

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Modeling and Control of an Ornithopter for Non-Equilibrium Maneuvers

Permalink

<https://escholarship.org/uc/item/22r0f1jr>

Author

Rose, Cameron Jarrel

Publication Date

2015

Peer reviewed|Thesis/dissertation

Modeling and Control of an Ornithopter for Non-Equilibrium Maneuvers

by

Cameron Jarrel Rose

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Engineering - Electrical Engineering and Computer Sciences

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Ronald S. Fearing, Chair

Professor Pieter Abbeel

Professor Robert Dudley

Fall 2015

Modeling and Control of an Ornithopter for Non-Equilibrium Maneuvers

Copyright 2015
by
Cameron Jarrel Rose

Abstract

Modeling and Control of an Ornithopter for Non-Equilibrium Maneuvers

by

Cameron Jarrel Rose

Doctor of Philosophy in Engineering - Electrical Engineering and Computer Sciences

University of California, Berkeley

Professor Ronald S. Fearing, Chair

Flapping-winged flight is very complex, and it is difficult to efficiently model the unsteady airflow and nonlinear dynamics for online control. While steady state flight is well understood, transitions between flight regimes are not readily modeled or controlled. Maneuverability in non-equilibrium flight, which birds and insects readily exhibit in nature, is necessary to operate in the types of cluttered environments that small-scale flapping-winged robots are best suited for. The advantages of flapping wings over quadrotors and fixed-wing fliers are realized in the ability to transition from forward flight to hover to gliding. Flight in the transitions between these regimes necessitates the development of novel modeling techniques and online control techniques to accurately complete these types of maneuvers.

In this thesis, methods for modeling and controlling the transitions between takeoff and diving maneuvers are developed for a flapping-winged micro aerial vehicle (MAV), the H²Bird. To transition into takeoff and steady state flight, a cooperative launching system is developed for the H²Bird by carrying it on the back of a 32 gram hexapedal millirobot, the VelociRoACH. The necessary initial velocity and pitch angle are determined for take off using force data collected in a wind tunnel, and the VelociRoACH is used to reach these initial conditions for successful launch. The models for the diving maneuver are generated using an automatic piece-wise affine identification technique. The flight conditions during the maneuver are segmented into separate regions and least-squares is used to estimate affine linear models for each modeling region. These models are used to compute the reachability sets for the recovery conditions for safe diving, and linear quadratic regulator controllers are used to maintain stable conditions before and after the dive. The data-driven automatic modeling techniques and controller design processes can be extended to additional flight maneuvers.

To my grandfather

For supporting me with words of wisdom and encouragement along the way, and always letting me know how proud of me he was. It meant a lot to me.

Contents

Contents	ii
List of Figures	iv
List of Tables	viii
1 Introduction and Background	1
1.1 Introduction	1
1.2 Background	4
2 Hardware and Robotic Platform	6
2.1 Ornithopter Platform	6
2.2 Aerodynamics	7
3 Comparison of Wind Tunnel Force Measurements with Free Flight	10
3.1 Flight Data Collection	10
3.2 Comparison of Data Sets	14
4 Takeoff and Flight Transition	21
4.1 Robotic Platforms and Behaviors	21
4.2 Control and Launching	23
4.3 Experimental Results and Discussion	24
5 Piece-wise Linear Modeling for Diving	35
5.1 Piece-wise Affine Modeling	35
5.2 Reachability Analysis	41
6 Online Control for Diving	44
6.1 Control Implementation	44
6.2 Experiments and Discussion	48
7 Conclusions	53

A Reachability Sets	55
A.1 Position Reachability Sets	55
A.2 Velocity and Pitch Angle Reachability Sets	58
Bibliography	60

List of Figures

1.1	The launch sequence of the H ² Bird from a cradle on the back of the VelociRoACH.	2
1.2	The H ² Bird dive sequence. Label 1 indicates the point at which the robot transitions from level flight to the unpowered dive, label 2 indicates the transition between the unpowered dive to the powered recovery, label 3 indicates the lowest point in the dive and the transition between the recovery back to level flight, and label 4 indicates flight to a new height.	2
2.1	The H ² Bird ornithopter [17].	6
2.2	The free body diagram for the wind tunnel data [4].	7
2.3	H ² Bird mounted to sensor	8
2.4	Aerodynamic horizontal force surface in world coordinates.	9
2.5	Aerodynamic vertical force surface in world coordinates.	9
2.6	Aerodynamic pitch moment surface.	9
2.7	Thrust (blue) and lift (green) forces as a function of the duty cycle.	9
2.8	Pitch moment as a function of the duty cycle.	9
3.1	Block diagram of the H ² Bird control system.	11
3.2	Diagram of the H ² Bird mounted to the force-torque sensor in the wind tunnel [36].	12
3.3	Free body diagram of the H ² Bird for the wind tunnel experiment data.	13
3.4	A sample plot of the vertical force surface measured in the wind tunnel as a function of angle of attack and wind speed.	14
3.5	The range of the change in pitch moment that the elevator can achieve for 80 percent duty cycle for different angles of attack and wind speeds. For example, at 40° angle of attack and 2.5 m/s wind speed, the elevator can affect a maximum change of about 1.1 N*m pitch moment through its entire range.	15
3.6	Side and top views of a sample flight path with start point (green square) and stop point (red circle) in the tracking space. The black bar indicates the target path.	16
3.7	The pitch, pitch velocity, elevator input, and velocity magnitude of the H ² Bird during one trial.	17
3.8	Block diagram of the estimation of equilibrium points from the wind tunnel data.	18

3.9	Equilibrium points measured in free flight (red squares) and equilibrium points predicted from the wind tunnel (blue triangles).	19
4.1	Free body diagram for the H ² Bird.	22
4.2	Net lift over a range of angles of attack and wind speeds at 16 Hz flap speed. The dashed black line indicates the line of zero net vertical force. Above the line are feasible conditions for takeoff and infeasible conditions are below.	22
4.3	VelociRoACH with cradle and H ² Bird ornithopter MAV (top), and launch sequence from left to right (bottom).	23
4.4	The launch cradle on top of the VelociRoACH, highlighted by the red dashed line. The carbon fiber spars through the back of the cradle are shown in the top-left inset.	24
4.5	Telemetry data at the start of running for a single launch trial.	25
4.6	Telemetry data around launch for a single launch trial. The red line indicates the launch point.	26
4.7	Launch experiments for varied running speeds and launch angles. The shaded region represented the wind tunnel predicted failure area, and the unshaded region represents the predicted success area. The red double triangles represent failures in the predicted success region.	27
4.8	The change in pitch angle 0.2 seconds post-launch for each tested velocity (left) and the elevator input at launch for each tested velocity (right). The red double triangles represent failures in the predicted success region.	28
4.9	The change in pitch angle vs. the elevator input for each trial. The red double triangles represent failures in the predicted success region.	29
4.10	Telemetry data for a single trial for the VelociRoACH running alone (left) and running with an inertial mass equivalent to the H ² Bird(right).	30
4.11	Telemetry data for a single trial for the VelociRoACH running with the passive H ² Bird (left) and running with the H ² Bird flapping at 5 Hz (right).	31
4.12	Roll velocity variance (left) and pitch velocity variance (right) for the VelociRoACH alone, with inertial mass, with passive H ² Bird, and with active H ² Bird.	32
4.13	Average running velocities (left) and average power consumed (right) at steady state for VelociRoACH alone, with inertial mass, with passive H ² Bird, and with active H ² Bird.	32
5.1	Free body diagram of the H ² Bird for the relevant state variables for the discrete-time models [4].	36
5.2	Position data for a single trial of the open loop H ² Bird diving experiment. The gray shading represents the unpowered portion of the dive. The black marker indicates the conditions at the lowest vertical position in the dive.	37

5.3	Telemetry data for a single trial of the open-loop H ² Bird diving experiment over the time period. The position is on the left and the horizontal velocity, vertical velocity, and pitch angle are on the right. The gray shading represents the unpowered portion of the dive. The black marker indicates the conditions at the lowest vertical position in the dive.	37
5.4	The means of the conditions for the lowest point in the open loop dive trials. The error bars represent one standard deviation above and below the mean for each state variable.	38
5.5	Graphical representation of the K-means segmentation of the dive (left) and escape (right) portions of the open loop experiments in the horizontal velocity, vertical velocity, and pitch angle space.	39
5.6	The terminal polytopes for the backwards reachability analysis. The set in the horizontal and vertical position space is on the left, and the set in the horizontal velocity, vertical velocity, and pitch angle space is on the right. The goal position is always zero and all other heights are relative to the goal.	42
5.7	The backwards reachable polytopes 16 timesteps, or 0.4 seconds, from the terminal set in the horizontal and vertical position space. Letters A - C in the left image correspond to the positions in the trajectory marked by letters A - C in the right image.	42
5.8	The backwards reachable polytopes 16 timesteps, or 0.4 seconds, from the terminal set in the horizontal velocity, vertical velocity, and pitch angle space.	43
6.1	Diagram of the control sections of the entire diving maneuver.	44
6.2	Controller block diagram of the diving maneuver.	45
6.3	Telemetry data for a single trial of the open-loop H ² Bird level flight. The position is on the left and the inputs to the wing motor and elevator are on the right.	46
6.4	The horizontal velocity, vertical velocity, and pitch angle for a single trial of the open-loop H ² Bird level flight.	47
6.5	Block diagram of the level flight controller.	47
6.6	Position data for a single trial of the H ² Bird diving experiment with the recovery controller active. The gray shading represents the unpowered portion of the dive. The black marker indicates the conditions at the lowest vertical position in the dive.	49
6.7	Telemetry data for a single trial of the H ² Bird diving experiment with the recovery controller active over the time period. The position is on the left and the horizontal velocity, vertical velocity, and pitch angle are on the right. The gray shading represents the unpowered portion of the dive. The black marker indicates the conditions at the lowest vertical position in the dive.	49
6.8	Number of successful trials (green) and failed trials (red) for each tested dive distance.	50
6.9	The total dive distance for each successful trial for the 1 meter and 1.5 meter desired dive distance set points.	51

6.10	The vertical velocity vs. the pitch angle at the time of recovery for successful (green circle) and failed (red square) trials.	51
6.11	The pitch angle at the top of the stall and the time to a downward pitch angle during the unpowered dive for successful (green circle) and failed (red square) trials.	52
A.1	The backward reachability polytopes in the position space 1 to 4 timesteps from the terminal set. The timestep used was 0.025 seconds.	55
A.2	The backward reachability polytopes in the position space 5 to 8 timesteps from the terminal set. The timestep used was 0.025 seconds.	56
A.3	The backward reachability polytopes in the position space 9 to 12 timesteps from the terminal set. The timestep used was 0.025 seconds.	56
A.4	The backward reachability polytopes in the position space 13 to 16 timesteps from the terminal set. The timestep used was 0.025 seconds.	57
A.5	The backward reachability polytopes in the velocity and pitch angle space 1 to 4 timesteps from the terminal set. The timestep used was 0.025 seconds.	58
A.6	The backward reachability polytopes in the velocity and pitch angle space 5 to 8 timesteps from the terminal set. The timestep used was 0.025 seconds.	58
A.7	The backward reachability polytopes in the velocity and pitch angle space 9 to 12 timesteps from the terminal set. The timestep used was 0.025 seconds.	59
A.8	The backward reachability polytopes in the velocity and pitch angle space 13 to 16 timesteps from the terminal set. The timestep used was 0.025 seconds.	59

List of Tables

3.1	Table of test parameters used for measurements collected in the wind tunnel. . .	11
3.2	Table of free flight measured equilibrium flight points and wind tunnel predicted thrust and lift forces and pitch moment. Ideally, the net thrust force, net lift force, and pitch moment should all be zero at the free-flight measured equilibrium points.	20
4.1	Table of percent increases and decreases of measured data for the active and passive H ² Bird cases over the case of the VelociRoACH running by itself.	30
4.2	Table of cost of transport for the VelociRoACH alone, the H ² Bird alone, and the active and passive cooperative cases.	34
5.1	Table of the Mahalanobis distance for 1 through 10 regions for the dive and escape portions of the maneuver. The Mahalanobis distance for the Null model is in column 1 for comparison.	41

Acknowledgments

I would like to thank my research advisor, Professor Ron Fearing, for going above and beyond his responsibilities as my PhD advisor. His insight, patience, and encouragement helped to make the research process a lot easier over the past six years.

I would also like to thank the other members on my PhD committee, Professor Robert Dudley and Professor Pieter Abbeel. They were both very helpful in providing feedback on my research plans, other perspectives outside of my particular area of robotics, and guidance. Professor Robert Dudley, was especially instrumental in the beginning of my PhD, allowing me to use the wind tunnel in his lab to complete all of my Master's work.

A huge thank you to the members of the Biomimetic Millisystems Lab for being excellent colleagues and making my experience at UC Berkeley interesting and worthwhile. I especially would like to thank Dr. Kevin Peterson, who provided invaluable advice and guidance when I first joined the lab. A big shout out to Austin Buchan, Duncan Haldane, Ethan Schaler, Jessica Lee, Andrew Pullin, Carlos Cesarez, Dr. Andrew Gillies, Dr. Nick Kohut, and Jane Esterline for making my time both inside and outside the lab enjoyable.

I would also like to thank my roommates, Dan Calderone, Ethan Schaler, and Jaakko Karras for the hilarious late-night conversations, and all of the friends I have made in and outside of Berkeley that have kept me sane.

Last, but not least, I would like to thank my family. My parents provided a lot of support for me during my time in Berkeley. They were also very understanding during the times that I did not call when I should have and missed countless events. They told me throughout the entire process that they were proud of me and provided the strength that I needed to finish.

Chapter 1

Introduction and Background

1.1 Introduction

Bio-inspired flapping-winged robotics is a rapidly growing field that is of interest to roboticians as well as biologists. In nature, birds and insects display high levels of maneuverability beyond the levels of current flapping-winged micro aerial vehicles (MAVs). For example, the Common Swift (*Apus apus*) is a remarkable glider, capable of sustained flight for up to a year at a time while roosting, migrating, and foraging and spends the vast majority of its life in the air [23]. The Peregrine Falcon (*Falco peregrinus*) is capable of high-performance flight, exceeding 88 m/s during dives and 40 m/s during horizontal flight [33]. Finally the hummingbird is able to achieve very high levels of maneuverability and flight performance. A study on Anna's hummingbird explored the effects of various sizes of vortex wakes at different wind speeds on the flight performance [29]. The hummingbirds are able to adjust their body posture, in addition to the flapping frequencies and amplitudes of their wings rapidly during flight to account for the turbulent flow. Maneuverability in nature is not only related to wing structure and wing movements, it is also dependent upon careful balancing of the forces and moments, using the body of the animal to achieve the levels of rotational agility observed [11]. The design of mechanisms to mimic those in nature is very difficult, especially regarding the weight and power constraints of robotic fliers. Flapping-winged fliers are often under-actuated due to these constraints. For this reason, it is important to design controllers for these systems to achieve the levels of actuation desired using minimal actuation. The main focus of this work is to determine how to design models and controllers for the transitional periods of non-equilibrium flight for an under-actuated flapping-winged MAV.

The aerodynamics of flapping-winged flight are nonlinear and complex, and are difficult to model. The flapping of the wings creates an unsteady airflow around the control surfaces of the flier, increasing the complexity of the aerodynamics associated with the surfaces [25][1][24][7]. An understanding of the behavior of these fliers in free flight is necessary for successful control. Modeling the aerodynamics and dynamics of flapping-winged fliers

is computationally intensive, so performing predictive control online is not possible on an autonomous millirobotic platform due to computational limitations. Therefore, a controller algorithm that can be updated at a high sampling rate is necessary.




Figure 1.1: The launch sequence of the H²Bird from a cradle on the back of the VelociRoACH.




Figure 1.2: The H²Bird dive sequence. Label 1 indicates the point at which the robot transitions from level flight to the unpowered dive, label 2 indicates the transition between the unpowered dive to the powered recovery, label 3 indicates the lowest point in the dive and the transition between the recovery back to level flight, and label 4 indicates flight to a new height.

While many researchers have developed various solutions for flapping flight mechanisms, there has been little work in controlling these fliers. The limitations in modeling and con-

trol of the unsteady aerodynamics prevent realization of the full maneuverable space of the fliers. In particular, non-equilibrium flight transitions are difficult to model and control, as they represent regions of flight in which the aerodynamics are the most complex. Some examples of non-equilibrium flight are exhibited in Figures 1.1 and 1.2. In Figure 1.1, a flapping-winged MAV, the H²Bird is being launched into the air by a robotic hexapod, the VelociRoACH. The H²Bird makes the transition from stationary, to moving forwards on top of the VelociRoACH, to flapping on top of the VelociRoACH, and finally to disengaging from the VelociRoACH into free flight. In the second scenario in Figure 1.2, the H²Bird transitions from forward, level flight to an unpowered dive at (1), to a powered recovery at (2), and then back to level flight at (3) and in (4). The transitions in the dive, the second scenario, are much different in that some foresight is necessary to decide when the recovery transition must occur to ensure safety and to minimize the final distance from the target height. Additionally, the level flight portions of the maneuver require a controller to maintain a certain height. In the takeoff scenario, the only requirements are that the H²Bird reach a certain speed before takeoff. If the VelociRoACH continues to run after this point, there are no adverse effects on performance. To achieve the types of maneuvers such as dives, sharp turns, and rapid transition between maneuvers similar to birds and insects in nature, advancements in control of flapping winged flight are necessary. This thesis presents modeling and control methods to complete transitions between different flight regimes using the two aforementioned scenarios as examples.

Millirobotic systems with mass on the order of 10 grams and centimeter-scale linear dimensions can provide advantages in field applications such as the navigation of collapsed buildings and other cluttered environments, over large-scale robots. In addition to the portability and lower mass of millirobots, they can also enter spaces that large-scale robots cannot. These spaces are typically found in emergency services applications such as earthquakes and collapsed buildings. Millirobots can also be constructed out of cheaper, lighter materials since the structures do not have to support as much mass as their larger counterparts. As robots provide benefits in the exploration of disaster environments dangerous to humans, reduced cost is important in the case of hardware loss. Due to the nature of these environments, however, it is important that the millirobots possess a high level of maneuverability to avoid obstacles in the confined space. High levels of maneuverability are dependent on both intelligent hardware design and effective controllers.

Flapping-winged millirobots for these types of applications require the following contributions presented in this work:

1. Determine the aerodynamic capacities of a flapping-winged millirobot.
2. Design a method of transitioning from docking to takeoff.
3. Design computationally tractable models for online control.
4. Demonstrate a method for control that can achieve maneuverability away from equilibrium.

1.2 Background

Dynamic and Aerodynamic Modeling

One method of understanding the flight behavior of these fliers involves modeling the wing motion during each wing stroke, e.g. using blade element theory [6][16][32]. Another method involves multi-body modeling to account for the changing mass distribution as the flier flaps its wings [5][13][28].

We desire a simplified representation of the aerodynamics which can be used for on-board model-based control in 10 gram scale fliers. The previous control strategies used for ornithopters of this scale were based upon Proportional-Integral-Differential (PID) control schemes for target tracking or height regulation [3][4]. Information about the aerodynamic interactions of the ornithopter can produce a more robust model that can be used in a more sophisticated control scheme such as Linear-Quadratic Regulator (LQR).

For this aerodynamic modeling, it is necessary to collect data that can be used to estimate the behavior of the MAV over time. Common approaches to develop these models include averaging the flight behavior over the wing beat period, and using linear low-dimensional models to predict the flight behavior over time. It has been shown that using time averaged aerodynamic data is a valid approximation over a given wing stroke [20][38]. Wind tunnels are often used to measure aerodynamic properties of robotic fliers. Although some of the degrees of freedom are constrained by the mounting mechanism, this method of aerodynamic force and moment measurement has been used by previous researchers for the purpose of developing aerodynamic models for simulation and control. A simulation of insect flight for the Robofly project uses aerodynamic models based upon wind tunnel measurements [10][37]. In addition, mounted sensor measurements have also been used to measure the aerodynamic properties of wings for modelling by Khan and Agrawal [19]. Lee and Han recently implemented a non-contact magnetic suspension and balance system to control the attitude of an experimental model [22]. This setup allows for data collection and controller verification using selected degrees of freedom.

