

AUTONOMOUS FLIGHTWORTHINESS DETERMINATION FOR MODULAR VERTICAL LIFT VEHICLES

A Dissertation
Presented to
The Academic Faculty

by

Jonathan Warner

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy in the
George W. Woodruff School of Mechanical Engineering

Georgia Institute of Technology
December 2018

COPYRIGHT © 2018 BY JONATHAN WARNER

AUTONOMOUS FLIGHTWORTHINESS DETERMINATION FOR MODULAR VERTICAL LIFT VEHICLES

Approved by:

Dr. Jonathan Rogers, Advisor
School of Mechanical Engineering
Georgia Institute of Technology

Dr. Eric Johnson
School of Aerospace Engineering
Pennsylvania State University

Dr. Jun Ueda
School of Mechanical Engineering
Georgia Institute of Technology

Dr. Patricio Vela
School of Electrical and Computer
Engineering
Georgia Institute of Technology

Dr. Aldo Ferri
School of Mechanical Engineering
Georgia Institute of Technology

Date Approved: September 19, 2018

To Mr. Vaughn

ACKNOWLEDGEMENTS

First, I would like to thank my advisor, Dr. Jonathan Rogers, for mentoring me through the PhD program and for being a source of continual support. I am deeply grateful for all the time that he devoted to guiding me through my graduate career. His enthusiasm, insight, and creativity towards engineering made researching an enjoyable and rewarding experience. My special gratitude goes to my thesis committee members, Dr. Jun Ueda, Dr. Aldo Ferri, Dr. Eric Johnson, and Dr. Patricio Vela for their guidance and assistance.

I would like to acknowledge my mother, Marcie Reynolds, my father, Bill Warner, my stepmother, Laura Warner, my sister, Kaleigh Stevens, and my brother, David Warner for their love and unwavering support. I want to thank my grandmother, Patricia Warner, and my grandfather, Bill Warner Sr., for taking a direct involvement in my life and in my education, I appreciate everything that was done on my behalf. I would like to express special appreciation to the Whittingtons for their constant support and encouragement throughout my time at Georgia Tech.

I want to thank my fellow Aggies, Austin Baty, Andrew Brown, Daniel Mann, and Cole Spencer for all the fun times, graduate school would not have been nearly as enjoyable without them. I want to thank my friends Justin King, Justin Wilbanks, and Lee Fowler for all the runs, discussions, and laughs. I would like to thank my lab mates for all the coffee breaks, off-topic debates, and memorable moments, it was a privilege working in the iREAL lab. Finally, I would like to express my gratitude to all my friends and family who encouraged me through the ups and downs, I cannot say thank you enough.

TABLE OF CONTENTS

ACKNOWLEDGEMENTS	iv
LIST OF TABLES	vi
LIST OF FIGURES	vii
LIST OF SYMBOLS AND ABBREVIATIONS	xi
SUMMARY	xiii
CHAPTER 1. INTRODUCTION	1
CHAPTER 2. COMPOSITE AIRCRAFT DYNAMIC MODEL	7
2.1 Rigid Body Dynamics	8
2.2 Vehicle Model	9
2.3 Ground Contact Model	12
CHAPTER 3. FLIGHTWORTHINESS ALGORITHM DESIGN	19
3.1 Spin-up Methodology	20
3.2 Excitation Axis Detection	22
3.3 Vehicle Position Estimation Algorithm	28
3.3.1 General Solution Strategy	29
3.3.2 Parallel Extension	34
3.4 Flightworthiness Determination	36
3.5 Bayesian Inference	38
CHAPTER 4. ERROR ANALYSIS	48
CHAPTER 5. SIMULATION RESULTS	53
5.1 Example Case Results	54
5.2 Monte Carlo Results	64
CHAPTER 6. EXPERIMENTAL SYSTEM DESCRIPTION	73
CHAPTER 7. EXPERIMENTAL RESULTS	79
7.1 Example Case Results	79
7.2 Monte Carlo Results	99
CHAPTER 8. CONCLUSION	105
APPENDIX A. Aircraft Weight Estimation	107
REFERENCES	109

LIST OF TABLES

Table 1	Flightworthiness Determination Trade Study Conclusions.	72
Table 2	Aircraft Weight and Inertia Properties.	100
Table 3	Flightworthiness Determination Trade Study Performance Metrics.	104

LIST OF FIGURES

Figure 1	Composite Aircraft Diagram and Coordinate system.	8
Figure 2	Simulated Excitation Event.	17
Figure 3	Sample Contact Model Validation Results.	18
Figure 4	Schematic of Flightworthiness Determination Algorithm.	20
Figure 5	Schematic of the Bayesian Update Procedure.	21
Figure 6	Example Contact Axis Excitation Angle.	25
Figure 7	Example Contact Axis Orientation - β .	25
Figure 8	Example Contact Axis Orientation - γ .	26
Figure 9	Example Contact Point Excitation Angle.	26
Figure 10	Example Contact Axis Orientation - β .	27
Figure 11	Example Contact Axis Orientation - γ .	27
Figure 12	Example Composite Aircraft Configuration and Contact Axis Definitions.	29
Figure 13	Example Actual Composite Aircraft Configuration (left) and Corresponding Slot Definitions (right). Each circle represents a vehicle location.	40
Figure 14	Example Equivalent Thrust Vector for given Hypothesis.	41
Figure 15	Example 2D Dynamic Transition FBD.	41
Figure 16	Example Likelihood Calculation given Normal Mass Center Location Distribution.	43
Figure 17	Example Placement History of Vehicles on Hypothetical Payload.	45
Figure 18	Example Modular Vehicle Configuration with Rectangular Contact Axes.	54
Figure 19	Rectangular Aircraft Excitation Axis History.	55
Figure 20	Vehicle Effectiveness Scores for Axes 1 (left) and 2 (right).	56

Figure 21	Vehicle Effectiveness Scores for Axes 3 (left) and 4 (right).	56
Figure 22	Example Vehicle Configuration and Position Estimates for Rectangular Contact Axes.	58
Figure 23	Estimated and Actual Vehicle Throttle Percentages Required to Hover for Rectangular Axes.	59
Figure 24	Example Vehicle Configuration and Position Estimates for Triangular Contact Axes.	60
Figure 25	Triangular Aircraft Excitation Axis History.	60
Figure 26	Estimated and Actual Vehicle Throttle Percentages Required to Hover for Triangular Axes.	61
Figure 27	Example Vehicle Configuration and Position Estimates for Pentagon Contact Axes.	62
Figure 28	Pentagon Aircraft Excitation Axis History.	63
Figure 29	Estimated and Actual Vehicle Throttle Percentages Required to Hover for Pentagon Axes.	63
Figure 30	Vehicle Position Estimation Error (left) and Weight Estimation Error (right) for Increasing Required Excitations per Axis.	66
Figure 31	Excessive Excitations (left) and Unsuccessful Simulations (right) for Increasing Required Excitations per Axis.	66
Figure 32	Vehicle Position Estimation Error (left) and Weight Estimation Error (right) for Increasing Aircraft Thrust to Weight Ratio.	68
Figure 33	Excessive Excitations (left) and Unsuccessful Simulations (right) for Increasing Aircraft Thrust to Weight Ratio.	68
Figure 34	Excessive Excitations (left) and Unsuccessful Simulations (right) for Increasing Mass Center Perturbations.	70
Figure 35	Vehicle Position Estimation Error (left) and Weight Estimation Error (right) for Increasing Mass Center Perturbations.	71
Figure 36	Modular Vehicle Design.	73
Figure 37	Experimental Payload Platforms.	74
Figure 38	Flightworthiness Determination Architecture Design.	75

Figure 39	Experimental Lift Vehicle Thrust-Throttle Mapping.	76
Figure 40	Composite Aircraft Endurance Experiment.	77
Figure 41	Sample Comparison between Simulated and Experimental Excitation.	78
Figure 42	Flightworthy (left) and Unflightworthy (right) Configurations of the Rectangle Aircraft.	79
Figure 43	Excitation Axis History for Flightworthy Rectanglar Aircraft.	80
Figure 44	Probabilistic Vehicle Placement History for Flightworthy Rectangle Aircraft.	82
Figure 45	Vehicle Position Estimates for Flightworthy Rectangle Aircraft.	83
Figure 46	Estimated and Actual Vehicle Hover Throttle Percentages for Flightworthy Rectangle Aircraft.	84
Figure 47	Excitation Axis History for Unflightworthy Rectanglar Aircraft.	85
Figure 48	Flightworthy (left) and Unflightworthy (right) Configurations of the Trianglar Aircraft.	86
Figure 49	Excitation Axis History for Flightworthy Trianglar Aircraft.	87
Figure 50	Probabilistic Vehicle Placement History for Flightworthy Triangle Aircraft (vehicles originate at origin at the start of the experiment).	88
Figure 51	Vehicle Position Estimates for Flightworthy Trianglar Aircraft.	89
Figure 52	Estimated and Actual Vehicle Hover Throttle Percentages for Flightworthy Triangle Aircraft.	90
Figure 53	Excitation Axis History for Unflightworthy Trianglar Aircraft.	90
Figure 54	Vehicle Position Estimates for Unflightworthy Triangle Aircraft.	92
Figure 55	Estimated and Actual Vehicle Hover Throttle Percentages for Unflightworthy Triangle Aircraft.	92
Figure 56	Flightworthy (left) and Unflightworthy (right) Configurations of the Pentagon Aircraft.	93
Figure 57	Excitation Axis History for Flightworthy Pentagon Aircraft.	94
Figure 58	Equivalent Thrust to Excitation Axis Mapping for Pentagon Payload.	95

Figure 59	Vehicle Position Estimates for Flightworthy Pentagon Aircraft.	96
Figure 60	Estimated and Actual Vehicle Hover Throttle Percentages for Flightworthy Pentagon Aircraft.	97
Figure 61	Excitation Axis History for Unflightworthy Pentagon Aircraft.	98
Figure 62	Vehicle Position Estimates for Unflightworthy Pentagon Aircraft.	98
Figure 63	Estimated and Actual Vehicle Hover Throttle Percentages for Unflightworthy Pentagon Aircraft.	99
Figure 64	Vehicle Position Estimates for Square Aircraft.	101
Figure 65	Vehicle Position Estimates for Triangle Aircraft.	102
Figure 66	Vehicle Position Estimates for Pentagon Aircraft.	103

LIST OF SYMBOLS AND ABBREVIATIONS

a_i	contact axis i for composite aircraft
$d_{cg,i}$	distance from contact axis i to mass center of aircraft in the \vec{J}_i direction
I	moment of inertia matrix of the composite aircraft about the composite system mass center
\vec{I}_i, \vec{J}_i	unit vectors in body-fixed coordinate system oriented about axis a_i
L_{ik}	distance from axis i to axis k in the \vec{J}_i direction
m	number of contact axes for composite aircraft
m_{ca}	mass of the composite aircraft
n	number of modular vehicles attached to the payload
L, M, N	total moment acting about the composite aircraft's mass center in the body reference frame
p_i	number of unique excitations about axis i
p, q, r	angular velocity components of composite aircraft in the body reference frame
$\vec{r}_{a/b}$	position vector from point b to point a in the body reference frame
T_j	total thrust force magnitude produced by vehicle j
T_H	transformation matrix from the inertia reference frame to the composite aircraft reference frame
u, v, w	translational velocity components of composite aircraft mass center in the body reference frame
W	weight of the composite aircraft
x, y, z	composite aircraft mass center location in the inertial frame
X, Y, Z	total force components acting on composite aircraft in the composite aircraft body frame

ϕ, θ, ψ Euler roll, pitch, and yaw angles

SUMMARY

Payload transportation via connected modular unmanned aerial vehicles (UAVs) is an emerging new area that offers unique advantages over other forms of aerial logistics. When considering rigidly attached modular vertical lift UAVs, differing payloads and vehicle attachment geometries have a significant effect on the composite aircraft's dynamic response during takeoff and stabilization. With no prior knowledge of payload parameters or vehicle attachment geometry, there is no inherent flightworthiness guarantee for a specific connected configuration. On-ground flightworthiness determination can be used to ensure acceptable performance during vehicle take-off or to prescribe changes to the vehicle attachment geometry if necessary. This work introduces an algorithm to determine flightworthiness while in partial ground contact by estimating the vehicle attachment positions and payload weight. The algorithm utilizes a probabilistic estimate of vehicle placement about the payload derived through a Bayesian learning technique to generate the necessary data to deterministically estimate the attached vehicles' positions. Following a description of the algorithm, simulation results are presented to illustrate the performance of the algorithm for a variety of modular aircraft configurations. The algorithm is experimentally validated through a series of tests using prototype modular vehicles.

CHAPTER 1. INTRODUCTION

Unmanned aerial vehicles (UAVs) have become increasingly popular for a variety of missions ranging from surveillance to payload transportation. The vast majority of UAV research and development has focused on the operation of single vehicles, oftentimes serving as a direct replacement for manned vehicles performing a similar mission. An exception to this is the emerging research area of swarming UAV operations. In this context UAVs are physically separate and interact through some type of cooperative control algorithm to collaboratively perform a mission. For the specific mission of payload manipulation and transportation, prior research has focused largely on use of single vehicles [1]. Nevertheless, numerous advantages may be obtained by performing payload transportation missions using cooperative teams of UAVs that attach to a payload either manually [2] or autonomously [3-6]. These advantages include scalability as well as portability considerations. For instance, to lift payloads of various weights, an operator may use a larger number of the same aircraft, rather than needing separate aircraft for different weight classes. This may ease logistical burden significantly. Likewise, a modular approach allows lifting capacity to be distributed across numerous smaller vehicles. In this manner, large lifting capabilities can exist in distributed, human-transportable packages. Due to these properties, the use of cooperative UAVs for search and rescue missions has been explored [7-9]. While this type of system provides clear advantages, numerous challenges arise including problems of control allocation [10] and of cooperative flight control in particular. The cooperative flight control problem has been investigated for a wide range of transportation scenarios [11-18] including cases where the

team of UAVs were unable to communicate with each other [19, 20]. This cooperative flight control problem involving physically connected vehicles becomes even more difficult in cases where the inertial characteristics of the payload are largely unknown. This is likely to be the case in practical scenarios, where teams of vehicles are deployed to carry payload items which have not been carefully surveyed and weighed.

Recently several authors have begun to explore this emerging field of cooperative UAV manipulation. Lindsey *et al.* [21] proposed a new model for construction of Special Cubic Structures using teams of quadrotors. Although Lindsey *et al.* did not address cooperative lifting, the research served as a proof-of-concept and as a starting point for future research into collaborative payload transportation and manipulation. Duffy and Samaritano [22] explored the scalability of payload transportation using modular, multi-vehicle propulsion and commented on the feasibility of the approach. A similar, small-scale approach to cooperative flight of modular vehicles was presented by Oung *et al* [23]. The approach of tethering the vertical lift vehicle(s) to a payload has been a topic of recent research [24-26]. Sreenath *et al* [27] introduced trajectory planning for cooperative transportation of a cable-suspended payload by multiple quadrotors.

When considering cooperative UAV manipulation, the in-air stabilization of differing payloads is dependent on the robustness of the active controller and/or the ability of an on-board observer to estimate the inertial parameters of the payload. Mellinger *et al* [28] successfully used multiple quadrotors to transport differing payload configurations. This cooperative control algorithm was developed under the assumption that all inertial parameters of the payload were known *a priori* and thus is unsuited for payloads with unknown weight and mass distribution. Transportation of a payload with uncertain mass

characteristics by a single quadrotor has been approached using several different control strategies. In subsequent work Mellinger *et al* [29] developed a method to obtain online estimates of a payload's weight and mass center location using a grasping quadrotor. The method requires the attached payload to achieve hover, then runs a least squares estimator on the static system's state vector to obtain the parametric corrections. By requiring the system to achieve hover, the approach makes the assumption that onboard control is robust enough to lift and stabilize the unknown payload and ignores the issue of flightworthiness determination. A similar approach was proposed by Abas *et al* [30] who utilized a full-state Unscented Kalman Filter to identify unknown system parameters of the quadrotor. Burri *et al* developed a maximum likelihood batch estimator to approximate the inertia matrix, mass center location, and certain aerodynamic parameters of a micro air vehicle [31]. Min *et al* [32] propose the use of Adaptive Robust Control (ARC) to compensate for the unknown weight of the payload, but make the assumption that the mass center of the system is a known parameter. This assumption limits the range of applicable payloads that the ARC can handle and adds the burden of prior payload characterization onto the (likely human) operator. Achtemlik *et al* [33] proposed using Model Reference Adaptive Control (MRAC) to handle large or complete parametric uncertainties of a quadrotor. The MRAC is designed to remain robust to a range of external disturbances, such as added weight or a shift in the quadrotor's mass center. Similarly, other adaptive and robust controllers have been proposed to remain robust to parameter uncertainties for single [34-36] and multiple [37-39] quadrotor(s).

All prior work above specifically addresses vehicles that have already attained stable flight – i.e., vehicles that have successfully performed a takeoff and achieved a stabilized

hover. While this prior work addresses parametric payload uncertainties, they ignore the issue of whether takeoff and stabilization from ground contact is even possible given a connected vehicle configuration. In many practical settings, it is likely that either a human operator attaches the vehicles to a payload, or the vehicles attach themselves. However, it is also likely that the payload inertial characteristics will not be measured or known precisely prior to flight, and thus there is no guarantee that a connected vehicle attachment geometry is even flightworthy. Prior work has investigated the topic of robust takeoff and landing [40, 41], but considers the ground surface as the only source of uncertainty. For the purposes of this work, a modular vehicle-payload aircraft (referred to hereafter as the “composite aircraft”) is deemed flightworthy if the following requirements are satisfied:

1. The total thrust to weight ratio of the composite aircraft exceeds 1.
2. The composite aircraft can achieve hover without exceeding user-defined throttle limits for any of the attached lift vehicles.

The first requirement is clearly required to achieve flight with vertical lift aircraft. The second requirement ensures that the aircraft can stabilize in air with available excess throttle to command pitching, rolling, and/or yawing moments. It is noted that for vertical lift aircraft satisfying the second flightworthiness requirement, the first flightworthiness requirement necessarily must be satisfied. This separation of requirements was purposefully introduced to divide the flightworthiness definition into considerations of available lifting capacity and of geometric placement of the vehicles about the composite aircraft. This approach does not directly solve for the inertia matrix of the composite aircraft and is intended purely as a method to assess whether the vehicle is capable of

achieving a stable hover, and to provide initial estimates of the aircraft weight and vehicle attachment positions. Once in-air, a parametric estimator such as the one proposed by Mellinger [29] can be used to refine parameter estimates and achieve desired flight performance. Note that this definition of flightworthiness is only meant to apply to modular vertical lift vehicles such as modular quadrotors or other similar aircraft.

The scenario investigated in this research is that of a set of modular vehicles rigidly attached to a payload with unknown inertial properties. The modular vehicles and payload combination forms the composite aircraft. Each modular vehicle does not know the positions of the other vehicles or itself with respect to the payload, but they are collectively tasked with determining whether the current configuration is flightworthy. The assumption that the vehicles do not know the relative locations of the other vehicles stems from the fact that the individual vehicles are envisioned to be equipped with only rudimentary sensing capabilities. When installed by a human operator (or when attaching autonomously) there is no assumption that the vehicles' relative positions are measured or sensed, since this may impose a substantial burden on the operator or autonomous vehicles. The composite aircraft is initialized in static ground contact. In the algorithm formulated here, the thrusts of each vehicle are assumed to be aligned and to act in the opposite direction to the payload weight force when the vehicle is at rest. Each vehicle thrust value is incrementally increased until one or more of the contact points loses ground contact and the system exhibits dynamic motion. The thrust distribution needed to induce dynamic excitation is then used to generate a probabilistic estimate of the effectiveness of individual vehicles at producing motion about the different ground contact axes. This probabilistic estimate is used to drive a series of excitations that generate the data needed to calculate

the weight of the aircraft and the positions of the attached vehicles. Finally, the estimated vehicle positions can be used to determine flightworthiness, or, if the vehicle is deemed unflightworthy, to suggest alternative vehicle layouts.

The thesis proceeds as follows. The composite aircraft dynamic model and the employed dynamic contact model are described. The flightworthiness algorithm is then discussed in detail, including the axis detection, Bayesian inference, vehicle position estimation, and flightworthiness determination algorithms. Simulation results are provided to illustrate the utility and performance of the algorithm, with trade studies that examine the effects of thrust to weight ratio, mass center position, and other algorithm parameters. Descriptions of the constructed test platforms are provided along with an overview of the testing procedure used for each experiment. A detailed error analysis of the algorithm is performed to better understand the sources of error in vehicle position estimates. Finally, a set of comprehensive experimental test results are presented to show the performance of the algorithm as a function of payload geometry, payload inertia, number of attached vehicles, and vehicle distribution about the payload.

CHAPTER 2. COMPOSITE AIRCRAFT DYNAMIC MODEL

Define a standard inertial reference frame with unit vectors $\vec{I}_I, \vec{J}_I,$ and \vec{K}_I and a body-fixed reference frame with unit vectors $\vec{I}_B, \vec{J}_B,$ and \vec{K}_B as shown in Fig. 1, which also shows an example configuration for the composite aircraft and contact axes. Here, the vehicles are attached to the top of a notional payload. Several assumptions are invoked throughout the remainder of this work. First, it is assumed that the number of attached vehicles is known *a priori*. Second, it is assumed that the vehicles are rigidly attached to the payload such that there is no relative motion between each vehicle and the payload. Third, it is assumed that the thrust vectors produced by the lift vehicles are known and parallel with the body \vec{K}_B direction. Fourth, it is assumed that the algorithm has real-time feedback of the orientation of the composite vehicle. Lastly, it is assumed that the payload has a set of contact points that form discrete contact axes. each of which may or may not lie at the same elevation (position along \vec{K}_I). Note that, while the method described below can be modified such that the thrust vectors do not have to be aligned with \vec{K}_B , such extensions are not addressed in this work.




