quan Quan. *Introduction to Multicopter Design and Control*, chapter 1, pages 1–27. Springer, 2017.

[2] Yefim Gordon. Vertical reality – the history of vtol. *Scribd*, January/February 2016.

[3] Cornu helicopter image, public domain, via wikimedia commons, accessed October 2025. URL `https://en.wikipedia.org/wiki/Cornu_helicopter`.

[4] Kaman k-225 image, public domain, via wikimedia commons, accessed October 2025. URL `https://en.wikipedia.org/wiki/Kaman_K-225`.

[5] Cl-84 tiltwing image, public domain, accessed August 2021. URL `https://vertipedia.vtol.org/aircraft/getAircraft/aircraftID/56`.

[6] Mv-22 mcas miramar 2014 image by fox 52, licensed under cc by-sa 4.0. license url: https://creativecommons.org/licenses/by-sa/4.0. cropped from the original, accessed October 2026. URL `https://commons.wikimedia.org/wiki/File:MV-22_mcas_Miramar_2014.JPG`.

[7] Christopher Silva, Wayne R. Johnson, Eduardo Solis, Michael D. Patterson, and Kevin R. Antcliff. Vtol urban air mobility concept vehicles for technology development. *2018 Aviation Technology, Integration, and Operations Conference*, 2018. doi: 10.2514/6.2018-3847. URL `https://arc.aiaa.org/doi/abs/10.2514/6.2018-3847`.

[8] Akshay R Kadhiresan and Michael J Duffy. Conceptual design and mission analysis for evtol urban air mobility flight vehicle configurations. In *AIAA aviation 2019 forum*, page 2873, 2019. doi: 10.2514/6.2019-2873.

[9] R. A. Danis, M. W. Green, J. L. Freeman, and D. W. Hall. Examining the conceptual design process for future hybrid-electric rotorcraft. Technical Report NASA/CR-2020-220460, NASA, 2020.

[10] Chana Anna Saias, Ioannis Goulos, Ioannis Roumeliotis, Vassilios Pachidis, and Marko Bacic. Preliminary design of hybrid-electric propulsion systems for emerging urban air mobility rotorcraft architectures. *Journal of Engineering for Gas Turbines and Power*, 143(11):111015, 10 2021. ISSN 0742-4795. doi: 10.1115/1.4052057.

[11] T. Edwards and G. Price. evtol passenger acceptance. Technical Report NASA/CR-2020-220460, NASA, 2020.

[12] Alessandro Bacchini and Enrico Cestino. Electric vtol configurations comparison. *MDPI Aerospace*, 6(26), 2019. doi: 10.3390/aerospace6030026.

[13] Airbus a3 vahana image by vertical flight society (vfs), licensed under cc by-sa 4.0. license url: https://creativecommons.org/licenses/by-sa/4.0. cropped from the original, accessed October 2026. URL `https://gallery.vtol.org/image/P0jan`.

[14] Rolls-royce evtol concept image, copyright rolls-royce, accessed Sep 2026. URL `https://www.rolls-royce.com/media/our-stories/discover/2018/blue-sky-thinking-rr-unveils-evtol-concept-at-farnborough-airshow.aspx`.

[15] J. P. Campbell. Status of v/stol research and development in the united states. *AIAA Journal of Aircraft*, 1, May–June 1964. doi: 10.2514/3.43576.

[16] Henry L. Kelley, John P. Reeder, and Robert A. Champine. Summary of a flight-test evaluation of the cl-84 tilt-wing v/stol aircraft. *NASA Report No. TM X-1914*, March 1970.

[17] J. Wolkovitch, C. W. Lamont, and D. W. Lochtie. Summary of flight-test results of the vz-2 tilt-wing aircraft. Technical Report NASA TN-D-989, NASA, 1962.

[18] D. C. Dugan, R. G. Erhart, and L. G. Schroers. The xv-15 tilt rotor research aircraft. Technical Report NASA TM-81244, NASA, 1980.

[19] H. L. Kelley. Transition and hovering flight characteristics of a tilt-duct vtol research aircraft. Technical Report NASA D-1491, NASA, 1962.

[20] W. F. Chana and T. M. Sullivan. The tilt wing advantage - for high speed vstol aircraft. *SAE Transactions Section 1: Journal of Aerospace*, 101:1535–1543, 1992.