An alternative method to using stationary wind tunnel data for modeling is collecting free flight telemetry data using a motion tracking system. Grauer et al. utilized this method to create a dynamic model for their "Slow Hawk" ornithopter [14]. The flight data collected was used to fit parameters to a multi-body dynamic model, and wind tunnel tests were used to determine the associated aerodynamic model. This method, however, involves the fitting of many model parameters. Faruque et al. also used system identification to generate linear models of the pitch and roll dynamics of their insect-inspired micro-flier [12]. Their models were of a lower complexity than Grauer et al., and they used the models to inform their on-board stabilizing avionics package.

Multi-modal Locomotion

Some researchers have pursued the design of robots that can traverse an environment using a variety of methods of locomotion. In particular, we are interested in a subset of these multi-modal robots that can navigate both terrestrial and aerial environments [21] [18]. Two such robots are the BOLT [30], designed by Peterson et. al, and the MALV II [2], by Bachmann et. al. Both robots use legs on the front of robotic fliers to overcome obstacles taller than the robot's height. Both robots can also transition from legged running to flying. The BOLT [30] uses small running legs to reach takeoff speed, but there is no way to run without flapping the wings, as the same motor drives both. Aerodynamic surfaces on legged robots can also provide benefits in terms of stability and running speed as demonstrated by Peterson et. al with Dash+ Wings [31] and Haldane et. al with the VelociRoACH [15]. In spite of these benefits and the versatility of multi-modal robots, it is difficult for multi-modal robots to excel at both modes with a single design, due to power and weight constraints.

Online Control

Researchers have implemented different methods of controlling various behaviors of flapping-winged MAVs. One approach that is often used is Proportional-Integral-Differential (PID) control for height regulation or path following. Baek et al. used on-board orientation estimation along with a camera to seek and fly towards an LED using their flapping-winged MAV, the iBird [4][3]. They predicted the position of the LED using a Kalman filter, and controlled the yaw and height of their flier using a PID controller.

Ma et al. use mixed-model based control methods for height regulation of a robotic fly, developed by Wood [40][26]. They use a PD controller to regulate the attitude of the flier, and calculate the angular reference to the attitude controller using the lateral position reference error. The altitude controller is designed using the linearized dynamics of the flier about hover. The controllers were decoupled to reduce the constraints on the sensitive attitude and lateral position controllers.

A vision-based approach is utilized by the researchers at the Delft University of Technology with their flapping flier, the Delfly II [9]. de Croon et al. use a camera mounted on the front of the Delfly to compute optical flow and texon histograms for texture detection [8]. Both methods are used to estimate the time to collision with an obstacle. A human controls the height, and open loop yaw inputs in the vision loop are used for obstacle avoidance. In this case, the vision inputs determine the necessary control, rather than a specific model or controller.

Another approach was used by Moore and Tedrake to control the perching behavior of a fixed wing aircraft [27]. The researchers used LQR Trees, a control method that Tedrake developed [39]. They designed a feedback controller using trajectory optimization and local linear feedback. The set of controllers over the local trajectory regions are computed within regions of attraction branching from the space of initial conditions. The regions of attraction are computed using time-varying Lyapunov functions.

Chapter 2

Hardware and Robotic Platform

2.1 Ornithopter Platform




Figure 2.1: The H²Bird ornithopter [17].

The robotic platform used is a flapping-winged MAV known as the H²Bird, shown in Figure 2.1 [17]. The H²Bird has a custom built carbon fiber frame, carbon fiber reinforced clap-fling wings, and carbon fiber reinforced tail, and uses the Silverlit i-Bird RC flier power train¹. The wingspan of the H²Bird is 26.5 cm, the length from front to back is 20 cm, and its mass is 13.6 grams. Yaw and pitch control are provided by a tail-mounted propeller and servo-controlled elevator, respectively. For control and sensing, the H²Bird uses an on-board ImageProc 2.4² controller that includes a 40 MIPS microprocessor, 6 DOF IMU, IEEE 802.15.4 radio, and motor drivers, powered by a 90 mAH lithium polymer battery [3].

The attitude estimation and control of the H²Bird are both performed on-board, and is computed at 300 Hz. To estimate the pose of the H²Bird, the angular rate values measured by

¹Silverlit Toys Manufactory Ltd.: i-Bird RC Flyer <http://www.silverlit-flyingclub.com/wingsmaster/>

²ImageProc 2.4:

https://github.com/biomimetics/imageproc_pcb

the on-board gyroscope are integrated over time. Separate proportional-integral-derivative (PID) controllers use the estimated pose and the desired pitch (θ) and yaw angles, provided by an external source, to determine the necessary control surface inputs to the elevator servo and tail propeller motor to achieve the desired pose. Another PID controller is used to regulate the flap frequency of the H²Bird, with the desired frequency provided by the external source, and flap frequency estimated by a Hall effect sensor on an output gear of the transmission, as inputs.

2.2 Aerodynamics

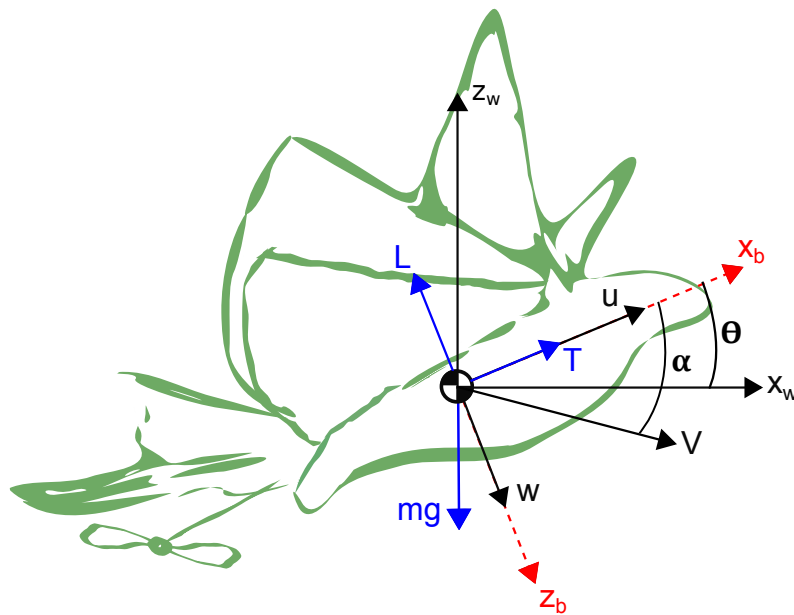



Figure 2.2: The free body diagram for the wind tunnel data [4].

To determine the aerodynamic capabilities of the lifting surfaces on the H²Bird, we collected data in a wind tunnel for various angles of attack, wind speeds, and flap frequencies. The dimensions of the wind tunnel working enclosure are 45.5 cm x 45.5 x 91.5 cm, and a picture of the robot in the enclosure can be found in Figure 3.2 for reference. The free body diagram of the H²Bird is in Figure 2.2. The red axes represent the body coordinate frame and the black axes represent the world frame. The lift and thrust indicated by the blue vectors and marked by L and T are expressed in the body frame. The pitch angle θ is the angle between the world frame and the body frame. The angle of attack α is the angle between the velocity vector, V , and the body frame. The total horizontal force and the total vertical force, referenced in Figures 2.4 and 2.5 are both expressed in the world frame.

The wind tunnel data were collected over a series of trials using an ATI Nano17 force-torque sensor. The ornithopter was affixed to an acrylic mount attached to the sensor and

Figure 2.3: H²Bird mounted to sensor

facing into the wind, as shown in Figure 3.2. The ornithopter was attached at approximately its center of mass, although this point fluctuates as the wings open and close. This fluctuation is minimal, however, so it was discounted. Each trial consisted of 3 seconds of data collected for wind speeds between 0 and 3 m/s in 0.5 m/s increments and longitudinal body axis angles relative to the horizontal, or angles of attack, between -70 and +90 degrees in 10 degree increments. Each of these trials was conducted with the wings of the robot closed and stationary. Each trial was averaged over the 3 seconds of collection, and only the thrust and lift forces in body coordinates and the pitch moment were stored. The results of the experiments are in Figures 2.4 through 2.6. Each force is expressed in the world reference frame with the weight of the robot absent from the measurements. The forces in the figures are measured with the wings closed and not flapping.

In addition to the stationary wing experiments, data were collected in still air for duty cycles between 50 percent and 100 percent in 10 percent increments using the same wind tunnel setup as the previous experiments. This range corresponds to wing flap speeds between 12 Hz and 20 Hz. These experiments required some damping around the connection of the robot to the acrylic mount with 4 mm of foam, as the oscillations from the wings moving created significant noise in the pitch moment signal. The data collected from the experiments were averaged over 5 seconds, and the thrust and lift forces in body coordinates and the pitch moment were stored. The results of the experiments are in Figures 2.7 and 2.8.




Figure 2.4: Aerodynamic horizontal force surface in world coordinates.




Figure 2.5: Aerodynamic vertical force surface in world coordinates.

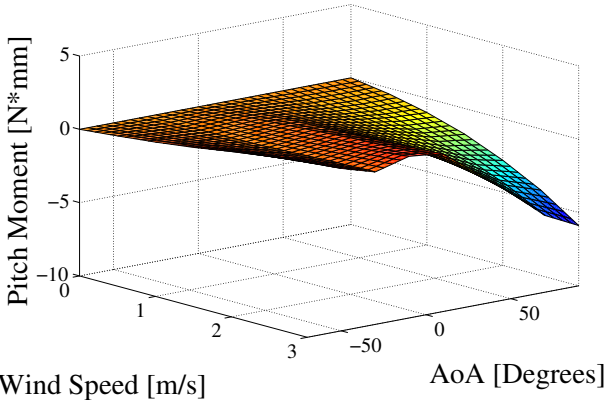


Figure 2.6: Aerodynamic pitch moment surface.

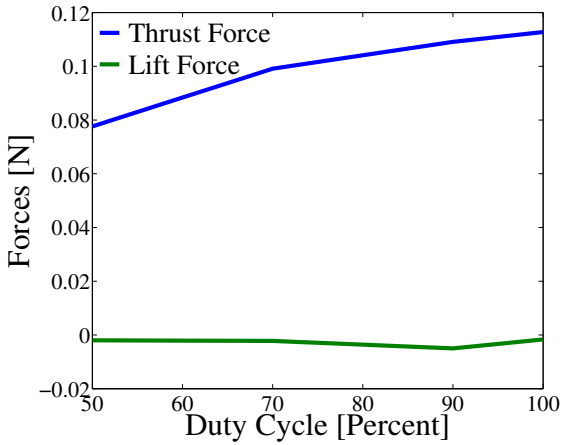


Figure 2.7: Thrust (blue) and lift (green) forces as a function of the duty cycle.

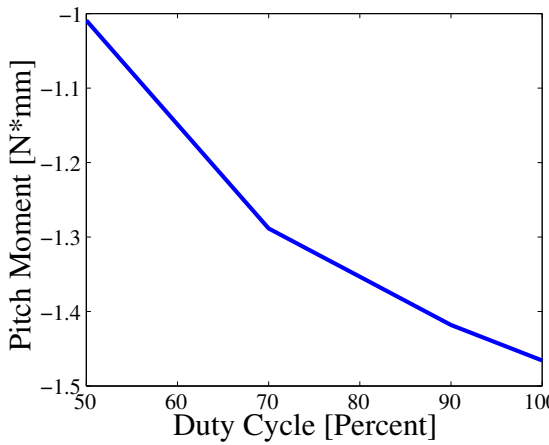


Figure 2.8: Pitch moment as a function of the duty cycle.

Chapter 3

Comparison of Wind Tunnel Force Measurements with Free Flight

3.1 Flight Data Collection

We desire a simplified representation of the aerodynamics which can be used for on-board model-based control for the H²Bird. The previous control strategies used for ornithopters of this scale were based upon Proportional-Integral-Differential (PID) control schemes for target tracking or height regulation [3] [4]. Information about the aerodynamic interactions of the ornithopter can produce a more robust model that can be used in more sophisticated control schemes such as Linear-Quadratic Regulator (LQR) or Model Predictive Control (MPC).

For this aerodynamic modeling, it is necessary to collect data that can be used to estimate the behavior of the MAV over time. We utilize two different methods to collect this data. For the first method, we use a wind tunnel to determine the forces and torques experienced by the H²Bird at different angles of attack, wind speeds, and flapping frequencies. For the second method, we use a motion tracking system to collect free-flight data at equilibrium. We compare both data sets under the same flight conditions to determine if the wind tunnel data can be used to accurately predict the free flight conditions of the ornithopter for simulation models.

Wind Tunnel

To determine the forces and moments that the ornithopter experiences in flight, wind tunnel data were collected over a series of trials using an ATI Nano17 force-torque sensor. The ornithopter was affixed to an acrylic mount attached to the sensor and facing into the air stream, as shown in Fig. 3.2. A 3 mm piece of foam was placed between the acrylic mount and the H²Bird to provide damping for the high frequency oscillations that the flapping of the wings causes in the pitch moment signal. The H²Bird was attached at approximately its center of mass, although this point fluctuates as the wings open and close. This fluctuation

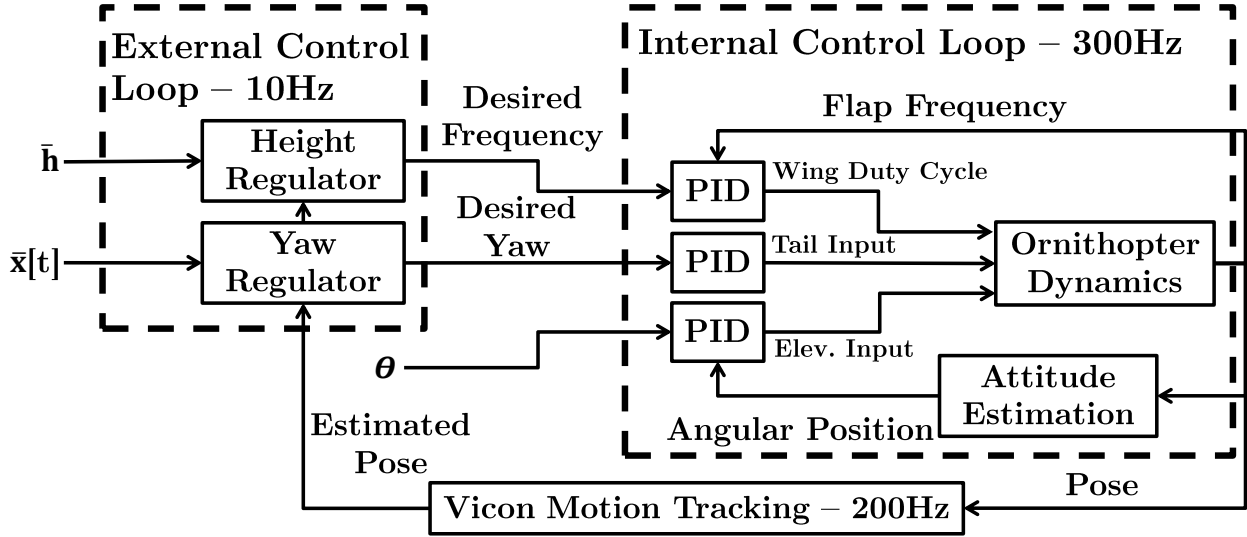


Figure 3.1: Block diagram of the H²Bird control system.

Variable	Test Parameters
Duty Cycle [Percent]	75,80,90,100
Angle of Attack [deg]	25,35,43,50,60
Wind Speed [m/s]	1.5,2.0,2.5
Elevator Deflection [deg]	-10,-6,0,8,19,30,35

Table 3.1: Table of test parameters used for measurements collected in the wind tunnel.

is minimal, however, so it was discounted in the data collection. For each trial, force and moment data were collected over a period of 7 seconds at all combinations of wing duty cycles, wind speeds, angles of attack, and elevator deflections shown in Table 3.1, a total of 420 data sets. The wing motor duty cycle corresponds to flapping frequencies between 14 and 20 Hz, and a positive elevator deflection angle corresponds to an upward pitch. Each trial was averaged over the 7 seconds.

The collected data form a data set with 4 inputs and 3 outputs. The outputs are the lift force (L) and thrust force (T) in body coordinates, and the pitch moment (M), which are dependent upon the inputs of the duty cycle, angle of attack (α), wind speed (V), and elevator deflection, the directions of which are shown in the free body diagram in Fig. 3.3. In the diagram, the pitch angle, θ , is the angle between the horizontal in world coordinates and the x-axis, x_b , of the body of the H²Bird, whereas the angle of attack, α , is the angle between the velocity vector and the x-axis of the body. This data set is used as a look-up table for the instantaneous forces and moments for a given pose. The collected data form a series of surfaces similar to Fig. 3.4, which is the surface for a wing duty cycle of 80 percent and elevator deflection of 24 degrees. There are similar surfaces for each duty cycle and




Figure 3.2: Diagram of the H²Bird mounted to the force-torque sensor in the wind tunnel [36].

elevator deflection, and linear interpolation is used to estimate values between the measured data points. Fig. 3.5 shows a representation of the aerodynamic effect of the elevator as a function of angle of attack for increasing wind speeds. Each data point is the amplitude of the pitch moment provided by the elevator for a given operating point. As expected, the range of moments increases with increasing wind speed. The plot illustrates the control authority of the elevator available to influence the pitch of the ornithopter.

Free Flight

The free flight experiments were conducted over variable amounts of time during approximately straight and level flight of the H²Bird. Before each experiment, the H²Bird was launched by hand and directed to follow the path shown in Fig. 3.6 using the external control loop of Fig. 3.1. This desired path allows the completion of several experiments during a flight, and ensures that there is a straight and level portion of flight time in which to record a data set. A Vicon motion tracking system¹ was used to track the position of the H²Bird at 200Hz, and this translational information was used with the desired reference trajectory, $x[t]$, shown in black in Fig. 3.6 as the input to a PID controller that computes the yaw angle necessary to maintain flight on the target path (black bar). A second PID controller was used to regulate the height of the H²Bird at a constant height input, h , of 1.5 meters by computing the necessary flap frequency to maintain level flight. Throughout the reference

¹Vicon Motion Systems: <http://www.vicon.com>




Figure 3.3: Free body diagram of the H²Bird for the wind tunnel experiment data.

path, the robot was directed to maintain a pitch angle of 35 degrees. The commanded angles and flap frequency were then transmitted to the robot at 10 Hz and the internal controllers on the H²Bird moved the robot to the correct pose. When the H²Bird reached the ends of the target path, it was directed to execute a 180 degree turn.

Each experiment consisted of a step from an initial pitch angle of 35 degrees to 50 degrees at 80 percent duty cycle. Each trial was conducted during the straight portion of the reference trajectory to minimize the effect of yaw and roll on the data. While straight and level flight was desired, some deviation occurred, although only trials with decidedly minimal disturbances were used in the data set. During the trials, the telemetry data, including the angular position, gyro values, and control motor inputs, were stored in the flash memory on the H²Bird at 300 Hz. Additionally, the translational position and velocity, angular position, directed angles, and commanded flap frequency were stored from the Vicon at 200 Hz. The data for one trial are shown in Fig. 3.7, where the pitch angle, pitch velocity, and elevator deflection are estimated on the H²Bird and the translational velocity is measured by the Vicon.




Figure 3.4: A sample plot of the vertical force surface measured in the wind tunnel as a function of angle of attack and wind speed.

3.2 Comparison of Data Sets

Equilibrium Point Estimation

We compare the free flight data set to the wind tunnel data set by determining the equilibrium flight points for both. The equilibrium points are flight conditions that satisfy the following criteria:

$$\begin{aligned}
 L_w &= mg \\
 T_w &= D \\
 \tau_P &= 0
 \end{aligned} \tag{3.1}$$

where

$$\begin{aligned}
 L_w &= T \sin \theta - L \cos \theta \\
 T_w &= T \cos \theta + L \sin \theta
 \end{aligned}$$

where m is the mass of the H²Bird, g is gravity, θ is the pitch angle, and τ_P is the total pitch moment. T is the thrust and L is the lift of the H²Bird in body coordinates, shown in Fig. 3.3. D is the drag force, which balances the thrust in world coordinates, T_w . Since we can only know the forces and moments in free flight at equilibrium flight conditions, we can only directly compare these points to analogous points in the wind tunnel data.




