

**LOAD ALLOCATION FOR OPTIMAL RISK MANAGEMENT IN
SYSTEMS WITH INCIPIENT FAILURE MODES**

A Ph.D Dissertation
Presented to
The Academic Faculty

by

Brian M. Bole

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy in the
Electrical and Computer Engineering

Georgia Institute of Technology
December 2013

Copyright © 2013 by Brian M. Bole

LOAD ALLOCATION FOR OPTIMAL RISK MANAGEMENT IN SYSTEMS WITH INCIPIENT FAILURE MODES

Approved by:

Dr David G. Taylor
School of Electrical and Computer
Engineering
Georgia Institute of Technology

Dr George Vachtsevanos, Advisor
School of Electrical and Computer
Engineering
Georgia Institute of Technology

Dr Thomas E Michaels
Dept. of ECE
Georgia Institute of Technology

Dr Thomas G Habetler
Dept. of ECE
Georgia Institute of Technology

Dr Kai Goebel
Director of NASA Ames Prognostics
Center of Excellence
NASA Ames Research Center

Date Approved: 11/14/2013

Work hard, dream big, and fail often

ACKNOWLEDGEMENTS

This work has been supported by multiple grants and collaborations with public and private organizations. The graduate funding and guidance provided by Dr. George Vachtsevanos and the Intelligent Control Systems Laboratory (ISCL) at Georgia Tech is gratefully acknowledged. The funding and project support provided by the United States National Aeronautics and Space Administration (NASA) under the Graduate Student Research Program (GSRP) was also a huge boon to this research effort. Thank you to Dr. Kai Goebel and the Prognostics Center of Excellence at NASA Ames for the support and guidance provided during my two years as a NASA GSRP fellow. The project support by AvSafe/SSAT and OCT/ACLO are also respectfully acknowledged.

TABLE OF CONTENTS

ACKNOWLEDGEMENTS	iv
LIST OF TABLES	ix
LIST OF FIGURES	x
SUMMARY	xiii
CHAPTER I: INTRODUCTION	1
1.1 A Motivating Discussion of the Prognostics-Based Decision Making Challenge	3
1.2 Identified Contributions of the Research Effort	6
1.2.1 Claim #1 - Description of Prognostics-Based Risk Management Using a Retrofit Nominal Load Modification Architecture	6
1.2.2 Claim #2 - Analysis of Supervisory Health Management Using Finite Horizon Prognostic Estimates	6
1.2.3 Claim #3 - Demonstration of the Explicit Incorporation of Future Demand Models into Health Management Policies	7
1.2.4 Claim #4 - Real-World Demonstration of Stochastic Component Health Depletion Modeling in terms of Uncertain Future Demand Estimates	7
1.3 Organization of the Dissertation	8
CHAPTER II: BACKGROUND	9
2.1 Origin and History of the Problem	9
2.1.1 Fault Risk Quantification Using Historical Data	9
2.1.2 Condition Based Maintenance and Control	10
2.1.3 Prognostics-Based Health Management	11
2.2 Uncertainty Representation in Prognostics	12
2.3 Risk Quantification	14
2.4 Component Load Allocation for Risk Management	16

CHAPTER III: A MODULAR METHODOLOGY FOR FAULT RISK ASSESSMENT AND MANAGEMENT	17
3.1 Introduction	17
3.2 A Component Load Allocation Control Structure	17
3.2.1 The Nominal Control Loop	17
3.2.2 Component Load Allocation	18
3.2.3 Low Level Component Load Allocation	21
3.3 Stochastic Modeling of Fault Growth in terms of Component Load and Net Output Control Effort Allocations	22
3.3.1 Performance Constraints	25
3.3.2 Prognostic Constraints	26
3.3.3 Verifying the Feasibility of the Constrained Load Allocation Problem	28
3.3.4 Strategic Definition of Performance Objectives and Risk Aversion Metrics	29
CHAPTER IV: AN EXPLORATION OF LOAD ALLOCATION AMONG SIMILAR DEGRADING COMPONENTS	31
4.1 Introduction	31
4.2 A Fault Growth Model	32
4.3 A Thermal Model	34
4.4 Application Example - A Triplex Redundant Electromechanical Actuator . .	35
4.4.1 Mechanical Dynamics Model	36
4.4.2 Performance and Prognostics Based Load Allocation Constraints . .	37
4.4.3 Quantification of Performance and Risk Metrics	39
4.4.4 Simulation Results	40
4.5 Application Example - A Skid-Steered Vehicle	49
4.5.1 Mechanical Dynamics Model	49
4.5.2 Performance and Prognostics Based Load Allocation Constraints . .	51
4.5.3 Quantification of Performance and Risk Metrics	53
4.5.4 Simulation Results	55

CHAPTER V: A GENERALIZED MDP FORMULATION OF PROGNOSTIC MODELING AND RISK BASED DECISION MAKING	64
5.1 Introduction	64
5.2 Markov Process Modeling of Fault Growth Dynamics	65
5.2.1 Incorporating a Stochastic Process Model for Environmental Loading Demand	67
5.3 Encoding Risk Aversion into a Markov Decision Process	68
5.3.1 Identification of Optimal Finite Horizon Control Policies with Dynamic Programming	69
5.4 Consideration of a Multi-Variable Stochastic System Example	70
5.4.1 Simulation Results	72
5.5 Summarizing Remarks	76
CHAPTER VI: SIL/HIL REPLICATION OF POWERTRAIN LOAD DYNAMICS AND ONLINE PROGNOSTICS FOR A UAV APPLICATION	77
6.1 Introduction	77
6.2 A General Software-in-the-Loop / Hardware-in-the-Loop Framework	79
6.2.1 Description of Inner-Loop Control Dynamics	80
6.2.2 Description of Outer-loop Failure Prognostics and Decision Making	84
6.3 Application Example - Battery Depletion Modeling on an Edge 540T UAV	84
6.3.1 Edge 540T Inner-Loop Controls	86
6.3.2 Battery Demand Modeling	87
6.3.3 An Equivalent Circuit Battery Model	92
6.3.4 Battery Charge Depletion Modeling	95
6.3.5 Prognostic Prediction	99
6.3.6 Online Controls Testing	104
6.3.7 Bounded Load Allocation	108
CHAPTER VII: CONCLUSIONS	110