Figure 1. Composite Aircraft Diagram and Coordinate system.

2.1 Rigid Body Dynamics

A six-degree-of-freedom dynamic model is used to simulate the position, orientation, and velocity states of the composite aircraft. The differential equations that govern the position and velocity states of the composite aircraft are provided in Eqs (1) and (2):

$$\begin{Bmatrix} \dot{x} \\ \dot{y} \\ \dot{z} \end{Bmatrix} = [T_H]^T \begin{Bmatrix} u \\ v \\ w \end{Bmatrix} \quad (1)$$

$$\begin{Bmatrix} \dot{u} \\ \dot{v} \\ \dot{w} \end{Bmatrix} = \frac{1}{m_{cv}} \begin{Bmatrix} X \\ Y \\ Z \end{Bmatrix} - \begin{Bmatrix} qw - rv \\ ru - pq \\ pv - qu \end{Bmatrix} \quad (2)$$

where

$$T_H = \begin{bmatrix} c_\theta c_\psi & c_\theta s_\psi & -s_\theta \\ s_\phi s_\theta c_\psi - c_\phi s_\psi & s_\phi s_\theta s_\psi + c_\phi c_\psi & s_\phi c_\theta \\ c_\phi s_\theta c_\psi + s_\phi s_\psi & c_\phi s_\theta s_\psi - s_\phi c_\psi & c_\phi c_\theta \end{bmatrix} \quad (3)$$

and $s_\alpha, c_\alpha, t_\alpha = \sin(\alpha), \cos(\alpha), \tan(\alpha)$. The total applied forces, X , Y , and Z , include contributions from the thrusts of the attached vehicles, the weight of the composite aircraft, and the normal and friction forces from the contact surfaces. The rotational dynamic equations for the composite aircraft are provided in Eqs (4) and (5):

$$\begin{Bmatrix} \dot{\phi} \\ \dot{\theta} \\ \dot{\psi} \end{Bmatrix} = \begin{bmatrix} 1 & s_\phi t_\theta & c_\phi t_\theta \\ 0 & c_\phi & -s_\phi \\ 0 & s_\phi / c_\theta & c_\phi / c_\theta \end{bmatrix} \begin{Bmatrix} p \\ q \\ r \end{Bmatrix} \quad (4)$$

$$\begin{Bmatrix} \dot{p} \\ \dot{q} \\ \dot{r} \end{Bmatrix} = [I]^{-1} \left[\begin{Bmatrix} L \\ M \\ N \end{Bmatrix} - \begin{bmatrix} 0 & -r & q \\ r & 0 & -p \\ -q & p & 0 \end{bmatrix} [I] \begin{Bmatrix} p \\ q \\ r \end{Bmatrix} \right] \quad (5)$$

where L , M , and N are the moments about the system's mass center. The external rolling and pitching moments include contributions from the vehicle thrusts, the frictional forces, and the normal forces acting on the composite aircraft.

2.2 Vehicle Model

The lift vehicles are modeled as lumped masses attached at defined locations about the payload. The rigid connection assumption is invoked, and thus the attached modular vehicles add to the total weight and inertia matrix of the payload and have a direct influence on the mass center location of the composite aircraft. The aerodynamic models used to express the external forces and moments produced by each vehicle are dependent on which

type of vertical lift vehicle is being used. If the attached vehicles are chosen to be rotorcraft, as was the case for all simulations in this thesis, then the thrust and torque generated by each vehicle can be derived from blade element theory [42] and are expressed as a function of the rotation rate of the rotor according to,

$$T = c_T \rho A \Omega^2 R^3 \quad (6)$$

$$Q = c_Q \rho A \Omega^2 R^3 \quad (7)$$

where c_T is the thrust coefficient of the rotor, c_Q is the torque coefficient, ρ is air density, A is the rotor disk area, Ω is the rotation rate of the rotor, and R is the radius of the rotor. Oftentimes for rotors powered by brushless DC motors, a first order transfer function is chosen to represent the dynamics between the desired rotor speed and its true value. In the context of this work, these rotor dynamics are neglected and it is assumed that the rotor speed can be changed quasi-statically. This results in a known relationship between the throttle effort of the vehicles and the resultant aerodynamic forces. As the composite vehicle remains in partial ground contact during the duration of the flightworthiness determination process, ground effect on the rotor wake could potentially be significant. While in ground effect, there is a reduction of induced airflow velocity through the rotors resulting in an increase in the thrust produced by the rotors. This increase can be modeled as a percent increase from the thrust produced out of ground contact [43] as

$$T_g = f_g T \quad (8)$$

where the percent increase, f_g , is a function of the rotor height above the ground, h , normalized by the radius of the rotor. This value is approximated by the Cheeseman and Benner model [44] as

$$f_g = \left[1 - \frac{1}{(4h/R)^2} \right]^{-1} \quad (9)$$

For the example aircraft analyzed in this work, the ground effect scaling term, f_g , was calculated to be below 1.01, only a 1% increase in thrust. As a result, ground effect was determined to have negligible impact and was not directly incorporated into the flightworthiness determination algorithm. Although the aerodynamic analysis above assumes the vertical lift vehicles to be rotorcraft, the only requirement is that the vehicles can produce a known external force that acts in the \vec{K}_B direction in response to a control command.

After the aerodynamics of the vehicle are fully defined, their contribution to the total external forces and moments can be expressed as

$$\begin{Bmatrix} X \\ Y \\ Z \end{Bmatrix}_V = \begin{Bmatrix} 0 \\ 0 \\ \sum_{i=1}^n T_i \end{Bmatrix} \quad (10)$$

and

$$\begin{Bmatrix} L \\ M \\ N \end{Bmatrix}_V = \sum_{i=1}^n \vec{r}_{v_i/cg} \times T_i \vec{K}_B + \sum_{i=1}^n Q_{v_i} \vec{K}_B \quad (11)$$

2.3 Ground Contact Model

When the system is in partial or full ground contact, the unknown normal forces generated at each contact point lead to an underdetermined dynamic system. Thus, a dynamic contact model must be utilized to determine the reactions. Three dynamic contact models that are commonly used in simulation are the soft contact model, a constraint-based contact model (sometimes known as a hard contact model), and an impulse-based contact model [45]. The soft contact model, such as the one presented by Goyal *et al* [46-47], uses an infinitesimal contact surface attached to each contact point via a stiff spring and damper. Due to the stiffness of the contact surface springs, the simulation timestep must be shrunk by orders of magnitude with respect to a typical rigid body code to capture the collision dynamics. For a composite vehicle that remains in partial ground contact for the duration of the flightworthiness determination process, this collision timestep is required for the entirety of the simulation. This approach was found to be prohibitively slow, especially when simulating a wide range of composite vehicle configurations. The constraint-based approach, such as the one presented by Trinkle *et al* [48], formulates the unknown contact point forces and accelerations into a linear complimentary problem [49]. The constraint-based approach aims to solve for the set of normal and frictional forces acting on each of the contact points. However, the method can struggle to find a unique solution when a contact point transitions into and out of ground contact. Finally, the impulse-based model proposed by Mirtich and Canny [50] solves for a set of impulses to accurately simulate

collisions and ground contact. This method produces solutions much faster than the soft contact model and does not suffer from the convergence issues of the root-finding process in the constraint-based method. The work presented here uses the impulse model described in [50] and expanded in [51-52] to resolve the ground-plane collisions that occur at each contact point. For each time step in the simulation, the model detects which contact points, if any, have collided with the ground-plane, then resolves the collisions iteratively. For a composite aircraft in ground contact, collisions are detected when the vertical position of a contact point of the aircraft falls below the defined ground-plane coordinate. When detected, a time of impact estimator is used to capture the instance of collision. After the collision detection phase, the contact model resolves each collision using a series of impulses applied to the composite vehicle. The impulses are computed using the impulse-momentum theorem and assumes an infinitesimal collision time, Poisson's hypothesis, and a Coulomb friction model.

When an impulse is applied to a contact point of the composite vehicle, the change in translational velocity of the mass center can be calculated as

$$\Delta \vec{v}_{cg} = \frac{1}{m_{ca}} \vec{p} \quad (12)$$

where $\Delta \vec{v}_{cg}$ is the change in translational velocity of the mass center before and after impact, and \vec{p} is the applied impulse. Similarly, the change in angular velocity of the composite vehicle can be computed as

$$\Delta \vec{\omega} = I^{-1} \left[\vec{r}_{c_i/cg} \times \vec{p} \right] \quad (13)$$

where $\Delta\vec{\omega}$ is the change in rotational velocity of the aircraft. These changes are related to the change in contact point velocity using rigid body kinematics as

$$\Delta\vec{v}_{c_i} = \Delta\vec{v}_{cg} + \Delta\vec{\omega} \times \vec{r}_{c_i/cg} \quad (14)$$

A relationship between the applied impulse and the change in contact point velocity can be obtained by plugging Eqns. (12) and (13) into Eq. (14) and is expressed as

$$\Delta\vec{v}_{c_i} = \left[\frac{1}{m} \mathbf{I} - \vec{r}_{c_i/cg} \times I^{-1} \times \vec{r}_{c_i/cg} \right] \vec{p} = M\vec{p} \quad (15)$$

where \mathbf{I} is the identity matrix. Equation (15) shows that the velocities of the contact points evolve over the duration of the applied impulse and, thus, the dynamics of the collision event must be considered. Differentiating each side of Eq. (15) with respect to the normal component of the impulse delivered to the contact point results in a differential equation describing the change in contact point velocity over the collision event.

$$v'_{c,i} = Mp' \quad (16)$$

Assuming Coulomb friction, Eq. (16) can be expressed as

$$\begin{Bmatrix} v'_{c_i,x} \\ v'_{c_i,y} \\ v'_{c_i,z} \end{Bmatrix} = M \begin{Bmatrix} -\mu \frac{v_{c_i,x}}{\sqrt{v_{c_i,x}^2 + v_{c_i,y}^2}} \\ -\mu \frac{v_{c_i,y}}{\sqrt{v_{c_i,x}^2 + v_{c_i,y}^2}} \\ 1 \end{Bmatrix} \quad (17)$$

Equation (17) is valid as long as the considered contact point is sliding relative to the ground plane. When collision occurs and $v_{c,z}$ is verified to be a negative value, Eq. (17) is numerically integrated over the growing impulse to the point that $v_{c,z}$ reaches zero. This event defines the point of maximum compression and the corresponding value of p_z is recorded. Poisson's hypothesis [53] is then used to calculate the final value of the normal impulse:

$$p_{z,f} = (1+e)p_z \quad (18)$$

Integration continues until the normal impulse reaches $p_{z,f}$, at which point the impulse is used to update the translational and rotational velocity of the composite vehicle according to Eqs. (12) and (13). During the integration, if the tangential velocities, $v_{c,x}$ and $v_{c,y}$, disappear, then sticking occurs and $v'_{c_i,x}$ and $v'_{c_i,y}$ are set to zero. While in the sticking mode, matrix M of Eq. (17) is inverted to solve for the relative magnitude of the normal and tangential impulse rates. As long as the static friction constraint from the Coulomb model is satisfied, sticking persists. Otherwise, sliding resumes along the unique sliding direction described in [50]. One drawback of this approach comes in the event of continuous static

contact. Since static forces do no work, the strategy of using Poisson's hypothesis from Eq. (18) with a non-unity e is no longer valid. Instead, static contact is approached using a series of micro-collision to reverse the initial collision velocity. This way, the micro-collisions are perfectly elastic and energy is conserved. The threshold for micro-collisions is based on the initial magnitude of $v_{c,z}$ and is defined as the velocity a resting object achieves as it falls some defined collision envelope, ε_c during one timestep.

$$v_e = \sqrt{2g\varepsilon_c} \quad (19)$$

When an impact is detected, the collision is classified as either a full collision or a micro-collision and the corresponding methodology is followed. If multiple collisions are detected during a given time-step, the collisions are resolved one at a time. After each collision resolution, the state of the composite vehicle is updated according to Eqs. (12) and (13) and the state of each contact point is recalculated. This process is performed iteratively until all collisions are resolved within an acceptable error tolerance at which point the simulation continues to the next timestep.

The described contact model, in conjunction with the rigid-body 6DOF equations of motion, is used to test a range of simulated payloads and attachment vehicles. Each payload is defined by a total weight, an inertia matrix, a mass center location, and a set of discrete points that define the payload's contact axes. Figure 2 shows a dynamic excitation experiment performed on a simulated composite aircraft with the described impulse-based contact model. In Fig. 2, each cylinder represents an attached modular vehicle location. As the attached vehicles increase their throttle effort from zero, the composite aircraft begins exciting about its south-west contact axes in Fig. 2. When excitation is detected, the

modular vehicles all set their throttle values to zero and the composite aircraft is returned to static equilibrium.




Figure 2. Simulated Excitation Event.

In order to verify the physical accuracy of the implemented contact model, a series of controlled tests were performed. Each test compared the state history of the full simulation to that of a simplified dynamic model whose motion was constrained. Figure 3 shows the results of one such test. For this case, the simulated composite aircraft increases the throttle levels of the attached vehicles resulting in the aircraft exciting about a contact axis of the payload. The simplified model was constructed by imposing the constraint that the vehicle excited purely about the same contact axis. By enforcing this constraint, the dynamics of the simplified model could be explicitly expressed without the need of a contact model according to,

$$I_{a_i} \ddot{\alpha} = \max \left(\sum_{j=1}^n \left(\vec{r}_{v_j/a_i} \times T_j \vec{K}_B \right) \cdot \vec{I}_i - \left(\vec{r}_{cg/a_i} \times W \vec{K}_I \right) \cdot \vec{I}_i, \quad 0 \right) \quad (20)$$

Figure 3 shows the excitation angle for the same composite aircraft using the full simulation and the simplified model. Figure 3 shows strong agreement between the two

simulation models and was used as one of several validation cases for the implemented contact model.




Figure 3. Sample Contact Model Validation Results.

The flightworthiness determination algorithm operates by measuring the thrust distribution needed to drive the composite aircraft to the transition between static and dynamic motion. Thus, the algorithm designed here operates only on measurements of the thrust required to initiate dynamic motion from static equilibrium and the corresponding excitation axis, and requires no other information for flightworthiness determination. The contact model described above is used only for simulation purposes and not within the algorithm itself.

CHAPTER 3. FLIGHTWORTHINESS ALGORITHM DESIGN

The flightworthiness algorithm is initialized with the assumptions that the total number of vehicles attached to the payload and the mapping between the vehicles' throttle and thrust are known *a priori*. Also, it is assumed that the algorithm has real-time feedback of the orientation of the composite aircraft so as to determine when and about which ground contact axis or point excitation occurs. An overview of the flightworthiness algorithm is provided in Fig. 4. The algorithm is divided into machine learning and static analysis subcomponents. The purpose of the machine learning subcomponent is to intelligently generate a thrust measurement data set which can be used to deterministically locate the modular vehicles about the composite aircraft and, thus, determine flightworthiness. This intelligent data acquisition is achieved by building a probabilistic model of the composite aircraft through a series of dynamic excitation observations. The probabilistic model is used to reduce the total control effort needed to generate the required data set for vehicle position estimation. Within the machine learning subcomponent, the algorithm performs an explore and exploit strategy. The purpose of the explore phase is to observe the different excitations that occur when the thrusts of the modular vehicles are increased from zero using a unique thrust distribution and to build a probabilistic model of the composite aircraft. These thrust experiments, where the vehicle thrust is increased at different rates for all vehicles, are referred to throughout this paper as spin-ups. During the exploration phase, the algorithm identifies and stores the orientations of observed contact axes. When the required number of contact axes are observed, the algorithm attempts to utilize the probabilistic model to

drive desired dynamic excitation. The algorithm commands excitation about the observed axes until the required amount of data is generated and vehicle position estimates can be generated. During the machine learning subcomponent, if the algorithm exceeds a set number of spin-up experiments the algorithm is terminated and no conclusions on flightworthiness can be made. Because this Bayesian process results in only probabilistic classifications, the resulting determination of the modular vehicle distribution about the payload may be subject to significant error. Thus, accurate determination of flightworthiness using the probabilistic model alone is not guaranteed.

The static analysis subcomponent takes the data collected by the machine learning subcomponent and uses a set of moment balance equations coupled with the geometric relationships to calculate the modular vehicles' positions with respect to the mass center of the composite aircraft.

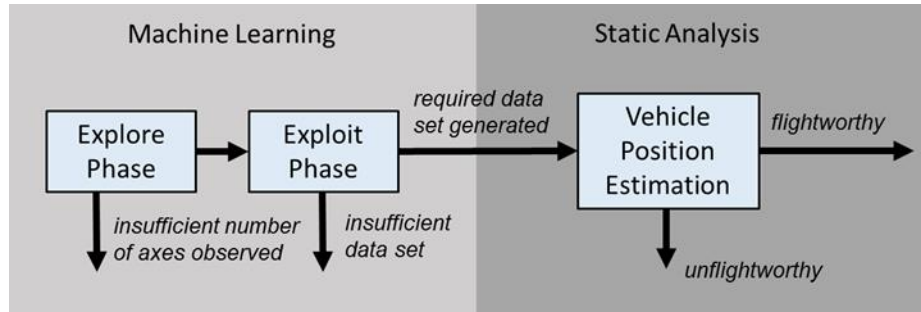


Figure 4. Schematic of Flightworthiness Determination Algorithm.

3.1 Spin-up Methodology

The entire algorithm is reliant on observations of discrete excitation experiments called spin-ups. During each experiment, the vehicle is initialized on the ground at rest. The modular vehicles' thrust values are increased continuously from zero in a prescribed way.

If dynamic excitation is observed, then the current vehicle thrust values and the direction of excitation are recorded and the vehicles are spun-down to rest to return the system to static equilibrium. Once recorded, the excitation direction is used to determine if the excitation was about a previously observed contact axis, a new contact axis, or about a contact point. If the observed excitation direction is classified as motion about a discrete contact axis, the recorded values are used to update the probabilistic model of the vehicle using Bayesian inference. If all vehicle thrusts saturate at their maximum value and no dynamic motion was observed, then the vehicles are spun down and the algorithm exits with a recommendation to add additional modular vehicles. The explore phase continually increases the throttle of the modular vehicles using a randomized thrust distribution, where the relative thrusts of each vehicle are randomized and this thrust ratio is held constant while the total thrust is increased. In contrast, the exploit phase utilizes the vehicles' axis effectiveness scores generated from the Bayesian model to maximize the probability of exciting about a commanded axis. An overview of this process is shown in Fig 5.




Figure 5. Schematic of the Bayesian Update Procedure.

The machine learning subcomponent in Fig. 5 is defined by either the explore phase or the exploit phase. The success of the algorithm is directly tied to how accurately the Bayesian update can estimate the ability of each attached vehicle to generate moments on

the composite aircraft. In order to build this probabilistic model, the algorithm must first be able to identify the contact axis about which motion occurred.

It is important to note that the flightworthiness algorithm takes measurements of the composite vehicle as it transitions from static equilibrium to dynamic excitation, thus each spin-up experiment must be initialized with the composite vehicle resting in ground contact. This procedure enables a quasi-static assumption to be made, simplifying the solution strategy. Potentially, a different strategy could be employed in which the thrusts can be continuously varied once the aircraft exhibits dynamic motion, and the vehicle responses could be recorded and analyzed to estimate model parameters. However, such an approach would require either angular acceleration measurements (which are difficult to obtain), or use of an estimator which would need to estimate a large number of model parameters simultaneously. Given the nonlinear dynamics of this system, observability would likely suffer in such a scheme and no convergence guarantees could be provided. Finally, from a safety perspective, the process of continuously changing the thrust profile while the vehicle is in partial ground contact would result in a higher probability of the vehicle potentially flipping over. As a result of these considerations, a quasi-static assumption is enforced here and only data that is recorded at the onset of dynamic excitation is used in the solution strategy.

3.2 Excitation Axis Detection

The flightworthiness determination algorithm is initialized with no prior knowledge of the payload geometry. However, the success of the algorithm is strictly dependent on observations of motion about discrete contact axes. When dynamic excitation occurs, it is

imperative to classify the excitation as motion about a previously observed contact axis, about a new contact axis, or about a contact point. During each experiment, the time histories of the Euler angles of the aircraft, ϕ , θ , and ψ , are analyzed. Euler's rotation theorem can be used to transform the roll, pitch, and yaw angles into a single rotation about the calculated Euler axis [54]. The single rotation, α , is calculated using

$$1 + 2\cos\alpha = \text{tr}[T_H] \quad (21)$$

where $\text{tr}[A]$ is the trace operator for matrix A . Using an onboard orientation sensor, the vehicle orientation is monitored during a spin-up and an excitation is identified when the vehicle orientation changes from rest by a defined excitation threshold angle. While the aircraft is exhibiting dynamic motion, the Euler axis is defined by β and γ and is computed as,

$$\beta = \tan^{-1}(l_{ZY}/l_{ZX}) \quad (22)$$

and

$$\gamma = \cos^{-1}(l_{ZZ}) \quad (23)$$

where

$$[T_H - \mathbf{I}] \begin{Bmatrix} l_{ZX} \\ l_{ZY} \\ l_{ZZ} \end{Bmatrix} = \{0\} \quad (24)$$

and \mathbf{I} is the identity matrix. The matrix in Eq. (24) is rank deficient and leads to a directional sign ambiguity for the resultant Euler axis orientation. The sign ambiguity can be resolved

using the realization that α is constrained to a positive value. The composite vehicle cannot rotate into the ground plane and thus, $\alpha \geq 0$. By constraining α , the resultant Euler axis is uniquely determined. When the composite aircraft excites purely about one of its contact axes, the computed Euler axis defines the contact axis of the aircraft. If the contact axis exists in the inertial $\vec{I}_I - \vec{J}_I$ plane, then γ is zero by definition and β defines the Euler axis. Otherwise, if the contact points lie at different heights, γ is a nonzero constant and β defines the projection of the Euler axis onto the inertial $\vec{I}_I - \vec{J}_I$ plane. While at rest, the attached vehicle thrusts act purely in the \vec{K}_I direction, therefore, the moments produced by the attached vehicles can be expressed using their planar distances to the $\vec{I}_I - \vec{J}_I$ projection of the contact axis. In this way, only the projection of the Euler axis (angle β) is important in the moment balance calculations used to generate the probabilistic model of the aircraft and to estimate the attached vehicle positions.