Figure 3.5: The range of the change in pitch moment that the elevator can achieve for 80 percent duty cycle for different angles of attack and wind speeds. For example, at 40° angle of attack and 2.5 m/s wind speed, the elevator can affect a maximum change of about 1.1 N*m pitch moment through its entire range.

To estimate the free flight equilibrium points, we computed the time averaged values for the velocity magnitude, pitch angle, angle of attack, and elevator deflection before and after the step from 35 to 50 degrees in pitch. The transitional portion during the step was not used in the analysis. We conducted 14 total trials, and the free flight equilibrium points are shown in the left half of Table 3.2. Each of these equilibrium flight conditions was then used as the input into the wind tunnel lookup tables to determine the total horizontal and vertical forces in world coordinates and the total pitch moment predicted by the wind tunnel at these flight conditions. The results are in the right half of Table 3.2, and correspond to the error between free flight and the wind tunnel data sets. Ideally, each net force and moment should be zero at equilibrium.

Determining the predicted equilibrium points from the wind tunnel data is a more complicated process, outlined in Fig. 3.8. For a given duty cycle, wind speed, elevator input, and angle of attack there is an associated lift force, thrust force, and pitch moment measured in the wind tunnel. We used each elevator deflection and the 80 percent duty cycle used in the free flight experiments to generate three dimensional surfaces for the net thrust force, net lift force, and net pitch moment, each dependent upon the wind speed and angle of attack. Each of these surfaces are similar to the one in Fig. 3.4. For a given free flight trial, we computed




Figure 3.6: Side and top views of a sample flight path with start point (green square) and stop point (red circle) in the tracking space. The black bar indicates the target path.

the level sets at zero for each of the lift, thrust, and pitch moment surfaces for the particular elevator deflection in the trial. We noticed that, in most of the data, the pitch moment never crosses zero, and therefore, no equilibrium is predicted to exist in pitch. We attribute this problem to an error in the sensor placement, due to the approximation of the center of mass of the ornithopter. While this approximation has minimal effect on the lift and thrust force values, it will affect the magnitude of the pitch moment data. To remedy this problem, we added the mean optimal offset, $\tau_{err} = 0.77 \text{ N}^*\text{mm}$, uniformly over the entire data set to shift the wind tunnel predicted pitch moments for each equilibrium point in free flight as close as possible to zero. This optimal offset is the mean of the pitch moment errors in the eighth column of Table 3.2, and the new values, τ' , are in the ninth column. After this shift, we needed to find the equilibrium points in angle of attack and wind speed space predicted by the wind tunnel, which corresponds to the G^{-1} block in Fig. 3.8. To do this, we solved the optimization problem using the level sets at zero pitch moment:

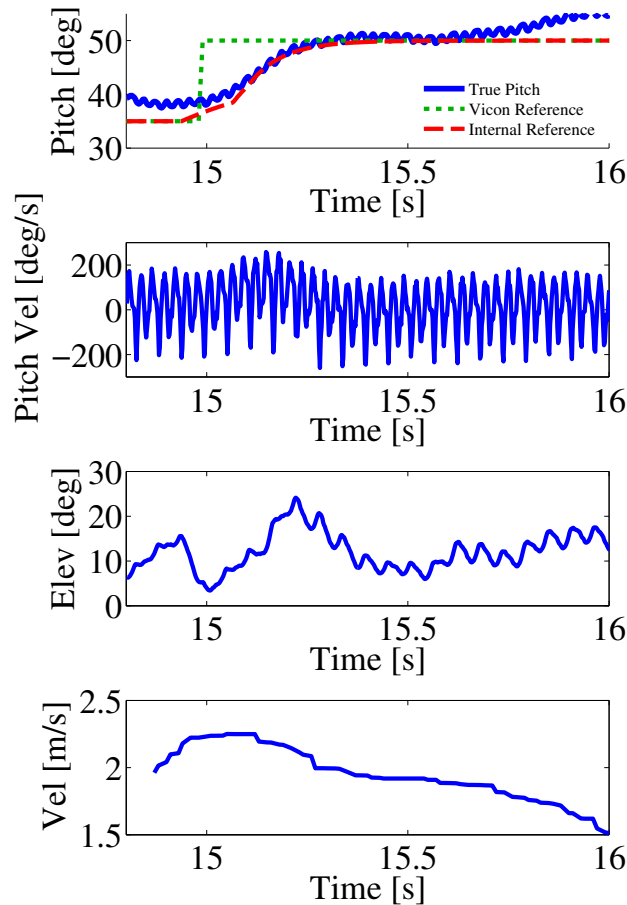


Figure 3.7: The pitch, pitch velocity, elevator input, and velocity magnitude of the H²Bird during one trial.

$$\begin{aligned}
 & \underset{\alpha, v}{\text{minimize}} && F_T(\alpha, v)^2 + F_L(\alpha, v)^2 \\
 & \text{subject to} && \tau' = 0
 \end{aligned} \tag{3.2}$$

where α is angle of attack, v is velocity, F_L is the net vertical force, F_T is the net horizontal force, and τ' is the sum of the pitch moment from the wind tunnel data and the moment offset. The optimal α and v are recorded as the wind tunnel predicted equilibrium point for the given elevator deflection. These optimal values are the angle of attack and velocity on the zero pitch moment level set that minimizes the net vertical and horizontal forces. If any of the zero level sets do not exist for each of the three surfaces, we determine that there is no predicted equilibrium point for that particular operating condition.

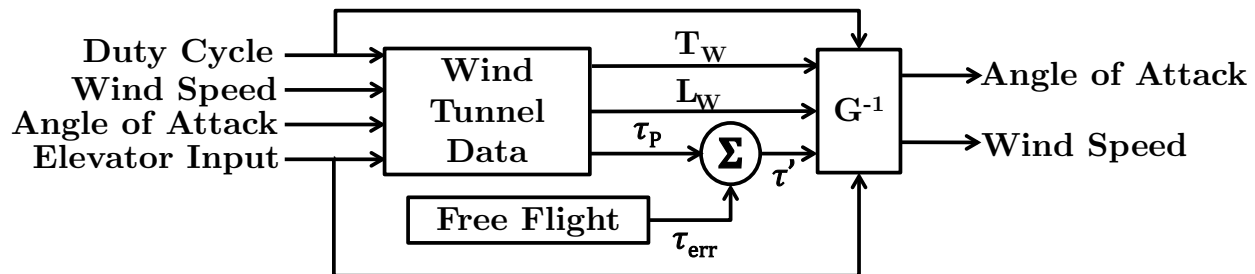


Figure 3.8: Block diagram of the estimation of equilibrium points from the wind tunnel data.

Comparison of Estimations

The end result of the aforementioned process is a set of equilibrium points in velocity and angle of attack for which the wind tunnel predicts a zero pitch moment and minimal net vertical and horizontal force. The results of the estimation are in Fig. 3.9, where the blue triangles represent the wind tunnel predicted equilibrium points for the analogous free flight equilibrium points, represented by red squares, at a particular elevator deflection. There are five free flight points for which the wind tunnel predicts the existence of an equilibrium, and they lie between -2 and 6 degrees elevator deflection.

As shown in Fig. 3.9, the wind tunnel predicted equilibrium points in velocity are much closer to the free flight measurements with an average error of 0.1 m/s, than the predicted equilibrium points in angle of attack, which have an average error of 15 degrees. Numerically, the source of this error is evident in Table 3.2 in the “Net Thrust Force” and “Net Lift Force” columns. Both columns represent the total force in their respective directions and should be zero at equilibrium. The table shows that the wind tunnel underestimates the net thrust force for a given free flight data point, while it overestimates the net lift force for a given data point. With an H²Bird weight of approximately 130 mN, the wind tunnel overestimates the net lift force by an average of 18 percent. Decreasing the angle of attack will decrease the drag and increase the lift caused by the airflow on the H²Bird, moving both the net thrust force and net lift force closer to zero, hence the underestimation of the equilibrium in angle of attack by 15 degrees.

While the error in the wind tunnel predicted equilibrium points is easily explained numerically, the physical reasons for the error are more complex. Since the ornithopter is fixed in the wind tunnel, it is not free to pitch up and down as it does in free flight. The forces and moments caused by these changes in pitch velocity are not captured in the wind tunnel measurements. Additionally, there are high frequency vibrations, caused by the interaction between the frame of the robot and the mounting mechanism from the flapping of the wings, that introduce noise in the wind tunnel measurements, but are not present in free flight.

An additional factor to consider in the differences between the wind tunnel collected data set and free flight is the existence of ground effect within the wind tunnel. Ground effect is the increased lift and decreased drag generated by the wings on an aircraft when it is in

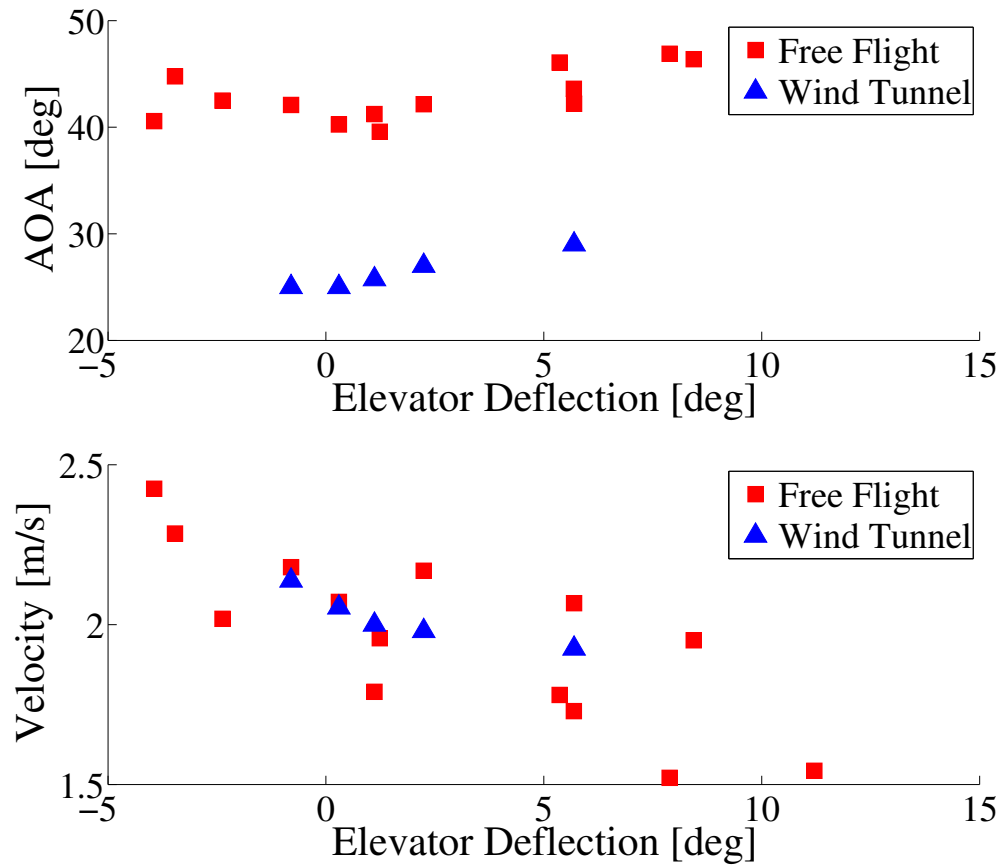


Figure 3.9: Equilibrium points measured in free flight (red squares) and equilibrium points predicted from the wind tunnel (blue triangles).

close proximity to a fixed surface. Ground effect is caused by interruption of formation of the vortices generated by the wings. Since the wind tunnel enclosure is only 1.6 times the wingspan of the H²Bird in width and height, ground effect could be affecting the aerodynamic force measurements in the wind tunnel.

Trial Number	Free Flight Measurements			Interpolated Net Wind Tunnel Values				
	AOA deg	Velocity Mag m/s	Elevator Input deg	Pitch deg	Net Thrust Force mN	Net Lift Force mN	Pitch Moment N*mm	Adj. Moment N*mm
1	41	2.4	-4	46	-57	48	-0.56	0.21
2	40	2.0	-1	58	-63	10	-0.76	0.01
3	42	2.2	2	42	-25	38	-0.50	0.27
4	44	1.7	6	53	-41	13	-0.80	-0.03
5	42	2.1	6	38	-12	31	-0.47	0.30
6	46	1.8	5	52	-48	11	-0.93	-0.16
7	42	2.2	-1	42	-28	39	-0.59	0.18
8	47	1.5	11	50	-37	0	-1.12	-0.35
9	45	2.3	-3	45	-49	47	-0.86	-0.09
10	46	2.0	8	52	-55	17	-0.72	0.05
11	41	1.8	1	42	-11	21	-0.88	-0.11
12	40	2.1	0	41	-24.	29	-0.69	0.08
13	43	2.0	-2	44	-26	30	-0.76	0.01
14	47	1.5	8	53	-43	-2	-1.16	-0.39
			Mean		-37	24	-0.77	0.00
			Max Error		63	48	1.16	0.39
			Min Error		11	0	0.47	0.01

Table 3.2: Table of free flight measured equilibrium flight points and wind tunnel predicted thrust and lift forces and pitch moment. Ideally, the net thrust force, net lift force, and pitch moment should all be zero at the free-flight measured equilibrium points.

Chapter 4

Takeoff and Flight Transition

4.1 Robotic Platforms and Behaviors

We experimented with two robotic platforms: the VelociRoACH, designed specifically for terrestrial locomotion, and the H²Bird, designed specifically for aerial locomotion.

H²Bird

To determine the region of initial conditions in angle of attack and velocity for which the H²Bird can take off, we conducted experimental trials in a 45.5 cm x 45.5 cm x 91.5 cm wind tunnel [35]. The free body diagram in Figure 4.1 shows the distribution of relevant forces for the H²Bird. In the free body diagram, T is the horizontal force in the H²Bird frame, L is the vertical force in the H²Bird frame, m is the mass, g is gravity, and α is the angle of attack. Using a 6DOF force and torque sensor we determined the total vertical force F_v in world coordinates. For each trial, we flapped the wings at the maximum speed of 16 Hz and measured the forces and moments for angles of attack from -60 to 90 degrees in 10 degree increments, and velocities from 0 to 2 m/s in 0.5 m/s increments. Using the force data, we computed the total vertical force, F_v , as follows:

$$F_v = T \sin \alpha + L \cos \alpha - mg \quad (4.1)$$

The results of the experiments are in Figure 4.2, where the dashed black line is the line of zero total vertical force. The operating points above the line are feasible takeoff conditions, and the operating points below the line are infeasible takeoff conditions. The aerodynamics of the H²Bird and the differences between wind tunnel and free flight collected data sets are discussed further in [35].

It is important to note that the H²Bird cannot produce enough lift to take off from the ground at rest. At 90 degrees pitch angle (sitting vertically, with all of the thrust in the upward direction), maximum flap speed, and zero velocity, the net vertical force ($T - mg$) on the H²Bird is approximately -10 mN. This net force indicates that the weight

of the H²Bird is 10 mN larger than its maximum thrust. Therefore, the forward velocity provided by the VelociRoACH is necessary for takeoff.




Figure 4.1: Free body diagram for the H²Bird.

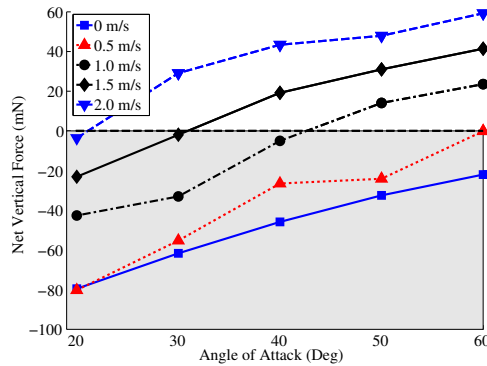


Figure 4.2: Net lift over a range of angles of attack and wind speeds at 16 Hz flap speed. The dashed black line indicates the line of zero net vertical force. Above the line are feasible conditions for takeoff and infeasible conditions are below.

VelociRoACH

The VelociRoACH [15] is a six-legged running robot weighing 30 grams (including 3.7 V, 120 mAh battery) that can run up to 2.7 m/s with a 24 Hz stride frequency. The VelociRoACH uses two 3.3 Ohm coreless, brushed DC motors in separate gear boxes for independently driving the legs, and magnetic Hall effect encoders to regulate the stride frequency and gait phase. On-board, the VelociRoACH has a micro-controller, the ImageProc 2.5¹, which holds a 40 MIPS microprocessor, 6 DOF IMU, IEEE 802.15.4 radio, and motor drivers.

VelociRoACH Dynamic Behavior

The VelociRoACH's gear ratio is 16:1 allowing a 4-24 Hz range of stride frequencies. The separate gear boxes enable differential steering for each set of legs. The fore and aft legs of the VelociRoACH are mechanically constrained 180 degrees out of phase from the middle leg for an alternating tripod gait. For straight-line running, a software controller is used to enforce a 180 degree offset phase between the right and left sets of legs, such that the middle leg of one side steps simultaneously with the fore and aft legs of the other side. For our experiments, yaw control using angular position sensing and steering was not used (e.g.

¹ImageProc 2.5:
https://github.com/biomimetics/imageproc_pcb

Pullin et. al [34]), so gait inconsistencies can result in yaw moments and deviations from straight running.

4.2 Control and Launching




Figure 4.3: VelociRoACH with cradle and H²Bird ornithopter MAV (top), and launch sequence from left to right (bottom).

During our launching experiments, each robot uses its own control scheme and mechanisms to govern its individual motion.

The H²Bird uses three PID controllers to control its pitch, yaw, and thrust. The pitch PID controller regulates the motion of the elevator mounted on the back of the tail and the yaw controller regulates the speed of a propeller mounted on the vertical stabilizer. The pitch and yaw controllers both use angular position estimates computed on-board by integrating the angular velocity measurements from the gyroscope. The thrust controller tracks a reference flap frequency using a Hall effect sensor mounted on one of the output gears on the wing gear box.

To bring the H²Bird to a desired launch velocity, a rigidly mounted cradle, shown in Figure 4.4 with the front and back annotated, was affixed to the top of the VelociRoACH, with the front and back of the cradle 7 cm apart. The cradle is constructed of 2.5 mm by




Figure 4.4: The launch cradle on top of the VelociRoACH, highlighted by the red dashed line. The carbon fiber spars through the back of the cradle are shown in the top-left inset.

1 mm flat carbon fiber spars for the support beams, and 2 mil PET for the sling in which the H²Bird sits. The mount is constructed such that the H²Bird sits at an angle of 25° in the cradle. This angle was selected to provide the H²Bird with a high enough initial pitch angle for takeoff, and to minimize the drag on the VelociRoACH as much as possible. The VelociRoACH has a maximum running velocity of 2.7 m/s without the H²Bird on top, and the minimum angle possible for takeoff at that speed is 18 degrees. Setting a cradle angle requiring the VelociRoACH to run at its maximum speed to take off is impractical, however, so we relaxed our velocity constraint by raising the cradle angle to 25°.

The cradle on the top of the VelociRoACH is 9 cm tall in the front and 6 cm tall in the back. We chose this height to prevent the tail of the H²Bird from scraping the ground as it pitches up before it takes off.

To prevent the H²Bird from sliding off of the cradle due to angular moments caused by the legs of the VelociRoACH making contact with the ground, we cut 5 mm x 1 mm slots in the back of the cradle at 25 mm apart, and affixed 2.5 mm by 1 mm flat carbon fiber spars to the back of the H²Bird. These spars are depicted in the top left inset of Figure 4.4. We observed that without the rods, the H²Bird tended to roll to the left or right and slide off of the cradle, so this modification constrains the rolling motion, thereby preventing the H²Bird from falling out of the cradle. The slots also helped to minimize some of the initial pitch motion as the VelociRoACH initially accelerates.

A picture of the full system, with the H²Bird sitting in the cradle on top of the VelociRoACH is in Figure 4.3.