APPENDIX A: RESOLVING ACTIVE REDUNDANCIES IN THE LOAD AL- LOCATION PROBLEM	113
APPENDIX B: THERMAL MODELING FOR DC MOTOR WINDING INSU- LATION	115
APPENDIX C: TRIPLEX REDUNDANT ELECTROMECHANICAL ACTU- ATOR MODELING INFORMATION	117
APPENDIX D: SKID-STEERED VEHICLE MODELING INFORMATION . .	119
APPENDIX E: DETAILED DESCRIPTION OF AN EDGE 540T SAMPLE FLIGHT	122
REFERENCES	125
VITA	134

LIST OF TABLES

Tables 4.1:	Parameter definitions for the EMA motor winding temperature dynamics model	41
Tables 4.2:	Parameter definitions for the skid-steered vehicle model	52
Tables 4.3:	Parameter definitions for the skid-steered vehicle motor winding temperature dynamics model	55
Tables 4.4:	Outcome evaluation costs and 98% confidence VaR estimates for motor state of health at end of mission, using three sample control routines	63
Tables 5.1:	Statistics of the accumulated state transition costs computed over repeated simulations of four sample control policies	75
Tables 6.1:	Symbol definitions for the software-in-the-loop / hardware-in-the-loop testing framework	80
Tables 6.2:	Parameter definitions for the equivalent circuit battery model	95
Tables C.1:	Parameter definitions for the electromechanical actuator model	118
Tables D.1:	Parameter definitions for the skid-steered vehicle model	121

LIST OF FIGURES

Figure 1.1:	A block diagram representation for outer-loop prognostics and risk-based decision making	2
Figure 1.2:	Measured (top) and predicted (bottom) distributions for the system failure time corresponding to two potential control policies	5
Figure 3.1:	Block diagram showing the assumed structure of nominally controlled systems	18
Figure 3.2:	A functional breakdown of control tasks in the proposed component load allocation framework	19
Figure 3.3:	A block diagram of component load assignment and inversion of embedded control logic	21
Figure 3.4:	A reduced order system model obtained using the assumptions: $\mathbf{u}_i = \mathbf{r}_i + \delta_i$ and $ \mathbf{u}_i - \tilde{\mathbf{r}}_i \leq \varepsilon$	22
Figure 3.5:	A discrete Markov process representation of fault growth modeling	23
Figure 3.6:	Evaluation of the prognostic constraint at several finite horizons for two sample component loading profiles denoted ‘High’ and ‘Low’	27
Figure 4.1:	Photograph and illustration of mechanical coupling in a triplex redundant electromechanical actuator	32
Figure 4.2:	An illustration of a skid-steered vehicle	32
Figure 4.3:	Addition of uncertainty to the insulation breakdown model	34
Figure 4.4:	An equivalent circuit model for motor winding temperature dynamics	35
Figure 4.5:	Illustration of the feasible load allocation domain for EMA motors under the performance constraint $0.8 \cdot T_c \leq T_o \leq T_c$	38
Figure 4.6:	Illustrations of finite horizon prognostic constraint evaluation using linear (a) and piecewise linear (b) interpolation from a state of health at end of mission constraint	39
Figure 4.7:	A sample EMA head speed profile command	40
Figure 4.8:	Simulation results for an EMA under nominal control	43
Figure 4.9:	Simulation results for an EMA following the minimum allowable performance path	45

Figure 4.10:	Plots of the parallel distribution cost function using $\rho \cdot T_c = 0.4\tilde{T}$ (a), $\rho \cdot T_c = 0.6\tilde{T}$ (b), and $\rho \cdot T_c = 0.7\tilde{T}$ (c)	46
Figure 4.11:	Simulation results using $\tau = 50$ sec, $\lambda = 25000$	48
Figure 4.12:	Simulation results using a piecewise linear constraint	50
Figure 4.13:	Visualization of motor torque contributions to system locomotion on a four wheeled skid-steered vehicle	51
Figure 4.14:	Sample skid-steered vehicle path command	55
Figure 4.15:	Simulation results for a skid-steered vehicle under nominal control	57
Figure 4.16:	Simulation results for a skid-steered vehicle following the minimum allowable performance path	59
Figure 4.17:	Results of a search for an optimal skid-steered vehicle load allocation policy over the repeated figure-8 path	60
Figure 4.18:	Simulation results for a skid-steered vehicle following the minimum allowable performance path	61
Figure 4.19:	Plots of the performance degradation penalty (a), the motor degradation penalty (b), and the net outcome evaluation metric (c) for control policies identified using $\lambda = [0, 3, 6, \dots, 160]$ and $\tau = [.4, 1.6, 2.8, \dots, 24.4]$ (optimal value found at $\lambda = 100$ and $\tau = 3.6s$)	62
Figure 5.1:	Time-series plots for 100 simulations of the random variables, w and ξ (top), and the product of w and ξ (bottom)	72
Figure 5.2:	Time-series data (left) and box plots (right) for 100 simulations of the example fault growth process with 80% deviation ($\rho = 0.2$) from nominal output demands (top), and no deviation ($\rho = 1$) from nominal output demands (bottom)	73
Figure 5.3:	Plots of the stochastic health degradation process over repeated simulations of the optimal control policies computed with (a) and without (b) model uncertainty	74
Figure 6.1:	Closed-loop control diagrams for online and offline flight testing	79
Figure 6.2:	Picture of an Edge 540T on runway	85
Figure 6.3:	Electrical and mechanical connections of the Edge 540T UAV powertrain	85
Figure 6.4:	Model fitting results from X-Plane flight load characterization tests	90