If the time history of the Euler axis remains near constant during the excitation, then the Euler axis direction is considered as a potential contact axis. If the potential contact axis is sufficiently close to an existing observed contact axis direction, then the excitation is grouped with this prior observed contact axis. Otherwise, the potential axis direction is saved until a second excitation with a near identical Euler axis direction is observed. This two-phase detection process is used to reject excitations that occur about a contact point, but that maintain a near constant Euler axis direction. Excitations with a varying Euler axis direction during a spin-up are classified as contact point excitations and are disregarded as they do not fit into the discrete contact axis framework of the algorithm. Figures 6, 7, and 8 show the time histories of α , β , and γ for a simulated composite vehicle performing a spin-up experiment. A detection threshold of $\pm 5^\circ$ was set for both β and γ over the duration

of the excitation. For this spin-up excitation, the aircraft excited about one of the discrete contact axes. Figure 7 shows that the orientation of the axis was measured as 90° with respect to the reference frame orientation of the attached IMU and remained nearly constant over the duration of dynamic excitation. Figure 8 shows the corresponding γ time history of the experiment. As both β and γ remained within the allowable detection threshold, the dynamic motion was classified as a contact axis excitation and the corresponding thrust values at the excitation detection angle of α were recorded.

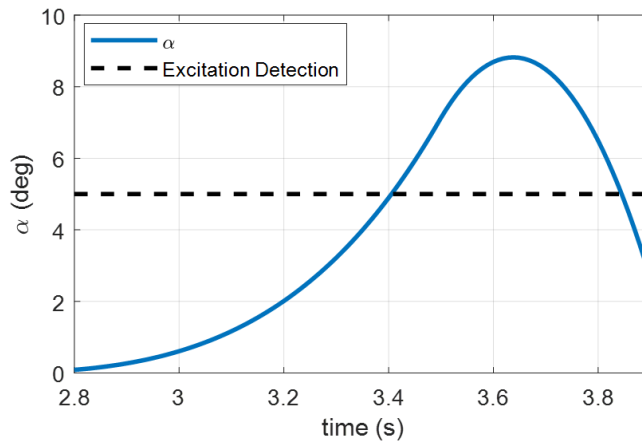


Figure 6. Example Contact Axis Excitation Angle.




Figure 7. Example Contact Axis Orientation - β .

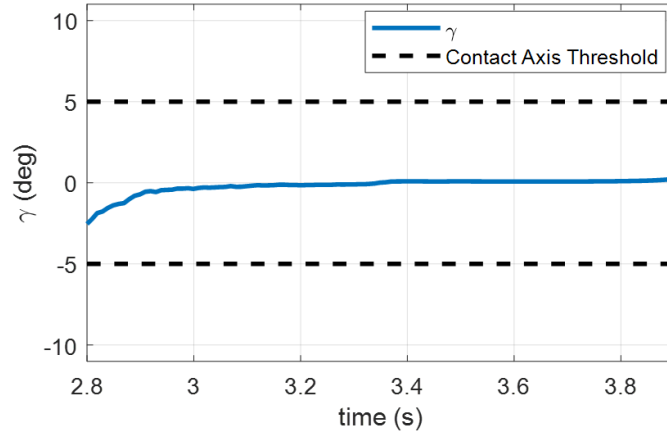


Figure 8. Example Contact Axis Orientation - γ .

Conversely, Figs. 9, 10, and 11 show the time histories of α , β , and γ for a simulated spin-up experiment where the composite vehicle excited about a contact point. Figures 10 and 11 show that, during the duration of the experiment, both β and γ varied significantly and violated the defined contact axis threshold. For the experiment, the excitation would not be classified as having occurred about a contact axis and the recorded thrust data would be discarded.

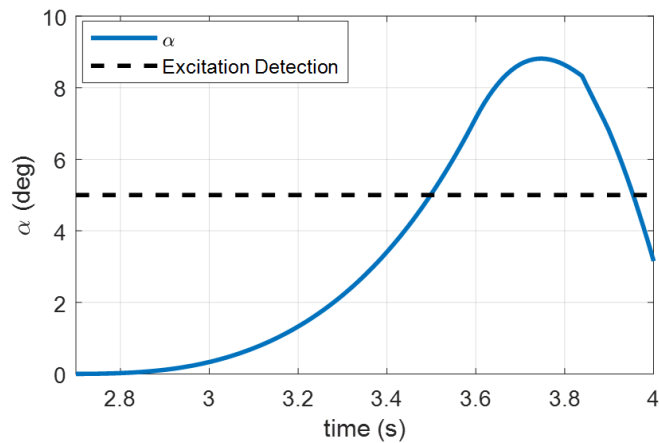


Figure 9. Example Contact Point Excitation Angle.




Figure 10. Example Contact Axis Orientation - β .

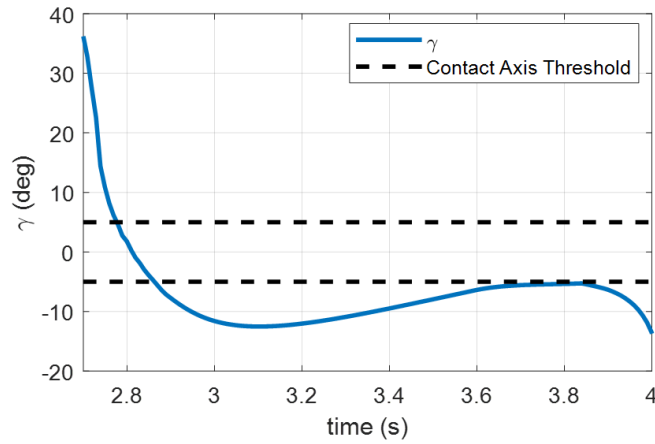


Figure 11. Example Contact Axis Orientation - γ .

It should be noted that this process of classifying excitations can result in a contact point excitation being classified as a contact axis. Depending on the selected contact axis threshold and the steadiness of the contact point excitation, false positives for contact axis detection can occur. The failsafe, discussed above, is the requirement for a contact axis to

be detected twice before being officially classified as a contact axis to be used in the algorithm. Although it is possible for a composite vehicle to excite about the same contact point twice with near identical, steady, axis orientations, this event was not observed over the course of many simulated and experimental tests. If a contact point did get classified as a contact axis, the algorithm would command excitations about the nonexistent axis. Most likely, this would result in the composite aircraft repeatedly failing to excite about the commanded axis, eventually reaching the maximum number of spin-ups allowed for flightworthiness determination. Potential solutions, such as contact axis de-classification, could prevent this from occurring but, due to the implausibility of the event, the added complexity to the algorithm was deemed unnecessary. Once an excitation direction is successfully classified as a contact axis, it can be used to command excitation in the exploit phase which drives the subsequent lift vehicle position estimation algorithm.

3.3 Vehicle Position Estimation Algorithm

The purpose of the lift vehicle position estimation algorithm is to deterministically calculate each lift vehicle's position on the composite aircraft with respect to the mass center location. This is accomplished by recording and utilizing the thrust distributions that were required to excite dynamic motion about differing contact axes of the composite aircraft. This section introduces the lift vehicle position estimation strategy for a composite aircraft with a general geometry that adheres to the assumptions listed in the dynamic model section. Figure 12 shows an example composite aircraft with 6 modular vehicles (circles) and a defined mass center location.




Figure 12. Example Composite Aircraft Configuration and Contact Axis

Definitions.

From Fig. 12, it is shown that the algorithm has observed excitations about axes i , $k1$, and $k2$ during the exploration phase. The orientations of these contact axes are defined by the angles θ_{k1} and θ_{k2} in the $\vec{I}_i - \vec{J}_i$ plane measured from the \vec{I}_i direction. The objective of the static analysis subcomponent is to approximate the \vec{I}_i and \vec{J}_i components of the position vectors of the modular vehicles with respect to the mass center location. This is accomplished by using the following solution strategy.

3.3.1 General Solution Strategy

At the instant when the composite vehicle transitions from static equilibrium to dynamic motion about one of the contact axes, the instantaneous moment balance equation about the contact axis can be written as

$$0 = -(\vec{r}_{cg/a_i} \cdot \vec{J}_i)W + \sum_{j=1}^n (\vec{r}_{v_j/a_i} \cdot \vec{J}_i)T_j \quad (25)$$

As additional experiments are performed and multiple excitations about the same axis occur, the moment balance equations can be accumulated in matrix-vector form according to

$$\begin{Bmatrix} W(\vec{r}_{cg/a_i} \cdot \vec{J}_i) \\ \vdots \\ W(\vec{r}_{cg/a_i} \cdot \vec{J}_i) \end{Bmatrix} = \begin{bmatrix} T_{11,i} & \cdots & T_{1n,i} \\ \vdots & \ddots & \vdots \\ T_{p_i 1,i} & \cdots & T_{p_i n,i} \end{bmatrix} \begin{Bmatrix} \vec{r}_{v_1/a_i} \cdot \vec{J}_i \\ \vdots \\ \vec{r}_{v_n/a_i} \cdot \vec{J}_i \end{Bmatrix} \quad (26)$$

where p_i is the total number of experiments about axis i and $T_{j,k,i}$ denotes the thrust of vehicle k used during spin-up experiment j which resulting in motion about axis i . At maximum, n linearly independent experiments can be performed about axis i . Equation (26) expresses the vehicle positions with respect to a contact axis, which is of little practical use for flightworthiness determination or flight control. Instead, Eq. (26) can be rewritten to express the vehicle positions with respect to the mass center location of the composite aircraft:

$$\begin{Bmatrix} \vec{r}_{v_1/cg} \cdot \vec{J}_i \\ \vdots \\ \vec{r}_{v_n/cg} \cdot \vec{J}_i \end{Bmatrix} = (\vec{r}_{cg/a_i} \cdot \vec{J}_i) \left(W \begin{bmatrix} T_{11,i} & \cdots & T_{1n,i} \\ \vdots & \ddots & \vdots \\ T_{p_i 1,i} & \cdots & T_{p_i n,i} \end{bmatrix}^+ \begin{Bmatrix} 1 \\ \vdots \\ 1 \end{Bmatrix} - \begin{Bmatrix} 1 \\ \vdots \\ 1 \end{Bmatrix} \right) \quad (27)$$

where $+$ is the pseudoinverse operator. From Eq. (27), there are $n+2$ total unknowns with only n linearly independent equations. Similarly, experiments about different axes each add $n+1$ unknowns while contributing, at maximum, n new equations. To correlate the excitations about differing axes, geometric relationships are used as follows. The positions of the modular vehicles expressed in the different axes reference frames can be equated using the transformation,

$$(\vec{r}_{v_j/cg} \cdot \vec{J}_i) \cos(\theta_k) - (\vec{r}_{v_j/cg} \cdot \vec{I}_i) \sin(\theta_k) = \vec{r}_{v_j/cg} \cdot \vec{J}_k \quad (28)$$

where θ_k is the angle between unit vectors \vec{I}_k and \vec{I}_i . For a single axis transformation, Eq. (28) introduces n additional unknowns $(\vec{r}_{v_j/cg} \cdot \vec{I}_i)$ and provides n additional equations. Each additional axis comparison, however, generates n equations without introducing any new unknowns. If axis i is designated as the primary axis, then Eqs. (27-28) can be used to relate excitations about the primary axis to excitations about two secondary axes, k_1 and k_2 .

$$\begin{Bmatrix} (\vec{r}_{cg/a_i} \cdot \vec{J}_i) \vec{f}(T_i, W) \\ \vec{0} \\ \vec{0} \end{Bmatrix} = \begin{bmatrix} \mathbf{I} & \mathbf{0} & \vec{0} & \vec{0} \\ -c\theta_{k_1} \mathbf{I} & s\theta_{k_1} \mathbf{I} & \vec{f}(T_{k_1}, W) & \vec{0} \\ -c\theta_{k_2} \mathbf{I} & s\theta_{k_2} \mathbf{I} & \vec{0} & \vec{f}(T_{k_2}, W) \end{bmatrix} \begin{Bmatrix} \mathbf{r}_{v/cg} \cdot \vec{J}_i \\ \mathbf{r}_{v/cg} \cdot \vec{I}_i \\ \vec{r}_{cg/a_{k_1}} \cdot \vec{J}_{k_1} \\ \vec{r}_{cg/a_{k_2}} \cdot \vec{J}_{k_2} \end{Bmatrix} \quad (29)$$

where

$$\vec{f}(T_i, W) = W \begin{bmatrix} T_{11,i} & \cdots & T_{1n,i} \\ \vdots & \ddots & \vdots \\ T_{p_l,i} & \cdots & T_{p_l n,i} \end{bmatrix}^+ \begin{Bmatrix} 1 \\ \vdots \\ 1 \end{Bmatrix} - \begin{Bmatrix} 1 \\ \vdots \\ 1 \end{Bmatrix} \quad (30)$$

Equation (30) contains the geometric relationships between the contact axes but cannot be used directly without knowledge of the composite aircraft weight or the mass center position from the primary axis. Assuming that the payload has a convex geometry, it is physically required that the mass center exists within the convex hull of the contact points and thus $\vec{r}_{cg/a_i} \cdot \vec{J}_i$ is a strictly positive constant denoted as $d_{cg,i}$. Without any additional prior information, the exact dimensions of the composite aircraft cannot be determined. Instead, $d_{cg,i}$ is set to a positive reference value, $\hat{d}_{cg,i}$, which acts as a scaling parameter. The scaling

factor, $\hat{d}_{cg,i}/d_{cg,i}$, is a strictly positive constant that scales the dimension of the composite aircraft without altering the lift vehicles' relative pitch and roll effectiveness about the composite aircraft mass center. With the inclusion of the scaling parameter, the total number of equations and unknowns can be expressed as a function of the number of modular vehicles and the number of axes about which excitation is observed:

$$\begin{aligned} \text{unknowns: } 2n + \bar{m} \\ \text{equations: } \bar{m} n \end{aligned} \tag{31}$$

where $\bar{m} \leq m$ is the number of contact axes used in the solution strategy. For three or more attached modular vehicles, excitations about three independent axes produce an equal or greater number of equations than unknowns. However even with the inclusion of the scaling parameter, there is still a nonlinear coupling between weight and the mass center positions from the secondary contact axes. In order to solve Eq. (29) for the vehicle positions, the weight of the composite aircraft must be determined. If the full nonlinear system of equations from Eq. (29) was shown to be globally convex, then a gradient based approach could be used to converge to the true weight and position parameters of the aircraft [55]. However, this system is not globally convex, which is easily shown as the trivial solution solves Eq. (29). Instead, a solution for the weight is found by reducing Eq. (29) into the compact form $\vec{b} = A\vec{x}$:

$$\left(c\theta_{k_2} s\theta_{k_1} - c\theta_{k_1} s\theta_{k_2} \right) \hat{d}_{cg,i} \vec{f}(T_i, W) = \begin{bmatrix} -\vec{f}(T_{k_1}, W) s\theta_{k_2} & \vec{f}(T_{k_2}, W) s\theta_{k_1} \end{bmatrix} \begin{Bmatrix} \vec{r}_{cg/a_{k_1}} \cdot \vec{J}_{k_1} \\ \vec{r}_{cg/a_{k_2}} \cdot \vec{J}_{k_2} \end{Bmatrix} \tag{32}$$

Equation (32) can be used to generate a least squares solution for the mass center location as a function of W . The squared error of this least squares fit can be expressed as

$$\vec{e}^T \vec{e} = \vec{b}^T \vec{b} - \vec{b}^T A (A^T A)^{-1} A^T \vec{b} \quad (33)$$

With perfect thrust measurements (i.e., measurements of the exact values needed to break static equilibrium), the correct weight value should generate zero error for the least squares estimate. The matrix A in Eq. (32) is known, and thus the inverse of matrix $A^T A \in \mathbb{R}^{2 \times 2}$ can be computed in closed form. Evaluating Eq. (33) it can be shown using a symbolic solver that the squared error term as a function of W can be expressed as,

$$\vec{e}^T \vec{e} = g(\theta_{k_1}, \theta_{k_2}) \frac{W^2 (c_1 W^2 + c_2 W + c_3)}{(c_4 W^2 + c_5 W + c_6)} = 0 \quad (34)$$

where

$$g(\theta_{k_1}, \theta_{k_2}) = \left(c(\theta_{k_1}) s(\theta_{k_2}) - c(\theta_{k_2}) s(\theta_{k_1}) \right)^2 \quad (35)$$

and where each c_i is a complex function of the axis angle and the measured thrust values. Expressions for these complex functions are provided in Appendix A. The trivial solution of zero weight is ignored. For a physical system, it is expected that a unique solution for weight exists, i.e. there are not two different values of W that result in the same exact thrust distributions for the same mass center and vehicle locations. Using the geometric relations of Eq. (28) it can be shown using a symbolic solver that the quadratic in the numerator of Eq. (34) has a unique solution $c_2^2 - 4c_1c_3 = 0$. Thus, the weight can be calculated as

$$W = \frac{-c_2}{2c_1} \quad (36)$$

and is dependent only on the measured thrust values. Plugging this solution into the dominator of Eq. (34) produces an expression that was verified symbolically to be non-zero – thus, Eq. (36) provides the unique solution for the composite vehicle weight. The identified W can then be substituted directly into Eq. (29) and the positions of the modular vehicles with respect to the mass center location can be calculated using a least squares estimate. With exact thrust measurements, these position estimates define the locations of the modular vehicles scaled by the term $\hat{d}_{cg,i}/d_{cg,i}$ with a minimum of $3n$ total spin-ups (n per unique contact axis) required. For a practical system, dynamic motion is detected at a non-zero excitation angle and, thus, the corresponding thrust measurements do not exactly reflect the thrust magnitudes needed to transition from static equilibrium. With additional excitations per axis, this error can be minimized in a least squares sense.

3.3.2 *Parallel Extension*

Equation (27) requires that the \bar{m} subset of contact axes used in the solution process are unique and linearly independent. As a pair of contact axes becomes closer to being parallel, the matrix of Eq. (29) becomes ill-conditioned and eventually loses rank, resulting in unobservable vehicle positions. Therefore, two parallel axes do not fit the required criteria and cannot be used in the general solution strategy. However, the parallel relationship between two contact axes can be utilized to couple the positions of the modular vehicles in the \vec{J}_i direction. This is accomplished by defining the distance between the two contact axes according to the constraint

$$\vec{r}_{v_j/a_i} \cdot \vec{J}_i + \vec{r}_{v_j/a_k} \cdot \vec{J}_k = L_{ik} \quad (37)$$

where L_{ik} is the constant distance between parallel axes i and k . Again, a positive constant, \hat{L}_{ik} , is introduced as a reference distance. Equation (25) can be used in combination with Eq. (37) to directly solve for the vehicle positions in a direction normal to the contact axes and scaled by defined length \hat{L}_{ik} .

$$\left\{ \begin{array}{c} \vec{0} \\ -\hat{L}_{ik} \sum_{j=1}^n T_{1j,k} \end{array} \right\} = \begin{bmatrix} \mathbf{T}_i & -\vec{1} & \vec{0} \\ -\mathbf{T}_k & \vec{1} & -\vec{L}_{ik} \end{bmatrix} \left\{ \begin{array}{c} \mathbf{r}_{v/a_i} \vec{J}_i \\ (\vec{r}_{cg/a_i} \cdot \vec{J}_i) W \\ W \end{array} \right\} \quad (38)$$

If p_i and p_k are the number of unique excitations about axes i and k respectively, then Eq. (38) can be used directly provided that $p_i + p_k \geq n + 2$ with $p_i \leq n$, and $p_k \leq n$. In order to calculate the vehicle positions in the \vec{I}_i direction, experiments about additional contact axes are required. If another pair of parallel contact axes i' and k' are observed, then Eq. (38) can be used with a new defined length $\hat{L}_{i'k'}$ to determine the vehicles positions in a direction normal to these new contact axes. If the two sets of parallel axes are not normal to each other, then the second set of position estimates can be transformed to the \vec{I}_i direction. The \vec{I}_i and \vec{J}_i components of the vehicle position estimates are decoupled and scaled by the two defined length constants, \hat{L}_{ik} and $\hat{L}_{i'k'}$. If the second set of contact axes are not parallel, then Eq. (29) can be reduced using the calculated weight and vehicle positions in the direction of \vec{J}_i . The simplified form can then be used to calculate the positions about the \vec{I}_i direction. The two presented solution strategies in this section and

the previous one are valid for any set of distinct, convex contact axes that form regular or irregular polygons.