[21] J. Rutherford and S. Bass. Advanced technology tilt wing study. *AIAA Guidance, Navigation, and Control Conference*, August 1992. doi: 10.2514/6.1992-4237.

[22] William F. Chana and T. M. Sullivan. The tilt wing advantage - for high speed vstol aircraft. *SAE Technical Paper*, 1992. doi: 10.4271/921911.

[23] Charles C. Crawford, Michael W. Heiges, and Mark Wasikowski. Analysis of tiltwing aircraft configuration potential. *SAE Transactions*, 105:1267–1280, 1996.

[24] Shamsheer S. Chauhan and Joaquim R. R. A. Martins. Tilt-wing evtol takeoff trajectory optimization. *AIAA Journal of Aircraft*, 57(1), January 2020. doi: 10.2514/1.C035476.

[25] M. O. McKinney, R. H. Kirby, and W. A. Newsom. Aerodynamic factors to be considered in the design of tilt-wing v/stol airplanes. *Annals New York Academy of Sciences*, 107:221–248, March 1963. doi: 10.1111/j.1749-6632.1963.tb13279.x.

[26] Leo Panish and Marko Bacic. A generalized full-envelope outer-loop feedback linearization control strategy for transition vtol aircraft. *AIAA Aviation Forum*, June 2023. doi: 10.2514/6.2023-4511.

[27] Leo Panish and Marko Bacic. Tiltwing evtol transition trajectory optimization. *AIAA Journal of Aircraft*, 62(1):81–93, Jan–Feb 2025. doi: 10.2514/1.C037862.

[28] Leo Panish and Marko Bacic. Transition trajectory optimization for a tiltwing vtol aircraft with leading-edge fluid injection active flow control. *AIAA SciTech Forum*, January 2022. doi: 10.2514/6.2022-1082.

[29] Leo Panish, Chris Nicholls, and Marko Bacic. Nonlinear dynamic inversion flight control of a tiltwing vtol aircraft. *AIAA SciTech Forum*, January 2023. doi: 10.2514/6.2023-1910.

[30] Eugene L. Duke, Robert F. Antonoiewicz, and Keith D. Krambeer. Derivation and definition of a linear aircraft model. *NASA Reference Publication 1207*, August 1988. Images reprinted courtesy of NASA.

[31] Mark Peters and Michael A. Konyak. The engineering analysis and design of the aircraft dynamic model for the faa target generation facility. *FAA Simulation Branch, Laboratory Services Division*, October 2012.

[32] Jan Roskam. *Airplane Flight Dynamics and Automatic Flight Controls, Part 1*, chapter 1–5, pages 3–412. DARcorporation, Kansas, 7 edition, 2018.

[33] B. W. McCormick, A. S. Aljabri, S. J. Jumper, and Z. N. Martinovic. The analysis of propellers including interaction effects. *Business Aircraft Meeting and Exposition, SAE Technical Paper 790576*, February 1979. doi: 10.4271/790576.

[34] John S. Carlton. *Marine Propellers and Propulsion*, chapter 5–8, pages 63–204. Elsevier Ltd., New York, 2 edition, 2007.

[35] David A. Caughey. Introduction to aircraft stability and control. *Sibley School of Mechanical and Aerospace Engineering, Cornell University, Ithaca, NY*, 2011. https://courses.cit.cornell.edu/mae5070/Caughey_2011_04.pdf.

[36] Duane T. Mcruer, Dunstan Graham, and Irving Ashkenas. *Aircraft Dynamics and Automatic Controls*, chapter 4, pages 220–238. Princeton University Press, Princeton, New Jersey, 2014.

[37] Richard E. Maine and Kenneth W. Iliff. Application of parameter estimation to aircraft stability and control. Technical Report NASA Reference Publication 1168, NASA, 1986.

[38] Richard C. Dorf and Robert H. Bishop. *Modern Control Systems*, chapter 7, 11, pages 443–552, 834–909. Prentice-Hall, 11 edition, 2008.

[39] Hassan K. Khalil. *Nonlinear Systems*, chapter 13, pages 505–550. Prentice-Hall, 3 edition, 2002.

[40] Jean-Jacques E. Slotine and Weiping Li. *Applied Nonlinear Control*, chapter 3–4, 6, pages 40–156, 207–275. Prentice-Hall, 1 edition, 1991.

[41] A. M. Lyapunov. The general problem of the stability of motion, (1892), in russian. translated from russian into french by edouard davaux. translated from french into english by a. t. fuller. *International Journal of Control*, 55(3):531–773, 1992.