Figure 6.5:	An equivalent circuit battery model	92
Figure 6.6:	Measured and fitted profiles for C_b and battery terminal voltage . . .	94
Figure 6.7:	Measured and fitted profiles for terminal battery voltage over a pulsed load test	95
Figure 6.8:	Plots of individual and net battery power consumption (top), and battery current load (bottom) measured over a sample flight	96
Figure 6.9:	Modeled and measured voltages of batteries B1 and B3 over a sample flight	98
Figure 6.10:	Estimated SOC for batteries B1 and B3 over a sample flight	98
Figure 6.11:	Measured and predicted net battery power consumption over steps 1-5 of the sample flight plan (annotated S1-S5 in the figure).	100
Figure 6.12:	Measured and predicted net battery power consumption over the sample flight plan.	101
Figure 6.13:	Measured battery power input to ESCs (Top) and observed ESC power ratio over a sample flight (bottom)	102
Figure 6.14:	Plots of estimated state of charge for B3 versus predicted state of charge at various points over the sample flight	103
Figure 6.15:	Predictions of battery B3 state of charge at $t = 846$ seconds is plotted at 60 second intervals over the sample flight.	104
Figure 6.16:	Operator displays captured near the beginning of a sample flight . . .	106
Figure 6.17:	Operator displays captured near the end of a sample flight	107
Figure 6.18:	Plots of the desired aircraft speed, acceleration, angel of climb, and altitude waypoints that define a sample mission	108
Figure B.1:	An equivalent circuit model for motor winding temperature dynamics	116
Figure D.1:	Visualization of motor torque contributions to system locomotion for a four wheeled skid-steered vehicle	120
Figure E.1:	Plots of throttle command, propeller RPM, and battery power draw measured over a sample flight	123
Figure E.2:	Battery current and voltage profiles recorded over a sample flight . . .	124

SUMMARY

The body of work described here investigates development and implementation challenges associated with a proposed paradigm for fault risk assessment and system health management. Health management actions are formulated here in terms of a value associated with improving system reliability, and a cost associated with inducing deviations from a system's nominal performance. The health management problem is considered to be complicated by the need to accommodate substantial uncertainties that are typically present in estimates of current component health states and predictions of future component deterioration.

Three simulated case study systems are considered here to highlight some of the fundamental challenges of formulating and solving an optimization on the space of available supervisory control actions in the described health management architecture. Repeated simulation studies on the three case-study systems are used to illustrate an empirical approach for tuning the conservatism of health management policies by way of adjusting risk assessment metrics in the proposed health management paradigm. The simulation studies also provide an opportunity to make illustrative comparisons between the supervisory control outcomes achievable if model uncertainty were reduced or eliminated.

The implementation and testing of a real-world prognostic system is also presented to illustrate model development challenges not directly addressed in the analysis of the simulated case study systems. Real-time battery charge depletion prediction for a small unmanned aerial vehicle is considered in the real-world case study. An architecture for offline testing of prognostics and decision making algorithms is explained to facilitate empirical tuning of risk assessment metrics and health management policies, as was demonstrated for the three simulated case study systems. The real-world case study also serves to provide further context to the role that unavoidable uncertainties play in the implementation and evaluation of failure risk assessment and management algorithms.

CHAPTER I

INTRODUCTION

Every physical system is subject to the continuous degradation of its constituent components as a function of time and usage. System designers can mitigate some of the risks posed by potential component failures through a more conservative specification of system redundancies, system operating policies, and system maintenance policies. However, system designers will also strive to avoid unnecessary sacrifices in performance and cost effectiveness incurred by overly cautious designs. Online estimates of the present health of system components (fault diagnostics) and prediction of their remaining useful life (failure prognostics) provide an opportunity for online health management actions that mitigate failure risks through responsive action rather than pre-planning.

The supervisory health management actions considered here are formulated in terms of a bounded set of adjustments that may be made to the component loads that would be allocated by a nominal control system. The component load terminology is used here as a stand-in for pressure, force, torque, or a wide variety of other stressors that drive component deterioration and system dynamics. The reduction of component loads will be considered to extend their likely useful lifetimes; however, this will typically come at the expense of degrading the nominal performance expected from an unmodified system.

A graphical introduction to the outer-loop supervisory health management paradigm to be described here is given in Figure 1.1. The architecture presented here is novel; however, it is posited that this functional breakdown is general enough to represent a wide swath of other prognostics-based risk management approaches.

The top shaded area in the figure denotes the inner-loop control dynamics of a system. A system operator or an imbedded controller is assumed to send control signals to powertrain

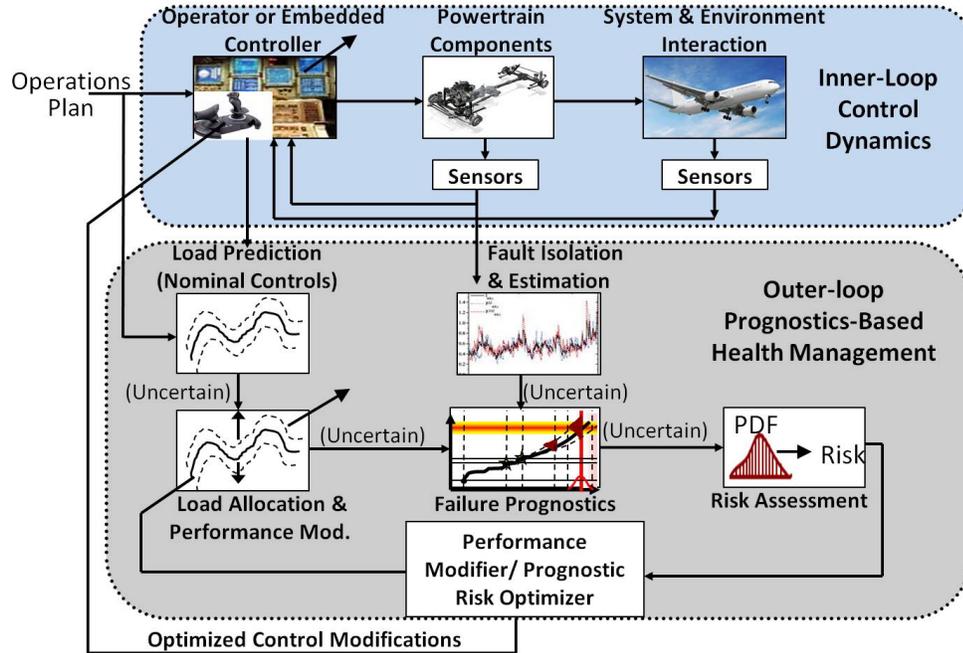


Figure 1.1: A block diagram representation for outer-loop prognostics and risk-based decision making

components based on a prescribed high level operations plan, the observed states of the controlled system, and the observed states of the system’s operating environment.