4.3 Experimental Results and Discussion

We conducted experiments launching the H²Bird from the cradle on top of the VelociRoACH to determine the feasibility and performance of the cooperative launch. We also

investigated the effect that the H²Bird has on the running performance of the VelociRoACH.

Cooperative Launching

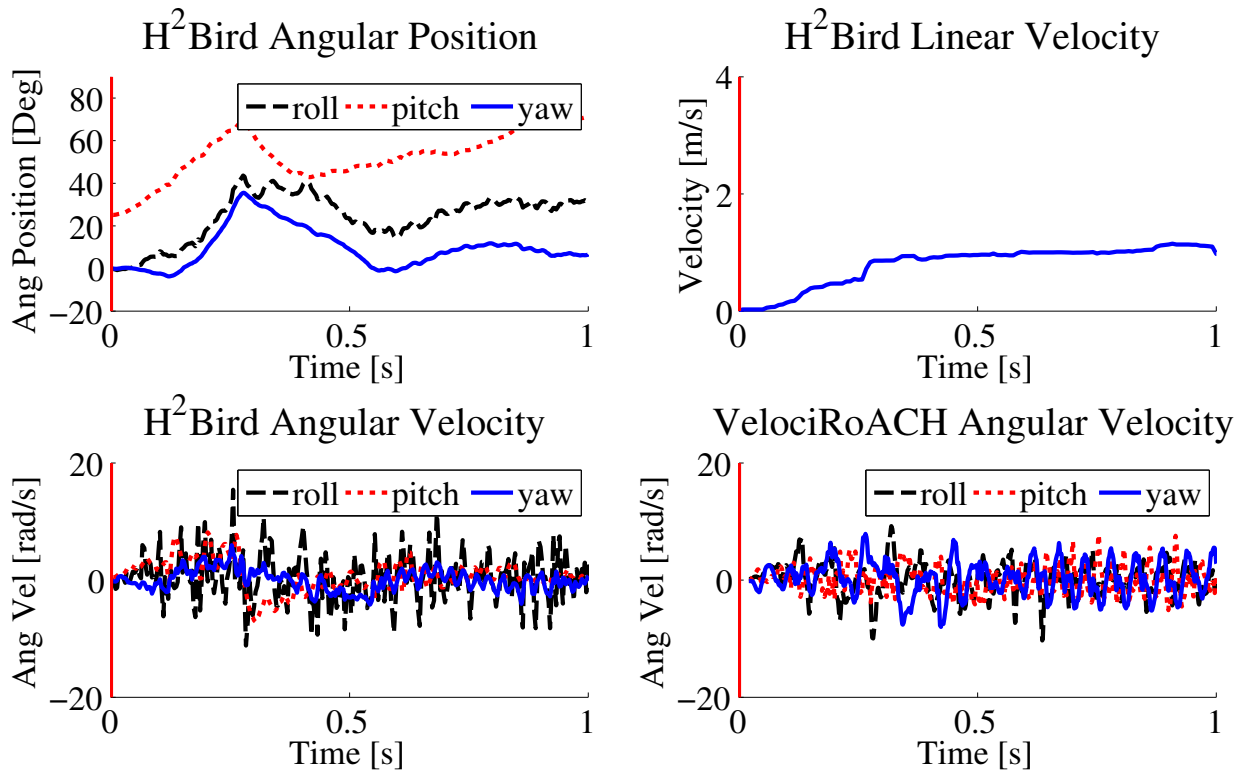


Figure 4.5: Telemetry data at the start of running for a single launch trial.

To test our cooperative launching system, we conducted experiments by attempting to launch the H²Bird under various conditions. Each trial consisted of the following steps, graphically depicted in the bottom portion of Figure 4.3:

1. The H²Bird is placed in the cradle at an angle of 25° and the VelociRoACH is at rest.
2. The VelociRoACH begins running at a predetermined stride frequency and the H²Bird pitches up initially due to the sudden forward acceleration.
3. The VelociRoACH reaches a steady state velocity and the H²Bird reaches a steady state pitch angle.
4. The H²Bird is given a launch command through a radio frequency (RF) link by the experimenter shortly after the steady state conditions are reached, and it detaches from the cradle.

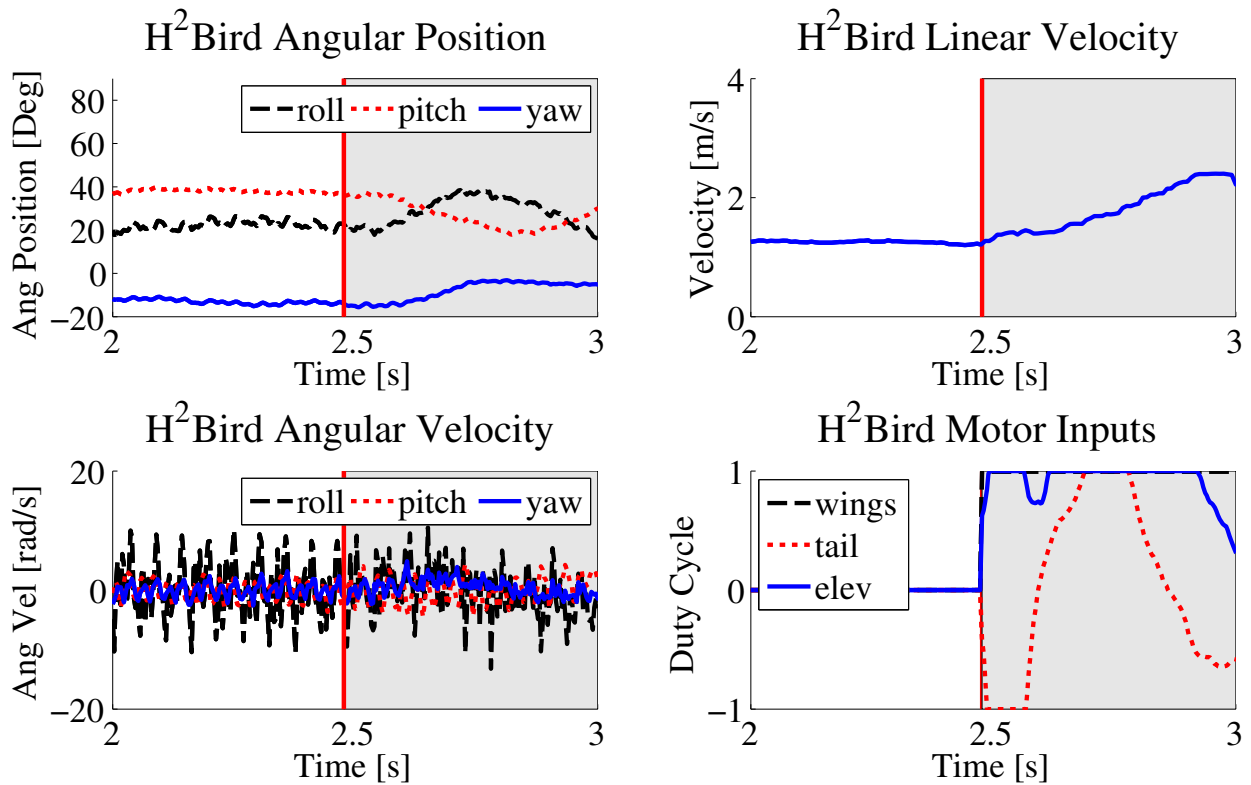


Figure 4.6: Telemetry data around launch for a single launch trial. The red line indicates the launch point.

A successful launch was defined as a launch in which the H²Bird did not touch the ground and ultimately traveled in an upwards direction after leaving the cradle. We did not have any requirements for the behavior of the VelociRoACH post-launch. We conducted trials at 16, 17, 18, and 20 Hz VelociRoACH stride frequencies in an effort to find the minimum stride frequency for successful launch. The H²Bird was free to pitch up in the cradle between 25 and 70 degrees from the horizontal. We did not constrain the pitching motion beyond the cradle properties previously mentioned. For each trial, we collected telemetry data from the H²Bird and VelociRoACH, and translational data from a Vicon² motion capture system.

The data collected for a typical trial are shown in Figures 4.5 and 4.6. The initial start-up transient is depicted in Figure 4.5. As shown in the figure, the H²Bird experiences an initial impulse increase in pitch angle of about 50 degrees, as well as roll oscillations of about 10 rad/s in magnitude. Both of these behaviors are due to the initial acceleration of the VelociRoACH. The behaviors 0.5 seconds before and after launch are in Figure 4.6. Right after launch, the H²Bird initially pitches down and there is an increase in velocity magnitude. The H²Bird angular controllers were on during the launch, which explains the

²Vicon Motion Systems: <http://www.vicon.com>

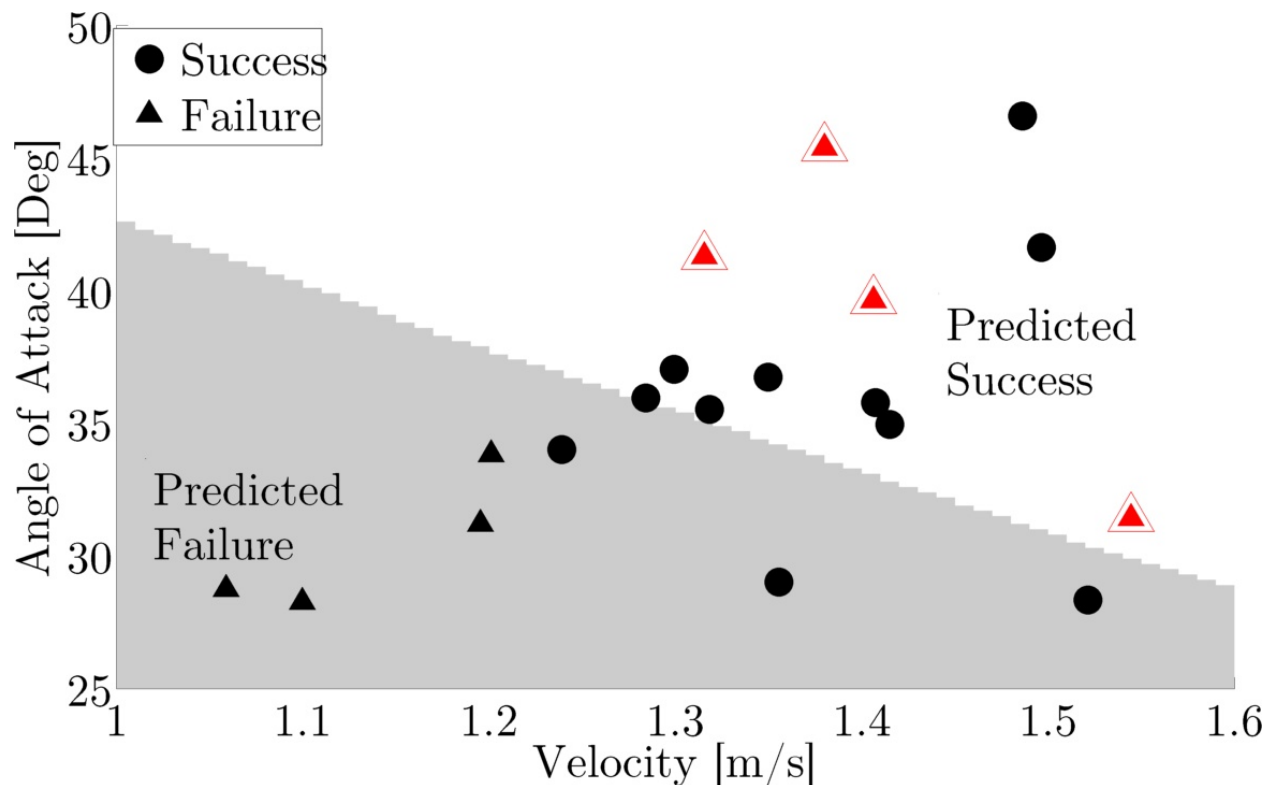


Figure 4.7: Launch experiments for varied running speeds and launch angles. The shaded region represented the wind tunnel predicted failure area, and the unshaded region represents the predicted success area. The red double triangles represent failures in the predicted success region.

changes in the duty cycle of the tail and elevator post-launch. Without these controllers, oscillations in pitch immediately after launch will cause the H²Bird to contact to the ground.

The results of the experiments are summarized in Figure 4.7, where the shaded region represents the area of initial conditions for which the wind tunnel data predicted failure, and the unshaded region represents the area of initial conditions for which the wind tunnel data predicted success. The circles represent experimental successes, and the triangles represent failures. The red triangles represent failures in the predicted success region.

We examined specific conditions of locomotion directly preceding and following the launch to determine why three launches succeeded in the wind tunnel predicted failure region and why four launches failed in the predicted success region.

There are several reasons for failure in the predicted success region of Figure 4.7. One failure mode that we observed is that the initial VelociRoACH acceleration could cause the H²Bird to rest in the stand in an unstable configuration. As shown on the right side of Figure 4.8, the failures in the predicted success region (indicated by red triangles) with less than 1.5 m/s takeoff velocity all had elevator input values less than 0.4. One trial at

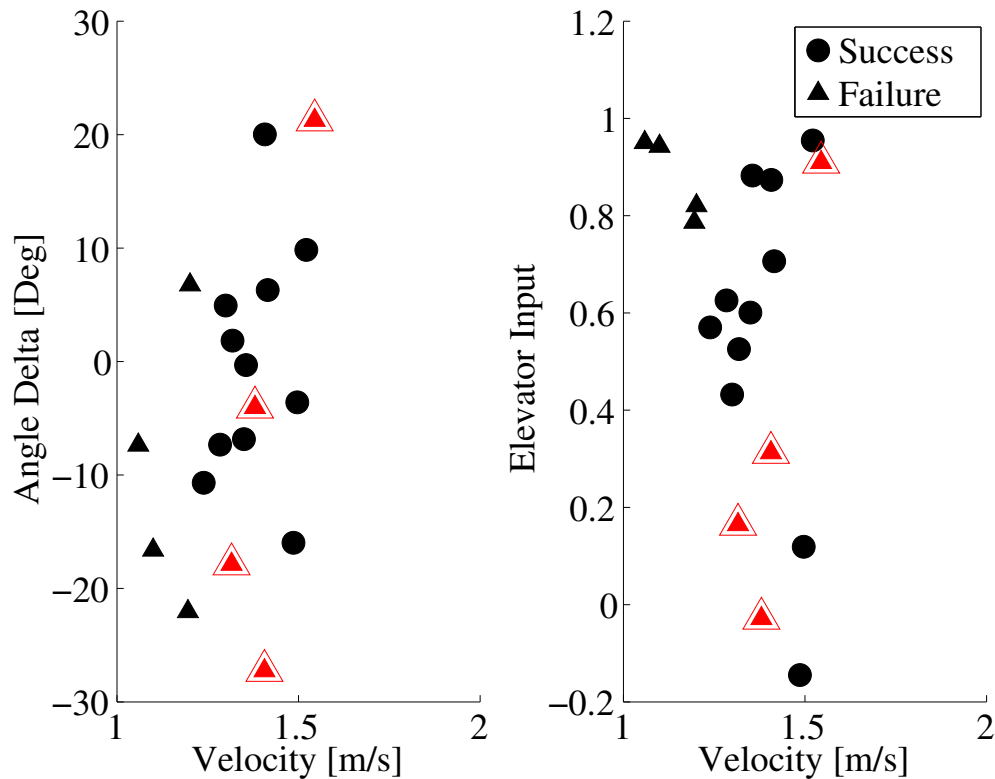


Figure 4.8: The change in pitch angle 0.2 seconds post-launch for each tested velocity (left) and the elevator input at launch for each tested velocity (right). The red double triangles represent failures in the predicted success region.

1.38 m/s had a negative elevator input, which indicates that the H²Bird tried to pitch down after takeoff. These low elevator input values are caused by a takeoff angle greater than 45 degrees. The pitch controller regulates the pitch angle to 45 degrees post-takeoff, so the pitch controller in combination with high takeoff angles due to a poor resting state in the cradle can cause failure.

Another failure mode that we observed is that the H²Bird sometimes became caught on the front of the stand. This occurrence manifests itself in the form of a drop in pitch immediately following launching. The left side of Figure 4.8 shows that in some failure cases, the H²Bird lost a significant amount of pitch angle post-launch. The two red triangles at approximately 1.3 m/s and 1.4 m/s takeoff velocities experienced losses in pitch angle of -18 degrees and -27 degrees, respectively. This large dip in pitch angle could indicate that the back of the H²Bird hit the front of the stand as it released from the VelociRoACH.

Examining both the elevator input and the change in the pitch angle in Figure 4.9, we see that three out of the four failure cases in the predicted success region exhibited a combination of both losses in pitch angle post-launch and low elevator set points.

The final failure point, at 1.54 m/s launch velocity exhibited both the highest launch

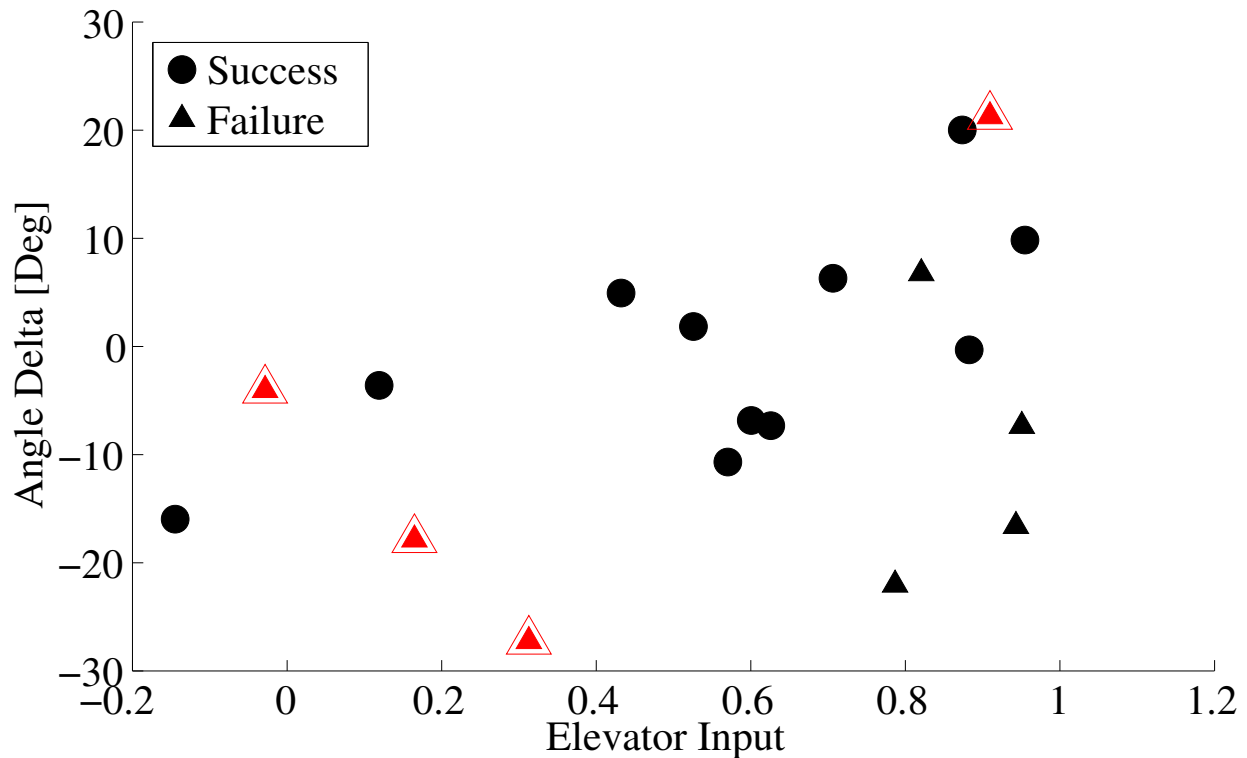


Figure 4.9: The change in pitch angle vs. the elevator input for each trial. The red double triangles represent failures in the predicted success region.

velocity and the largest increase in pitch angle post-launch. The combination of these two circumstances can cause the H²Bird to stall, experience a reduction in velocity and lift, and fall to the ground.

Although the wind tunnel can provide some idea as to the forces and moments that are experienced by the H²Bird in free flight, the fact that the robot is rigidly mounted can cause some inconsistencies between what is happening in the wind tunnel and in unconstrained flight. These inconsistencies are detailed further in [35] and could explain the existence of successful trials in the predicted failure region.