3.4 Flightworthiness Determination

Once vehicle position estimates have been calculated using Eq. (29) or Eq. (38), they can be used to determine whether or not the current composite aircraft is flightworthy. For a composite aircraft to be classified as flightworthy, it must have a total thrust-to-weight ratio greater than one and must be able to maintain an in-air hover with all vehicles below a user-defined throttle limit. The thrust-to-weight ratio can be directly determined from the weight estimate and the known thrust capabilities of each modular vehicle. The throttle required to hover can be computed by commanding a total thrust of W and zero roll and pitch moments:

$$\begin{Bmatrix} W \\ 0 \\ 0 \end{Bmatrix} = \begin{bmatrix} 1 & \cdots & 1 \\ \vec{r}_{v_1/cg} \cdot \vec{I}_i & \cdots & \vec{r}_{v_n/cg} \cdot \vec{I}_i \\ \vec{r}_{v_1/cg} \cdot \vec{J}_i & \cdots & \vec{r}_{v_n/cg} \cdot \vec{J}_i \end{bmatrix} \begin{Bmatrix} T_1 \\ \vdots \\ T_n \end{Bmatrix} \quad (39)$$

If $n > 3$, the minimum Euclidean norm solution to Eq. (39) can be used to uniquely solve for the thrust required to hover through use of the pseudoinverse. The scaling term(s) presented in the vehicle position estimation section have no impact on these calculations due to the fact that zero moments are being commanded and that each position estimate in a given direction is scaled by the same value. The position estimates in the 2nd row and 3rd row of Eq. (39) are each scaled by the same nonzero value and, thus, can be factored out. Each vehicle's maximum thrust limit can be used to determine the vehicles' throttle percentages at hover. It is important to note that Eq. (39) considers only pitching and rolling

moment balances for the composite aircraft and does not explicitly command a yawing moment balance. Although yaw control is necessary for general flight, the form of the external yawing moments is not explicitly defined for the general definition of the composite aircraft. If each vehicle is assumed to be a single rotor, then, depending on the rotor's direction of rotation, each vehicle would produce a positive or negative yawing moment that could be directly mapped to the rotation rate of the propeller. However, if the vehicles are assumed to be standard quadrotors, then each vehicle could satisfy a commanded thrust and yaw moment simultaneously (within rotor saturation limits) by constraining the throttles of the clockwise and counterclockwise rotors of the quadrotor. With a defined vehicle platform, the actuation matrix of Eq. (39) could be augmented to include the defined yawing moment contributions for the individual vehicles. In the context of this work, however, the vehicle platform is not assumed, and a yawing moment balance is not explicitly commanded. If all hover throttles are within user-defined cutoffs, then the composite aircraft is deemed flightworthy. If any throttle values are outside of the cutoffs, then the current vehicle configuration is determined to be unflightworthy and vehicles should be rearranged or added before flightworthiness determination is attempted again. When determining flightworthiness, the thrust-to-weight ratio of the aircraft dictates the mean throttle percentage of the modular vehicles while the distribution of vehicles about the payload affects the variance of throttle percentages. As the total thrust-to-weight ratio decreases, the vehicle placement plays an increasingly significant role in achieving a flightworthy aircraft.

3.5 Bayesian Inference

The vehicle position estimation algorithm presented in the preceding sections require linearly independent excitations about each commanded contact axis to adequately populate the thrust matrix from Eq (26). Until a minimum of n excitations are observed about a given axis, the inverse/pseudoinverse of the thrust matrix cannot be computed and that particular axis cannot be used in the solution strategy. Therefore, it is important to efficiently generate the required set of excitations for the set of commanded axes. One naïve method to generate excitations is to simply randomize the relative thrust distributions during each spin-up experiment. However, this is likely to lead to excessive excitations about other axes which are not used in the solution set. Such excitations are detrimental in that they do not provide any usable information and use unnecessary power and time. In order to minimize wasted control effort, an intelligent procedure for commanding excitation is needed.

The machine learning component of the flightworthiness determination algorithm, which may be viewed as an unsupervised learning process, forms a probabilistic estimate of each modular vehicle's effectiveness in generating moments about the observed contact axes and assists in generating the dynamic excitation needed for vehicle position estimation. The machine learning is performed through Bayesian inference which uses observations of composite aircraft excitations and the corresponding throttle levels of the modular vehicles. This information is used to compute a probabilistic likelihood of each vehicle's effectiveness at producing moments about the observed contact axes. The inference begins when an initial contact axis is observed and continues for each excitation about current or newly observed contact axes.

Bayesian inference is a method used to update the probability of a set of hypotheses as information is gathered [56-58]. These likelihood estimates can be expressed with Bayes' theorem as,

$$P(H_l | a_i) = \frac{P(a_i | H_l)P(H_l)}{P(a_i)} \quad (40)$$

where the posterior likelihood of hypothesis H_l being correct given that event a_i occurred is equal to the product of the probability that event a_i occurred given hypothesis H_l and the prior probability of H_l , $P(H_l)$, all divided by the total probability $P(a_i)$. For this application, the “events” are discrete excitations about one of the contact axes. For instance, an example event may be recorded as “the vehicle began to roll about contact axis a_i for the applied thrust distribution”. The different hypotheses vary the vehicle placement locations about the payload. The hypotheses are formed by constructing a hypothetical payload containing a series of vehicle locations evenly distributed in a circular path on the payload perimeter. Each hypothesis places one of the modular vehicles into a corresponding location (or slot) on the hypothetical composite aircraft. When a new contact axis is observed, a corresponding axis is drawn on the hypothetical payload tangent to the circular path. Figure 13 shows an example slot layout for a six-vehicle composite aircraft that has observed excitations about axes a_i and a_k . The actual vehicle placement on the composite aircraft is shown at left, whereas the slot location definitions are shown at right.




Figure 13. Example Actual Composite Aircraft Configuration (left) and Corresponding Slot Definitions (right). Each circle represents a vehicle location.

For the example case shown in Fig. 13, there are a total of $6!$ vehicle-slot permutations (or hypotheses). It is noted that the hypothetical composite aircraft, such as the one shown in the right of Fig. 5, is constructed with no prior knowledge of the actual composite aircraft's geometry or dimensions. The dimensions of the hypothetical composite aircraft are of no consequence to the algorithm and thus the radius of the circular path is initialized to a unit length. The objective of the hypothetical composite aircraft is to construct vehicle slot locations with relative distances from each contact axis so as to differentiate vehicle effort when excitation occurs. When excitation does occur, the equivalent thrust vector of the modular vehicles is

$$T_{eq} = \sum_{j=1}^n T_j \quad (41)$$

where T_j is the thrust produced by the vehicle placed into slot j for a given hypothesis. This equivalent thrust vector acts at a perpendicular distance, r_{eq} , from contact axis i , according to

$$r_{eq/a_i} = \frac{\sum_{j=1}^n T_j (\vec{r}_{s_j/a_i} \cdot \vec{J}_{a_i})}{\sum_{j=1}^n T_j} \quad (42)$$

Figure 14 shows an example case where the composite aircraft excites about axis a_i and a certain hypothesis places the modular vehicles into slots that result in an equivalent thrust vector positioned at a distance of r_{eq} from a_i .




Figure 14. Example Equivalent Thrust Vector for given Hypothesis.

A similar equivalent thrust position can be computed for each observed contact axes. An inequality constraint on the relationship between the equivalent thrust position and the position of the mass center from the axis of excitation can be established by analyzing the transition from static equilibrium to dynamic excitation. A 2D representation of the transition event is shown in Fig. 15.




Figure 15. Example 2D Dynamic Transition FBD.

In Figure 15, N_{a_i} is the normal force acting at contact axis a_i and r_{cg} is the perpendicular distance from contact axis a_i to the mass center location of the composite aircraft. At the instant before dynamic equation, the aircraft is in static equilibrium. Therefore, a simple force balance ($T_{eq} + N_{a_i} = W$) and moment balance ($r_{eq}T_{eq} = r_{cg}W$) can be written to relate the external forces and their corresponding positions

$$N_{a_i} = T_{eq} \left(\frac{r_{eq}}{r_{cg}} - 1 \right) \quad (43)$$

If the contact axes are assumed to be non-sticking, then the associated normal forces, if present, are constrained to positive values ($N_{a_i} > 0$). As each vehicle is constrained to producing a positive thrust in the \vec{K}_B direction ($T_{eq} > 0$), the following constraint on r_{eq} is established:

$$r_{eq} > r_{cg} \quad (44)$$

Due to the requirement that the equivalent thrust vector be located past the mass center in a direction normal to the axis of excitation for motion to occur, the probability of excitation for each contact axis can be computed as the probability that this constraint is satisfied. If the mass center of the composite aircraft is assumed to be a random variable that is distributed about the hypothetical payload according to a Gaussian distribution with mean located at the payload's geometric center, this probability can be approximated using the cumulative distribution function according to,

$$P(a_i | H_l) = \frac{1}{\sigma\sqrt{2\pi}} \int_{-\infty}^{r_{eq}} e^{-\frac{(t-\mu)^2}{2\sigma^2}} dt \quad (45)$$

where the standard deviation σ and mean μ are chosen such that the diameter of the hypothetical payload spans three standard deviations of the mean. A 2D visualization of this calculation is shown in Fig. 16.




Figure 16. Example Likelihood Calculation given Normal Mass Center Location Distribution.

As more information becomes available, i.e. excitations about current and new contact axes occur, the posterior probabilities can be updated as,

$$P(H_l | a_i, a_k) = \frac{P(a_k | H_l)P(a_i | H_l)P(H_l)}{P(a_k)P(a_i)} \quad (46)$$

where a_k is a different event (dynamic excitation) that is observed.

When implemented, the posterior likelihood estimates will favor hypotheses which order the modular vehicles from lowest to highest throttle percentage in a direction normal from the axis of excitation, thus maximizing the expression of Eq. (42). This strategy of maximizing r_{eq} can produce misleading hypotheses, depending on the magnitude of T_{eq} . The required torque needed to excite dynamic motion about a given contact axis remains

constant for a given composite aircraft configuration. As the observed value of T_{eq} increases between experiments, the corresponding r_{eq} on the actual payload must decrease to maintain this required torque value. Thus, maximizing the r_{eq} of the hypothetical payload for relatively large values of T_{eq} would skew the posterior probabilities of Eq. (46) towards misleading hypothesis. Therefore, weighting experiments that exhibit a lower equivalent thrust over experiments that exhibit a larger equivalent thrust improves the Bayesian model of the composite aircraft and ensures greater accuracy in the Bayesian learning process. One method of weighting likelihood estimates has been presented by Newton and Raftery [59]. Incorporating this weighting factor, the posterior probability is computed as,

$$P(H_l | a_i, a_k) = \frac{P(a_k | H_l)^{w_k} P(a_i | H_l)^{w_i} P(H_l)}{P(a_k)P(a_i)} \quad (47)$$

where the weight for a particular experiment, w_i , takes on a value between zero and one and is inversely proportional to the equivalent thrust value of that experiment. As new spin-up experiments occur, these weighting parameters are normalized according to the range of equivalent thrust values observed. Each contact axis maintains a set of weights that are dependent only on the thrust values observed during excitation about it. By decoupling the different axes weights, the Bayesian inference remains robust to non-geocentric mass center locations and uneven vehicle distributions.

Each hypothesis places each vehicle into one of the potential slots. By marginalizing the posterior probabilities, each vehicle is assigned a probability of being located in one of the available slots. This marginalization takes the total set of hypotheses and generates an $n \times n$ vehicle-slot likelihood estimate matrix, \mathbf{L} . The position vector of each slot about the

hypothetical payload can be expressed in the body fixed reference frame, $\mathbf{P}_s \in \mathbb{R}^{n \times 2}$. With \mathbf{L} and \mathbf{P}_s , each vehicle is assigned a probabilistic placement on the hypothetical payload that is updated after each excitation experiment.

$$\mathbf{P}_v = \mathbf{L}\mathbf{P}_s \quad (48)$$

Initially, the probabilistic vehicle position estimates are all located at the center of the hypothetical payload as each vehicle is equally likely to be located in any of the defined slot positions. As more information is observed, these probabilistic estimates grow and shift about the hypothetical payload.

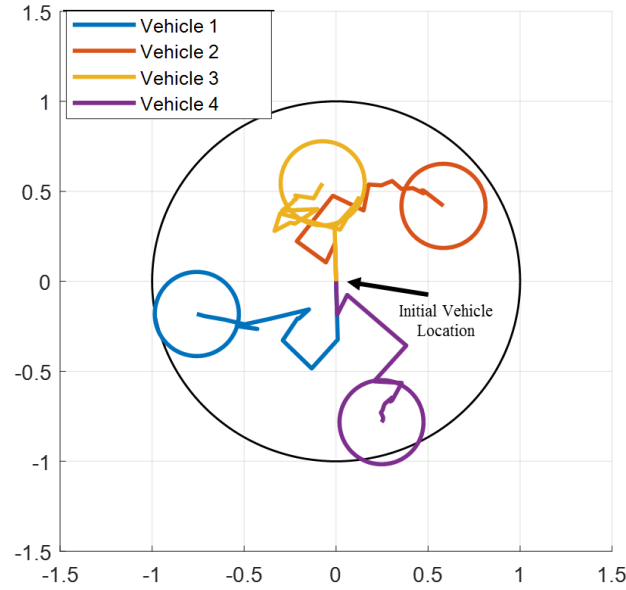


Figure 17. Example Placement History of Vehicles on Hypothetical Payload.

Figure 17 shows an example scenario where the flightworthiness algorithm is performed on a composite aircraft comprised of 4 attached vehicles. This figure shows the time evolution of the vehicle placements over the duration of the flightworthiness

algorithm. The circles in the figure represent the estimate that was reached at the end of the algorithm while the solid lines trace the placement history. The position estimates are spatially-weighted averages of the likelihoods for each slot location for a given vehicle. When the model is initialized, each vehicle has equal likelihood of being located in any of the defined slot positions. This equal likelihood results in the vehicle position estimates originating from the center of the hypothetical payload. As excitation thrust data is observed, the likelihood estimates update, and the position estimates shift according to the new likelihood averages. The Bayesian estimates are purely probabilistic and can vary significantly as new data is observed for reasons previously discussed.

While the probabilistic vehicle placements, \mathbf{P}_v effectively illustrate the Bayesian updating process, an alternate metric is used to command the thrust distribution when excitation about a particular axis is desired. This is accomplished by expressing each slot location by a perpendicular distance from each of the contact axes, $\mathbf{R}(s_j, a_i)$, thus, each vehicle can be assigned an axis effectiveness score, \mathbf{S} , that is weighted by the vehicle-slot probabilities.

$$\mathbf{S} = \mathbf{L}\mathbf{R} \quad (49)$$

where $\mathbf{S} \in \mathbb{R}^{n \times m'}$. In this way, each vehicle is assigned axes effectiveness scores proportional to its probabilistic relative distance from each contact axis. These probabilistic scores are used to drive excitations in the exploit phase by scaling the nominal throttle rate by the vehicle's relative effectiveness score. One drawback of this approach is observed when multiple excitations about the same axis are commanded in the exploit phase. For these cases, the probabilistic model between the multiple excitations may be largely

similar, thus producing near identical effectiveness scores and commanded throttle rates. However, to use Eq. (29) from the vehicle position estimation algorithm, the excitations about a contact axis must be linearly independent. To ensure unique excitations, a user-defined level of noise is added to perturb the throttle rates generated from the Bayesian model. The noise is a random variable defined by a uniform distribution. After the throttle rate perturbations are introduced, the vehicle throttle rates are re-normalized according to the highest throttle rate percentage. By varying the commanded throttle rates, the thrust matrix in Eq. (26) can better span the space of throttle rates, thus avoiding ill-conditioned data sets. By increasing the condition number of the thrust data sets, the vehicle position estimates are more robust to sources of error. These error sources are quantified in detail in the next section. The cost of these throttle rate perturbations comes in the form of commanded excitation reliability. As the introduced noise is increased in magnitude, the throttle rates approach purely random values (resulting in the explore phase strategy) and the algorithm no longer has a meaningful method to drive dynamic excitations. Therefore, this strategy introduces a tradeoff between the condition of the measured thrust matrix, and the accuracy of the probabilistic model in driving desire excitations. For this work, the throttle rate noise was chosen using a uniform distribution between 0% and 30%. This range was selected as it was found to consistently produce well-conditioned thrust data sets while only suffering a slight loss in commanded excitation accuracy. Ultimately, the goal of the machine learning component of the flightworthiness algorithm is to minimize the number of spin-ups and total control effort needed to generate the required, linearly independent, data for the vehicle position estimation algorithm.

CHAPTER 4. ERROR ANALYSIS

This chapter seeks to analyze and quantify the sources of error in the vehicle position estimation scheme described above, which has a direct effect on the accuracy of the overall flightworthiness determination. Three major sources of error have been identified which contribute to error in the resulting vehicle position estimates. These sources of error are:

1. A mismatch between the measured thrusts needed to create excitation during a spin-up, and the actual thrusts needed to break static equilibrium with the ground.
2. A mismatch between the thrust-throttle mapping used to convert the throttle percentages to thrust values in the flightworthiness algorithm, and the actual thrust-throttle mapping.
3. A mismatch between the weight estimate used in Eqs. (1)-(4) and the actual weight of the composite aircraft.

The goal of this section is to quantify the relationship between error in the vehicle position estimates and the three sources of error outlined above. This relationship highlights important trends in algorithm performance as shown in the subsequent experimental results.

With reference to the first source of error (item 1 above), when computing the vehicle position estimates using the static moment balance in Eq. (25), it is assumed that the thrust values about a given axis are recorded at the instant that the composite vehicle transitions from static equilibrium to dynamic excitation. Practically speaking, this transition must be

detected at some non-zero excitation angle α_D . During each spin-up experiment, the attached vehicles linearly increase their throttle levels according to the commanded throttle rate distribution. At the instant of excitation detection, the throttle levels are higher than those needed to break static equilibrium. The violation of this static assumption introduces error into the position estimates. In order to quantify the effects of this violation, consider the moment balance equation for excitation about axis a_i at the exact time t_s at which the vehicle thrusts balance the moment of the composite vehicle weight about the contact axis,

$$0 = -(\vec{r}_{cg/a_i} \cdot \vec{J}_i)W + t_s \sum_{j=1}^n (\vec{r}_{v_j/a_i} \cdot \vec{J}_i) \Delta T_j \quad (50)$$

where ΔT_j is the throttle rate of vehicle j . After p spin-up experiments about a particular axis, the moment balance equations can be accumulated into a matrix-vector form:

$$\left\{ \begin{array}{c} \frac{W(\vec{r}_{cg/a_i} \cdot \vec{J}_i)}{t_{s_1}} \\ \vdots \\ \frac{W(\vec{r}_{cg/a_i} \cdot \vec{J}_i)}{t_{s_p}} \end{array} \right\} = \begin{bmatrix} \Delta T_{11,i} & \cdots & \Delta T_{1n,i} \\ \vdots & \ddots & \vdots \\ \Delta T_{p1,i} & \cdots & \Delta T_{pn,i} \end{bmatrix} \left\{ \begin{array}{c} \vec{r}_{v_1/a_i} \cdot \vec{J}_i \\ \vdots \\ \vec{r}_{v_n/a_i} \cdot \vec{J}_i \end{array} \right\} = A\vec{x} \quad (51)$$

Let the left hand side of Eq. (51) be denoted as \vec{b} . After the vehicle breaks static equilibrium, the rotational dynamics of the composite vehicle can be described using a small angle approximation as,

$$I_{a_i} \ddot{\alpha}_i = -(\vec{r}_{cg/a_i} \cdot \vec{J}_i)W + (t_s + t_D) \sum_{j=1}^n (\vec{r}_{v_j/a_i} \cdot \vec{J}_i) \Delta T_j \quad (52)$$

where I_{ai} is the inertia about axis a_i and t_D is the amount of time that the composite vehicle spends in dynamic motion (after leaving static equilibrium). Integrating Eq. (52) twice and solving for t_D yields,

$$t_D = \sqrt[3]{\frac{6I_{a_i}\alpha_D}{\sum (\vec{r}_{v_j/a_i} \cdot \vec{J}_i)\Delta T_j}} \quad (53)$$

where α_D is the excitation detection threshold angle. With these quantities defined, the static expression of Eq. (50) used to estimate vehicle positions can be rewritten in terms of the total spin-up time and estimated parameters as,

$$\left\{ \begin{array}{c} \frac{\bar{W}(\vec{r}_{cg/a_i} \cdot \vec{J}_i)}{t_{s_i} + t_{D_i}} \\ \vdots \\ \frac{\bar{W}(\vec{r}_{cg/a_i} \cdot \vec{J}_i)}{t_{s_p} + t_{D_p}} \end{array} \right\} = \left[\begin{array}{ccc} \Delta \bar{T}_{11,i} & \cdots & \Delta \bar{T}_{1n,i} \\ \vdots & \ddots & \vdots \\ \Delta \bar{T}_{p1,i} & \cdots & \Delta \bar{T}_{pn,i} \end{array} \right] \left\{ \begin{array}{c} \vec{r}_{v_1/a_i} \cdot \vec{J}_i \\ \vdots \\ \vec{r}_{v_n/a_i} \cdot \vec{J}_i \end{array} \right\} = \bar{A} \hat{\vec{x}} \quad (54)$$

where \bar{W} is the estimated aircraft weight and $\Delta \bar{T}$ is the predicted vehicle thrust rate using the assumed thrust-throttle mapping. If the thrust-throttle relationship of each vehicle is exactly known, then $\bar{A} = A$ from Eq. (54) and the only differences between Eqs. (51) and (54) are from the nonzero t_D values and from error in the estimated weight value, \bar{W} . Practically, the thrust-throttle mapping cannot be exactly known and thus \bar{A} will differ from A as well. The extent of this error can be minimized through extensive testing or potentially through the use of embedded force sensors [60]. Letting the left-hand side of Eq. (54) be denoted as $\vec{\bar{b}}$, the quantity $\delta \vec{b} = \vec{\bar{b}} - \vec{b}$ captures the error due to the violation of the static assumption and error in the weight estimate. The quantity $\delta A = \bar{A} - A$ contains the error

due to an imperfect thrust model. In estimating the vehicles' positions, the expected error $\delta\vec{x} = \vec{\bar{x}} - \vec{x}$ can be bounded by the following inequality

$$\frac{\|\delta\vec{x}\|}{\|\vec{x}\|} \leq \kappa(A) \left(\frac{\|\delta A\|}{\|A\|} + \frac{\|\delta\vec{b}\|}{\|\vec{b}\|} \right) \quad (55)$$

given that

$$\|\delta A\| < \frac{1}{\|A^{-1}\|} \quad (56)$$

where $\kappa(A) = \|A\|\|A^+\|$ is the condition number of A [61]. A key takeaway from Eq. (55) is that the degree to which errors in the weight estimate, thrust-throttle mapping, and time delay to detect excitation affect error in the resulting vehicle position calculation $\delta\vec{x}$ is dependent on $\kappa(A)$. The condition number of A is determined by how linearly independent the thrust distributions are that cause excitation in each spin-up experiment about axis a_i . If each of the thrust distributions used in Eq. (54) are nearly orthogonal, then $\kappa(A)$ is small and the algorithm will be robust to the three main sources of error. If each of the thrust distributions are similar, the algorithm sensitivity to these error sources will be higher.