[42] Alberto Isidori. *Nonlinear Control Systems*, chapter 4,5, pages 137–292. Springer-Verlag, 3 edition, 1995.

[43] Tor A. Johansen and Thor I. Fossen. Control allocation - a survey. *Elsevier Automatica*, 49:1087–1103, 2013.

[44] G. H. Golub and C. F. van Loan. *Matrix Computations*, chapter 5, pages 233–302. North Oxford Academic Press, 4 edition, 2013.

[45] D. Enns. Control allocation approaches. *AIAA Guidance, Navigation, and Control Conference and Exhibit*, pages 98–108, 1998.

[46] T. I. Fossen and S. I. Sagatun. Adaptive control of nonlinear systems: a case study of underwater robotic systems. *Journal of Robotic Systems*, 8:393–412, 1991.

[47] K. A. Bordignon and W. C. Durham. Closed-form solution to constrained control allocation problem. *Journal of Guidance, Control, and Dynamics*, 18: 1000–1007, 1995.

[48] W. C. Durham. Constrained control allocation. *Journal of Guidance, Control, and Dynamics*, 16:717–725, 1993.

[49] Sigurd Skogestad and Ian Postlethwaite. *Multivariable Feedback Control: Analysis and Design*, chapter 3, pages 67–118. Wiley, 2 edition, 2005.

[50] W. C. Durham. Attainable moments for the constrained control allocation problem. *Journal of Guidance, Control, and Dynamics*, 17:1371–1373, 1994.

[51] W. C. Durham. Constrained control allocation: Three-moment problem. *Journal of Guidance, Control, and Dynamics*, 17:330–336, 1994.

[52] W. C. Durham. Efficient, near-optimal control allocation. *Journal of Guidance, Control, and Dynamics*, 22:369–372, 1999.

[53] J. A. M. Petersen and M. Bodson. Fast implementation of direct allocation with extension to coplanar controls. *Journal of Guidance, Control, and Dynamics*, 25:464–473, 2002.

[54] M. Bodson. Evaluation of optimization methods for control allocation. *Journal of Guidance, Control, and Dynamics*, 25:703–711, 2002.

[55] J. Virnig and D. Bodden. Multivariable control allocation and control law conditioning when control effectors limit. *AIAA Guidance, Navigation, and Control Conference*, 1994. doi: 10.2514/6.1994-3609.

[56] J. Shi, W. Zhang, G. Li, and X. Liu. Research on allocation efficiency of the redistributed pseudo inverse algorithm. *Science China Information Sciences*, 53:271–277, 2010.

[57] J. M. Buffington and D. F. Enns. Lyapunov stability analysis of daisy chain control allocation. *Journal of Guidance, Control, and Dynamics*, 19:1226–1230, 1996.

[58] T. A. Johansen, T. I. Fossen, and S. P. Berge. Constrained nonlinear control allocation with singularity avoidance using sequential quadratic programming. *IEEE Transactions on Control Systems Technology*, 12:211–216, 2004.

[59] V. L. Poonamalle, S. Yurkovich, A. Serrani, D. B. Doman, and M. W. Oppenheimer. A nonlinear programming approach for control allocation. *American Control Conference*, pages 1689–1694, 2004.

[60] K. Levenberg. A method for the solution of certain problems in least-squares. *Quarterly Applied Mathematics 2*, pages 164–168, 1944.

[61] D. Marquardt. An algorithm for least-squares estimation of nonlinear parameters. *SIAM Journal Applied Mathematics*, 11:431–441, 1963.

[62] J. J. Moré. The levenberg-marquardt algorithm: Implementation and theory. *Numerical Analysis, ed. G. A. Watson, Lecture Notes in Mathematics 630, Springer Verlag*, pages 105–116, 1977.

[63] Daniel Liberzon. *Calculus of Variations and Optimal Control Theory: A Concise Introduction*, chapter 6, pages 180–199. Princeton University Press, 1 edition, 2012.

[64] Anil V. Rao. A survey of numerical methods for optimal control. *Advances in the Astronautical Sciences*, (preprint) 09–334, January 2010.

[65] John T. Betts. Survey of numerical methods for trajectory optimization. *Journal of Guidance, Control, and Dynamics*, 21(2):193–207, Mar–Apr 1998. doi: 10.2514/2.4231.