Cooperative Running

To examine the effects of transporting the H²Bird on top of the VelociRoACH, we conducted experiments running the VelociRoACH at a stride frequency of 17 Hz. For each trial, we started the robot from rest and recorded the angular velocities, linear velocity, battery voltage, motor duty cycles, and motor back electromotive force (back-EMF) for 5 seconds of running. We collected this data at 17 Hz stride frequency for the VelociRoACH running alone, running with an inertial mass sitting on top, running with the H²Bird passively

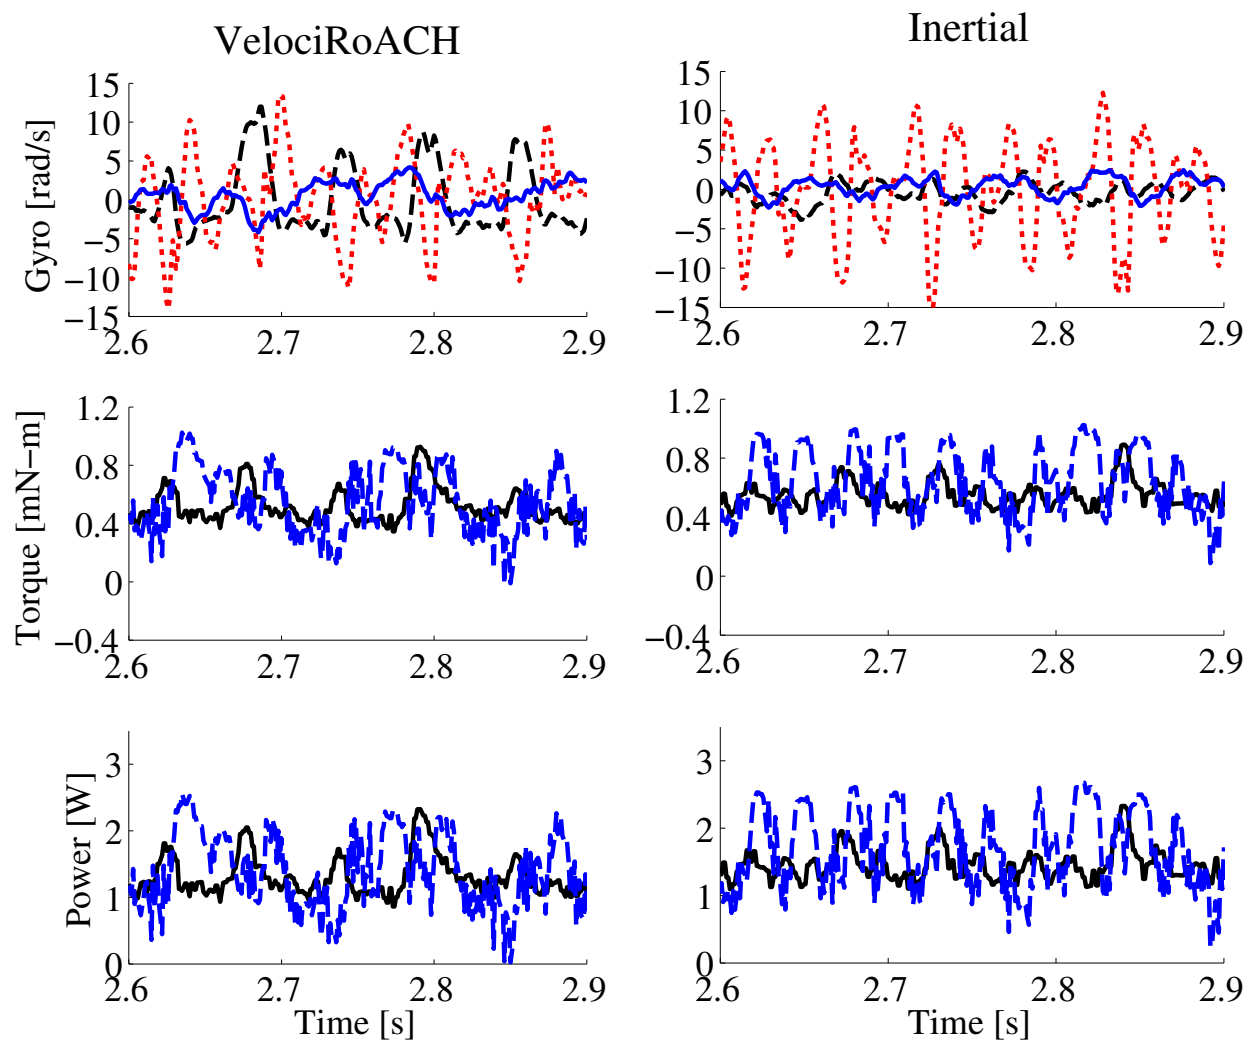


Figure 4.10: Telemetry data for a single trial for the VelociRoACH running alone (left) and running with an inertial mass equivalent to the H²Bird(right).

Data	Passive [% Change]	Active [% Change]
Average Power	+24.5	+18.1
Roll Velocity Variance	-91.1	-95.0
Pitch Velocity Variance	-80.0	-90.5
Yaw Velocity Variance	-25.3	-41.5
Average Velocity	-8.4	+12.7

Table 4.1: Table of percent increases and decreases of measured data for the active and passive H²Bird cases over the case of the VelociRoACH running by itself.

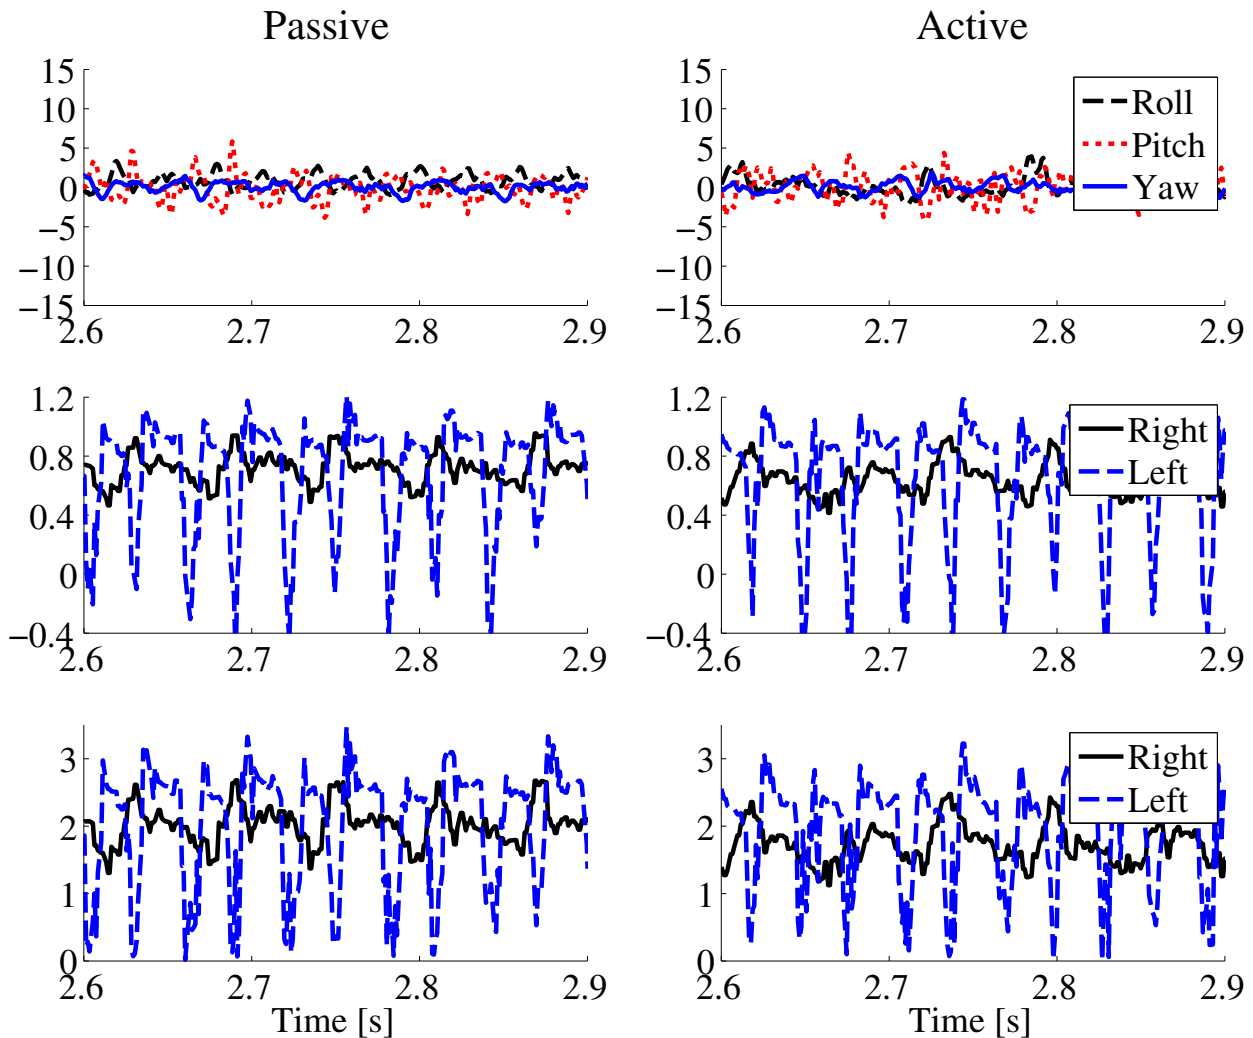


Figure 4.11: Telemetry data for a single trial for the VelociRoACH running with the passive H²Bird (left) and running with the H²Bird flapping at 5 Hz (right).

sitting in the cradle, and running with the H²Bird with the yaw controller active and the wings flapping at 5 Hz. We chose a 5 Hz flap frequency because it was just low enough for the H²Bird to remain stationary sitting on the stand. The inertial mass had approximately the same mass and inertia tensor as the H²Bird.

The data for a typical trial are shown in Figures 4.10 and 4.11, with the case of the VelociRoACH alone, the VelociRoACH with the inertial mass, the passive H²Bird case, and the active H²Bird case from left to right. Examining the top row of plots, it is clear that the magnitude of the pitch velocity is reduced dramatically from the 'alone' case to the 'passive' and 'active' cases. The middle row of plots show that the presence of the H²Bird induces a periodic spike in the motor torque for the VelociRoACH that is not as prominent in the

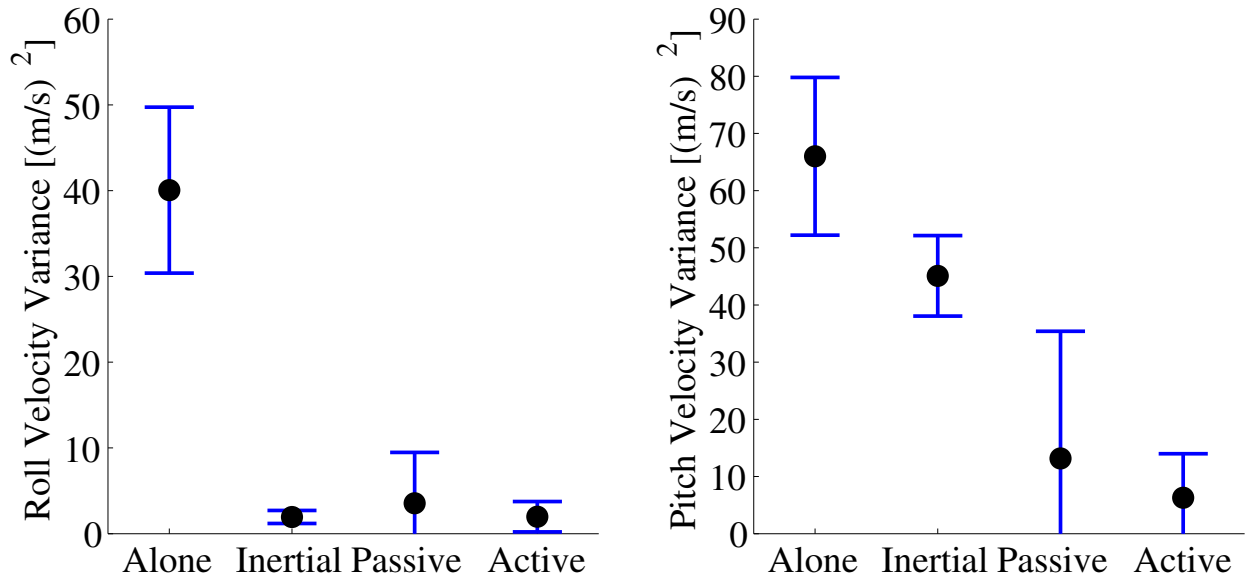


Figure 4.12: Roll velocity variance (left) and pitch velocity variance (right) for the VelociRoACH alone, with inertial mass, with passive H²Bird, and with active H²Bird.

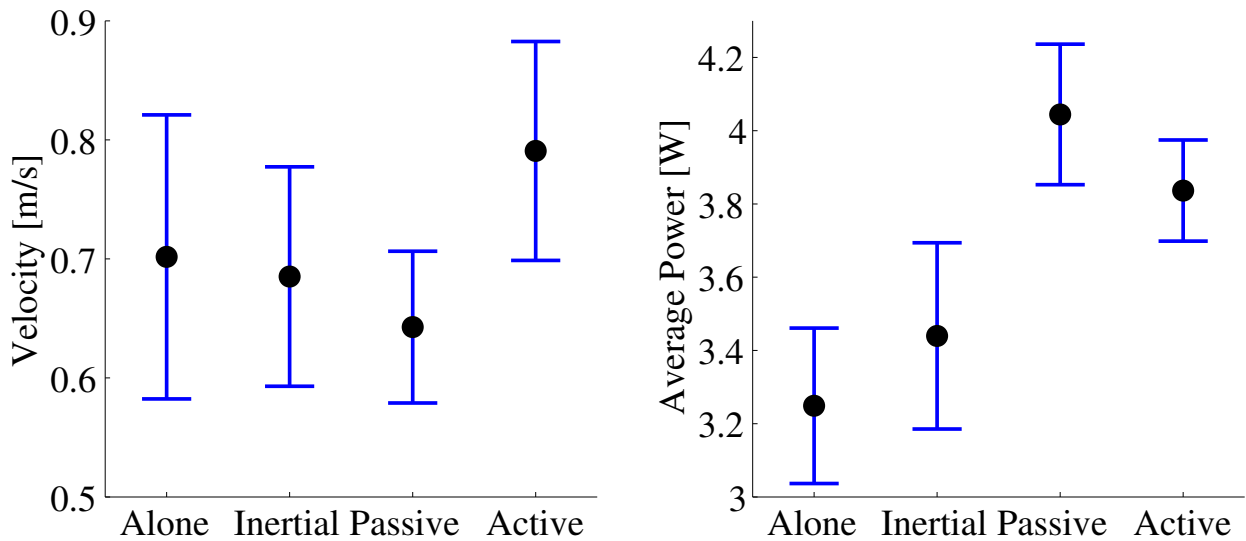


Figure 4.13: Average running velocities (left) and average power consumed (right) at steady state for VelociRoACH alone, with inertial mass, with passive H²Bird, and with active H²Bird.

'inertial' case. The increase in the magnitude of the peaks in the pitch velocity and the motor torque from the 'inertial' case to the 'passive' case indicates that the running gait of the VelociRoACH is not only affected by the change in inertia provided by the H²Bird, but

the wings as well.

To determine the energetic penalty to the the VelociRoACH's running performance while carrying the H²Bird, we computed the average power, $Power_i$, consumed by each motor over the 5 second trials at 17 Hz stride frequency, using the following equation:

$$Power_i = |DC_i| * V_{Batt} * \frac{(V_{Batt} - BEMF_i)}{R_i} \quad (4.2)$$

where $Power_i$ is the power input to motor i , DC_i is the duty cycle, $BEMF_i$ is the back-EMF, R_i is the motor resistance, and V_{Batt} is the battery voltage. The energetic data are presented in the right plot of Figure 4.13. Each point represents the average sum of the power going into each motor for the 6 trials for each experiment set, and the error bars are one standard deviation above and below the mean.

From Figure 4.13, the VelociRoACH motors consume more power for the passive and active trials with the H²Bird on top, than they do for the VelociRoACH alone or with the inertial mass. This result is caused by the increased mass and higher center of gravity from the H²Bird sitting on top of the VelociRoACH, as well as the drag provided by the wings. These additional forces require the motors to produce more torque as the feet contact the ground to keep the stride frequency at 17 Hz. Although carrying the H²Bird is detrimental to the VelociRoACH in terms of power consumption, there are several benefits.

One benefit comes in the form of the angular velocities of the VelociRoACH as it runs. Figure 4.12 depicts the effect of the H²Bird on the variance of the roll and pitch velocities experienced during running. Each point represents the mean variance in roll (left) and pitch (right) velocities for the four cases over 6 trials, and the error bars are one standard deviation above and below the mean. The VelociRoACH experienced a 91.1 percent reduction in the roll velocity variance simply by placing the H²Bird in the stand, and a 95.0 percent reduction in the variance for the active H²Bird. For pitch velocity variance, the VelociRoACH experienced an 80.0 percent reduction for the passive H²Bird, and a 90.5 percent reduction for the active case.

While the active H²Bird case had little benefit in terms of rotational damping over the passive case, and requires more total energy overall for locomotion, the benefit to the VelociRoACH comes in the form of running speed. Figure 4.13 shows the change in the average running velocity for each experimental case with the error bars representing one standard deviation above and below the mean of the 6 trials. The passive H²Bird case results in a 8.4 percent decrease in average running velocity at steady state. This reduction is caused by the added drag force from the wings of the H²Bird. By actively flapping the wings at 5 Hz, however, we measured a 12.7 percent increase in running velocity. The flapping provides a forward thrust force to counteract the drag force caused by the wings. The lift provided by the wings can reduce the ground contact forces as the VelociRoACH runs, reducing the torque required by the motors and enabling faster running.

All of the aforementioned benefits and detriments are summarized in Table 4.1. Positive percentages indicate an increase in a particular measure over the VelociRoACH alone case, while negative percentages indicate a decrease.

We calculated the cost of transport for the H²Bird in flight, for the VelociRoACH running alone, and for both cooperative cases according to the following equation:

$$COT = \frac{P}{mgv} = \frac{W}{mgd} \quad (4.3)$$

where P is average power consumed, m is mass, g is gravity, v is velocity, W is work, and d is distance. The calculated cost of transport for the VelociRoACH alone, the H²Bird alone, and our two cooperative cases are in Table 4.2. Placing the H²Bird on top of the VelociRoACH decreases the cost of transport of the VelociRoACH by approximately 16 percent. This decrease in the cost of transport would be useful in a situation where the VelociRoACH and the H²Bird had to both reach a point 80 meters away and the H²Bird had to fly 20 meters in the air, where the VelociRoACH cannot reach. In one case, both robots travel the 80 meters separately, and then the H²Bird continues the last 20 meters. In a second case, the VelociRoACH carries the H²Bird for the first 80 meters, then the H²Bird is launched and flies 20 meters. The second case consumes 25 percent less energy than the first. In situations such as these, cooperative locomotion would be more efficient than independent locomotion.

	VelociRoACH	H²Bird	Passive	Active
Cost of Transport	8.1	10.1	6.8	6.6

Table 4.2: Table of cost of transport for the VelociRoACH alone, the H²Bird alone, and the active and passive cooperative cases.

Chapter 5

Piece-wise Linear Modeling for Diving

5.1 Piece-wise Affine Modeling

The diving behaviors of the H²Bird were modeled using a piece-wise affine discrete time model with a 0.025 second time step. The state data were segmented into diving and recovery sections, and each section was segmented into several linear models using K-means and least squares regression. From these linear models, we used reachability analysis to determine when to switch control behaviors within the maneuver to reach the goal height.

Data Collection

The diving data were collected in free flight, using a Vicon motion capture system to collect translational position and velocity data and pitch angle data from the H²Bird during the dive. Each experiment consisted of:

1. Release the H²Bird by hand from one side of the tracking space.
2. Wait for the initial release transient flight to stabilize.
3. Manually begin the unpowered dive portion of the maneuver.
4. Manually begin the powered recovery portion of the maneuver.