In comparing the results of the experimental tests detailed below with analogous results from simulation, it was found that vehicle position estimates were relatively insensitive to error in the thrust-throttle mapping. Thus, to more fully explore the algorithm error characteristics, consider a case involving perfect knowledge of the thrust-throttle mapping, and error only in the weight estimate and the thrust values needed to break static

equilibrium (items 1 and 3 above). The error associated with spin-up j due to these factors can be expressed as,

$$\delta b_j = b_j - \bar{b}_j = \frac{(\vec{r}_{cg/a_i} \cdot \vec{J}_i)(t_{s,j}\delta W - t_{D,j}W)}{t_{s,j}^2 + t_{s,j}t_{D,j}} \quad (57)$$

where $\delta W = \bar{W} - W$. Rewriting this quantity in the form $\delta b_j = k_j b_j$ where k_j is a constant yields,

$$k_j = \frac{\delta W(\vec{r}_{cg/a_i} \cdot \vec{J}_i) - (6I_{a_i}\alpha_D)^{1/3} \left(\sum (\vec{r}_{v_j/a_i} \cdot \vec{J}_i) \Delta T_j \right)^{2/3}}{W(\vec{r}_{cg/a_i} \cdot \vec{J}_i) + (6I_{a_i}\alpha_D)^{1/3} \left(\sum (\vec{r}_{v_j/a_i} \cdot \vec{J}_i) \Delta T_j \right)^{2/3}} \quad (58)$$

where (58) is computed by plugging in Eqs. (50) and (53) into Eq. (57). If the weight estimation error, δW , is negligible, then k_j is a strictly negative value that is bounded by 0 and 1. Equations (58) and (55) show that vehicle position errors grow as the detection angle α_D , the inertia about the excitation axis, and the number of attached vehicles increase. Conversely, increasing the total weight of the composite vehicle while the other parameters remain constant decreases the value of k_j , thereby decreasing position estimate errors. As the composite aircraft increases in scale, the relative increase between the weight of the aircraft and the aircraft's inertia about its contact axes can be used to predict the error behavior using Eqs. (55) and (58). If the weight estimation error, δW , is non-negligible then a constant bias is introduced to each element in the k vector.

CHAPTER 5. SIMULATION RESULTS

The proposed algorithm has been implemented on a simulation model of a composite aircraft. The flightworthiness algorithm is initialized assuming no prior knowledge of vehicle layout, composite aircraft mass center location, total weight, or contact axis geometry. Each vehicle's thrust is limited by a defined maximum value and the number of vehicles in any given simulation is defined by a constant n . Thrust values for a successful spin-up were recorded at an excitation detection angle threshold of $\alpha = 5$ degrees. The attached vehicles were assumed to each occupy 1/12 of the total attachment area of the payload and were placed within the convex hull of the contact axes (all payloads have convex shapes). The vehicles were placed semi-randomly about the perimeter of the payload with the requirement that no two vehicles occupy the same space. The algorithm was tested on payloads with contact axes that formed rectangles, triangles, and pentagons. Note that the effects of sensor noise or other random disturbances that may occur during excitation are not considered here. However, it should be noted that the excitation detection angle threshold α should be set high enough so that noisy measurements when the vehicle is at rest on all contact points do not mistakenly trigger the threshold. In light of typical IMU pitch and roll root-mean-square error values on the order of a few degrees, $\alpha = 5$ deg is chosen for the simulation studies that follow.

The first section of results presents several example cases to highlight the various stages and outputs of the flightworthiness determination algorithm. The second section quantifies the algorithm's performance by presenting Monte Carlo simulations in which several parameters are varied. Unless stated otherwise, all simulations are initialized with a mass

center location at the geometric center of the composite aircraft (although this is varied in a Monte Carlo simulation). The accuracy of the estimates as well as the number of spin-ups required to achieve the vehicle position estimates and an overall flightworthiness evaluation are used as metrics to quantify the algorithm's performance.

5.1 Example Case Results

Three example cases are presented with $n = 6$ attached vehicles and a total thrust to weight ratio of 2.0 at the maximum thrust saturation limit (where this is calculated from the total thrust generated from all vehicles). The three example cases consider payloads with contact axes that form a rectangle, a triangle, and a pentagon, respectively. Each case commands excitations for the minimum number of spin-ups required to estimate the vehicles' positions as defined in the vehicle position estimation section. A hover throttle threshold of 70% is set to determine flightworthiness.

The first example case was performed on a rectangular payload with a vehicle configuration shown in Fig. 18. The axes are labeled in the order that they were observed during the exploration phase.




Figure 18. Example Modular Vehicle Configuration with Rectangular Contact Axes.

Figure 19 shows the excitation results for the first example case. The solid line in Fig. 19 shows the axis about which the composite aircraft excited during each of the vehicle spin-ups. The dashed black line shows the points in the algorithm where the composite aircraft was commanded to excite about a certain axis. These commands occur during the exploitation phase after the algorithm observed all four axes twice. The algorithm labels the axes in the order that they are observed. Spin-up excitations end when the required amount of data has been generated and the vehicle position estimation algorithm has enough data to calculate the vehicle weight and position estimates.

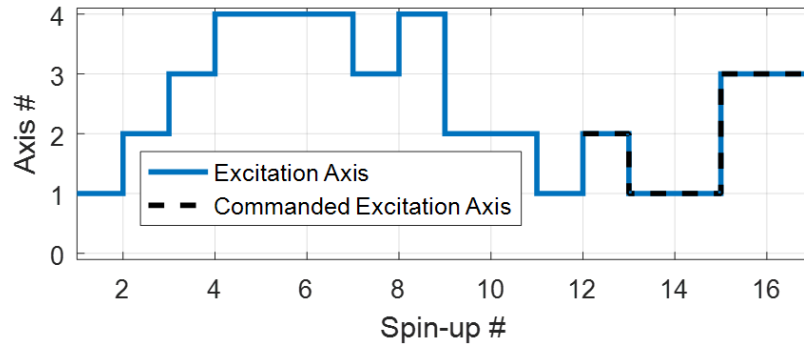


Figure 19. Rectangular Aircraft Excitation Axis History.

For the first 11 spin-ups the algorithm performed the exploration phase. During this phase, the vehicles' commanded thrust ratios were randomly selected and the resulting excitation axis and thrust distributions were used to propagate the Bayesian likelihood estimates. For this example case, the algorithm recognized that two of the contact axes were parallel and, as a result, required that four axes total be observed before progressing. After spin-up 11, the algorithm entered the exploitation phase and utilized the Bayesian estimates to command excitation about the 4 discovered axes. Figure 19 shows that the

composite aircraft was able to successfully excite about all commanded axes. After the minimum number of spin-ups per axis (4) was observed, the algorithm was allowed to progress to the deterministic vehicle position estimation algorithm. In this case, the algorithm was able to generate the required data in 16 spin-ups, which is the minimum number possible for a rectangular geometry.




Figure 20. Vehicle Effectiveness Scores for Axes 1 (left) and 2 (right).




Figure 21. Vehicle Effectiveness Scores for Axes 3 (left) and 4 (right).

Figures 20 and 21 show the evolutions of the vehicles' axis effectiveness scores as the algorithm progresses. Each axis effectiveness score is initialized when the axis is first observed and is updated at each spin-up. Due to the pairs of parallel axes, the effectiveness scores are mirrored for opposite axes and only update when excitation occurs about a non-

perpendicular contact axis. Figures 20 and 21 show that the algorithm was able to produce an accurate effectiveness distribution for the majority of the vehicles. Figure 21 shows that for axis 3, vehicle 5 was given a high effectiveness score despite vehicle 5's small moment arm about axis 3. This is not unexpected as there is no guarantee that the Bayesian likelihood estimates accurately represent the physical vehicle placements. However, with sufficient exploration they can be used to command excitation and to drive the vehicle position estimation algorithm. After 16 spin-ups, the recorded thrust values and their associated axes of excitation were sent to the vehicle position estimation algorithm. Equation (38) was used for both pairs of parallel contact axes and two separate scaling dimensions were introduced. The accuracy of the weight and vehicle position estimates was determined by using the correct scaling parameters and computing relative errors for the composite aircraft weight and each vehicle position in the \vec{J}_i and \vec{I}_i directions. For position estimates in each direction, the median of all relative errors is taken as a representative value. For this example case, the weight was estimated with 10% error and the vehicle positions in the \vec{J}_i and \vec{I}_i directions were estimated with a median error of 37% and 30%, respectively. Figure 22 shows the position estimates of the vehicles displayed alongside the actual vehicle configuration, where the dashed circles denote the estimated locations.




Figure 22. Example Vehicle Configuration and Position Estimates for Rectangular Contact Axes.

In this case, the algorithm was able to estimate the vehicle positions about the composite aircraft with reasonable accuracy. The estimation errors can be attributed to the fact that the vehicle thrusts are continually increasing to the point of excitation detection, which occurs at a nonzero excitation angle. By the time that the thrust values are recorded, they have surpassed the required amount needed to exit static equilibrium and, therefore, no longer exactly satisfy Eq. (25). The extent and form of this error is discussed in detail in the previous Error Analysis section. The vehicle position estimates were substituted into Eq. (39) to determine the vehicle throttle levels required to achieve a stable hover. The minimum norm throttle solution for the estimated and actual vehicle locations is shown in Fig. 23. In this case, because both the estimated and actual vehicle positions are distributed somewhat evenly about the mass center, the required throttle values to hover are relatively even across all vehicles. With an available thrust to weight ratio of 2, both the estimated and actual aircraft would be able to achieve hover with all vehicles operating well below the defined throttle limit. The flightworthiness determination algorithm was able to estimate the required throttle percentages of the actual aircraft with a root mean square

(RMS) error of 6.41% throttle. With these results, the algorithm classifies the composite aircraft as flightworthy.

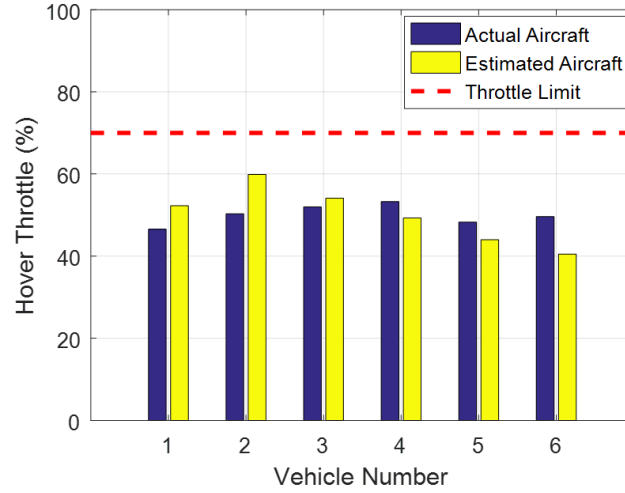


Figure 23. Estimated and Actual Vehicle Throttle Percentages Required to Hover for Rectangular Axes.

The second example case was performed using a payload with triangular contact axes and a vehicle distribution and mass center location as shown below in Fig. 24. The spin-up excitations are shown in Fig. 25. For triangular contact axes, the algorithm must discover each axis during the exploration phase in order to progress. During the exploration, two spin-ups were classified as contact point excitations and the corresponding thrust information was discarded (these are denoted as an excitation about axis 0). Figure 25 shows that 18 spin-up iterations were required before axis 3 was observed a second time. After the exploration phase, the exploit phase was able to successfully command the composite aircraft about the remaining required axes. For this example, 23 spin-up iterations were required in order to generate the necessary data. The

minimum possible number of spin-ups for 3 independent contact axes and 6 modular vehicles is 18, thus, a total of 5 excess spin-ups were performed for this example case.




Figure 24. Example Vehicle Configuration and Position Estimates for Triangular Contact Axes.

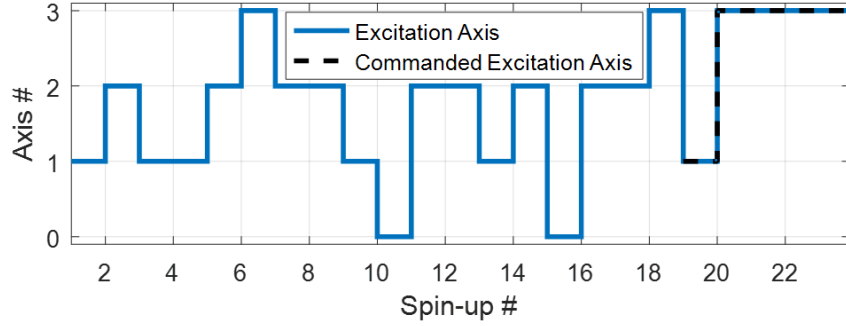


Figure 25. Triangular Aircraft Excitation Axis History.

The three contact axes shown in Fig. 24 are linearly independent and, thus, Eq. (29) was used to estimate the weight of the composite aircraft and the positions of the modular vehicles. For this example case, Eq. (29) resulted in a weight error of 16% and a median position error in the \vec{J}_i and \vec{I}_i directions of 19% and 40% respectively. Figure 24 shows the position estimates of the modular vehicles (dashed lines). The disproportionate amount of error in the \vec{I}_i direction for this case occurs due to the defined direction of the scaling

term used in calculation of the position estimates. From Eq. (29), this scaling term is used to directly solve for positions in the \vec{J}_i direction. Conversely, due to the non-orthogonality of the contact axes, the position estimates along the \vec{I}_i direction are dependent on the \vec{J}_i estimates which, themselves, contain errors from the imperfect thrust measurements. This dependency leads to compounded error for estimates in the \vec{I}_i direction.




Figure 26. Estimated and Actual Vehicle Throttle Percentages Required to Hover for Triangular Axes.

Figure 26 shows the flightworthiness results for this example case. It can be seen that all throttle percentages are below the threshold value and so the composite aircraft is deemed flightworthy. Even with an estimation error in the \vec{I}_i direction of 40%, the estimated throttle percentages matched those of the actual system with an RMS error of 3.34% throttle. The estimated vehicle positions in the \vec{I}_i direction are all closer to the mass center than the actual values, and thus the estimated throttle percentages are still reasonably accurate. If the threshold were defined at 60% instead of 70%, the estimated aircraft would

be deemed flightworthy despite the fact that vehicle 1 would be required to run at the throttle cutoff. While the estimated throttle percentages are always subject to some error and thus it is possible to produce false positives or false negatives for flightworthiness, this can be mitigated by using a threshold throttle cutoff well below 100% and/or a higher total thrust to weight ratio.

The third example case was performed on a payload with five distinct contact axes and a vehicle distribution shown in Fig. 27. This geometry exhibits an uneven distribution of vehicles about the mass center location. Although the payload has five contact axes, the algorithm only requires excitations about three to estimate the vehicle positions. The algorithm explores until three axes have been observed, then progresses to the exploitation phase without checking for the existence of additional axes.




Figure 27. Example Vehicle Configuration and Position Estimates for Pentagon Contact Axes.

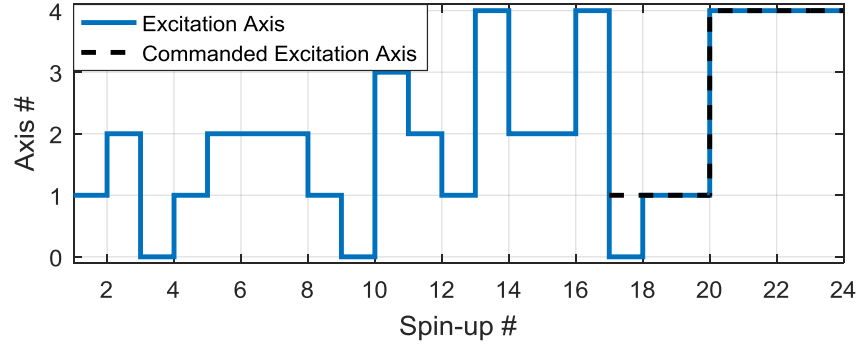


Figure 28. Pentagon Aircraft Excitation Axis History.

The spin-up excitations for this case are shown in Fig. 28. The algorithm observed excitation about four of the five contact axes. During the exploit phase, the algorithm commanded excitation about axis 1 and detected motion about a contact point. The machine learning subcomponent was able to generate the required set of thrust data about axes 1, 2, and 4 in 23 spin-up iterations, resulting in an excess of 8 spin-ups. For this example case, Eq. (29) resulted in a weight error of 12% and a median position error in the \vec{J}_i and \vec{I}_i directions of 15% and 26% respectively.

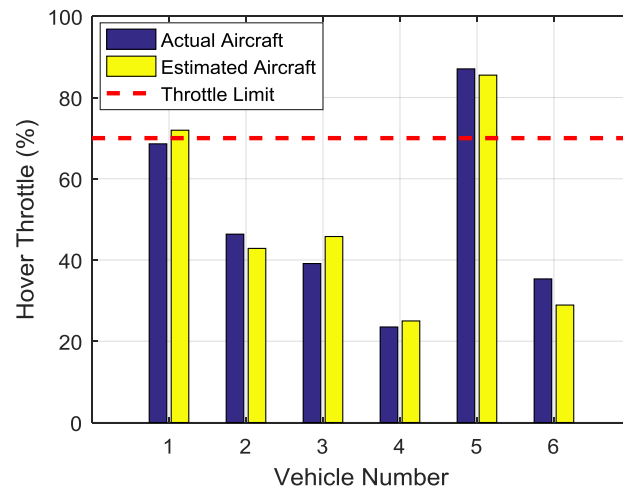


Figure 29. Estimated and Actual Vehicle Throttle Percentages Required to Hover for Pentagon Axes.

Figures 27 and 29 show the position estimates of the modular vehicles and the estimated throttle percentages to maintain a hover, respectively. Figure 29 shows that the algorithm was able to successfully determine that the modular vehicle configuration was not flightworthy. Due to the uneven distribution of vehicles, vehicle 5 would be required to maintain over 80% throttle to establish hover. Figure 29 shows that, even with a substantial amount of position estimation error, the flightworthiness determination algorithm was able to produce hover throttle estimates with 4.36% RMS error.

5.2 Monte Carlo Results

Due to the limited number of assumptions, there is a large set of parameters that can potentially be varied to alter the algorithm's behavior. This section analyzes several parameters that have dominant, meaningful, and coupled effects on the algorithm's performance. The investigated set of parameters includes the required spin-ups per axis, the overall thrust-to-weight ratio of the composite aircraft, and the mass center location with respect to the geometric center of the contact points. These studies are performed using payloads with contact axes that form rectangles, triangles, and pentagons. Each trade study is performed with varying values of n between 4 and 8. Changing the number of modular vehicles directly affects the algorithm's performance due to attachment restrictions on fixed area payloads and the geometry of the contact axes. To account for this coupling, the results are averaged together for all values of n to produce a representative metric independent of the number of attached vehicles. For each combination of parameters, 200 simulations are performed in which the vehicle locations and the direction in which the mass center is shifted are randomized. After each simulation, the accuracy of the weight and positions estimates is computed. Also, the number of spin-

up excitations performed in each simulation is recorded and compared to the minimum required value, resulting in an excessive spin-up metric. Finally, the number of simulations that hit the maximum allowed number of spin-ups while attempting to generate the required set of data was recorded (denoted as a failed case). The estimation accuracy quantifies the performance of the vehicle position estimation algorithm while the number of excessive spin-ups and the percentage of failed cases quantify the performance of the Bayesian inference algorithm.

For the first set of Monte Carlo simulations the number of required excitations per axis was varied from 1.0 to 2.0 times the minimum number required to sufficiently populate Eqns. (29) and (38). The simulations were run with a thrust-to-weight ratio of 2. Figure 30 shows the weight and vehicle position estimation errors as a function of the required spin-up multiplier, for all three contact axis geometries. These results show that for each composite aircraft the vehicle position estimates can be considerably improved by requiring additional excitations per axis. For payloads with rectangular contact axes, requiring an additional 25% beyond the minimum resulted in a 170% improvement in position estimates. Each set of thrust measurements contains error due to the assumption that the thrust values are recorded at the transition from static equilibrium to dynamic motion, which cannot happen practically. With additional excitations per axis, these errors can be minimized in a least squares sense. Figure 30 also shows that the weight estimate for the rectangular payload benefits from the additional spin-ups whereas the weight estimates for pentagon and triangle payloads remain relatively accurate and constant.

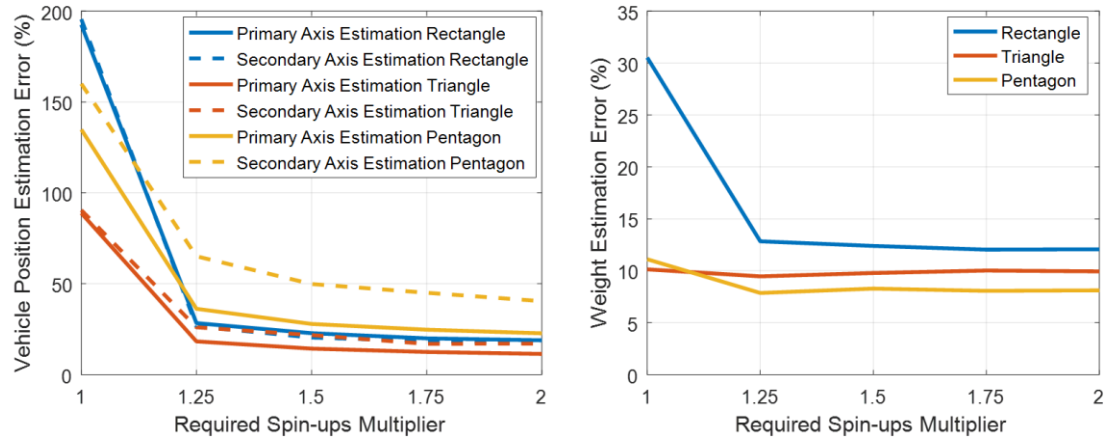


Figure 30. Vehicle Position Estimation Error (left) and Weight Estimation Error (right) for Increasing Required Excitations per Axis.