The bottom shaded area in the figure denotes the functional elements in the proposed outer-loop formulation of prognostics-based health management. The ‘Fault Isolation and Estimation’ functional block is shown to output estimates of current component health states, based on available sensor measurements from vehicle powertrain components. The ‘Load Prediction’ block represents the translation of a high-level operations plan into estimates of anticipated current and future loads that a given nominal control system would place on system components. The production of component failure time predictions based on estimates of current component health states and predictions of future component loads is represented by the ‘Failure Prognostics’ block. The translation of prognostic estimates into a value representing relative aversion to predicted outcome distributions is represented by the ‘Risk Assessment’ block. The prognostics-based decision making problem is then considered to be represented by a search for optimal modifications to a given nominal compo-

ment load allocation policy. Optimal load allocation policy modifications will be evaluated in terms of a trade-off between assessed failure risk and adherence to the system's nominal input-output behavior.

This dissertation takes the position that the 'Risk Assessment' and performance modification blocks in the depicted outer-loop health management architecture are the least mature, and thus the most deserving of our attention. The approach taken in much of the research described here has been to replace all other blocks in the functional breakdown shown in Figure 1.1 with compact models that highlight various sources of model and estimation uncertainties. This approach is intended to highlight fundamental formulation challenges associated with the 'Risk Assessment' and performance modification blocks in the proposed architecture. Several sources of model and estimation uncertainties are identified in Figure 1.1 with the 'Uncertain' label.

Section 1.1 of this chapter motivates the challenge of optimizing health management policies based on prognostic estimates. Section 1.2 summarizes the identified contributions of research described here. Section 1.3 briefly describes the information presented in each chapter of the dissertation.

1.1 A Motivating Discussion of the Prognostics-Based Decision Making Challenge

Risk management processes may be viewed as inherent to almost every decision that we make. However, arriving at a consensus quantification for the risk assessments that guide our decisions would almost always be the subject of much debate.

The following three questions are central to how the failure risk assessment and prognostics-based decision making tasks are addressed in this document.

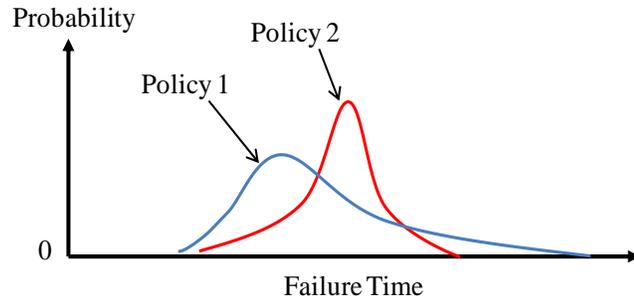
1. Given two probability distributions over the space of possible future outcomes, how would the preference for one distribution be established?

2. What is the relationship between the distribution of control outcomes that would be observed after repeated run-to-failure testing, and the outcome distributions predicted using prognostic modeling information?
3. How is preference evaluated for predictions on the space of future control outcomes?

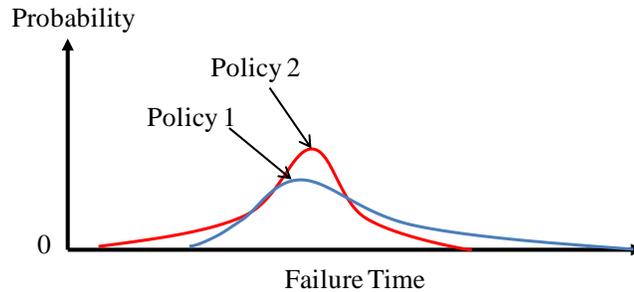
Consider these questions as they pertain to the sample failure time distributions illustrated in Figure 1.2. Figure 1.2a depicts distributions of component failure times corresponding to repeated run-to-failure trials of two sample control policies. Figure 1.2b depicts failure time predictions produced by a prognostic routine for the same two sample control policies. Figure 1.2a illustrates the fact that results of repeated run-to-failure trials of a given control policy are generally expected to show some variability corresponding to manufacturing variations and uncontrolled environmental variations. Figure 1.2b illustrates the fact that unmodeled dynamics and errors in the system models used by prognostic algorithms are generally expected to result in prognostic estimates that differ substantially from the failure time distributions observed in repeated run-to-failure testing.

First, consider question 1: Given knowledge of the failure time distributions observed over repeated trials of two sample control policies, how would a preference be established? One can imagine an application dependent process for weighting the relative aversion or preference for the failure time distribution's mean, median, mode, upper tail, lower tail, etc. Different system experts may come up with a different set of metrics, and a different decision on the preferred control policy, but let us assume that a satisfactory consensus could be reached for any particular application.

Question 2 asks how prognostic model uncertainty should be assessed in practice. The inherent variability expected in repeated run-to-failure trials, and the high cost of conducting repeated run-to-failure trials to verify the accuracy of prognostic estimates will typically complicate this problem. The question must however be addressed in the explanation of every implementation of prognostic prediction. The view taken in this research, is that the



(a) Sample failure time distributions given by repeated run-to-failure trials



(b) Sample failure time probability distributions estimated by a prognostic routine

Figure 1.2: Measured (top) and predicted (bottom) distributions for the system failure time corresponding to two potential control policies

measure of an effective prognostic estimator lies in desirability of the control outcomes achieved by a decision making policy that acts it. This view posits that a viable prognostic estimator is one that would enable better supervisory control actions than would otherwise be made.