During the unpowered dive portion of the maneuver, the H²Bird stops flapping, and the wings remain pressed together. The angle of the wings is not directly controlled; the force of the air interacting with the wings forces them to the “closed” position. The elevator is commanded to a neutral position of 20 degrees and the yaw controller remains active. Since the maneuver is within the saggital plane, we did not model the yaw motion for these experiments.

During the powered recovery portion of the maneuver, the wings are flapped at a maximum frequency of approximately 20 Hz, and the elevator is commended to its maximum

deflection of 60 degrees. The 160 mAh LiPo battery was replaced after every 3 trials, to eliminate the effects of battery drain on the experiments.

A sample data set from the open-loop diving trials is in Figures 5.2 and 5.3. The gray shading represents the unpowered dive portion of the maneuver, and everything after the gray shading is considered the recovery part of the maneuver. The black marker indicates the lowest point of the maneuver, which is the final dive height. Hereafter, all references to the beginning and final heights of the dive will represent the height at the beginning of the unpowered portion and the lowest point during the recovery portion, respectively. Figure 5.4 illustrates the means of the terminal conditions of the maneuver for all of the trials. The error bars are one standard deviation above and below the means.

We did not specifically seek to model ground effect during the diving trials, although there were dives that came very close to the ground. As a result ground effect could be aiding the recovery of the H²Bird from some of the dives, as ground effect increases lift and decreases drag on an airfoil as it comes into close proximity with a fixed surface. This phenomenon could have some effect on the end result of the modeling that we used to design our controllers.




Figure 5.1: Free body diagram of the H²Bird for the relevant state variables for the discrete-time models [4].




Figure 5.2: Position data for a single trial of the open loop H²Bird diving experiment. The gray shading represents the unpowered portion of the dive. The black marker indicates the conditions at the lowest vertical position in the dive.

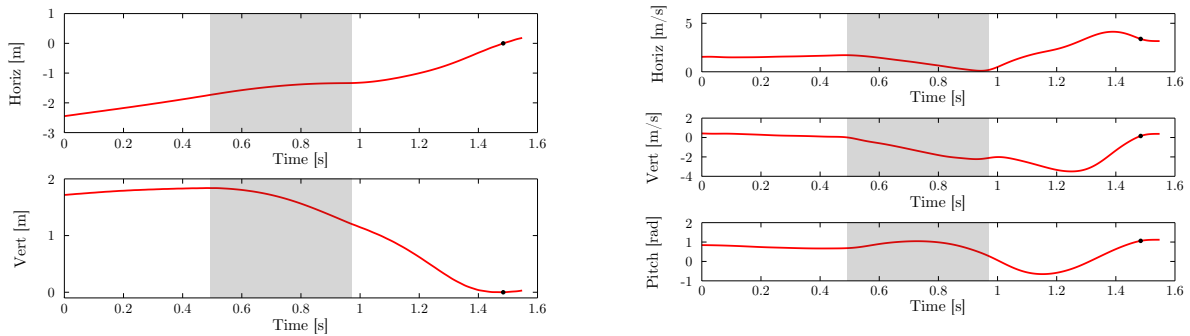


Figure 5.3: Telemetry data for a single trial of the open-loop H²Bird diving experiment over the time period. The position is on the left and the horizontal velocity, vertical velocity, and pitch angle are on the right. The gray shading represents the unpowered portion of the dive. The black marker indicates the conditions at the lowest vertical position in the dive.

Data Segmentation

The flight dynamics of the H²Bird are nonlinear, and the aerodynamics of flapping-winged flight are very complex. We decided to segment the flight trajectory data into several regions and linearize about the regions to work around the numerical difficulties of the aerodynamics. In our models, we represented the state of the H²Bird in the form of Equation 5.1.




Figure 5.4: The means of the conditions for the lowest point in the open loop dive trials. The error bars represent one standard deviation above and below the mean for each state variable.

$$\mathbf{X} = [x \quad z \quad \dot{x} \quad \dot{z} \quad \theta] \quad (5.1)$$

In the vector, x is the forward position in world coordinates, z is the vertical position in world coordinates, θ is the pitch angle, \dot{x} is the horizontal velocity in world coordinates, and \dot{z} is the vertical velocity in world coordinates. Figure 5.1 is a diagram of the relevant state variables and their associated directions.

For the purposes of the models, we did not want the model chosen for a particular data segmentation region to be dependent upon the translational position of the robot, so we used our segmentation routine on only a portion of the state vector:

$$\mathbf{X}_p = [\dot{x} \quad \dot{z} \quad \theta] \quad (5.2)$$

We used k-means clustering, which partitions our data into k clusters to segment the data into regions for linearization. To segment the data we first randomly selected an initial set of k region centers m_1^1, \dots, m_k^1 from all of our observed state data for all of our diving trials. We then complete an assignment step in which we assign each data point x_p to the closest region center, according to Equation 5.3, where S_i^t is the i th region cluster at time t .

$$S_i^t = \{x_p \ni \|x_p - m_i^t\|^2 \leq \|x_p - m_j^t\|^2 \forall j, 1 \leq j \leq k\} \quad (5.3)$$

After the assignment step, we recalculate the the region centers according to Equation 5.4. The centers of the regions are recalculated as the mean of the regions of clustered data points.

$$m_i^{t+1} = \frac{1}{|S_i^t|} \sum_{x_j \in S_i^t} x_j \quad (5.4)$$

The update and assignment steps are run iteratively 100 times. To account for bad seeding of the initial randomly selected region centers, we run the entire k-means algorithm 10 times and select the regions that minimize the sum of the squared distances from each data point to its associated region cluster:

$$\arg \min_S \sum_{i=1}^k \sum_{x \in S_i} \|x - \mu_i\|^2 \quad (5.5)$$

The results of the k-means segmentation of the two portions of the maneuver are graphically represented in Figure 5.5, with the model regions for the dive on the left and the regions for the recovery on the right. Each color represents a different region of the state space described in Equation 5.2. We tested from one to ten regions, however only the three-region case is represented in Figure 5.5 for illustrative purposes.




Figure 5.5: Graphical representation of the K-means segmentation of the dive (left) and escape (right) portions of the open loop experiments in the horizontal velocity, vertical velocity, and pitch angle space.

Piece-wise Affine Model Generation

Using the segmented model regions of similar dynamic conditions generated using k-means, we used least squares to fit linear models to each region. The piece-wise linear discrete-time models that we fit to the data in each region are of the form:

$$x_{t+1} = Ax_t + Bu_t + f \quad (5.6)$$

where $x_{t+1} \in \mathbb{R}^{5 \times 1}$ is the state vector described by Equation 5.1 at the next time step, $x_t \in \mathbb{R}^{5 \times 1}$ is the state vector at the current time step, $u_t \in \mathbb{R}^{2 \times 1}$ is the input vector, $A \in \mathbb{R}^{5 \times 5}$ is the state evolution matrix, $B \in \mathbb{R}^{5 \times 2}$ is the input matrix, and $f \in \mathbb{R}^{5 \times 1}$ is an affine portion of the state evolution equation. For the purposes of our design, we used a time step of 0.025 seconds.

Since the first two variables of our state vector are the x and z position, we know that the state update equations for that portion of the state space should be:

$$\begin{aligned} x_{t+1} &= x_t + T_s * \dot{x}_t \\ z_{t+1} &= z_t + T_s * \dot{z}_t \end{aligned} \quad (5.7)$$

where T_s is our time step. Therefore, we are only fitting the model for the reduced state vector in Equation 5.2, the same state vector that we used for our k-means segmentation. This ensures that our segmentation and models are independent of the position.

To calculate the model parameters A_i , B_i , and f_i for each model region i , we solved the following least squares optimization problem:

$$\mathbf{X}_{t+1} = [A \ B \ f]_i \begin{bmatrix} \mathbf{X} \\ \mathbf{u} \\ \mathbf{1} \end{bmatrix}_t = \beta_i \bar{\mathbf{X}}_t \quad (5.8)$$

β_i is the unknown model parameter matrix for all of the points in region S_i that satisfy Equation 5.3. $\mathbf{X}_{t+1} \in \mathbb{R}^{3 \times N}$ is the vector of all N next states in the region i , collected into the observation matrix of the least squares problem. $\mathbf{X} \in \mathbb{R}^{3 \times N}$ is the vector of all N current states in region i , $\mathbf{u} \in \mathbb{R}^{2 \times N}$ is the vector of all N inputs in region i , and $\mathbf{1} \in \mathbb{R}^{1 \times N}$ is a vector of N ones, collected into the data matrix $\bar{\mathbf{X}}_t$ of the least squares problem. We solved the least squares problem for each region i in each the dive and recovery regions for one- to ten-region k-means segmentations. The final learned discrete time piece-wise affine model equation is of the following template for each region i :

$$x_{t+1} = \begin{bmatrix} 1 & 0 & T_s & 0 & 0 \\ 0 & 1 & 0 & T_s & 0 \\ 0 & 0 & a_{11}^i & a_{12}^i & a_{13}^i \\ 0 & 0 & a_{21}^i & a_{22}^i & a_{23}^i \\ 0 & 0 & a_{31}^i & a_{32}^i & a_{33}^i \end{bmatrix} \begin{bmatrix} x \\ z \\ \dot{x} \\ \dot{z} \\ \theta \end{bmatrix}_t + \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ b_{11}^i & b_{12}^i \\ b_{21}^i & b_{22}^i \\ b_{31}^i & b_{32}^i \end{bmatrix} \begin{bmatrix} u_1 \\ u_2 \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ f_1^i \\ f_2^i \\ f_3^i \end{bmatrix} \quad (5.9)$$

To determine the appropriate number of regions to model the motion of the H²Bird, we used the Mahalanobis distance metric:

$$\sqrt{\frac{\sum_{i=1}^N (\beta_i \bar{\mathbf{X}}_t - \mathbf{X}_{t+1})^\top * \Sigma_x * (\beta_i \bar{\mathbf{X}}_t - \mathbf{X}_{t+1})}{N}} \quad (5.10)$$

N is the number of data points in the particular model region i and Σ_x is the covariance metrics of the state measurements. We used the Mahalanobis distance as our metric because it relates vector quantities in different spaces based on the distribution of values in a signal. We used 80 percent of our measured data as the training set for our models and the other 20 percent as the validation set to determine the validity of our models. The distance values for the validation sets are in Table 5.1. The distance metric for the ‘‘Null’’ model is in the

first column of the table for comparison with the other fitted models. The “Null” model is defined as:

$$\mathbf{X}_{t+1} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \mathbf{X}_t \quad (5.11)$$

For ease of computation, we selected the number of regions that provided an appreciable decrease in error from a smaller number of regions.

# of Regions	Null	1	2	3	4	5	6	7	8	9	10
Dive	0.113	0.0309	0.0188	0.0208	0.0205	0.0215	0.0212	0.0209	0.0205	0.0218	0.0209
Escape	0.501	0.113	0.0859	0.0586	0.0540	0.0516	0.0470	0.0440	0.0418	0.0413	0.0411

Table 5.1: Table of the Mahalanobis distance for 1 through 10 regions for the dive and escape portions of the maneuver. The Mahalanobis distance for the Null model is in column 1 for comparison.

5.2 Reachability Analysis

To determine when to switch between the uncontrolled dive to the bang-bang recovery controller, we computed the backward reachable sets over a fixed time horizon from our goal state. We defined the set of goal states as the region of observed terminal conditions at the lowest vertical point in our maneuver. This point is indicated by a black marker in Figures 5.2 and 5.3. Figure 5.4 shows the mean horizontal and vertical displacement, the mean horizontal and vertical velocity, and mean pitch angle of the terminal conditions for all of our trials. The error bars indicate one standard deviation above and below the terminal condition. To conduct the reachability analysis, we defined the set of possible goal states as the minimum and maximum of the observed terminal velocities and pitch. We define the goal position as zero for our controller, so we defined the goal terminal condition as 0.1 meters above and below zero for the horizontal and vertical position. An equality constraint at zero for the goal position constrained our problem too much and did not provide the result that we needed for the analysis. The projection of the five dimensional goal region on the translational position space is on the left in Figure 5.6, and the projection on the translational velocity and pitch angle space is on the right.

We define the backwards reachable set as:

$$B(x_f, U, t) = \{x_0 \in X : \exists u \in U, \exists t \in [0, t] \text{ s.t. } x(t) = x_f\} \quad (5.12)$$

The backwards reachable set is the set of all states x_0 such that there exists an input $u \in U$ that can drive the collection of states x_0 into the collection of states x_f in t time. For our analysis, we use a time step of 0.025 seconds and find the backwards reachable set for our goal region in Figure 5.6 for 16 time steps, or 0.4 seconds. The total collection of all of the polytopes for the reachable sets for each step in time is the total set of states that can




Figure 5.6: The terminal polytopes for the backwards reachability analysis. The set in the horizontal and vertical position space is on the left, and the set in the horizontal velocity, vertical velocity, and pitch angle space is on the right. The goal position is always zero and all other heights are relative to the goal.

reach our target height in $t \leq 0.4s$ time. We use this set boundary as a guard condition for switching control schemes from the dive portion to the recovery portion of the maneuver. The total reachable set within 0.4 seconds for our goal set is shown in Figures 5.7 and 5.8. When the trajectory of the H²Bird during the dive passes into this set, the control switches into the recovery phase. The evolution of the reachable sets can be found in Appendix A.




Figure 5.7: The backwards reachable polytopes 16 timesteps, or 0.4 seconds, from the terminal set in the horizontal and vertical position space. Letters A - C in the left image correspond to the positions in the trajectory marked by letters A - C in the right image.




Figure 5.8: The backwards reachable polytopes 16 timesteps, or 0.4 seconds, from the terminal set in the horizontal velocity, vertical velocity, and pitch angle space.

Chapter 6

Online Control for Diving

6.1 Control Implementation

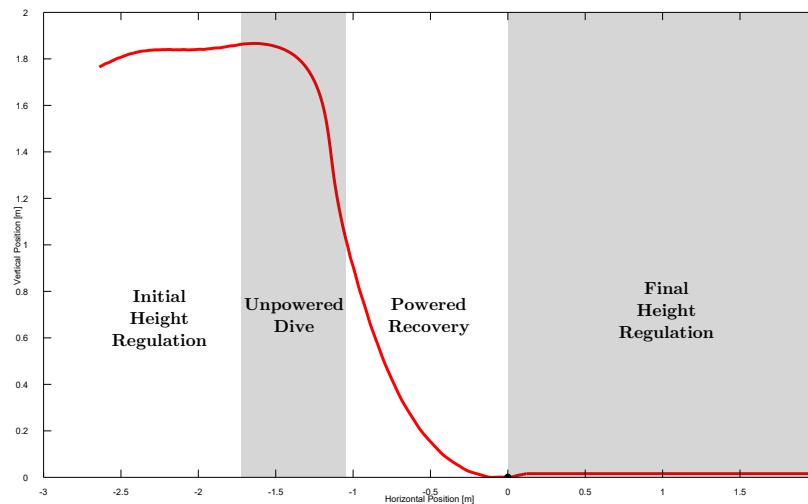


Figure 6.1: Diagram of the control sections of the entire diving maneuver.

The controllers for the diving maneuver are sectioned into three distinct pieces: the uncontrolled dive, the powered recovery, and initial and final height regulation. The location within the dive of each section of the control is graphically depicted in Figure 6.1. During the beginning of the maneuver, the H²Bird uses a height regulation Linear Quadratic Regulator (LQR) controller to maintain a predetermined initial height. The H²Bird then begins the unpowered dive portion and recovers from the dive at a time dictated by our models described in Chapter 5. Finally, the final height is maintained using the same LQR controller as the initial height regulation section. A block diagram of the control implementation is in

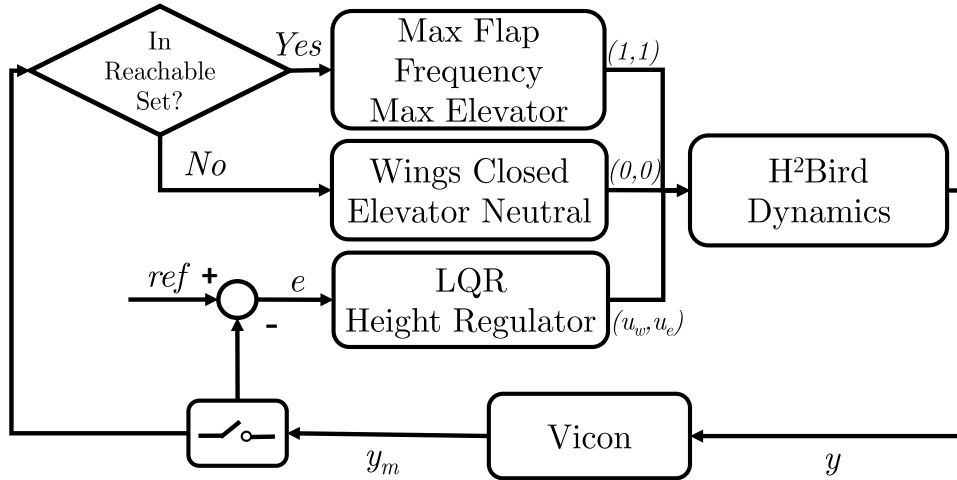


Figure 6.2: Controller block diagram of the diving maneuver.

Figure 6.2. The controller is divided into the three sections described previously. The controller switches from the initial height regulation when the dive is initiated by the user. The transition from the unpowered dive to the recovery controller occurs when the estimated state, y_m , of the H²Bird is within the reachability polytope computed in Chapter 5. The controller transitions back to the LQR height regulator for the final height set point when the H²Bird stops descending, or the vertical velocity, v_z , is greater than zero. Each controller section along with the transition conditions are described in the following subsections.

Uncontrolled Dive

The dive portion of the maneuver is uncontrolled; the wings stop flapping and the elevator is commanded to a neutral position of 20 degrees upwards deflection. The position at which the dive begins is marked as the starting position of the maneuver.

Recovery Control

To determine the transition point between the uncontrolled dive and recovery portions of the maneuver, we used the set generated by the reachability analysis as our hybrid guard condition. The reachability set for our desired goal set is numerically represented by a collection of N polytopes of the form $A_i x_t \leq b_i$, where x_t is the current state, $A_i \in \mathbb{R}^{M \times 5}$, and $b_i \in \mathbb{R}^{M \times 1}$. M depends on the number of faces in the polytopes and varies polytope to polytope. If the current state is within the reachability set, the control scheme is switched to the recovery controller. This check can be formally stated as:

$$\text{if } \exists \mathcal{P}_i \mid \mathcal{P}_i = A_i x \leq b_i, A_i x_t - b_i \leq 0 \Rightarrow \text{recover} \quad (6.1)$$

\mathcal{P}_i is the i th polytope of the reachable set.

The recovery controller is a bang-bang controller that stays active until the vertical velocity becomes positive. During the recovery, the wings are flapped at maximum frequency and the elevator is at its full upward deflection. The moment the vertical velocity becomes positive is recorded as the final dive height, and the controller switches to the height regulator.

Powered Height Regulation

The powered height regulation controller is a Linear Quadratic Regulator (LQR) controller to stabilize the system, and an integral controller for height reference tracking. To generate the controllers, we use a model designed using the same techniques as the diving and recovery models.




Figure 6.3: Telemetry data for a single trial of the open-loop H²Bird level flight. The position is on the left and the inputs to the wing motor and elevator are on the right.

Flight data were collected using a Vicon motion capture system and the on-board telemetry system of the H²Bird. A sample data set is in Figures 6.3 and 6.4. Each data set was collected for the H²Bird at varying flap frequencies and elevator deflections to determine how the height, forward velocity, vertical velocity, and pitch angle change with respect to the inputs. The data were then used to generate a single linear model of the dynamics of the H²Bird around equilibrium flight conditions. Only one model was generated because it was determined that multiple models did not provide appreciable increases in accuracy over a single model. The generated model is of the form:

$$x_{t+1} = \begin{bmatrix} 1 & 0 & T_s \\ 0 & a_{11}^i & a_{12}^i \\ 0 & a_{21}^i & a_{22}^i \end{bmatrix} \begin{bmatrix} z \\ \dot{x} \\ \dot{z} \end{bmatrix}_t + \begin{bmatrix} 0 & 0 \\ b_{11}^i & b_{12}^i \\ b_{21}^i & b_{22}^i \end{bmatrix} \begin{bmatrix} u_1 \\ \theta_d \end{bmatrix}_t \quad (6.2)$$




Figure 6.4: The horizontal velocity, vertical velocity, and pitch angle for a single trial of the open-loop H²Bird level flight.