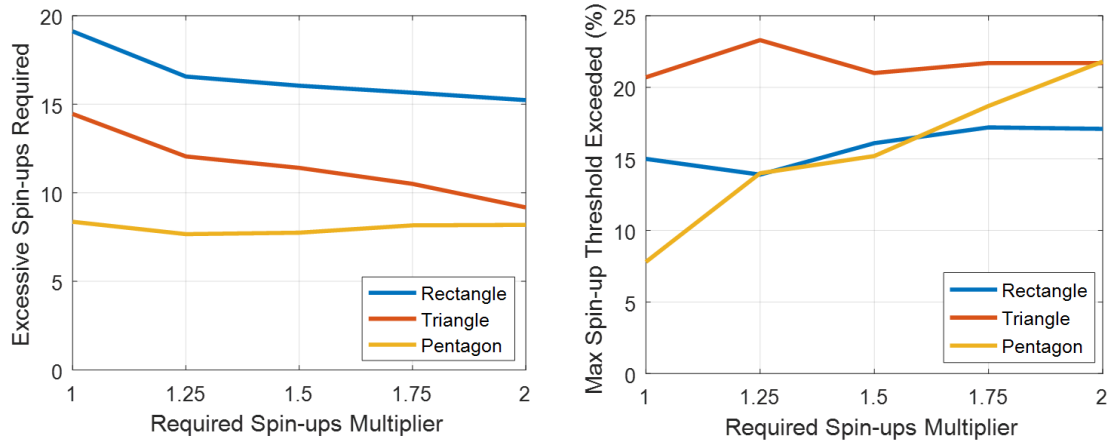


Figure 31. Excessive Excitations (left) and Unsuccessful Simulations (right) for Increasing Required Excitations per Axis.

Figure 31 shows the average excessive spin-ups required and the percentage of the 200 simulations that exceeded the max spin-up threshold. It is clear that the additional required spin-ups has a minimal effect on the average number of excessive spin-ups. If the explore phase can successfully observe the required number of contact axes (3 or 4), then commanding excitations using the axis effectiveness scores yields a high success rate.

Therefore, commanding additional excitation in simulations that already identified all required axes carries little risk of failure. The total number of required spin-ups increases, but the number of excessive spin-ups remains nearly constant. The spin-up multiplier adds to the minimum number of excitations required to perform the vehicle position estimation algorithm, but greatly improves the resulting position estimates. As the multiplier increases past 1.25, the improvements plateau while the actuation effort (spin-ups) continues to increase. Therefore, a spin-up multiplier of 1.25 was chosen as the baseline to run the following scenarios.

The second set of trade studies varied the thrust-to-weight ratio of the composite aircraft by altering the saturation thrust of each attached vehicle. Again, payloads with all three contact axis geometries were simulated with the number of attached vehicles ranging from 4 to 8. The resulting weight and position estimation errors are shown in Fig. 32. Note that the vehicle position estimates improve with an increase in the available thrust. For payloads with five contact axes, increasing the thrust to weight ratio from 1.25 to 1.5 improves position estimations about the primary and secondary directions by 14% and 25%, respectively. As the available thrust decreases, there is an increasing probability that a vehicle, or set of vehicles, hits maximum throttle during a spin-up. Accurate position estimates are reliant on a well-conditioned thrust history matrix from Eq. (26). If a set of vehicles saturate for every excitation about a given axis, they no longer provide linearly independent data and the resultant thrust matrix becomes ill-conditioned. Once the available thrust increases to the point that saturation is avoided, there is no benefit, in the context of this algorithm, to further increases in the lifting capabilities of the attached vehicles.

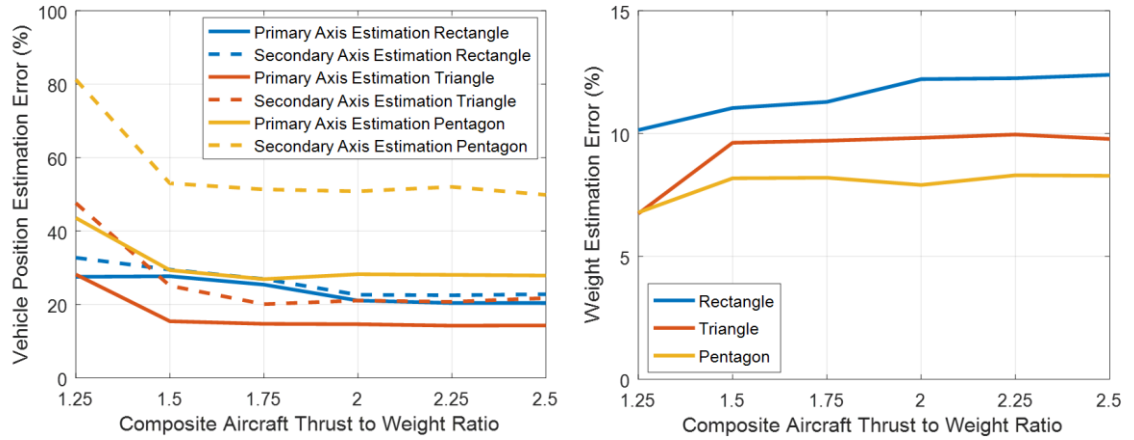


Figure 32. Vehicle Position Estimation Error (left) and Weight Estimation Error (right) for Increasing Aircraft Thrust to Weight Ratio.




Figure 33. Excessive Excitations (left) and Unsuccessful Simulations (right) for Increasing Aircraft Thrust to Weight Ratio.

The available thrust also has a significant effect on the machine learning algorithm's ability to generate the required set of data. Figure 33 shows the mean excessive spin-ups and the percentage of simulations that hit the max spin-up threshold. Note that, for rectangle and triangle contact geometries, there is significant coupling between the available thrust and the percentage of failed simulations. For triangular and rectangular

contact axes, each contact axis must be excited about in order to progress to the vehicle position estimation phase. Due to the randomized vehicle placement about the payload, there are simulations with an uneven distribution of vehicles. These uneven distributions, in combination with limited thrust, can result in scenarios where it is highly unlikely, or even impossible, for the aircraft to excite about each axis. A payload with contact axes that form a pentagon must excite only about three of the five possible axes and, thus, is more robust to limited thrust and uneven vehicle distributions. Figures 32 and 33 show that insufficient available thrust degrades the performance of both the vehicle position estimation algorithm as well as the machine learning subcomponent. For all results metrics, however, there is no noticeable improvement in performance at thrust-to-weight ratios greater than 2.

The final set of simulations investigated the sensitivity of the algorithm's performance to shifts in the mass center of the composite aircraft. The distance from the geometric center of the contact points to the nearest contact axis was used as a maximum length for the mass center perturbation. Perturbations were made at increasing percentages of this maximum distance in a direction that was randomized for each simulation. This set of simulations was performed with a required spin-up multiplier of 1.25 and with a total thrust to weight ratio of 2.0.

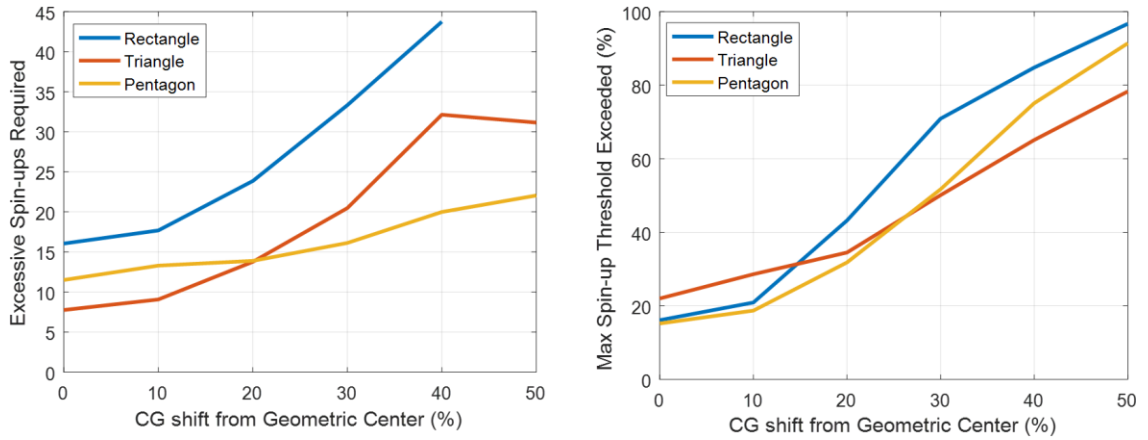


Figure 34. Excessive Excitations (left) and Unsuccessful Simulations (right) for Increasing Mass Center Perturbations.

Figure 34 shows the machine learning subcomponent’s performance with increasing mass center shifts. This figure shows the dominant role that the mass center location plays in allowing the algorithm to excite about the required number of contact axis. At a mass center shift magnitude of 50%, all three payload geometries reach the maximum number of spin-ups in over 80% of simulations. For payloads with rectangular contact axes, simulations run with 7 and 8 attached vehicles hit the spin-up threshold in 100% of simulations. At 80% failed simulations, only 40 successful simulation runs are being averaged to generate the excessive spin-ups plot in Fig 34. With so few simulations, there is a large amount of variance in the results. As the mass center is shifted from the geometric center of the contact points, the vehicle distribution about the mass center location becomes increasingly skewed. Vehicles that are closer to the mass center location lose effectiveness at producing moments about the contact axes and are largely dominated by the vehicles with an increased distance from the mass center. Because the exploration phase uses random thrust distributions, it becomes increasingly difficult to observe excitations about

the contact axes furthest from the shifted mass center location. The machine learning algorithm is effective for mass center shifts between 0% and 10%, but begins to deteriorate as the shifts become larger.

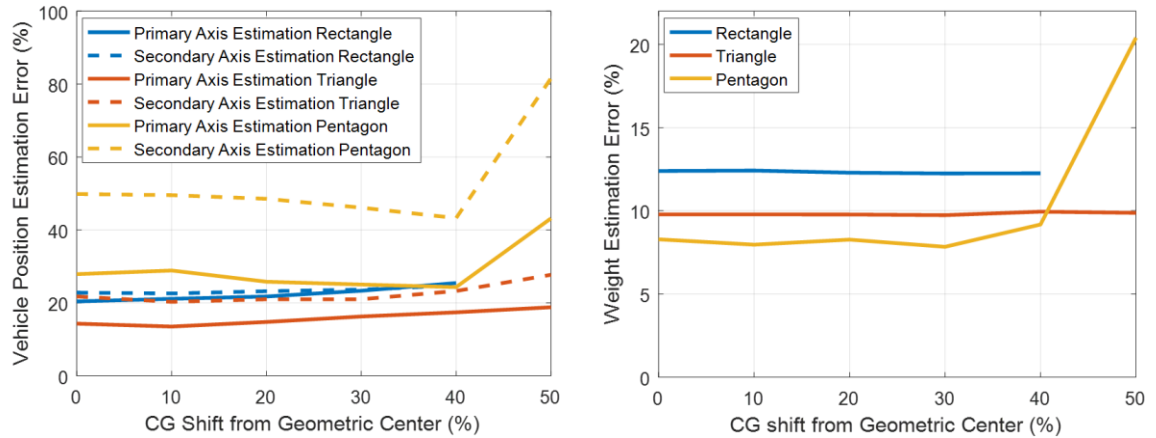


Figure 35. Vehicle Position Estimation Error (left) and Weight Estimation Error (right) for Increasing Mass Center Perturbations.

Figure 35 shows the accuracy of the vehicle position estimation phase as a function of the mass center shifts. These results show that, for both position and weight estimates, mass center shifts have a mostly insignificant effect on the estimation error. With adequate available thrust, moving the mass center location has negligible effect on how well-conditioned the thrust matrices about each axis are. If the machine learning subcomponent is able to generate the required excitations about the contact axes, the vehicle position estimation phase is able to successfully utilize the data regardless of the associated mass center shift. Unlike the thrust-to-weight trade study, perturbations to the mass center location only meaningfully affect the machine learning component of the flightworthiness determination algorithm. Figures 34 and 35 show that for simulations performed with a thrust to weight ratio of 2.0, a required spin-up multiplier of 1.25, and mass center

perturbations within 10%, the algorithm is able to perform the required excitations for approximately 80% of the simulations run and is able to do so with an average of approximately 15 excessive spin-ups. With these generated sets of data, the algorithm is able to approximate the weight of the composite aircraft and the attached vehicle positions within 13-20% of their true value, excluding the pentagon secondary axis position estimates.

Table 1. Flightworthiness Determination Trade Study Conclusions.

	Decreases Estimation Error (Vehicle Position Estimation)	Improves Data Generation (Machine Learning Phase)
Increasing Required Spin-ups	X	
Increasing Available Thrust	X	X
Decreasing CG Perturbations		X

The results of all three trade studies are summarized in Table 1, showing the impact that the three varied quantities have on the performance of the algorithm. In order to accurately estimate the weight and vehicle positions, and to do so using an acceptable amount of control effort, it is imperative that the composite aircraft has a sufficient amount of total thrust and that the mass center location is somewhat near the geometric center of the contact points. Also, the quality of the generated estimates can be greatly improved by requiring additional excitations about the commanded contact axes. These results were critical in driving component selection and composite aircraft construction for the experimental portion of this work, detailed in the next section.

CHAPTER 6. EXPERIMENTAL SYSTEM DESCRIPTION

To evaluate experimental performance of the flightworthiness determination algorithm, a set of unique payloads and modular vertical lift vehicles were constructed. This section details the design of the experimental setup and provides a brief description of the associated processing elements.




Figure 36. Modular Vehicle Design.

Figure 36 shows the modular vehicle design and its corresponding elements constructed for experimental testing. A set of 6 identical vehicles were built. Each lift vehicle can attach to the payloads via a clip mechanism. Extension arms were used to offset the attachment point from the rotor blades such that the downwash on the composite vehicle was minimized. Each vehicle houses its own power supply, microcontroller, radio, and rotor. As such, each vehicle is entirely modular. The lift vehicles each weigh 3.6 N and are capable of producing a gross thrust of approximately 11 N at maximum throttle. The

attached microcontrollers were used to provide throttle commands to their corresponding rotor and to communicate with the separate onboard CPU attached to the payload.



Figure 37. Experimental Payload Platforms.

Acrylic payloads of differing weights and geometries were constructed with mounting surfaces along their parameters so as to interface with the modular vehicles, shown in Fig. 37. The payloads had contact axes forming a square, a triangle, and a pentagon and weighed 9.3, 11.7, and 14.9 N, respectively. The extended mounting surfaces along the payload perimeter provide flexibility in vehicle attachment locations. Figure 1 shows an example configuration where six of the modular vehicles have been attached to the square payload. An inertial measurement unit (IMU) coupled to a microcontroller (MCU) is placed on the top face of the payload. The IMU contains a three-axis gyro and provides real time feedback on the orientation of the composite aircraft. The IMU-MCU package is tasked with running each spin-up experiment while actively monitoring the orientation of the composite vehicle.

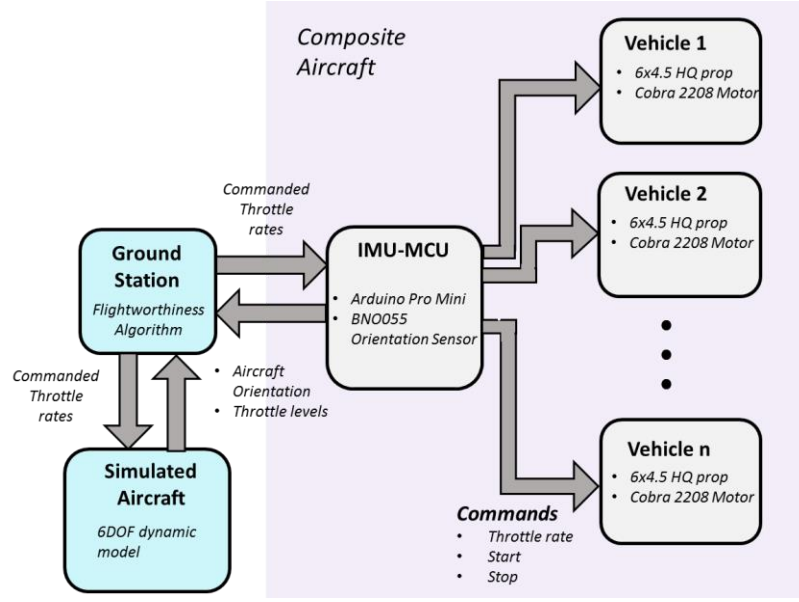


Figure 38. Flightworthiness Determination Architecture Design.

Figure 38 shows a diagram of the different experimental elements and the corresponding data flow. The described flightworthiness algorithm is implemented on a laptop computer (ground station). For each spin-up, the ground station sends the computed throttle rates to the IMU-MCU attached to the payload. The IMU-MCU sends the throttle rates to the corresponding vehicles and initiates the spin-up experiment. While the vehicles are increasing their thrusts according to the prescribed throttle rates, the IMU-MCU is monitoring the orientation of the composite vehicle. Once the IMU-MCU observes excitation, a stop command is sent, and the current throttle levels of each vehicle is reported back to the ground station to be processed in the flightworthiness algorithm.

To calculate the weight of the aircraft and the positions of the attached vehicles from Eqs. (29) and (38), the recorded throttle percentages of the attached vehicles must be converted to thrust values. The thrust-throttle relationship was obtained experimentally by

attaching the individual vehicles to a Futek LCF300 loadcell and recording the produced force across the range of throttle percentages. Figure 39 shows the results of the thrust-throttle experiment for a particular vehicle when tested at high and low battery charge. The results generated at full battery charge and near empty battery charge were averaged to generate the predictive model. Figure 39 was used to provide a direct mapping between throttle percentage and generated thrust.




Figure 39. Experimental Lift Vehicle Thrust-Throttle Mapping.

For modular lift vehicles, such as the one shown in Fig. 36, there is a noteworthy tradeoff between the lift vehicles' physical footprint on the composite aircraft, the weight of the vehicles, the lifting capacity of each vehicle, and the endurance of the composite aircraft. When selecting a power supply for the vehicles, it was desired that the aircraft determine flightworthiness while retaining sufficient battery charge for the subsequent stabilization and transportation of the payload. These considerations drove component selection for the lift vehicles and the overall aircraft design. Figure 40 shows an experiment in which a composite aircraft, with a thrust to weight ratio of 2 at maximum throttle, was

commanded to perform 100 consecutive spin-up excitations. Each spin-up randomly assigned throttle percentages to the attached vehicles. At various points in the 100 spin-ups, the resting voltage of each vehicle battery was measured, averaged, and then plotted as a function of the corresponding spin-up number.




Figure 40. Composite Aircraft Endurance Experiment.

Figure 40 shows that the aircraft was able to complete the 100 spin-ups without significantly depleting the onboard lithium polymer batteries. A maximum spin-up limit of 100 was enforced for this algorithm and Fig. 40 shows that, given a thrust to weight ratio of 2, the composite aircraft was able to reach the limit with power to spare. What's more, for the composite aircraft tested, flightworthiness determination typically required 20-30 spin-ups. Figure 40 shows that for the range of 20 to 30 spin-ups, the composite vehicle still retains ~80% battery charge. These results affirmed and finalized the component selection for the modular vehicles.

In the experimental studies below, results are compared to a 6-degree-of-freedom model of the composite vehicle to further analyze the algorithm's performance. The model parameters of the experimental system, including payload weight, inertia, and thrust-throttle mapping, were measured and provided to the simulation for comparison studies. Several comparison cases were performed to verify accuracy of the simulation model with respect to the physical aircraft. Figure 41 shows one such experiment in which the time histories of the excitation angle, $\alpha(t)$, for the simulated and experimental systems are plotted against each other.




Figure 41. Sample Comparison between Simulated and Experimental Excitation.

Figure 41 shows strong agreement and enables a direct comparison between experimental and simulation results to be described below. It is important to note that the thrust-throttle relationship and the number of attached vehicles are the only pieces of information required to determine flightworthiness for a given composite vehicle configuration.

CHAPTER 7. EXPERIMENTAL RESULTS

7.1 Example Case Results

For each of the three payloads, two example cases are presented, one flightworthy and one unflightworthy, in order to demonstrate the algorithm's performance for differing payload geometries and inertial properties. For the flightworthy cases, the number of attached vehicles was selected to maintain a near-constant thrust to weight ratio for the aircraft. A hover throttle threshold of 75% was set to determine flightworthiness.

The first set of results utilized the square payload and a set of three modular vehicles, shown on the left of Fig. 42. The combined weight of the aircraft was 20.0 N yielding a thrust-to-weight ratio of 1.7 at maximum throttle.




Figure 42. Flightworthy (left) and Unflightworthy (right) Configurations of the Rectangle Aircraft.

Figure 43 shows the excitation history for the first example (flightworthy) case. The axes are labeled in the order that they are observed with the 0 axis representing spin-ups

that excited about a contact point instead of a contact axis. The solid blue line shows the axis about which the composite vehicle excited for each of the spin-ups. The dashed red line shows the axes that the simulated composite vehicle excited about using the same commanded throttle rates, and the thick, dashed black line shows the points in the algorithm where the composite vehicle was commanded to excite about a particular axis (exploit phase).




Figure 43. Excitation Axis History for Flightworthy Rectangular Aircraft.

For the first 11 spin-ups the algorithm was in the exploration phase and commanded random vehicle thrust ratios. From these random thrust ratios, the resulting excitation axis and thrust distributions were used to refine the probabilistic model of the aircraft. For this example case, the algorithm determined that two of the contact axes were parallel and, as a result, required that four total axes be observed before progressing to the exploit phase. After spin-up 11, the algorithm entered the exploitation phase and leveraged the Bayesian estimates to command excitation about the four discovered axes. With the exception of spin-up 7, the simulated composite vehicle directly matched the excitations

of the physical system. For the 7th spin-up the simulated system classified the excitation as occurring about a contact point. The discrepancy occurs due to uncertainty in the thrust throttle relationship and unmodeled flexibility of the composite aircraft. In this case, the algorithm was able to generate the required data in 16 spin-ups.