[66] Miriam E. Dennis, William W. Hager, and Anil V. Rao. Computational method for optimal guidance and control using adaptive gaussian quadrature collocation. *Journal of Guidance, Control, and Dynamics*, 42(9):2026–2041, September 2019. doi: 10.2514/1.G003943.

[67] Divya Garg, Michael Patterson, William W. Hager, Anil V. Rao, David A. Benson, and Geoffrey T. Huntington. A unified framework for the numerical solution of optimal control problems using pseudospectral methods. *Automatica*, 46:1843–1851, August 2010.

[68] Brian Jackson. Exploring quadrature rules in direct collocation methods for trajectory optimization. *Principles of Optimal Control*, (AA 203), June 2018.

[69] Christopher L. Darby, William W. Hager, and Anil V. Rao. A preliminary analysis of a variable-order approach to solving optimal control problems using pseudospectral methods. *AIAA Guidance, Navigation, and Control Conference*, 2–5 August 2010.

[70] Michael A. Patterson and Anil V. Rao. Gpops-ii: A matlab software for solving multiple-phase optimal control problems using hp-adaptive gaussian quadrature collocation methods and sparse nonlinear programming. *ACM Trans. Math. Softw.*, 41(1), oct 2014. ISSN 0098-3500. doi: 10.1145/2558904.

[71] Michael A. Patterson, William W. Hager, and Anil V. Rao. A *ph* mesh refinement method for optimal control. *Optim. Control Appl. Meth.*, 36:398–421, 2015. doi: 10.1002/oca.2114.

[72] Martin Doff-Sotta, Mark Cannon, and Marko Bacic. Fast optimal trajectory generation for a tiltwing vtol aircraft with application to urban air mobility. In *2022 American Control Conference (ACC)*, pages 4036–4041, 2022. doi: 10.23919/ACC53348.2022.9867852.

[73] Martin Doff-Sotta and Mark Cannon. Difference of convex functions in robust tube nonlinear mpc. In *2022 IEEE 61st Conference on Decision and Control (CDC)*, pages 3044–3050, 2022. doi: 10.1109/CDC51059.2022.9993390.

[74] Martin Doff-Sotta, Mark Cannon, and Marko Bacic. Robust trajectory optimisation for transitions of tiltwing vtol aircraft. In *2023 Conference on Control Technology and Application (CCTA)*, pages 425–431, 2023. doi: 10.1109/CCTA54093.2023.10253282.

[75] Martin Doff-Sotta, Mark Cannon, and Marko Bacic. Data-driven robust model predictive control of tiltwing vertical take-off and landing aircraft. *AIAA Journal of Guidance, Control, and Dynamics*, 48(1):203–211, January 2025. doi: 10.2514/1.G008315.

[76] M. Cannon, J. Buerger, B. Kouvaritakis, and S. Rakovic. Robust tubes in nonlinear model predictive control. *IEEE Transactions on Automatic Control*, 56(8):1942–1947, 2011. doi: 10.1109/TAC.2011.2135190.

[77] Mohamed Gad-el Hak, Andrew Pollard, and Jean-Paul Bonnet. *Flow Control, Fundamentals and Practices*, chapter 1, pages 1–108. Springer-Verlag, 1998.

[78] Linda D. Kral. Active flow control technology. *ASME Fluids Engineering Division Technical Brief*, January 2000.

[79] Robert J. Englar. Circulation control for high lift and drag generation on stol aircraft. *Journal of Aircraft*, 12(5):457–463, May 1975. doi: 10.2514/3.59824.

[80] John L. Loth and Michael Boasson. Circulation controlled stol wing optimization. *Journal of Aircraft*, 21(2):128–134, February 1984. doi: 10.2514/3.48235.

[81] Robert J. Englar, Rodney A. Hemmerly, W. Horace Moore, Vladimir Seredinsky, Walter Valckenaere, and John A. Jackson. Design of the circulation control wing stol demonstrator aircraft. *Journal of Aircraft*, 18(1):51–58, January 1981. doi: 10.2514/3.57463.

[82] Robert J. Englar, Lynn A. Trobaugh, and Rodney A. Hemmerly. Stol potential of the circulation control wing for high-performance aircraft. *Journal of Aircraft*, 15(3):175–181, March 1978. doi: 10.2514/3.58337.

[83] Robert J. Englar. Circulation control pneumatic aerodynamics: Blown force and moment augmentation and modification - past, present and future. *Fluids 2000 Conference and Exhibit*, June 2000. doi: 10.2514/6.2000-2541.