Question 3 asks how should preference be evaluated over a set of possible prognostic predictions. This question is notably similar to Question 1, which asked for the specification of preference for measured distributions over the space of possible control outcomes. While it is reasonable to imagine that system experts would be able to express a preference for measured distributions, as posed in Question 1, it is much more difficult to imagine how a consensus answer would be generated for Question 3. The difficulty arises because in addition to the system experts evaluating their relative aversion or preference for the predicted outcome distributions, the system experts might also want to consider how much trust they have in the accuracy of the prognostic routines used. The problem of develop-

ing prognostics-based preference evaluation metrics is addressed here by demonstrating a means of tuning prognostics-based risk assessments to achieve desirable control outcomes

1.2 Identified Contributions of the Research Effort

1.2.1 Claim #1 - Description of Prognostics-Based Risk Management Using a Retrofit Nominal Load Modification Architecture

A retrofit architecture for enacting modifications to a nominal component load allocation policy is described in Chapter 3. Stochastic models of component deterioration as a function of component loads are used to evaluate constraints on the risk that component fault modes will exceed specified thresholds at a given prognostic horizon. Constraints on maximum tolerances for induced deviations from nominal system input-output performance are also used to evaluate constraints on supervisory load allocation policies. Strategies for tuning prognostic risk aversion metrics and for verifying the feasibility of prognostic constraints are described in the context of the retrofit architecture. Much of the architecture description covered in Chapter 3 was previously published in [8, 12, 15].

1.2.2 Claim #2 - Analysis of Supervisory Health Management Using Finite Horizon Prognostic Estimates

Risk assessment metrics are used to weight a relative aversion to the probability of component failure against the relative aversion to supervisory control actions that decrease failure probabilities at the expense of system performance. Presently, nearly all published explanations of prognostics-based health management systems consider remaining useful life (RUL) estimates in risk assessments. However, it is posited here that if future system demand and fault model uncertainties are high, it may be more useful to assess fault predictions over short time-scales where estimation uncertainty is smaller than it would be for RUL estimates. Chapter 4 explores the supervisory control behaviors resulting from manipulating certain variables in the risk assessment metrics introduced in Chapter 3, in-

cluding the length of the prognostic horizon over which control preferences are evaluated. Empirical results are provided for repeated simulation studies on two example systems with multiple degrading effectors. The described simulation studies are used to develop novel insights into the general problem of tuning the risk assessment metrics to achieve a desirable trade-off between system performance and system reliability. Analysis of the prognostics-based component load allocation problem on the multi-effector systems described in Chapter 4 was previously published in [12, 8, 17].

1.2.3 Claim #3 - Demonstration of the Explicit Incorporation of Future Demand Models into Health Management Policies

Chapter 5 describes the incorporation of stochastic models for future system demands into a Markov decision process formulation of the supervisory health management problem. Another simulated case study example is presented to demonstrate the incorporation of stochastic process models for future system loading and fault growth physics into a representative Markov decision process. Dynamic programming is shown to identify an optimal load allocation policy for the example stochastic system. Repeated simulations are used to compare the optimal health management control outcomes calculated with and without uncertainty in models of the sample fault growth process and future loading demands. Much of the material describing the Markov decision process formulation of the health management problem and the illustrative case study example system was previously published in [10, 11].

1.2.4 Claim #4 - Real-World Demonstration of Stochastic Component Health Depletion Modeling in terms of Uncertain Future Demand Estimates

Chapter 6 describes a framework for offline software-in-the-loop and hardware-in-the-loop testing of battery charge depletion prediction and charge management routines for an electric aircraft. The dynamic loads on aircraft powertrain components are identified using a real-time simulation of airframe dynamics running on a flight simulator, and an inner-loop

flight control policy (executed by either an autopilot routine or a human pilot). A stochastic model is created to predict the depletion of battery charge based on estimates of the current battery charge and a defined vehicle flight plan. The analysis presented in Chapter 6 is drawn mostly from two papers, [13, 22]. Development of battery charge management for a similar application, a hybrid electric vehicle, was presented in [9].

1.3 Organization of the Dissertation

- Chapter 2 gives background on the problem of prognostics-based risk assessment and the use of supervisory control policies to manage fault risk.
- A modular methodology for fault risk assessment and management is introduced in Chapter 3.
- A high level exploration of the load allocation problem for systems with multiple connected effectors at different states of health is given for two sample applications in Chapter 4. The sample applications considered are an electromechanical actuator (increasingly used in aerospace applications) and a four wheeled skid-steered vehicle.
- The identification of a component load allocation policy that optimizes a given stochastic fault growth model and future demand model is formulated as a Markov decision process in Chapter 5.
- Chapter 6 considers software-in-the-loop and hardware-in-the-loop testing of outer-loop prognostics and supervisory decision making routines on a real-world system. The real-world failure mode considered is battery charge depletion on a small battery powered aircraft.
- Chapter 7 gives concluding remarks.

CHAPTER II

BACKGROUND

Continuous improvements in fault diagnostic and fault growth prognostic technologies have spawned a growing research interest in the use of diagnostic and prognostic information to improve system control and maintenance. In general, the application of an analytical approach to the implementation and evaluation of prognostics-based decision making techniques will be complicated by the potential for high uncertainty in estimating the future effects of available supervisory control actions. There is also a need to define a computationally tractable space of present and future supervisory decisions to be optimized over.

The origin and history of the prognostics-based decision making (PDM) is described briefly in Section 2.1. Background on the representation of uncertainty in failure prognostic estimates is provided in Section 2.2. Risk quantification metrics are discussed in Section 2.3. Finally, Section 2.3 introduces the use of component load allocation for risk management.