Figure 44 shows where the probabilistic model placed the vehicles on the hypothetical payload over the course of the 16 spin-ups. Initially, the vehicles are equally likely to be located in any of the locations defined about the hypothetical payload and, thus, originate at the geometric center. As the experiment progresses and excitations about different contact axes are observed, the probabilistic vehicle positions evolve to the final estimate. The solid lines in Fig. 44 show the progression of these estimates and the circles denote the final probabilistic estimates. The estimated positions in Fig. 9 show a strong agreement with the physical system and were successfully utilized to command excitation as desired in the exploit phase. Due to the small number of attached rotors, when excitation occurred the Bayesian updating process was able to easily identify which vehicles were most effective at producing moments about the differing axes.




Figure 44. Probabilistic Vehicle Placement History for Flightworthy Rectangle Aircraft.

For a rectangular payload, all four axes must be observed and excited about in order to calculate the vehicle positions about the mass center of the aircraft. Figure 44 shows the contact axes that were observed by the algorithm. Due to imperfect sensor data and minor payload flexibility, the observed axes were 1.9° , 87° , 179.8° , and 273.3° instead of the expected 0° , 90° , 180° , and 270° . Nevertheless, these errors fell within acceptable bounds and the algorithm was able to determine that the sets of axes were parallel (leading to use of the parallel axes solution scheme in Section 3.3.2. Following execution of the 16 spin-ups, the recorded thrust values and their associated axes of excitation were used to estimate the system weight and the vehicle positions. In this case, the weight of the aircraft was computed with a relative error of 15.22% and 7.4% for the physical and simulated systems, respectively. Figure 45 shows the physical and simulated position estimates of the vehicles scaled by the correct physical dimensions, as well as the actual positions about the mass

center. For this case, both the experimental and simulated position estimates closely match the actual configuration.




Figure 45. Vehicle Position Estimates for Flightworthy Rectangle Aircraft.

To determine flightworthiness, the weight estimate and vehicle position estimates were substituted into Eq. (39) to calculate the throttle levels required to achieve a stable hover. The resulting minimum-norm throttle levels generated from the experimental and simulated aircraft were compared to the true throttle levels required for hover. These results are shown in Fig. 46. For this case, the near-even distribution of vehicles about the mass center results in relatively uniform throttle levels across the attached vehicles. Figure 46 also shows that, due to the accurate weight and vehicle position estimates, the estimated throttle levels closely match the actual values. The RMS error of the experimentally-derived and simulated throttle levels were 9.65% and 4.67% throttle, respectively. With a

predefined cutoff of 75% throttle, the algorithm classified this composite vehicle as flightworthy.




Figure 46. Estimated and Actual Vehicle Hover Throttle Percentages for Flightworthy Rectangle Aircraft.

The second example used the same payload and vehicles as the first example, but reconfigured the vehicles as shown in the right of Fig. 42. Due to the non-negligible weight of the vehicles, this reconfiguration shifted the mass center location of the aircraft. Technically this configuration is flightworthy as, even with the shift, the required throttle levels to achieve hover are 68%, 45%, and 68%, respectively. However, while in ground contact the distribution of vehicles about the contact axes, as well as about the mass center location, play a critical role in determining flightworthiness. If the mass center and/or vehicles are located at disproportionate distances from each contact axis, then the probability of randomly exciting about all of the necessary contact axes decreases

significantly. The spin-up excitations for this example are shown in Fig. 12. For a rectangular payload, the algorithm must observe all four axes before progressing to the exploit phase. To classify an excitation direction as a contact axis, the aircraft must excite about the same contact axis twice (in order to filter out excitations that occur about a contact point). Figure 47 shows that the algorithm was able to excite about all four axes once but was unable to randomly excite about a fourth axis over the course of 25 spin-ups. Note in Fig. 47 that both the physical and simulated aircraft excited about a fifth axis despite the payload only having four contact axes. This axis actually represents excitation about a contact point. Given the maximum exploration spin-up limit of 25 spin-ups, the algorithm ended with a recommendation to add vehicles or to redistribute the three currently attached.




Figure 47. Excitation Axis History for Unflightworthy Rectangular Aircraft.

The third and fourth example cases use the triangular payload. The triangle payload has a greater weight compared to the square payload, and thus additional vehicles were attached for the experimental trials.




Figure 48. Flightworthy (left) and Unflightworthy (right) Configurations of the Trianglar Aircraft.

The third example configuration is shown on the left of Fig. 48 and had a composite vehicle weight of 29.7 N and thrust-to-weight ratio of 1.8 at maximum throttle. Figure 49 shows the excitation history of the experimental and simulated systems for this flightworthy configuration. Note that there were multiple instances of disagreement between the experimental system and the simulated aircraft (spin-ups 5, 8, 10, 18, and 19). For these excitations, the experimental system detected that the excitations were not purely about a contact axes while, in simulation, contact axis excitations were observed. Despite these instances of disagreement, there was never a case where the experimental system and the payload excited about two different contact axes. The algorithm was able to generate

the required excitations in a total of 26 spin-ups, with five of the spin-ups not providing any useful data.




Figure 49. Excitation Axis History for Flightworthy Triangular Aircraft.

Figure 50 shows the evolution of the probabilistic model over the course of the excitation. With a larger number of attached vehicles, the possible thrust distributions that lead to excitation about a given axis significantly increases. The higher dimensionality in thrust distributions can result in the propagation of misleading probabilistic placements of the individual vehicles about the hypothetical payload. Figure 51 shows that in this case, the probabilistic model did not yield vehicle placements close to the actual locations. This is not unexpected as there is no guarantee that the Bayesian likelihood estimates accurately represent the physical system; however, they were successfully used to command excitation in the exploit phase.




Figure 50. Probabilistic Vehicle Placement History for Flightworthy Triangle Aircraft (vehicles originate at origin at the start of the experiment).

The three observed contact axes shown in Fig. 50 for the flightworthy triangular aircraft are linearly independent, and thus Eq. (29) was used to estimate the weight of the composite aircraft and the positions of the modular vehicles. For this example, Eq. (29) resulted in a weight error of 14% and 6% for the experimental and simulated cases, respectively. Figure 51 shows the experimental and simulated position estimates of the modular vehicles. Again, the algorithm showed strong agreement between the estimated and actual positions.




Figure 51. Vehicle Position Estimates for Flightworthy Triangular Aircraft.

Figure 52 shows the estimated hover throttle values for this example case. For this aircraft, the estimated throttle percentages matched those of the actual system with an RMS error of 8.29% and 4.56% throttle for the experimental and simulated systems, respectively. All throttle percentages are below the threshold value and so the composite aircraft is deemed flightworthy.




Figure 52. Estimated and Actual Vehicle Hover Throttle Percentages for Flightworthy Triangle Aircraft.

For the fourth example case, the previous triangular aircraft was altered by removing one of the vehicles. This configuration is shown in the right of Fig. 48. With four attached vehicles, the total thrust to weight ratio for the aircraft decreased to 1.7.




Figure 53. Excitation Axis History for Unflightworthy Triangular Aircraft.

Figure 53 shows the excitation history of the experimental and simulated aircraft. Ten spin-ups were required to observe each of the three axes twice. Even with a single vehicle removed, the aircraft was able to sufficiently vary the equivalent thrust vector to excite about all contact axes. Despite multiple excitations about contact points (spin-ups 3, 7, 20, and 21) the algorithm was able to gather the required axis excitation data. The resulting weight errors were 11% and 9% for the experimental and simulated systems, respectively, with vehicle position estimates shown in Fig. 54. Figure 55 shows the computed required hover throttle levels, where it is evident that Vehicle 4 exceeds the acceptable threshold and the configuration is deemed unflightworthy. As expected, the removal of one of the vehicles had a large skewing effect on the throttle levels required to achieve stable hover, and even though the thrust-to-weight ratio is acceptable the arrangement of the vehicles about the mass center is highly uneven, negatively affecting the distribution of thrust required.




Figure 54. Vehicle Position Estimates for Unflightworthy Triangle Aircraft.




Figure 55. Estimated and Actual Vehicle Hover Throttle Percentages for Unflightworthy Triangle Aircraft.

The final two example cases were performed using the pentagon-shaped payload. The pentagon payload was designed to be the heaviest of the three payloads, and thus all 6 vehicles were used for the flightworthy configuration shown on the left of Fig. 56 (resulting in a thrust to weight ratio of 1.84).




Figure 56. Flightworthy (left) and Unflightworthy (right) Configurations of the Pentagon Aircraft.

The spin-up excitations for the experimental and simulated aircraft are shown in Fig. 57. Although the payload has five contact axes, the algorithm only requires excitations only about three to complete the flightworthiness evaluation. The algorithm explores until three axes have been identified, then progresses to the exploitation phase without checking for the existence of additional axes. Figure 57 shows a significant percentage of excitations about a contact point (approximately 1/3) which provided no useful information to the algorithm. As demonstrated by this case, as the number of contact axes increases it becomes increasingly difficult to command excitation about a particular desired axis. Also, with more contact axes, there is a greater probability that the thrust vector will be located

in a region that leads to excitation about a contact point. Figure 58 illustrates this point using a pentagon payload with a geocentric mass center location. The color-coded regions show the areas in which the equivalent thrust vector has the greatest moment arm with respect to the corresponding color-coded contact axis. As the equivalent thrust vector approaches the boundary between two of the colored regions, there is an increasing probability that the aircraft will excite about a contact point as opposed to one of the contact axes. For an N -sided polygon, as the number of contact axes N grows, the highlighted regions become smaller and there are more boundaries between regions, leading to a higher likelihood of contact point excitation and a smaller area to place the equivalent thrust to achieve excitation about a desired axis.




Figure 57. Excitation Axis History for Flightworthy Pentagon Aircraft.




Figure 58. Equivalent Thrust to Excitation Axis Mapping for Pentagon Payload.

This experiment resulted in weight estimation errors of 39.89% and 2.2% for the experimental and simulated systems, respectively. Figure 59 shows the estimated positions of the attached vehicles about the mass center location of the aircraft. While the estimated positions still clearly correspond to the actual locations, there is larger error with respect to the true position. The increasing weight and position estimate errors can be readily explained by the increasing inertia of the payload and the larger number of attached vehicles, which increase the system's sensitivity to model and parameter uncertainty as shown by Eq. (58) in the previous Error Analysis section.




Figure 59. Vehicle Position Estimates for Flightworthy Pentagon Aircraft.

Figure 60 shows the results of the flightworthiness analysis of Eq. (39). The large overpredictions of throttle levels for Vehicles 2, 4, and 5 can be explained by the sizeable overprediction in the aircraft's weight. Also, as the error in estimated vehicle positions grows, there is larger disagreement between the true required throttle values and the estimated throttle required. The RMS errors for the experimental and simulated systems were 21.6% and 9.13%, respectively. Despite these errors, all vehicle hover throttles were correctly estimated below the 75% throttle limit and, thus, the aircraft was deemed flightworthy.




Figure 60. Estimated and Actual Vehicle Hover Throttle Percentages for Flightworthy Pentagon Aircraft.

The final example case used the payload and vehicle distribution from the previous example, except two vehicles were removed yielding a thrust to weight ratio of 1.53. This configuration is shown on the right side of Fig. 56. The history of axis excitations is shown in Fig. 61. Compared to the previous example case, this experiment observed a different set of three axes to use in the flightworthiness algorithm. This experiment yielded weight estimation errors of 8% and 7% for the experimental and simulated vehicles, respectively. The estimated vehicle positions shown are shown in Fig. 62, with the flightworthiness analysis shown in Fig. 63. Despite the four vehicles being evenly placed about the mass center of the aircraft, the low overall thrust-to-weight ratio of the aircraft leads to higher-than-acceptable throttle percentages. This particular aircraft was classified as unflightworthy with a recommendation to use additional vehicles.




Figure 61. Excitation Axis History for Unflightworthy Pentagon Aircraft.




Figure 62. Vehicle Position Estimates for Unflightworthy Pentagon Aircraft.




Figure 63. Estimated and Actual Vehicle Hover Throttle Percentages for Unflightworthy Pentagon Aircraft.

7.2 Monte Carlo Results

Due to the random nature of the exploration phase and the significant role that the probabilistic model plays in the flightworthiness algorithm, a set of Monte Carlo results are generated to measure the expected performance of the algorithm. The three flightworthy configurations presented in the previous section were chosen for these studies as each has different inertial properties and numbers of attached vehicles. As discussed in the error analysis section, these parameters affect the weight and vehicle position error sensitivity to uncertainty. Table 2 lists the numbers of attached vehicles, the total mass, and the average contact axis inertia for the three aircraft.

Table 2. Aircraft Weight and Inertia Properties.

	Rectangular Payload	Triangular Payload	Pentagon Payload
Attached Vehicles	3	5	6
Total Weight (N)	20.9	30.5	37.4
Average Contact Axis Inertia (kg m ²)	0.20	0.29	0.731

For each aircraft, 15 flightworthiness determination experiments were performed. Each of the 15 experiments used the same aircraft but randomized the throttle distributions used in the exploration phase. The metrics used to quantify the results were the excessive number of spin-ups required to determine flightworthiness and the RMS hover throttle errors. The excessive spin-up metric quantifies the degree to which the machine learning approach was able to generate the required thrust data set, while the RMS throttle error measures the accuracy of the static analysis approach in determining flightworthiness. These metrics are compared to results generated from the simulated model using the same throttle rate commands.

Figures 64-66 and Table 2 show the vehicle position estimates from the series of experiments compared to simulation results. In each figure, the actual vehicle positions are denoted by solid circles and the estimated positions by the +’s. The dashed circles show the constant radius areas which capture 93% (approximately 2 standard deviations) of the estimates. The mean error (accuracy) of the estimates, shown as the distance between the centers of the solid and dashed circles, is a result of the constant bias

introduced by weight estimation error. The distribution, or precision, of the estimates around this mean are a result of the varying throttle rates of the vehicles used in each spin-up experiment and grow with the aircraft inertia about the contact axes and the number of attached vehicles. As seen by the increasing radii of the dashed circles and the increasing distance between the centers of the two circles in Figs. 64-66, there is a deterioration of estimation accuracy and precision as the payload size and number of attached vehicles is increased.




Figure 64. Vehicle Position Estimates for Square Aircraft.

The results in Fig. 64 and Table 3 show that, for the rectangular payload, the algorithm was able to fairly accurately place the vehicles over the course of the 15 experiments. The relatively low inertia of the rectangular payload reduces sensitivity to sources of error (as shown in Eq. (58)), resulting in good position estimates and favorable correlation with simulation results as shown in Table 3. Figure 65 and Table 3 show the position estimation results for the second aircraft. For this configuration, the mean position estimation error is

over twice as large as for the rectangular payload. As discussed in the error analysis section, a larger number of vehicles increases sensitivity to sources of uncertainty which, in turn, results in a larger weight estimation error. From Eq. (58), the weight estimation error acts as a source of constant bias and can be expected to increase the mean positioning error value. Note that the similar precision of the estimates (radii of the dashed lines) between the rectangle and triangle cases in this section can be explained by the proportionate increase in aircraft weight and contact axis inertia. From Eq. (58), increasing aircraft weight reduces sensitivity to uncertainty, while increasing contact axis inertia increases it. From the rectangular aircraft to the triangular aircraft both these quantities increase by approximately 50%, effectively negating the effects of one another in the context of vehicle position estimation error.




Figure 65. Vehicle Position Estimates for Triangle Aircraft.




Figure 66. Vehicle Position Estimates for Pentagon Aircraft.

As expected, the pentagon aircraft, which has the largest inertia and number of attached vehicles, exhibited the poorest estimation accuracy. The results for this case are shown in Fig. 66 and Table 3. Although the weight increase from the triangle aircraft to the pentagon aircraft was approximately 20%, the average contact axis inertia increased by 150%. This increase, coupled with an additional lift vehicle, resulted in reduced accuracy of both precision and accuracy of the vehicle position estimates.

For all aircraft, the vehicle position estimates, along with the aircraft weight estimates, were used to compute the hover throttle RMS error shown in Table 3. As expected, the RMS throttle error of the pentagon vehicle suffers as a result of the poor weight and vehicle position estimates. Compared to the experimental systems, the simulated aircraft had perfect knowledge of the vehicle thrust and was able to mitigate a substantial portion of

the weight and hover throttle error. Overall, Table 3 shows the general trends and expected consequences of increasing the inertia of the payload and adding additional lift vehicles. Table 3 shows that, for all aircraft, the algorithm was able to determine flightworthiness using on average five or less excessive spin-ups. Once again, this highlights the utility of the probabilistic model in driving contact axis excitations.

Table 3. Flightworthiness Determination Trade Study Performance Metrics.

	Rectangular Payload (exp / sim)	Triangular Payload (exp / sim)	Pentagon Payload (exp / sim)
Weight Error (%)	9.3 / 9.6	14.3 / 7.9	28.7 / 9.8
RMS Throttle Error (%)	9.9 / 6.5	11.8 / 7.4	16.7 / 10.6
Position Estimate Mean Error (m)	0.06 / 0.01	0.12 / 0.07	0.20 / 0.13
Position Estimate 2 σ Radius (m)	0.11 / 0.06	0.13 / 0.09	0.30 / 0.33
Excessive Spin-ups	4.1 / -	3.7 / -	5.3 / -

CHAPTER 8. CONCLUSION

This work introduced an algorithm to determine flightworthiness of a composite aircraft with unknown inertial properties and attached vehicle locations. The flightworthiness algorithm first attempts to discover and excite about discrete contact axes of the composite aircraft through a machine learning component that provides axis effectiveness scores for each attached vehicle. Next, the algorithm utilizes the thrust data collected during the excitations to estimate the weight of the aircraft and the positions of the attached vehicles. Using these estimates, the algorithm determines whether or not the current vehicle configuration is flightworthy. The algorithm was evaluated using a dynamic simulation of the composite aircraft. The simulations incorporated a ground contact model to accurately capture the dynamic motion of the aircraft throughout the flightworthiness determination algorithm. Through Monte Carlo simulations, the algorithm's performance was shown to depend on the number of successful excitations to mitigate thrust errors, a sufficient thrust-to-weight ratio to avoid vehicle saturation, and a reasonably centralized mass center location to enable excitation about the required number of contact axes. With a sufficiently high thrust-to-weight ratio and mass center location within certain bounds, the algorithm was shown to accurately estimate the weight and vehicle positions in a reasonable number of spin-ups such that the algorithm can execute in an efficient manner. A practical implementation of the algorithm is demonstrated using a set of constructed payloads and prototype vertical lift vehicles. The performance of the algorithm was demonstrated for a range of flightworthy and unflightworthy composite aircraft. The differing sources of estimation error are analyzed mathematically and explored through a series of

flightworthiness experiments. Overall, the algorithm is shown to be a useful tool in determining flightworthiness of a composite aircraft and in providing an initial estimate of aircraft weight and vehicle locations to be used in subsequent flight control strategies.

Currently, the output of the algorithm is a simple yes or no response regarding the flightworthiness of the composite aircraft. Future work could augment the current algorithm to directly recommend a direction to add or rearrange vertical lift vehicles in scenarios where the composite aircraft was found to be unflightworthy. Additionally, future work could investigate a method to ensure that all vehicles are operating as expected and to detect if a vehicle has failed. This could potentially be accomplished by analyzing the vehicle position estimates and the associated error magnitudes in the least-squares calculation. The algorithm, as described, was designed to be ran on any composite aircraft that adheres to the listed assumptions. Future work could consider a specific class of payloads, constrained in scale and/or in geometry, and analyze the corresponding performance of the flightworthiness algorithm. By enforcing constraints on potential payloads, additional studies could be conducted to explicitly define the performance bounds of the algorithm as a function of the various user-defined parameters. Finally, future work could integrate this flightworthiness algorithm with an on-board flight controller. The output of the algorithm, the flightworthiness determination and the associated aircraft weight and vehicle position estimates, could be directly fed to the flight controller as initial parametric estimates to be refined in-air. This way, full autonomy would be achieved through takeoff and hover.