The main difference between the previous models is that we no longer care about the horizontal position, so it was omitted from the state vector. Additionally, we planned to run the level flight controller at 10 Hz, which is too slow to directly control the elevator deflection. For this reason, we use a commanded pitch set-point θ_d as our second input.

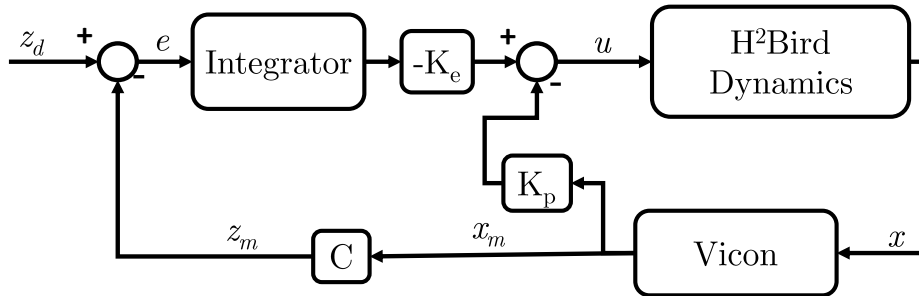


Figure 6.5: Block diagram of the level flight controller.

The block diagram for the controller is in Figure 6.5. In the diagram, z_d is the desired height set-point, z_m is the measured height, x_m is the measured full state, and u is a vector with the wing motor duty cycle and a pitch reference. The controller combines a gain matrix K_p , used to stabilize the system, and an integrator gain K_e , for height reference tracking. Both controllers were computed using the Linear Quadratic Regulator (LQR) optimal control algorithm with an infinite time horizon. The controller gains are as follows:

$$\begin{aligned} K_p &= \begin{bmatrix} -3.3452 & 0.0833 & -0.3116 \\ 1.7272 & 0.0297 & 0.1908 \end{bmatrix} \\ K_e &= \begin{bmatrix} 2.2754 \\ -1.0025 \end{bmatrix} \end{aligned} \tag{6.3}$$

The closed loop poles of the entire system, including the integrator, are:

$$p = \begin{bmatrix} 0.0662 \\ 0.9344 + 0.0417i \\ 0.9344 - 0.0417i \\ 0.7168 \end{bmatrix} \tag{6.4}$$

6.2 Experiments and Discussion

We conducted experiments on the dive recovery control scheme. Unfortunately, the available tracking space in the Vicon room was not large enough to test both controllers in sequence, although the distinct transition requirements between them can be easily implemented with a larger space.

Recovery Control

To determine the effectiveness of the recovery controller, we conducted experiments using various desired dive heights. Each trial consisted of releasing the H²Bird by hand, then triggering a dive to a particular setpoint. The rest of the maneuver is autonomous, and the current yaw controller remained active to prevent lateral movement. A successful dive was counted as one that did not contact the ground, and a failed dive was counted as one that did. We did not have any requirements on post-dive behavior, although the H²Bird ascended at the end of each trial. We conducted trials for target dive distances of 1.0, 1.5, and 2.0 meters. For each trial, we released the H²Bird from approximately 2.6 meters above the ground with a 30 degree pitch angle setpoint and waited for the initial transient at the start of flight to end. During each trial, we collected the motor inputs from the H²Bird and translational position, translational velocity, and pitch angle from the Vicon motion capture system.

The data for a typical trial are presented in Figures 6.6 and 6.7. The gray shaded region indicates the unpowered dive portion of the experiment, and the recovery trigger point is




Figure 6.6: Position data for a single trial of the H²Bird diving experiment with the recovery controller active. The gray shading represents the unpowered portion of the dive. The black marker indicates the conditions at the lowest vertical position in the dive.

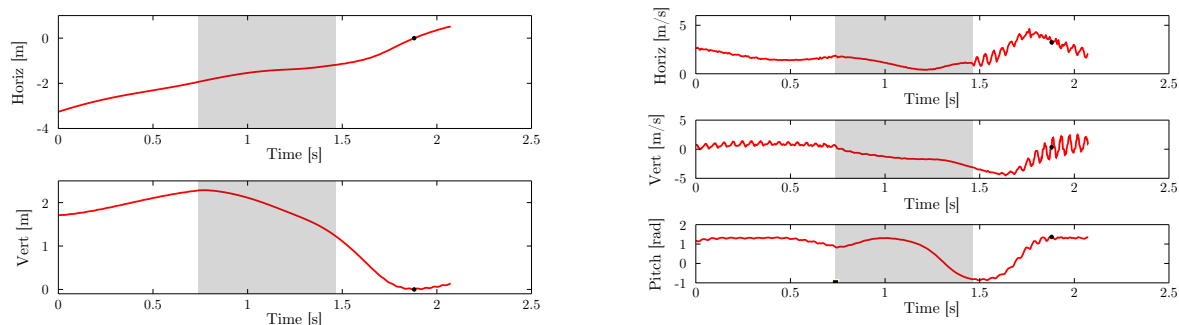


Figure 6.7: Telemetry data for a single trial of the H²Bird diving experiment with the recovery controller active over the time period. The position is on the left and the horizontal velocity, vertical velocity, and pitch angle are on the right. The gray shading represents the unpowered portion of the dive. The black marker indicates the conditions at the lowest vertical position in the dive.

at the rightmost edge of the gray region. The black marker indicates the conditions of the lowest vertical point in the trajectory. The velocity in the parts of the trajectory when the wings are flapping is noisy because the velocity is a time derivative of the tracked position. The Vicon only provides angular position in axis-angle format and the translational position of the tracked body by default. Since the Vicon data was streaming directly to the control

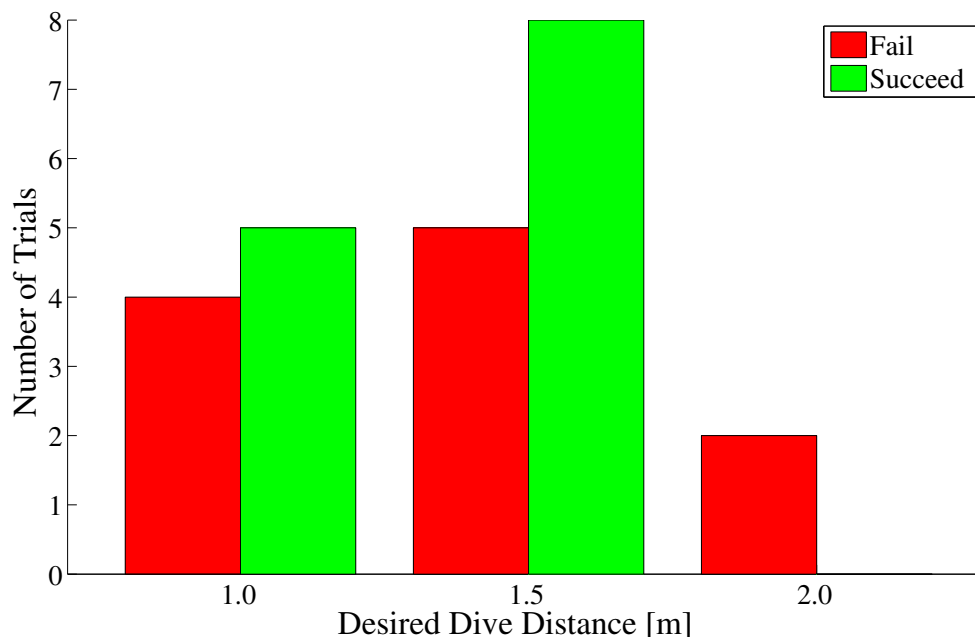


Figure 6.8: Number of successful trials (green) and failed trials (red) for each tested dive distance.

program, we could not take advantage of the built-in post-processing provided by the tracking software.

We conducted 24 total dive trials: 9 at a 1.0 meter desired dive depth, 13 at a 1.5 meter depth, and 2 at a 2.0 meter depth. The results of the dives are summarized in Figure 6.8. The green bars represent successful dives and the red bars represent failed dives. For a 1.0 meter target dive distance, we observed a success rate of approximately 55 percent, for a 1.5 meter target distance we observed a 61 percent success rate, and for 2.0 meters we observed no successful dives. We did not attempt more than two trials for the target distance of 2.0 meters because the robot showed no signs of success at that set-point and would hit the ground at a high rate of speed, eventually resulting in damage.

The total dive distance from the point that the dive is initiated (the wings stop flapping) to the lowest vertical point in the maneuver is shown in Figure 6.9 for the 1.0 meter and 1.5 meter desired distance set-points. Unfortunately, the Vicon tracking for some of the trials was of low quality in the end parts of 1 dive at 1.0 meters and 1 dive at 1.5 meters, so these trials were omitted. Although the H²Bird managed to enter the backward reachable set in each of these trials and successfully recover, the mean dive distance was 2.2 meters for both 1.0 and 1.5 meter set-point. The reason for this can be explained by the vertical height data in Figure 5.4. As shown in the figure, the mean dive distance for our initial trials used to train our models was 1.8 meters. The lack of diversity in the dive distances that the model was based upon resulted in a model that consistently yields dive distances around 2.2 meters.

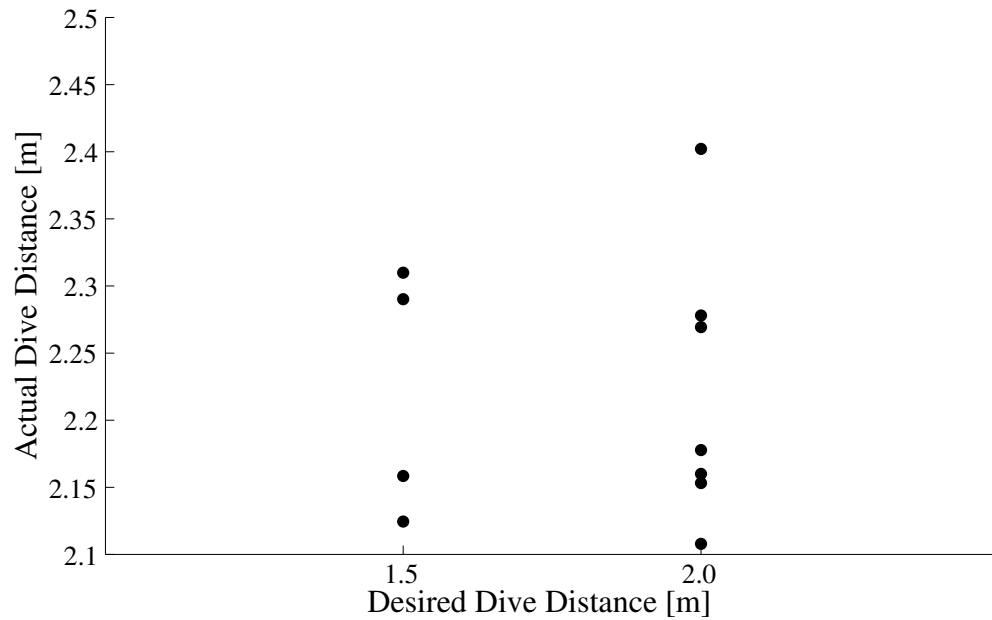


Figure 6.9: The total dive distance for each successful trial for the 1 meter and 1.5 meter desired dive distance set points.

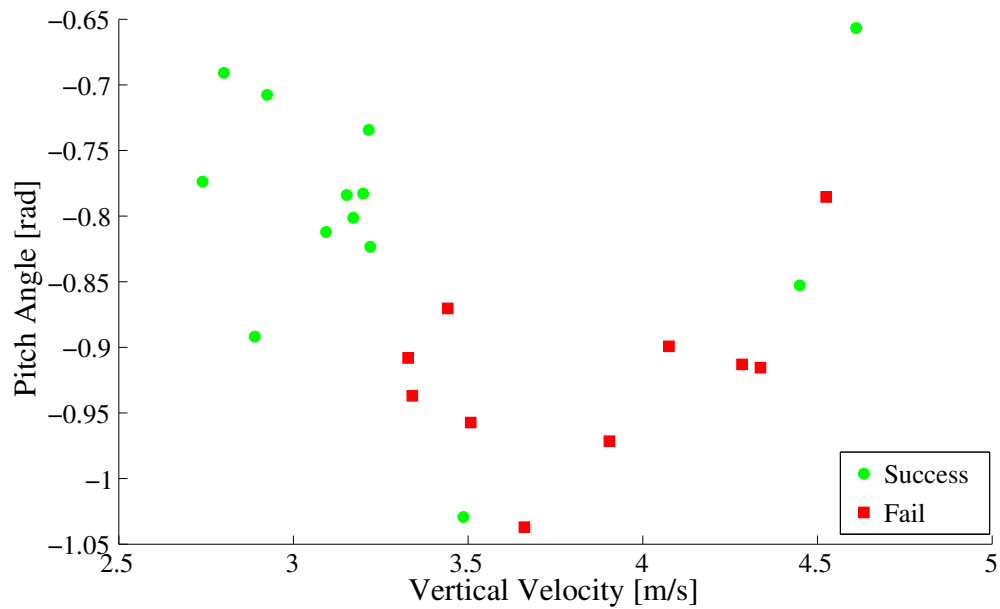


Figure 6.10: The vertical velocity vs. the pitch angle at the time of recovery for successful (green circle) and failed (red square) trials.

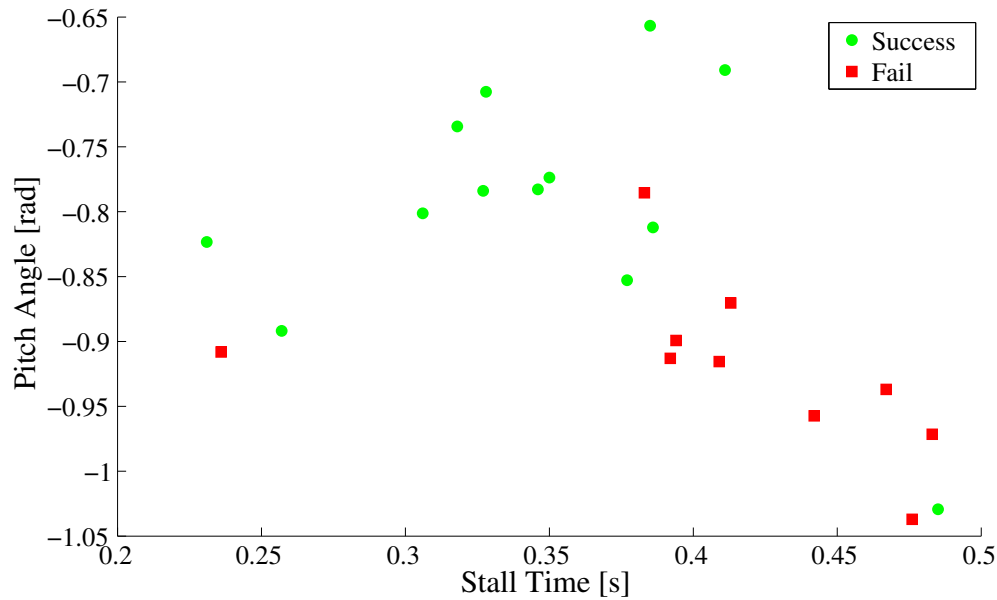


Figure 6.11: The pitch angle at the top of the stall and the time to a downward pitch angle during the unpowered dive for successful (green circle) and failed (red square) trials.

To hypothesize why some of the trials resulted in failed dives, we examined the pose of the H²Bird at the beginning of the dive, at the maximum pitch angle during the stall, and at the point that recovery is initiated. The combinations of variables that provided the most prominent clues as to why a particular trial failed are in Figures 6.10 and 6.11. Figure 6.10 shows the vertical velocity and the pitch angle at the recovery point when the wings start flapping for successful (green circle) and unsuccessful (red square) trials. All of the failed trials exhibit both a high vertical velocity and high downward pitch angle at the time of recovery. These conditions are not ideal because when the wings begin flapping, they will induce both a downward pitch moment and increase in velocity, which further exacerbates the problem. When this happens, the H²Bird is likely to lose enough height before reaching a positive pitch angle, and crash into the ground. These recovery conditions are caused by the properties of the stalling behavior shown in Figure 6.11. On the x axis is the amount of time spent in a stall in the unpowered dive, and on the y axis is the pitch angle at the recovery point. The stall time is the time elapsed between the maximum pitch angle during the stall and the first occurrence of a negative pitch angle. On average, the failed trials exhibit higher stall times, which leads to increased vertical velocities at the point of recovery.

Chapter 7

Conclusions

In this thesis, we detailed methods for designing controllers and models to complete non-equilibrium maneuvers for a flapping-winged micro aerial vehicle.

We compared the predicted behavior of an ornithopter from a wind tunnel measured data set to measured free flight equilibrium conditions. As a comparison metric, we determined the equilibrium velocity magnitude and angle of attack predicted by the wind tunnel data set for given free flight measured elevator deflections. The wind tunnel equilibrium was represented as the angle of attack and velocity on the pitch moment level set at zero that minimizes the net horizontal and vertical forces measured in the wind tunnel. We found that the wind tunnel underestimates the angle of attack observed in free flight at equilibrium by approximately 15 degrees, whereas the error between the equilibrium velocities between the two data sets is approximately 0.1 m/s for an average flight speed of 2.0 m/s. Although the wind tunnel can provide some information about non-equilibrium flight, the information was not sufficient for online control of the flier.

We demonstrated a method for launching a flapping-winged MAV, the H²Bird, using a legged hexapod, the VelociRoACH. We found that it is possible to reliably launch the H²Bird from atop the VelociRoACH for successful flight, provided the legged robot reaches an appropriate minimum velocity of 1.2 m/s. A failure mode for velocities greater than this minimum is improper positioning in the carrying cradle as a result of pitch velocity impulses as the VelociRoACH initially accelerates to run. Additionally, the H²Bird catching on the stand can cause a downward pitch immediately after launch that can cause the launch to fail.

Although the H²Bird causes the VelociRoACH motors to consume approximately 18.1 to 24.5 percent more power, we found that the H²Bird can have some beneficial effects on the running performance of the VelociRoACH. Just by simply resting on top of the VelociRoACH, the H²Bird can reduce the variance of the roll and pitch velocities by about 80 and 90 percent, respectively. This pitch and roll damping can allow the legged robot to run more stably. Although the H²Bird reduces the average running velocity at 17 Hz stride frequency by 8.4 percent, by flapping the wings at 5 Hz, the average velocity can be increased by 12.7 percent. We make no claims about the efficiency of the running gait for the VelociRoACH by carrying

the H²Bird. There are better ways to provide rotational damping, but the ability to carry a robot with an advantageous mode of transportation with minimal losses is an important result.

To design computationally tractable models for online control during a diving maneuver, we fit linear models around clusters of flight conditions from a data set of sample dives. We segmented the data sets into the clusters of similar flight poses using k-means, and fit the models using linear least squares. We then used the models to compute the backward reachable set from a goal set of previously experienced dive end poses. Using the backward reachable set, we determined when to begin recovery in the dives by checking our current pose and simulated poses 0.1 seconds ahead for inclusion in the reachable set. Using this control method, we were able to complete 2.2 meter dives at a success rate of 60 percent. The majority of the failures can be attributed to extended stall conditions that result in dangerous vertical velocities and pitch angles during recovery.

Future Work

Linear piece-wise affine modeling of segments of flight conditions within a maneuver has proven to be an effective method for determining transition points between hybrid controllers. In the future, this method could be extended to incorporate models of chunks of additional maneuvers. These sections of flight could form a library of maneuver segments that could be stitched together to form flight patterns and trajectories that were not previously experienced by the robot. The controllers for each segment could be stored on-board the robot to enable autonomous navigation or obstacle avoidance by picking maneuvers applicable to an observed situation. Since the models are linear, the computational overhead for on-board look-ahead simulation or computation of feedback controllers is reduced compared to complex nonlinear models.