2.1 Origin and History of the Problem

Fault analysis and fault risk mitigation strategies have become deeply ingrained in the design processes for many systems, and a wide variety of standardized methodologies have emerged over the past several decades.

2.1.1 Fault Risk Quantification Using Historical Data

The fault analysis process typically starts with the identification of potential failure modes and the quantification of severity and likelihood for each potential failure mode, based on expert knowledge and historical data. The failure modes, effects, and criticality analysis (FMECA) is one of the most widely known and applied *a priori* fault analysis methods. It was developed in the early 1970's by NASA for aerospace and defense [48], and is

currently widely applied in industrial automation [34], automotive [45], and aerospace [69] industries.

Fault Tree Analysis (FTA), Event Tree Analysis (ETA), Six Sigma (6σ), and Reliability Block Diagrams (RBD) are examples of other popular techniques for fault analysis and mitigation based on historical failure rates. Such techniques will continue to play an ever more prominent role in the design of hardware redundancy specifications and contingency management policies.

2.1.2 Condition Based Maintenance and Control

In addition to the growing trend towards utilizing historical fault data to manage failure risks, there is also a growing push to utilize technologies for online fault identification and fault growth prediction to improve the operation and maintenance of systems. The implementation of online anomaly detection and diagnostic routines are enabling the increased development and use of condition based maintenance and control (CBMC) policies. Reductions in maintenance costs and failure rates resulting from the implementation of CBMC are quantified for several systems in [25]. Automated control recovery after a known fault mode is detected has been shown using Pseudo-inverse [43, 19], model predictive control (MPC) [56], and H_2 and H_∞ robust control theory [28].

CBMC techniques can yield dramatic improvements in risk management for incipient fault modes, to the extent that they are able to detect faults before they become catastrophic and adjust maintenance and control policies accordingly. Real-time diagnostics-based techniques for CBMC are rapidly growing in popularity in industrial, automotive, and aerospace applications. Further improvements in performance and safety are expected if the diagnostic information used by CBMC routines is supplemented with prognostic estimates of fault growth as a function of future use.

2.1.3 Prognostics-Based Health Management

Several recent government initiatives are dedicated to the analytical development of diagnostics and prognostics enhanced control techniques, such as the Integrated Vehicle Health Management (IVHM) Project (initiated in 2008) [89], the Integrated Resilient Aircraft Control (IRAC) Project (initiated in 2007) [51], and the Aircraft Aging and Durability (AAD) project [97] (initiated in 2008).

Predicting the evolution of future system states may be performed in real-time by propagating input uncertainty, model uncertainty, and state uncertainty forward in time until failure conditions are reached. Component failure time estimates may be reported in terms of a probability distribution or in terms of a range of values falling within a given confidence interval. Particle filtering [3], extended Kalman filtering [67], and Markov modeling [38] are examples of predictive filtering techniques that have been used to propagate current state and model uncertainties forward in time.

The development and use of prognostic information is typically a challenging proposition due to significant uncertainty that is generally unavoidable in predicting future fault growth. The models used to estimate fault dynamics are typically split into two main camps; one is data driven modeling [36, 82], and the other is physics based modeling [7]. Although, an ideal approach would likely make use of both physics of failure models and data driven model adaptation techniques, as described in [72, 23, 67].

The use of stochastic estimates of future loading in prognostic predictions is described in [76, 90]. Because outer-loop supervisory control actions will effect component load allocations, outer-loop supervisory control policies could also be factored into the computation of fault growth predictions. The incorporation of outer-loop control policies into fault growth predictions is described in [10].

2.2 Uncertainty Representation in Prognostics

The eventual failure of components within a system is considered in this document in terms of component fault modes that will grow in severity until they cross a threshold, after which they are considered no longer viable. Fault magnitudes are assumed to be represented by a real number corresponding to a measurable physical property such as crack length, spall width, or pitting depth; although, in many cases, faults cannot be directly measured *in situ* and diagnostic routines are needed to approximate current fault magnitudes based on the secondary effects observed in available sensor measurements.

Consider the dynamics of a particular component failure mode to be defined by the following generic nonlinear mapping:

$$\dot{\gamma}_l = f(\gamma_l, u_l, \xi_l) \quad (2.1)$$

where γ_l is a random variable representing a fault magnitude for the l^{th} component in a system, u_l represents a loading profile applied to component l , and ξ_l is a random variable representing uncertainty in this physics of failure model.

The component loading variable, u_l , may represent pressure, force, torque, or a wide variety of other stressors that drive component deterioration. Component loading profiles are assumed to be dictated partly by the dynamics of the system's operating environment, which may not be entirely predictable, and partly by the supervisory health management control actions taken in response to online observations of environmental states and component fault magnitudes.

Uncertainty in estimating the dynamics of component degradation as a function of the loads or stresses applied to a specimen is often included as a noise term in lumped parameter and data driven fault growth prediction models [72, 63]. A bounded environmental disturbance term is typically used to incorporate modeling of future exogenous system stresses into the analysis of robust and reconfigurable control design techniques, such as H_∞ control, L_1

control, and gain scheduling [104, 102].

Sensor noise and feature mapping uncertainties will often result in significant diagnostic uncertainty, and it is common practice for diagnostic estimates to be reported in terms of a probability distribution over the potential fault magnitudes that could correspond to a given set of observations. A recursive expression for the Bayesian belief in the current state of a particular fault mode, given a set of sequential observations is written as;

$$p(\gamma(k) | z_l(k)) = \alpha p(z_l(k) | \gamma(k)) \cdot \int p(\gamma(k) | \gamma(k-1), u_l(k-1)) \cdot p(\gamma(k-1) | z_l(k-1)) d\gamma \quad (2.2)$$

where z_l represents observations of a component fault magnitude, $p(\gamma(k) | \gamma(k-1), u_l(k-1))$ represents uncertain fault growth physics modeling, $p(z_l(k) | \gamma(k))$ represents an observation model, and α is a normalizing constant. The recursive nature of this expression makes it computationally expensive to compute, opening the door for approximate Bayesian methods and machine learning to be used for fault magnitude estimation [26].