APPENDIX A. AIRCRAFT WEIGHT ESTIMATION

In order to use Eq. (29) and obtain a least squares estimate of the vehicle placements about the mass center of the aircraft, the weight of the aircraft must be calculated. From the Vehicle Placement Algorithm section, this is accomplished by explicitly writing out the squared error term of the estimate as a function of the aircraft weight. This way, the weight of the aircraft can be solved for as the value that minimizes the squared error term (Eqns. (32-36)). In computing the error squared term, the terms on the righthand side of Eq. (33) can be expressed by

$$A^T A = \begin{bmatrix} s^2(\theta_{k_2})(T_{k_1 k_1} W^2 - 2T_{k_1} W + n) & -s(\theta_{k_2})s(\theta_{k_1})(T_{k_1 k_2} W^2 - (T_{k_1} + T_{k_2})W + n) \\ -s(\theta_{k_2})s(\theta_{k_1})(T_{k_1 k_2} W^2 - (T_{k_1} + T_{k_2})W + n) & s^2(\theta_{k_1})(T_{k_2 k_2} W^2 - 2T_{k_2} W + n) \end{bmatrix} \quad (A1)$$

$$A^T \vec{b} = \begin{Bmatrix} (c(\theta_{k_1})s^2(\theta_{k_2}) - c(\theta_{k_2})s(\theta_{k_2})s(\theta_{k_1})) (T_{ik_1} W^2 - (T_i + T_{k_1})W + n) \\ (c(\theta_{k_2})s^2(\theta_{k_1}) - c(\theta_{k_1})s(\theta_{k_2})s(\theta_{k_1})) (T_{ik_2} W^2 - (T_i + T_{k_2})W + n) \end{Bmatrix} \quad (A2)$$

$$\vec{b}^T \vec{b} = (c(\theta_{k_2})s(\theta_{k_1}) - c(\theta_{k_1})s(\theta_{k_2}))^2 (T_{ii} W^2 - 2T_i W + n) \quad (A3)$$

where

$$T_a = \sum_{i=1}^{p_a} \sum_{j=1}^n \begin{bmatrix} T_{11,a} & \cdots & T_{1n,a} \\ \vdots & \ddots & \vdots \\ T_{p_a 1,a} & \cdots & T_{p_a n,a} \end{bmatrix}^+ \quad (A4)$$

and

$$\mathbb{T}_{ab} = \sum_{i=1}^n \sum_{j=1}^n \left\{ \left(\begin{bmatrix} T_{11,a} & \cdots & T_{1n,a} \\ \vdots & \ddots & \vdots \\ T_{p_a 1,a} & \cdots & T_{p_a n,a} \end{bmatrix}^+ \right)^T \begin{bmatrix} T_{11,b} & \cdots & T_{1n,b} \\ \vdots & \ddots & \vdots \\ T_{p_b 1,b} & \cdots & T_{p_b n,b} \end{bmatrix}^+ \right\} \quad (\text{A5})$$

are scalar values containing the summation of thrusts used to excite about axes a and b . By plugging these expressions into Eq. (33) and using a symbolic solver, the error squares term as a function of W is expressed by Eq. (34). The complex functions, $c_1 - c_6$, are defined as

$$c_1 = -\mathbb{T}_{ii} \mathbb{T}_{k_1 k_1} \mathbb{T}_{k_2 k_2}^2 + \mathbb{T}_{ii}^2 \mathbb{T}_{k_1 k_2}^2 + \mathbb{T}_{ik_1}^2 \mathbb{T}_{k_2 k_2} - 2\mathbb{T}_{ik_1} \mathbb{T}_{ik_2} \mathbb{T}_{k_1 k_2} + \mathbb{T}_{ik_2}^2 \mathbb{T}_{k_1 k_1} \quad (\text{A6})$$

$$\begin{aligned} c_2 = & 2\mathbb{T}_i \mathbb{T}_{ik_1} \mathbb{T}_{k_1 k_2} - 2\mathbb{T}_i \mathbb{T}_{ik_1} \mathbb{T}_{k_2 k_2} - 2\mathbb{T}_i \mathbb{T}_{ik_2} \mathbb{T}_{k_1 k_1} + 2\mathbb{T}_i \mathbb{T}_{ik_2} \mathbb{T}_{k_1 k_2} + 2\mathbb{T}_i \mathbb{T}_{k_1 k_1} \mathbb{T}_{k_2 k_2} - 2\mathbb{T}_i \mathbb{T}_{k_1 k_2}^2 \dots \\ & \dots - 2\mathbb{T}_{ii} \mathbb{T}_{k_1} \mathbb{T}_{k_1 k_2} + 2\mathbb{T}_{ii} \mathbb{T}_{k_1} \mathbb{T}_{k_2 k_2} + 2\mathbb{T}_{ii} \mathbb{T}_{k_1 k_1} \mathbb{T}_{k_2} - 2\mathbb{T}_{ii} \mathbb{T}_{k_1 k_2} \mathbb{T}_{k_2} - 2\mathbb{T}_{ik_1}^2 \mathbb{T}_{k_2} + 2\mathbb{T}_{ik_1} \mathbb{T}_{ik_2} \mathbb{T}_{k_1} \dots \\ & \dots + 2\mathbb{T}_{ik_1} \mathbb{T}_{ik_2} \mathbb{T}_{k_2} - 2\mathbb{T}_{ik_1} \mathbb{T}_{k_1} \mathbb{T}_{k_2 k_2} + 2\mathbb{T}_{i2} \mathbb{T}_{k_1 k_2} \mathbb{T}_{k_2} - 2\mathbb{T}_{ik_2}^2 \mathbb{T}_{k_1} + 2\mathbb{T}_{ik_2} \mathbb{T}_{k_1} \mathbb{T}_{k_1 k_2} - 2\mathbb{T}_{ik_2} \mathbb{T}_{k_1 k_1} \mathbb{T}_{k_2} \end{aligned} \quad (\text{A7})$$

$$\begin{aligned} c_3 = & \mathbb{T}_i^2 \mathbb{T}_{k_1 k_1} - 2\mathbb{T}_i^2 \mathbb{T}_{k_1 k_2} + \mathbb{T}_i^2 \mathbb{T}_{k_2 k_2} - 2\mathbb{T}_i \mathbb{T}_{ik_1} \mathbb{T}_{k_1} + 2\mathbb{T}_i \mathbb{T}_{ik_1} \mathbb{T}_{k_2} + 2\mathbb{T}_i \mathbb{T}_{ik_2} \mathbb{T}_{k_1} \dots \\ & \dots - 2\mathbb{T}_i \mathbb{T}_{k_1} \mathbb{T}_{k_2 k_2} - 2\mathbb{T}_i \mathbb{T}_{k_1 k_1} \mathbb{T}_{k_2} + 2\mathbb{T}_i \mathbb{T}_{k_1 k_2} \mathbb{T}_{k_2} + \mathbb{T}_{ii} \mathbb{T}_{k_1}^2 - 2\mathbb{T}_{ii} \mathbb{T}_{k_1} \mathbb{T}_{k_2} - \mathbb{T}_{ii} \mathbb{T}_{k_1 k_1} n + 2\mathbb{T}_{ii} \mathbb{T}_{k_1 k_2} n \dots \\ & \dots + 2\mathbb{T}_i \mathbb{T}_{k_1} \mathbb{T}_{k_1 k_2} + \mathbb{T}_{ii} \mathbb{T}_{k_2}^2 - \mathbb{T}_{ii} \mathbb{T}_{k_2 k_2} n + \mathbb{T}_{ik_1}^2 n - 2\mathbb{T}_{ik_1} \mathbb{T}_{ik_2} n + 2\mathbb{T}_{ik_1} \mathbb{T}_{k_1} \mathbb{T}_{k_2} - 2\mathbb{T}_{ik_1} \mathbb{T}_{k_1 k_2} n \dots \\ & \dots - 2\mathbb{T}_i \mathbb{T}_{ik_2} \mathbb{T}_{k_2} - 2\mathbb{T}_{ik_1} \mathbb{T}_{k_2}^2 + 2\mathbb{T}_{ik_1} \mathbb{T}_{k_2 k_2} n + \mathbb{T}_{ik_2}^2 n - 2\mathbb{T}_{ik_2} \mathbb{T}_{k_1}^2 + 2\mathbb{T}_{ik_2} \mathbb{T}_{k_1} \mathbb{T}_{k_2} + 2\mathbb{T}_{ik_2} \mathbb{T}_{k_1 k_1} n \dots \\ & \dots - 2\mathbb{T}_{ik_2} \mathbb{T}_{k_1 k_2} n + \mathbb{T}_{k_1}^2 \mathbb{T}_{k_2 k_2} - 2\mathbb{T}_{k_1} \mathbb{T}_{k_1 k_2} \mathbb{T}_{k_2} + \mathbb{T}_{k_1 k_1} \mathbb{T}_{k_2}^2 - \mathbb{T}_{k_1 k_1} \mathbb{T}_{k_2 k_2} n + \mathbb{T}_{k_1 k_2}^2 n \end{aligned} \quad (\text{A8})$$

$$c_4 = -\mathbb{T}_{k_1 k_1} \mathbb{T}_{k_2 k_2} + \mathbb{T}_{k_1 k_2}^2 \quad (\text{A9})$$

$$c_5 = (-2\mathbb{T}_{k_1} - 2\mathbb{T}_{k_2}) \mathbb{T}_{k_1 k_2} + 2\mathbb{T}_{k_2 k_2} \mathbb{T}_{k_1} + 2\mathbb{T}_{k_2} \mathbb{T}_{k_1 k_1} \quad (\text{A10})$$

$$c_6 = 2\mathbb{T}_{k_2 k_2} n + \mathbb{T}_{k_1}^2 - 2\mathbb{T}_{k_1} \mathbb{T}_{k_2} + \mathbb{T}_{k_2}^2 - n(\mathbb{T}_{k_2 k_2} + \mathbb{T}_{k_1 k_1}) \quad (\text{A11})$$

and are used directly in Eq. (36) to compute the weight of the composite aircraft.

REFERENCES

- [1] Pounds, P., Bersak, D., and Dollar, A., "Stability of small-scale UAV helicopters and quadrotors with added payload mass under PID control," *Autonomous Robots*, Vol. 33, No. 1-2, Aug. 2012, pp. 129-142.
- [2] Mellinger, D., Michael, N., Shomin, M., and Kumar, V., "Recent Advances in Quadrotor Capabilities," In *Robotics and Automation (ICRA)*, 2011 IEEE International Conference on, IEEE., 2011 pp. 2964-2965.
- [3] Baabaie, R., and Ehyae, F., "Robust Optimal Motion Planning Approach to Cooperative Grasping and Transporting using Multiple UAVs based on SDRE", *Transaction of the Institute of Measurement and Control*, Vol. 39, No. 9, 2016, pp. 1391-1408.
- [4] Gawel, A., Kamel, M., Novkovic, T., Widauer, J., Schindler, D., von Altishofen, B.P., Siegwart, R. and Nieto, J., "Aerial picking and delivery of magnetic objects with mavs.", In *Robotics and Automation (ICRA)*, 2017 IEEE International Conference on, pp. 5746-5752.
- [5] Papayanopoulos, J. F., "Autonomous UAV Precision Item Pickup." Master's Thesis., Georgia Institute of Technology, 2017.

- [6] Bähnemann, R., Schindler, D., Kamel, M., Siegwart, R., and Nieto, J., “A decentralized multi-agent unmanned aerial system to search, pick up, and relocate objects”, IEEE International Symposium on Safety, Security and Rescue Robotics (SSRR) 2017 p. 8088150.

- [7] Al-Kaff, A., Moreno, F.M., de la Escalera, A. and Armingol, J.M., 2017, “Intelligent vehicle for search, rescue and transportation purposes.” In *Safety, Security and Rescue Robotics (SSRR), 2017 IEEE International Symposium on*, 2017 pp. 110-115.

- [8] Bernard, M., Kondak, K., Maza, I. and Ollero, A., “Autonomous transportation and deployment with aerial robots for search and rescue missions.” *Journal of Field Robotics* 28, No. 6, 2011, pp. 914-931.

- [9] Beebe, M. K., and Ret., C. R., "Unmanned Aircraft Systems for Casualty Evacuation-What Needs to be Done." In *Proceedings of the NATO STO-MP-HFM-231 Symposium, Beyond Time and Space, Orlando, FL*. 2013.

- [10] Durham, W.C., “Constrained control allocation.”, *Journal of Guidance, control, and Dynamics*, 16, no. 4, 1993, pp.717-725.

- [11] Palunko, I., Cruz, P. and Fierro, R., “Agile load transportation: Safe and efficient load manipulation with aerial robots.”, *IEEE robotics & automation magazine* 19, No. 3, 2012, pp.69-79.

- [12] Maza, I., Kondak, K., Bernard, M. and Ollero, A., “Multi-UAV cooperation and control for load transportation and deployment.” In *Selected papers from the 2nd International Symposium on UAVs, Reno, Nevada, USA June 8–10, 2009*, pp. 417-449.
- [13] Gimenez, J., Gandolfo, D. C., Salinas, L. R., Rosales, C., and Carelli, R., “Multi-objective control for cooperative payload transport with rotorcraft UAVs”, *ISA transactions*, 2018.
- [14] Pereira, P. O., and Dimarogonas, D. V., “Collaborative transportation of a bar by two aerial vehicles with attitude inner loop and experimental validation”, *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, Melbourne, VIC, 2017, pp. 1815-1820.
- [15] Tan, Y., Lai, S., Wang, K., and Chen, B., “Cooperative Control of Multiple Unmanned Aerial Systems for Heavy Duty Carrying”, *Annual Reviews in Control*, 2018.
- [16] Six, D., Chriette, A., Briot, S., and Martinet, P., “Dynamic Modeling and Trajectory Tracking Controller of a Novel Flying Parallel Robot”, *IFAC-PapersOnLine*, Vol. 50, No. 1, 2017, pp. 2241-2246.
- [17] Alonso-Mora, J., Baker, S., and Rus, D., “Multi-robot formation control and object transport in dynamic environments via constrained optimization”, *The International Journal of Robotics Research* 36, no. 9, 2017, pp. 1000-1021.

- [18] Estévez, J., Lopez-Guede, J., and Grana, M., “Particle Swarm Optimization Quadrotor Control for Cooperative Aerial Transportation of Deformable Linear Objects”, *Cybernetics and Systems*, 2016, Vol. 47, pp. 4-16.
- [19] Wang, Z., and Schwager, M., “Multi-robot Manipulation without communication using only local measurements”, *2015 54th IEEE Conference on Decision and Control (CDC)*, Osaka, 2015, pp. 380-385.
- [20] Gassner, M., Cieslewski, T., and Scaramuzza, D., “Dynamic collaboration without communication: Vision-based cable-suspended load transport with two quadrotors”, *2017 IEEE International Conference on Robotics and Automation (ICRA)*, Singapore, 2017, pp. 5196-5202.
- [21] Lindsey Q., Mellinger, D., and Kumar V., “Construction with Quadrotor Teams,” *Autonomous Robots*, published Oct. 2012; Vol. 33, No. 3, pp. 323-336.
- [22] Duffy, M.J., Samaritano, T.C., “The LIFT! Project – Modular, Electric Vertical Lift System with Ground Power Tether,” *Applied Aerodynamics Conference*, AIAA, Dallas, TX, June 2015.
- [23] Oung, R., Bourgault, F., Donovan, M., D’Andrea, R., “The Distributed Flight Array,” *2010 IEEE International Conference on Robotics and Automation*, Anchorage, Alaska, 3-8 May, 2010.

- [24] Michael, N., Fink, J., and Kumar, V., "Cooperative Manipulation and Transportation with Aerial Robots," *Autonomous Robots*, 30(1), 2011, pp. 73-86.
- [25] Montemayor, G. and Wen, J.T., "Decentralized Collaborative Load Transport by Multiple Robots," In *Robotics and Automation, 2005. Proceedings of the 2005 IEEE International Conference on*, IEEE, 2005, pp. 372-377.
- [26] Dai, S., Lee, T. and Bernstein, D.S., "Adaptive Control of a Quadrotor UAV Transporting a Cable-Suspended Load with Unknown Mass," In *Decision and Control (CDC), 2014 IEEE 53rd Annual Conference on*, IEEE., 2014, pp. 6149-6154.
- [27] Sreenath, K., Michael, N., and Kumar, V., "Dynamics, Control and Planning for Cooperative Manipulation of Payloads Suspended by Cables from Multiple Quadrotor Robots," In *Robotics: Science and Systems*, 2013.
- [28] Mellinger, D., Shomin, M., Michael, N., and Kumar, V., "Cooperative Grasping and Transport using Multiple Quadrotors," *Distributed Autonomous Robotic Systems*, Lausanne, Switzerland, 2010, Vol. 83, pp. 545- 558.
- [29] Mellinger, D., Lindsey Q., Shomin M., and Kumar V., "Design, Modeling, Estimation and Control for Aerial Grasping and Manipulation," *Int. Conf. on Intelligent Robots and Systems*, IEEE, 2011, San Francisco, CA, USA, pp. 2668-2673.

- [30] Abas, N., Legowo, A., and Akmeliawati, R., "Parameter Identification of an Autonomous Quadrotor," *Int. Conf. on Mechatronics*, IEEE, Kuala Lumpur, Malaysia, 2011, pp. 1-8.
- [31] Burri, M., Nikolic, J., Oleynikova, H., Achtelik, M.W. and Siegwart, R., "Maximum likelihood parameter identification for MAVs.", In *Robotics and Automation (ICRA), 2016 IEEE International Conference on*, pp. 4297-4303.
- [32] Min, B.C., Hong, J.H., and Matson E.T., "Adaptive Robust Control (ARC) for an Altitude Control of a Quadrotor Type UAV Carrying an Unknown Payloads," *Int. Conf. on Control, Automation and Systems*, IEEE, Gyeonggi-do, South Korea, 2011, pp. 1147-1151,
- [33] Achtelik, M., Bierling, T., Wang, J., Hocht, L., and Holzapfel, F., "Adaptive Control of a Quadcopter in the Presence of large/complete Parameter Uncertainties," *AIAA Infotech@Aerospace* 2011, March 2011.
- [34] Rashid, M.I. and Akhtar, S., Adaptive Control of a Quadrotor with Unknown Model Parameters. In *Applied Sciences and Technology (IBCAST), 2012 9th International Bhurban Conference on*, IEEE., 2012, pp. 8-14.
- [35] Xu, R. and Ozguner, U., Sliding Mode Control of a Quadrotor Helicopter. In *Decision and Control, 2006 45th IEEE Conference on*. IEEE., 2006, pp. 4957-4962.

- [36] Lee, D., Kim, H.J., and Sastry, S., "Feedback Linearization vs. Adaptive Sliding Mode Control for a Quadrotor Helicopter," *International Journal of control, Automation and systems* 7, No.3, 2009, pp. 419-428.
- [37] Tagliabue, A., Kamel, M., Nieto, J., and Siegwart, R., "Robust Collaborative Object Transportation Using Multiple MAVs," arXiv:1711.08753, 2017.
- [38] Anderson, Brady. "Adaptive Control for Modular Vertical Lift Air Vehicles." Master's Thesis., Georgia Institute of Technology, 2017.
- [39] Lee, H.I., Lee, B.Y., Yoo, D.W., Moon, G.H. and Tahk, M.J., "Dynamics Modeling and Robust controller design of the Multi-UAV transportation system.", In *29th Congress of the International Council of the Aeronautical Sciences.*, 2014.
- [40] Cabecinhas, D., Naldi, R., Marconi, L., Silvestre, C. and Cunha, R., "Robust take-off and landing for a quadrotor vehicle.", In *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, 2010, pp. 1630-1635.
- [41] Roberto, N., Marconi, L., and Gentili, L., "Robust Takeoff and Landing for a Class of Aerial Robots", In *Decision and Control Proceedings of the 48th IEEE Conference on*, 2009, pp. 3436-3441.

- [42] Bouabdallah, S. and Siegwart, R.Y., "Full control of a quadrotor," *International Conference on Intelligent Robots and Systems*, 2007, pp. 153-158.
- [43] Bernard, D.D.C., Giurato, M., Riccardi, F. and Lovera, M., "Ground Effect Analysis for a Quadrotor Platform," In *Advances in Aerospace Guidance, Navigation and Control*, 2018, pp. 351-367.
- [44] Cheeseman, I.C., Bennett, W.E., "The Effect of Ground on a Helicopter Rotor in Forward Flight," Technical Report, Aeronautical Research Council, 1957.
- [45] Kraus, P., Kumar, V., and Dupont, P., "Analysis of Frictional Contact Models for Dynamic Simulation," *Robotics and Automation, 1998. Proceedings. 1998 IEEE International Conference on*. Vol. 2. IEEE, 1998, pp. 976-981.
- [46] Goyal, S., Pinson E.N., and Sinden F.W., "Simulation of Dynamics of Interacting Rigid Bodies Including Friction I: General problem and contact model," *Engineering with computers* 10, No. 3 ,1994, pp. 162-174.
- [47] Goyal, S., Pinson, E.N., and Sinden, F.W., "Simulation of Dynamics of Interacting Rigid Bodies Including Friction II: Software System Design and Implementation," *Engineering with computers* 10, No. 3, 1994, pp. 175-195.

- [48] Trinkle, J., Pang, J-S., Sudarsky, S., and Lo, G., "On Dynamic Multi-Rigid-Body Contact Problems with Coulomb Friction," *ZAMM-Journal of Applied Mathematics and Mechanics/Zeitschrift für Angewandte Mathematik und Mechanik* 77, no. 4, 1997, pp. 267-279.

- [49] Kojima, M., Megiddo, N., Noma, T., & Yoshise, A., *A Unified Approach to Interior Point Algorithms for Linear Complementarity Problems*, Vol. 538, Springer Science & Business Media, 1991.

- [50] Mirtich, B., and Canny, J., "Impulse-Based Simulation of Rigid Bodies," In *Proceedings of the 1995 symposium on Interactive 3D graphics*, New York, NY, 1995, pp. 181-ff.

- [51] Mirtich, B., and Canny, J., *Impulse-Based Dynamic Simulation*. California: Computer Science Division (EECS), University of California, 1994, pp. 407-418, <https://people.eecs.berkeley.edu/~jfc/papers/94/ibds94.pdf>.

- [52] Bender, J., and Schmitt, A., "Constraint-Based Collision and Contact Handling using Impulses," *Proceedings of the 19th international conference on computer animation and social agents*. 2006, pp. 3-11.

- [53] Poisson, S.D., *Mechanics*, Longmans, London, 1817.

- [54] Ginsberg, J., “Relative Motion,” *Engineering Dynamics*, Cambridge: Cambridge University Press, 2007, pp. 106-112

- [55] Hestenes, M.R., *Calculus of variations and optimal control theory*. Wiley, 1996.

- [56] Kyburg, H.E., “Bayesian and non-Bayesian Evidential Updating,” *Artif. Intell.*, Vol. 31, 1987, pp. 271-293.

- [57] Jensen, Finn V. *An Introduction to Bayesian Networks*. Vol. 210. London: UCL press, 1996.

- [58] Bernardo, J.M. and Smith, A.F., *Bayesian Theory*. Vol. 405. John Wiley & Sons, 2009.

- [59] Newton, M.A., and Raftery, A.E., “Approximate Bayesian Inference with the Weighted Likelihood Bootstrap,” *J. Roy. Statistical Soc. (series B)*, 1994, Vol. 56, pp. 3-26, <http://www.jstor.org/stable/2346025>.

- [60] Davis, E., and Pounds, P., “Direct Sensing of Thrust and Velocity for a Quadrotor Rotor Array”, *IEEE Robotics and Automation Letters*, vol. 2, no. 3, 2017, pp. 1360-1366.

- [61] Demmel, J., “Linear Equation Solving,” *Applied Numerical Linear Algebra*, 1st ed., SIAM, Philadelphia, PA, 1997, pp. 31–63.