[84] B. G. McLachlan. Study of a circulation control airfoil with leading/trailing-edge blowing. *Journal of Aircraft*, 26(9):817–821, September 1989. doi: 10.2514/3.45846.

[85] J. Ottensoser. Two-dimensional subsonic evaluation of a 15-percent thick circulation control airfoil with slots at both leading and trailing edges. *David Taylor Naval Ship R and D Center, Rept. 4456*, July 1974.

[86] Sigurd Skogestad and Ian Postlethwaite. *Multivariable Feedback Control: Analysis and Design*, chapter 7–9, pages 259–382. Wiley, 2 edition, 2005.

[87] D. James, A. Manohar, D. Fillingham, and M. Bacic. Vertical take-off and landing aircraft. *European Patent Office*, 2019. Patent Number EP3594107A1.

[88] E. Cetinsoy, S. Dikyar, C. Hancer, K. T. Oner, E. Sirimoglu, M. Unel, and M. F. Aksit. Design and construction of a novel quad tilt-wing uav. *Mechatronics*, 22, Issue 6:723–745, 2012. doi: 10.1016/j.mechatronics.2012.03.003.

[89] Kai Masuda and Kenji Uchiyama. Robust control design for quad tilt-wing uav. *Aerospace*, 5(1, Paper 2), 2018. doi: 10.3390/aerospace5010017.

[90] David Rohr, Thomas Stastny, Sebastian Verling, and Roland Siegwart. Attitude and cruise control of a vtol tiltwing uav. *IEEE Robotics and Automation Letters*, 4, No. 3, July 2019. doi: 10.1109/LRA.2019.2914340.

[91] Fpv compute thrust to weight–a comprehensive analysis of calculating and utilizing thrust-to-weight ratio for fpv drones, accessed November 2025. URL https://rcdrone.top/blogs/articles/fpv-compute-thrust-to-weight?srsltid=AfmBOoobGFTBLuJQzTzL0In832-UBLh_wnJgdRGt379eyguv7T-pwOLY.

[92] Drone thrust testing, accessed November 2025. URL https://www.halfchrome.com/drone-thrust-testing/.

[93] Internal data provided by rolls-royce plc., October 2020.

[94] Daisuke Kubo and Shinji Suzuki. Tail-sitter vertical takeoff and landing unmanned aerial vehicle: Transitional flight analysis. *AIAA Journal of Aircraft*, 45(1), January 2008. doi: 10.2514/1.30122.

[95] Adnan Maqsood and Tiauw Hiong Go. Optimization of transition maneuvers through aerodynamic vectoring. *Aerospace Science and Technology*, 23(1):363–371, 2012. doi: 10.1016/j.ast.2011.09.004.

[96] Sebastian Verling, Thomas Stastny, Gregory Battig, Kostas Alexis, and Roland Siegwart. Model-based transition optimization for a vtol tailsitter. *IEEE International Conference on Robotics and Automation (ICRA)*, May–June 2017. doi: 10.1109/icra.2017.7989454.

[97] Daniel Milz and Gertjan Looye. Tilt-wing control design for a unified control concept. *AIAA SciTech Forum*, January 2022. doi: 10.2514/6.2022-1084.

[98] Marc May, Daniel Milz, and Gertjan Looye. Dynamic modeling and analysis of tilt-wing electric vertical take-off and landing vehicles. *AIAA SciTech Forum*, January 2022. doi: 10.2514/6.2022-0263.

[99] Rodrigo Sorbilli Cardoso de Sousa and Roberto Gil Annes da Silva. Nasa-crm lateral-directional stability derivatives comparison between semi-empirical, vortex lattice and panel method. *32nd Congress of the International Council of the Aeronautical Sciences*, September 2021.

[100] Larry L. Erickson. Panel methods—an introduction. *NASA Technical Paper 2995*, 1990.

[101] J. Holsten, P. Hartmann, and D Moormann. Modeling of tiltwing aircraft dynamics as linear system. *6th EASN International Conference on Innovation in European Aeronautics Research*, 2016.

[102] Daniel Milz, Marc May, Sophie F. Armanini, and Gertjan Looye. Stability and robustness of a tandem tilt-wing vtol with dynamic inversion control laws. *AIAA SciTech Forum*, January 2026. doi: 10.2514/6.2026-2070.