The work in this thesis could be extended to modeling and control for landing the H²Bird on top of the VelociRoACH, both stationary and running. An interesting problem would be to determine the feasibility of predicting the recovery point for a dive or controlled descent in which the goal state is time-varying.

Appendix A

Reachability Sets

A.1 Position Reachability Sets




Figure A.1: The backward reachability polytopes in the position space 1 to 4 timesteps from the terminal set. The timestep used was 0.025 seconds.




Figure A.2: The backward reachability polytopes in the position space 5 to 8 timesteps from the terminal set. The timestep used was 0.025 seconds.




Figure A.3: The backward reachability polytopes in the position space 9 to 12 timesteps from the terminal set. The timestep used was 0.025 seconds.

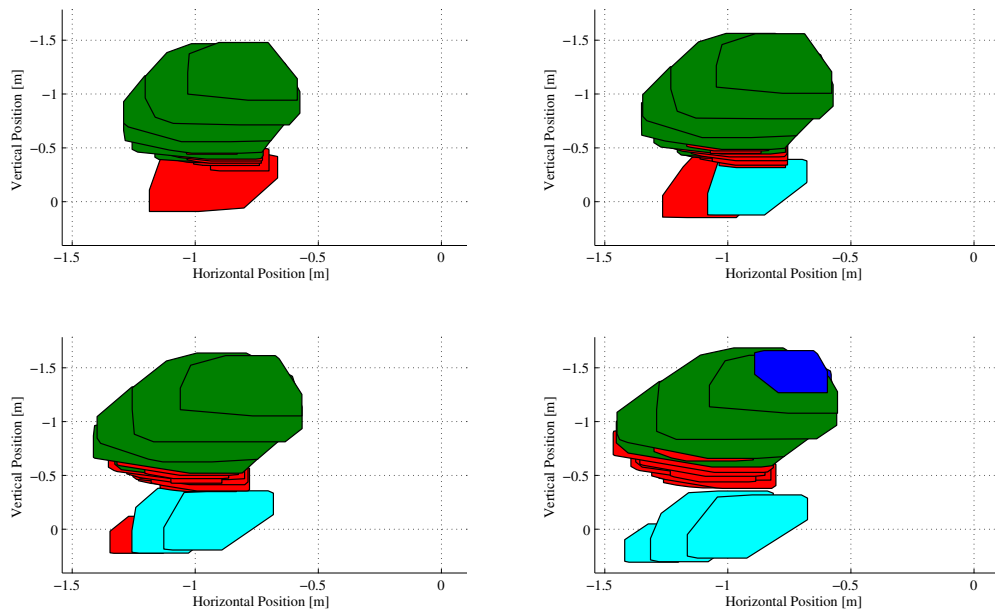


Figure A.4: The backward reachability polytopes in the position space 13 to 16 timesteps from the terminal set. The timestep used was 0.025 seconds.

A.2 Velocity and Pitch Angle Reachability Sets

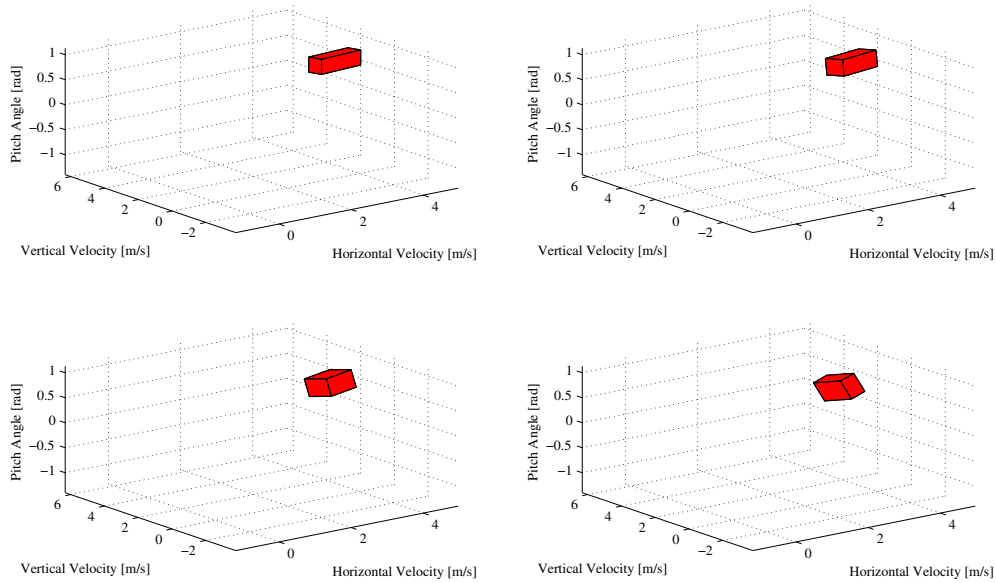


Figure A.5: The backward reachability polytopes in the velocity and pitch angle space 1 to 4 timesteps from the terminal set. The timestep used was 0.025 seconds.

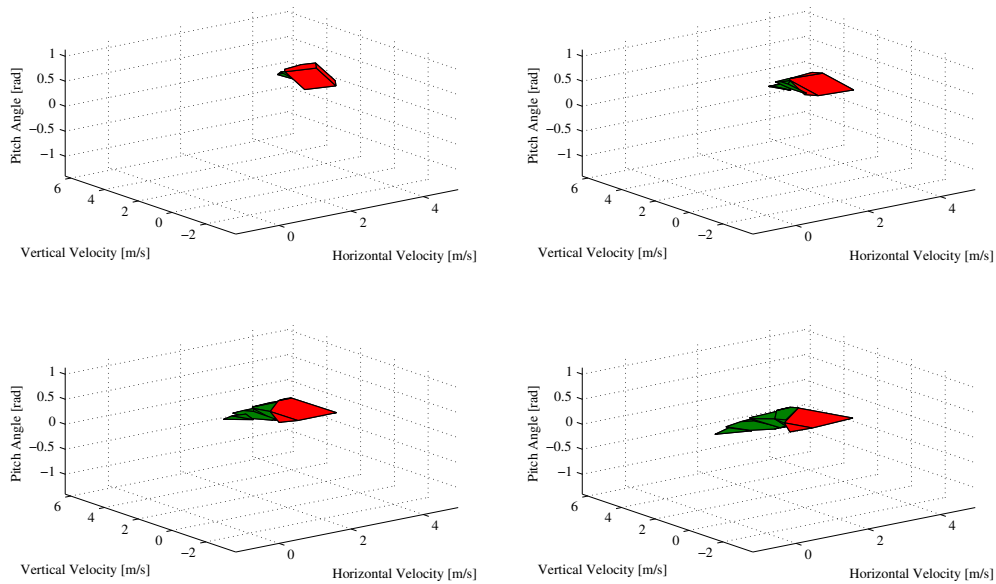


Figure A.6: The backward reachability polytopes in the velocity and pitch angle space 5 to 8 timesteps from the terminal set. The timestep used was 0.025 seconds.

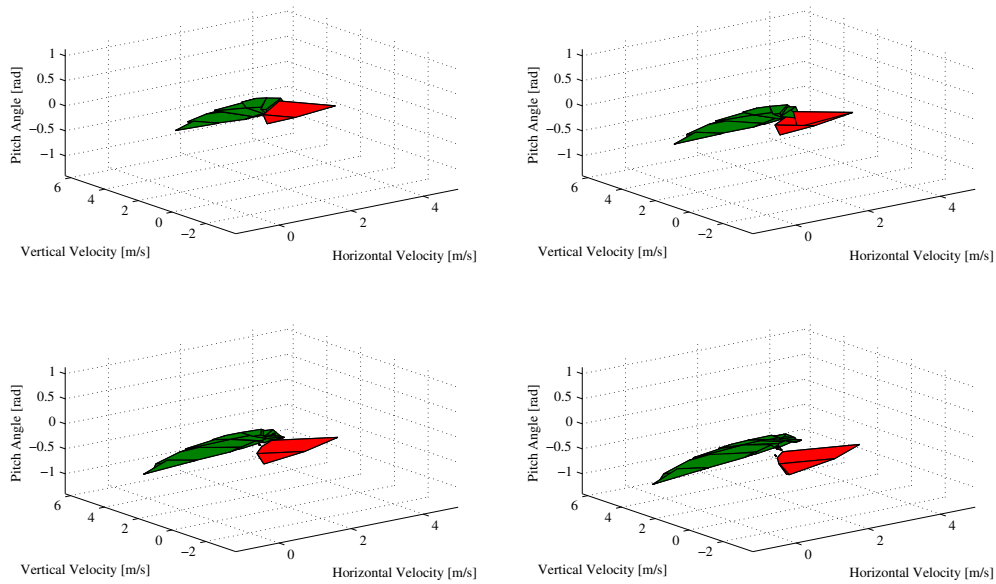


Figure A.7: The backward reachability polytopes in the velocity and pitch angle space 9 to 12 timesteps from the terminal set. The timestep used was 0.025 seconds.

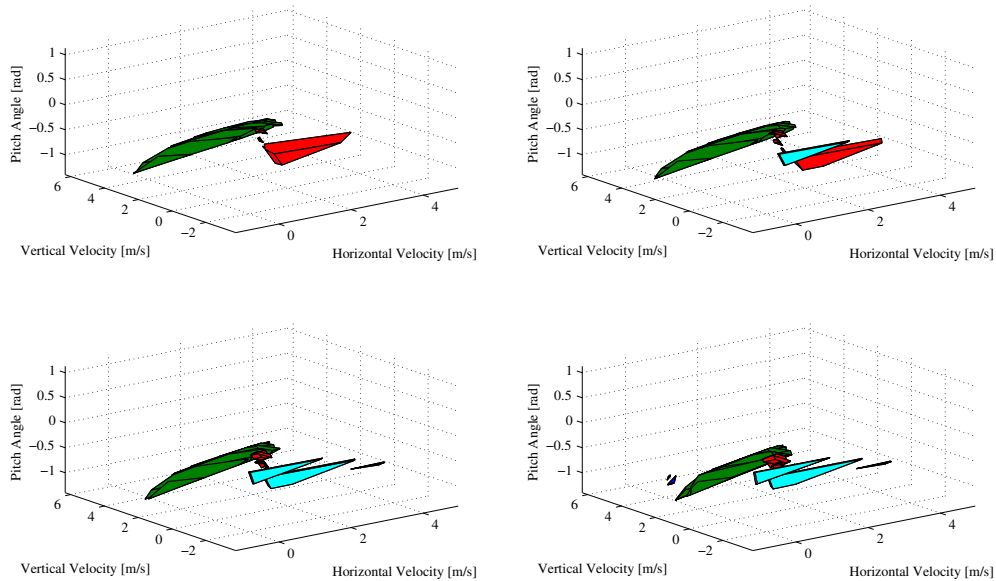


Figure A.8: The backward reachability polytopes in the velocity and pitch angle space 13 to 16 timesteps from the terminal set. The timestep used was 0.025 seconds.

Bibliography

- [1] R. Ames, O. Wong, and N. Komerath. “On the flowfield and forces generated by a flapping rectangular wing at low Reynolds number.” In: *Progress in Astronautics and Aeronautics* 195 (2001), pp. 287–305.
- [2] R.J. Bachmann et al. “A biologically inspired micro-vehicle capable of aerial and terrestrial locomotion”. In: *Mechanism and Machine Theory* 44.3 (2009). Special Issue on Bio-Inspired Mechanism Engineering, pp. 513–526.
- [3] S.S. Baek, F.L. Garcia Bermudez, and R.S. Fearing. “Flight Control for Target Seeking by 13 gram Ornithopter.” In: *IEEE Intl. Conf. on Intelligent Robots and Systems*. 2011, pp. 2674–2681.
- [4] S.S. Baek and R.S. Fearing. “Flight forces and altitude regulation of 12 gram I-Bird”. In: *IEEE Intl. Conf. on Biomedical Robotics and Biomechatronics*. 2010, pp. 454–460.
- [5] M. A. Bolender. “Rigid Multi-Body Equations-of-Motion for Flapping Wing MAVs using Kanés Equations”. In: *AIAA Guidance, Navigation, and Control Conference*. 2009.
- [6] B. Cheng and X. Deng. “Translational and Rotational Damping of Flapping Flight and Its Dynamics and Stability at Hovering”. In: *IEEE Transactions on Robotics* Vol. 27.5 (2011), pp. 849–864.
- [7] G. C. H. E. de Croon et al. “Design, aerodynamics, and vision-based control of the DelFly”. In: *International Journal of Micro Air Vehicles* Vol. 1.2 (June 2009), pp. 71–97.
- [8] G.C.H.E. de Croon et al. “Design, aerodynamics, and vision-based control of the DelFly”. In: *International Journal of Micro Air Vehicles* 1.2 (2009), pp. 71–97. DOI: 10.1260/175682909789498288. eprint: <http://dx.doi.org/10.1260/175682909789498288>. URL: <http://dx.doi.org/10.1260/175682909789498288>.
- [9] K.M.E. De Clercq et al. “Aerodynamic Experiments on DelFly II: Unsteady Lift Enhancement”. In: *International Journal of Micro Air Vehicles* 1.4 (2009), pp. 255–262. DOI: 10.1260/175682909790291465. eprint: <http://mav.sagepub.com/content/1/4/255.full.pdf+html>. URL: <http://mav.sagepub.com/content/1/4/255.abstract>.

- [10] W. B. Dickson et al. “An integrative model of insect flight control”. In: *44th American Institute of Aeronautics and Astronautics Sciences Meeting*. 2006.
- [11] R. Dudley. “Mechanisms and Implications of Animal Flight Maneuverability”. In: *Integrative and Comparative Biology* 42.1 (2002), pp. 135–140. DOI: 10.1093/icb/42.1.135. eprint: <http://icb.oxfordjournals.org/content/42/1/135.full.pdf+html>. URL: <http://icb.oxfordjournals.org/content/42/1/135.abstract>.
- [12] I. Faruque, P.D. Samuel, and J. Humbert. “Moment Generation of Stabilizing Axes for Insect-Inspired Flapping Wing Flight”. In: *AIAA Atmospheric Flight Mechanics Conference* (2014).
- [13] J. Grauer and J. Hubbard. “Multibody model of an ornithopter”. In: *Journal of Guidance, Control, and Dynamics*. Vol. 32. 5. 2009, pp. 1675–1679.
- [14] J. Grauer et al. “System identification of an ornithopter aerodynamics model”. In: *AIAA Atmospheric Flight Mechanics Conference*. 2010.
- [15] D.W. Haldane et al. “Animal-inspired design and aerodynamic stabilization of a hexapedal millirobot.” In: *Intl. Conf. on Robotics and Automation*. 2013.
- [16] JH Han, JY Lee, and DK Kim. “Ornithopter Modeling for Flight Simulation”. In: *Intl. Conf. on Control, Automation and Systems*. 2008, pp. 1773–1777.
- [17] R. Julian et al. “Cooperative Control for Window Traversal with an Ornithopter MAV and Lightweight Ground Station”. In: *12th Intl. Conf. on Autonomous Agents and Multiagent Systems, 2013. Proceedings*. 2013.
- [18] A. Kalantari and M. Spenko. “Design and experimental validation of HyTAQ, a Hybrid Terrestrial and Aerial Quadrotor.” In: *Intl. Conf. on Robotics and Automation*. IEEE, 2013, pp. 4445–4450.
- [19] Z. A. Khan and S. K. Agrawal. “Force and Moment Characterization of Flapping Wings for Micro Air Vehicle Application”. In: *American Control Conference*. 2005.
- [20] Z.A. Khan and S.K. Agrawal. “Control of Longitudinal Flight Dynamics of a Flapping-Wing Micro Air Vehicle Using Time-Averaged Model and Differential Flatness Based Controller”. In: *American Control Conference*. 2007, pp. 5284–5289.
- [21] A. Kossett and N. Papanikolopoulos. “A robust miniature robot design for land/air hybrid locomotion.” In: *Intl. Conf. on Robotics and Automation*. IEEE, 2011, pp. 4595–4600.
- [22] DK Lee and JH Han. “Flight Controller Design of a Flapping-Wing MAV in a Magnetically Levitated Environment”. In: *Intl. Conf. on Robotics and Automation*. 2013.
- [23] D. Lentink et al. “How swifts control their glide performance with morphing wings”. In: *Nature* 446 (), pp. 1082–1085.
- [24] D. Lentink et al. “Vortex interactions with flapping wings and fins can be unpredictable”. In: *Biology Letters* 6 (3 2010), pp. 394–397.

- [25] H. Liu and K. Kawachi. “Leading-edge vortices of flapping and rotary wings at low Reynolds number.” In: *Progress in Astronautics and Aeronautics* 195 (2001), pp. 275–285.
- [26] K.Y. Ma et al. “Controlled Flight of a Biologically Inspired, Insect-Scale Robot”. In: *Science* 340.6132 (2013), pp. 603–607. DOI: 10.1126/science.1231806. eprint: <http://www.sciencemag.org/content/340/6132/603.full.pdf>. URL: <http://www.sciencemag.org/content/340/6132/603.abstract>.
- [27] J. Moore and R. Tedrake. “Control synthesis and verification for a perching UAV using LQR-Trees”. In: *Decision and Control (CDC), 2012 IEEE 51st Annual Conference on*. 2012, pp. 3707–3714. DOI: 10.1109/CDC.2012.6425852.
- [28] C. Orlowski and A. Girard. “Modeling and simulation of nonlinear dynamics of flapping wing micro air vehicles”. In: *AIAA Journal* 49.5 (2011), pp. 969–981.
- [29] V.M. Ortega-Jimenez et al. “Into turbulent air: size-dependent effects of von Kármán vortex streets on hummingbird flight kinematics and energetics”. In: *Proceedings of the Royal Society of London B: Biological Sciences* 281.1783 (2014). ISSN: 0962-8452. DOI: 10.1098/rspb.2014.0180.
- [30] K. Peterson and R.S. Fearing. “Experimental dynamics of wing assisted running for a bipedal ornithopter.” In: *Intl. Conf. on Intelligent Robotics and Systems*. 2011.
- [31] K. Peterson et al. “A wing-assisted running robot and implications for avian flight evolution”. In: *Bioinspiration and Biomimetics* 6.4 (2011), p. 118102.
- [32] A. T. Pfeiffer et al. “Ornithopter Flight Simulation Based on Flexible Multi-Body Dynamics”. In: *Journal of Bionic Engineering* Vol. 7.1 (2010), pp. 102–111.
- [33] B. Ponitz et al. “Diving-Flight Aerodynamics of a Peregrine Falcon (*Falco peregrinus*)”. In: *PLOS One* 9 (2), pp. 1–13.
- [34] A.O. Pullin et al. “Dynamic turning of 13 cm robot comparing tail and differential drive.” In: *Intl. Conf. on Robotics and Automation*. IEEE, 2012, pp. 5086–5093.
- [35] C. Rose and R.S. Fearing. “Comparison of ornithopter wind tunnel force measurements with free flight”. In: *Intl. Conf. on Robotics and Automation*. IEEE, 2014.
- [36] C. Rose and R.S. Fearing. *Flight Simulation of an Ornithopter*. MS Thesis. EECS Department, University of California, Berkeley, 2013. URL: <http://www.eecs.berkeley.edu/Pubs/TechRpts/2013/EECS-2013-60.html>.
- [37] S.P. Sane and M.H. Dickinson. “The aerodynamic effects of wing rotation and a revised quasi-steady model of flapping flight”. In: *Journal of Experimental Biology* Vol. 205.8 (2002), pp. 1087–1096.
- [38] L. Schenato, D. Campolo, and S. Sastry. “Controllability issues in flapping flight for biomimetic micro aerial vehicles (MAVs)”. In: *42nd IEEE Conf. on Decision and Control*. Vol. 6. 2003, pp. 6441–6447.

- [39] R. Tedrake. “LQR-Trees: Feedback motion planning on sparse randomized trees”. In: *Papers of the fifth annual Robotics: Science and Systems conference*. 2009.
- [40] R.J. Wood. “The first takeoff of a biologically inspired at-scale robotic insect”. In: *IEEE Trans. Robot* (2008).