The application of analytical processes for uncertainty representation and management have been described in many publications on fault diagnostic classification [32, 52, 60] and failure prognostic modeling [78, 64, 77, 91]. Prudent methods for obtaining diagnostic and prognostic probability density functions (pdfs) should be selected on a case-by-case basis. Particle Filtering [4, 64, 63, 93] and Bayesian Reasoning [72, 84] are commonly used approaches for estimating fault magnitudes and predicting future growth based on available measurements and physical modeling.

Fault growth predictions may be generically represented as:

$$\Pr[(\gamma(t_p + \tau) = \gamma | u_l(t))], t \in [t_p..t_p + \tau] \quad (2.3)$$

where t_p denotes the time at which a prediction is made, and τ denotes the length of the

horizon to be predicted over.

Prognostic pdfs are represented in this document using the following more compact form of Eqn. 2.3:

$$p(\mathcal{Y}(t + \tau) | u_l(t)) \quad (2.4)$$

The time period until the hazard threshold is reached is referred to as the remaining useful life (RUL) or the time-to-failure (TTF). RUL and TTF probability distributions are represented here as

$$\Pr[\mathcal{Y}(t_p + \tau) = \zeta_l | u_l(t)], \quad \forall \tau \in [0..∞], t \in [t_p..t_p + \tau] \quad (2.5)$$

where ζ_l denotes the hazard threshold for component l , and τ is a random variable representing RUL.

2.3 Risk Quantification

The application of analytical processes for uncertainty representation and management has been described in many publications on fault diagnostic classification and failure prognostic modeling, as cited in the previous section. However, evaluating the effectiveness of tools for making intelligent control decisions based on available diagnostic and prognostic information is a much more complex proposition, due to the lack of certainty and repeatability inherent to such problems.

The specification of appropriate metrics for assessing the risk posed by probabilistic prognostic predictions of future component health deterioration may generally be considered independently from the problem of building prognostic models. In practice, some form of scenario analysis could be used to derive and validate risk assessment metrics through empirical studies wherein people familiar with the targeted systems would quantify the relative value of preserving nominal system performance and the relative aversion to the risks posed by potential fault scenarios [2, 80].

As was described in Chapter 1, Section 1.1, the translation of outcome evaluation metrics into metrics for evaluating the relative preference for various possible fault scenarios is complicated by the fact that the trust in prognostic outputs must also be incorporated. The utilization of strategic analytics to redesign a system's mission based on diagnostic and prognostic estimates is discussed at a high level in [29]. A discussion of the formulation of convex risk metrics to enforce asymptotic safety and asymptotic precision limits is found in [33]. Quantitative metrics for assessing the performance of prognostic algorithms are compared in [78].

Candidate risk metrics may be drawn from the growing body of publications offering discussions of health management systems (HMS), or integrated systems health management (ISHM) architectures. Although, nearly all current studies in this area consider only end of life predictions in risk calculations and ignore data regarding short term fault growth, which is probably not ideal in many cases.

Literature on risk management in finance and actuarial science contains a rich array of tools that may be adapted to analyze prognostic information. Value at risk (VaR) is an example of a prolific financial risk management tool that may be utilized in analyzing fault prognostic risk. VaR metrics are commonly used to judge the financial risk of investments [30]. The VaR metric has been applied to a wide variety of risk management problems in areas such as industrial systems management [81] and agricultural resource planning [95]. Chapter 4 provides several examples of VaR metrics for assessing risk and enforcing prognostic constraints at a chosen prognostic horizon.

Expected utility (EU) is perhaps the most prolific tool for risk based decision making in the fields of decision theory and Bayesian reasoning. It is essentially a way of encoding risk into an expected value calculation. The reader interested in understanding EU theory is encouraged to start by considering the famous St. Petersburg paradox [75], which originally motivated mathematicians to formulate EU theory. The reader is then encouraged to look

up explanations for the three fundamental axioms of EU theory: ordering, independence, and continuity. The review papers by Schoemaker [79] and Camerer [20] describe many of the key empirical studies, dating back from the early 1950s, that have played a particularly important role in shaping EU theory.

2.4 Component Load Allocation for Risk Management

The load allocated to individual components at each control time-step by a given control policy is assumed to be a primary driver of component degradation. The aggregate of individual component loads is also assumed to drive a system's mechanical dynamics.

Practical applications of control allocation based on fault information are already found in aerospace [31, 85, 101, 55], automotive vehicles [40], and marine vehicles [46]. Optimal load allocation to extend the RUL of a single component is discussed in [37, 101]. A survey of efficient methods for determining the optimal control allocation for general linear and nonlinear systems is discussed in [62]. Proof of the equivalence of this type of control allocation and optimal control is given in [39] for nonlinear systems with precise modeling and a quadratic cost function.

A modularized approach to the control allocation for risk management problem is described in the next chapter.

CHAPTER III

A MODULAR METHODOLOGY FOR FAULT RISK ASSESSMENT AND MANAGEMENT

3.1 Introduction

A modular architecture for fault risk assessment and management is described in this chapter. The potential benefits of the proposed approach, as well as its fundamental limitations, are discussed.

The following general assumptions are made regarding the data and modeling available to the control routines.

Assumption 3.1. A sufficiently accurate dynamic model of the system is known at all times. Where a definition of sufficient model accuracy is provided in Section 3.2.3.

Assumption 3.2. Fault magnitudes will increase monotonically with component load.

Assumption 3.3. Diagnostic and prognostic estimates for each component are given in terms of probability density functions (pdfs).

3.2 A Component Load Allocation Control Structure

This section introduces the component load allocation for risk management control structure in terms of a nominal control loop, a component load allocation loop, and a low-level load enacting controller.