[103] J. Wolkovitch, C. W. Lamont, and D. W. Lochtie. Fundamental limitations on v/stol terminal guidance due to aircraft characteristics. Technical Report NASA CR-1901, NASA, 1971.

[104] J. Freudenberg and D. Looze. Right half plane poles and zeros and design tradeoffs in feedback systems. *IEEE Transactions on Automatic Control*, 30 (6):555–565, 1985. doi: 10.1109/TAC.1985.1104004.

[105] Priyank Pradeep and Peng Wei. Energy optimal speed profile for arrival of tandem tilt-wing evtol aircraft with rta constraint. *Proceedings of IEEE/CSAA Guidance, Navigation and Control Conference (GNCC)*, pages 1–6, 2018. doi: 10.1109/GNCC42960.2018.9018748.

[106] Shuan-Tai Yeh, Guirong Yan, and Xiaosong Du. Inverse machine learning prediction for optimal tilt-wing evtol takeoff trajectory. *AIAA Aviation Forum*, June 2023. doi: 10.2514/6.2023-3593.

[107] Marc May, Daniel Milz, Sophie Armanini, and Gertjan Looye. Transition strategies for tilt-wing aircraft. *AIAA Journal of Guidance, Control, and Dynamics*, 48(10), October 2025. doi: 10.2514/1.G008257.

[108] Marc May, Daniel Milz, Sophie F. Armanini, and Gertjan Looye. Impact of failure on the tilt-wing evtol backward transition. *AIAA Journal of Aircraft*, Article in Advance, 2025. doi: 10.2514/1.C038391.

[109] R. H. Stone and G. Clarke. Optimization of transition manoeuvres for a tail-sitter unmanned air vehicle. *Australian Aerospace International Congress [online paper], Paper 105*, March 2001.

[110] M. Jabbal, S. C. Liddle, and W. J. Crowther. Active flow control systems architectures for civil transport aircraft. *Journal of Aircraft*, 47(6), Nov-Dec 2010. doi: 10.2514/1.C000237.

[111] Louis N. Cattafesta and Mark Sheplak. Actuators for active flow control. *Annual Review of Fluid Mechanics*, 2011. doi: 10.1146/annurev-fluid-122109-160634.

[112] John A. Ekaterinaris. Prediction of active flow control performance on airfoils and wings. *Aerospace Science and Technology*, 8(5):401–410, 2004. doi: 10.1016/j.ast.2004.02.003.

[113] Daniel P. Raymer. *Aircraft Design: A Conceptual Approach*, chapter 12, pages 303–364. AIAA Education Series, 6 edition, 2018.

[114] Gpops-ii reference manual, accessed December 2023. URL `https://gpops2.com/resources/gpops2UsersGuide.pdf`.

[115] Jeffrey J. Dickeson, Oguzhan Cifdaloz, David W. Miles, Paul M. Koziol, Valana L. Wells, and Armando A. Rodriguez. Robust h_∞ gain-scheduled conversion for a tilt-wing rotorcraft. *45th IEEE Conference on Decision and Control*, December 2006. doi: 10.1109/CDC.2006.377256.

[116] Jeffrey J. Dickeson, David Miles, Oguzhan Cifdaloz, Valana L. Wells, and Armando A. Rodriguez. Robust lpv h-infinity gain-scheduled hover-to-cruise conversion for a tilt-wing rotorcraft in the presence of cg variations. *46th IEEE Conference on Decision and Control*, December 2007. doi: 10.1109/CDC.2007.4435028.

[117] Masayuki Sato and Koji Muraoka. Flight test verification of flight controller for quad tilt wing unmanned aerial vehicle. *AIAA Guidane, Navigation, and Control Conference*, August 2013. doi: 10.2514/6.2013-5100.

[118] Masayuki Sato and Koji Muraoka. Flight controller design and demonstration of quad-tilt-wing unmanned aerial vehicle. *Journal of Guidance, Control, and Dynamics*, May 2014. doi: 10.2514/1.G000263.

[119] Hironori Totoki, Yoshimasa Ochi, Masayuki Sato, and Koji Muraoka. Design of gain scheduled stability and control augmentation system for quad-tilt-wing uav. *AIAA SciTech Forum*, January 2015. doi: 10.2514/6.2015-0606.

[120] Anh Tuan Tran, Noboru Sakamoto, Masayuki Sato, and Koji Muraoka. Control augmentation system design for quad-tilt-wing unmanned aerial vehicle via robust output regulation method. *IEEE Transactions on Aerospace and Electric Systems*, 53, No. 1, February 2017. doi: 10.1109/TAES.2017.2650618.

[121] Philipp Hartmann, Carsten Meyer, and Dieter Moormann. Unified approach for velocity control and flight state transition of unmanned tiltwing aircraft. *AIAA SciTech Forum*, January 2016. doi: 10.2514/6.2016-2101.

[122] Philipp Hartmann, Carsten Meyer, and Dieter Moormann. Unified velocity control and flight state transition of unmanned tilt-wing aircraft. *AIAA Journal of Guidance, Control, and Dynamics*, November 2016. doi: 10.2514/1.G002168.

[123] Anthony Comer, Aashutosh Aman Mishra, and Imon Chakraborty. Total energy flight control architecture optimization for a tilt-wing aircraft. *AIAA Aviation Forum*, June 2023. doi: 10.2514/6.2023-4510.

[124] Guillaume J. J. Ducard and Mike Allenspach. Review of designs and flight control techniques of hybrid and convertible vtol uavs. *Aerospace Science and Technology*, 118, November 2021. doi: 10.1016/j.ast.2021.107035.

[125] Kai Masuda and Kenji Uchiyama. Flight controller design using μ-synthesis for quad tilt-wing uav. *AIAA SciTech Forum*, January 2019. doi: 10.2514/6.2019-1918.

[126] Johannes Autenrieb, Hyo-Sang Shin, and Marko Bacic. Development of a neural network-based adaptive nonlinear dynamic inversion controller for a tilt-wing vtol aircraft. *The 2019 International Workshop on Research, Education and Development on Unmanned Aerial Systems*, November 2019. doi: 10.1109/RED-UAS47371.2019.8999700.

[127] Daniel Milz, Marc May, and Gertjan Looye. Dynamic inversion-based control concept for transformational tilt-wing evtols. *AIAA SciTech Forum*, January 2024. doi: 10.2514/6.2024-1290.

[128] Daniel Milz, Marc May, and Gertjan Looye. Tandem tilt-wing control design based on sensory nonlinear dynamic inversion. *AIAA Aviation Forum*, Jul–Aug 2024. doi: 10.2514/6.2024-4418.

[129] Daniel Milz and Marc May. Dynamic inversion-based control design for multi-phase tandem tilt-wing flight. *AIAA SciTech Forum*, January 2025. doi: 10.2514/6.2025-2698.

[130] Michael G. Safonov and Michael Athans. Gain and phase margin for multiloop lqg regulators. *IEEE Transactions on Automatic Control*, AC-22(2):173–179, April 1977. doi: 10.1109/TAC.1977.1101470.

[131] John C. Doyle. Guaranteed margins for lqg regulators. *IEEE Transactions on Automatic Control*, AC-23(4), August 1978.

[132] John C. Doyle and Gunter Stein. Robustness with observers. *IEEE Transactions on Automatic Control*, AC-24(4), August 1979.

[133] John C. Doyle, Roy S. Smith, and Dale F. Enns. Control of plants with input saturation nonlinearities. *IEEE American Control Conference*, June 1987. doi: 10.23919/ACC.1987.4789464.

[134] Original f-16 line drawing from http://www.af.mil/art/, public domain, via wikimedia commons, accessed August 2025. URL `https://commons.wikimedia.org/wiki/File:General_Dynamics_F-16_Fighting_Falcon_3-view_line_drawing.svg`.

[135] Yih Tang Yeo and Hugh H. T. Liu. Transition control of a tilt-rotor vtol uav. *AIAA SciTech Forum*, January 2018. doi: 10–.2514/6.2018-1848.

[136] Hao Yang and Rafael M. Morales. Robust full envelope flight control design for an evtol vehicle. *AIAA SciTech Forum*, January 2021. doi: 10.2514/6.2021-0254.

[137] Anthony Comer, Aashutosh Mishra, and Imon Chakraborty. Extension of a full envelope flight control system to ctol/stol capabilities on an urban air mobility vectored-thrust vtol vehicle. *AIAA SciTech Forum*, January 2025. doi: 10.2514/6.2025-0658.

[138] Philipp Hartmann, Marten Schutt, and Dieter Moormann. Control of departure and approach maneuvers of tiltwing vtol aircraft. *AIAA SciTech Forum*, January 2017. doi: 10.2514/6.2017-1914.

[139] David Rohr, Matthias Studiger, Thomas Stastny, Nicholas R. J. Lawrance, and Roland Siegwart. Nonlinear model predictive velocity control of a vtol tiltwing uav. *IEEE Robotics and Automation Letters*, 6, No. 3, July 2021. doi: 10.1109/LRA.2021.3084888.

[140] Mike Allenspach and Guillaume Jacques Joseph Ducard. Nonlinear model predictive control and guidance for a propeller-tilting hybrid unmanned air vehicle. *Automatica*, 132, October 2021. doi: 10.1016/j.automatica.2021.109790.

[141] Thomas Lombaerts, John Kaneshige, Stefan Schuet, Gordon Hardy, Bimal Aponso, and Kimberlee Shish. Dynamic inversion based full envelope flight control for an evtol vehicle using a unified framework. *AIAA SciTech Forum*, January 2020. doi: 10.2514/6.2020-1619.

[142] Metehan Yayla, Ali Turker Kutay, Murat Senipek, and Osman Gungor. An adaptive flight controller design for a tilt-prop fixed wing uav for all flight modes. *AIAA SciTech Forum*, January 2020. doi: 10.2514/6.2020-0593.

[143] Stefan Raab, Jiannan Zhang, Pranav Bhardwaj, and Florian Holzapfel. Proposal of a unified control strategy for vertical take-off and landing transition aircraft configurations. *AIAA Aviation Forum*, June 2018. doi: 10.2514/6.2018-3478.

[144] Salahudden Salahudden, Henil Agrawal, Aditya Karnam, Anirban Roy, and Abhishek Singh. Flight dynamics and control of tricopter vtol and tilt rotor transition to fixed wing. *AIAA Aviation Forum*, Jul–Aug 2024. doi: 10.2514/6.2024-4513.

[145] Michael A. Henson and Dale E. Seborg. *Nonlinear Process Control*, chapter 4, pages 149–232. Prentice-Hall, 1 edition, 1997.

[146] Michael A. Henson and Dale E. Seborg. Input-output linearization of general nonlinear processes. *AIChE Journal*, 36(11):1753–1757, November 1990.

[147] Francis J. Doyle III and Michael A. Henson. *Nonlinear Process Control*, chapter 3, pages 111–148. Prentice-Hall, 1 edition, 1997.

[148] Jacob Reiner, Gary J. Balas, and William L. Garrard. Robust dynamic inversion for control of highly maneuverable aircraft. *Journal of Guidance, Control, and Dynamics*, 18(1):18–24, January–February 1995. doi: 10.2514/3.56651.

[149] Jacob Reiner, Gary J. Balas, and William L. Garrard. Flight control design using robust dynamic inversion and time-scale separation. *Automatica*, 32(11): 1493–1504, 1996.

[150] Dale Enns, Dan Bugajski, Russ Hendrick, and Gunter Stein. Dynamic inversion: an evolving methodology for flight control design. *International Journal of Control*, 59(1):71–91, 1994. doi: 10.1080/00207179408923070.

[151] Amir Shakouri. On the prescribed-time attractivity and frozen-time eigenvalues of linear time-varying systems. *Automatica*, 140, 2022. doi: 10.1016/j.automatica.2022.110173.

[152] Francesco Amato, Giovanni Celentano, and Franco Garofalo. New sufficient conditions for the stability of slowly varying linear systems. *IEEE Transactions on Automatic Control*, 1993. doi: 10.1109/9.237657.

[153] S. Ray and J. Zhu. A case study of frozen-time eigenvalues in the stability analysis for periodic linear systems. *[1991 Proceedings] The Twenty-Third Southeastern Symposium on System Theory*, 1991. doi: 10.1109/SSST.1991.138601.

[154] Lawrence Markus and Hidehiko Yamabe. Global stability criteria for differential systems. *Osaka Math. J.*, 12:305–317, 1960. doi: 10.18910/9397.

[155] Petter Lundstrom, Sigurd Skogestad, and Zi-Qin Wang. Performance weight selection for h-infinity and μ-control methods. *Trans Inst MC*, 13(5), 1991.

[156] John E. Bibel and D. Stephen Malyevac. Guidelines for the selection of weighting functions for h-infinity control. *Naval Surface Warfare Center, Dahlgren Division, Weapons Systems Department*, NSWCDD/MP-92/43, AD-A251 781 (92-16379), January 1992.