3.2.1 The Nominal Control Loop

Figure 3.1 shows a block diagram for a system under nominal control. Here, each system effector is assumed to have an embedded controller that adjusts component input signals,

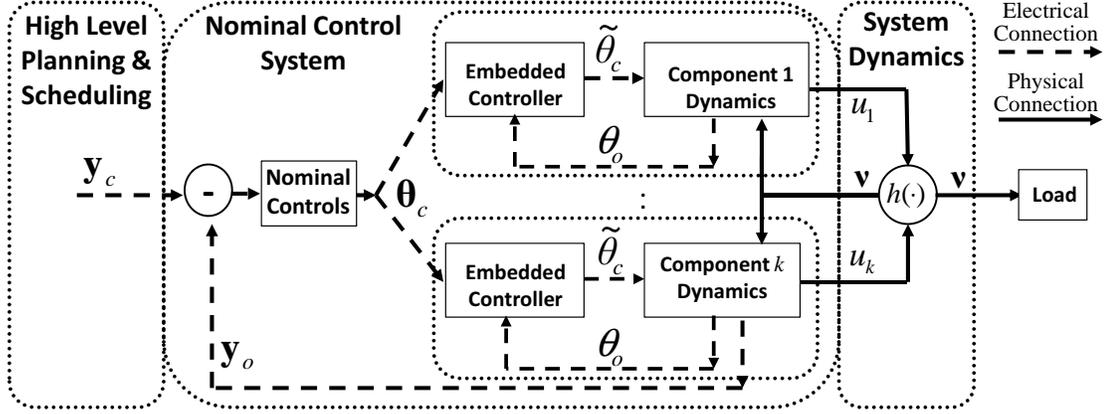


Figure 3.1: Block diagram showing the assumed structure of nominally controlled systems

$\tilde{\theta}_c$, based on the difference between component set-point commands, θ_c , and component state feedback signals, θ_o . A reference system state vector, designated y_c , is assumed to be specified by a high level planning and scheduling routine. The difference between the desired state vector and an observed state vector, designated y_o , is assumed to be acted on by some nominal control logic that maps $(y_c - y_o) \rightarrow \theta_c$.

The load exerted on or by component i at a given time is represented by u_i . The net output force exerted by the system on its environment is represented by the variable, v , and the mapping between u_i and v is represented by $h(\cdot)$. The physical mapping represented by $h(\cdot)$ is assumed to be known for a system of interest; however, it is allowed to be dynamic and nonlinear.

3.2.2 Component Load Allocation

A block diagram designating separable control tasks in a load allocation for risk management controller is given in Figure 3.2. The optimization problem that is the focus of this research is formulated as a middle layer in the control architecture illustrated in Figure 3.2.

The first step in the load allocation control structure shown in Figure 3.2 is the inversion of nominal control dynamics to determine the net control effort that would be output by the nominally controlled system at each control time-step. After a fault has been detected,

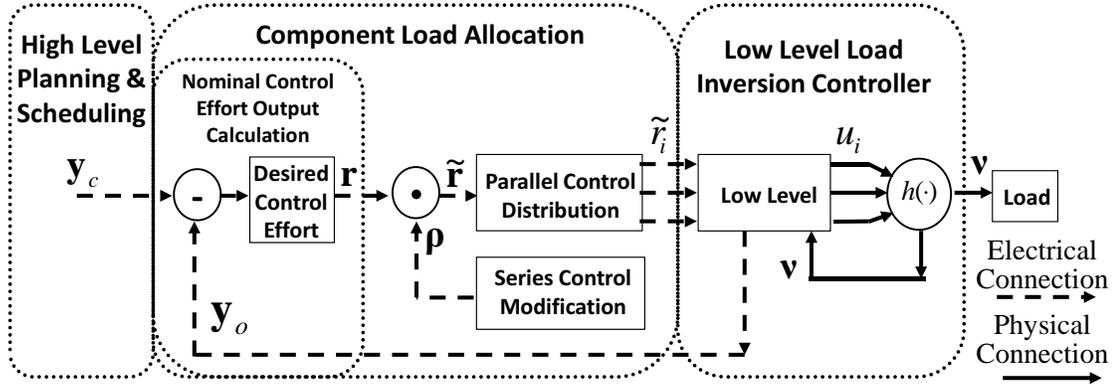


Figure 3.2: A functional breakdown of control tasks in the proposed component load allocation framework

the state observations and control commands that would have been input to the ‘Nominal Controls’ block, shown in Figure 3.1, are rerouted to the ‘Desired Control Effort’ block, shown in Figure 3.2. The output of the ‘Desired Control Effort’ block is an estimate of the forces that would have been exerted by a system on its environment if no fault had occurred and the nominal control laws were still in effect.

The vector, r , in Figure 3.2 is used to represent the current estimate of a nominally controlled system’s output force response to a given set of inputs from high level planning and scheduling routines. This vector is viewed as the reference performance to be delivered by component load allocation routines in each of the system’s output degrees of freedom. Individual component loads are then allocated by a risk management policy that will attempt to minimize induced deviations from a system’s nominal input-output performance, while also lowering the loads on degrading components in order to increase safety.

An advantage of having both system kinematics and fault growth dynamics formulated in terms of component loading is that load allocations between functionally redundant components may be separated from the top level supervisory control problem and handled by a nested optimization routine. A mathematical description of the nested optimization problem for systems with functionally redundant components is given in Appendix A.

The optimal load allocation for risk management problem is represented in Figure 3.2 by

two functional blocks. The ‘Series Control Modification’ block represents the problem of determining the degree to which the faulty system will deviate from the output force that would have been exerted by a healthy system in each of the system’s output degrees of freedom at each control time-step. The ‘Parallel Control Distribution’ block represents the problem of distributing load among active redundancies in a system while delivering a specified net output force from the composite system.

The vector, ρ , in Figure 3.2 represents the relative deviation between the net output control effort that would have been exerted by a system on its environment if a nominal control law were used, and the net output control effort that is actually exerted by a given health management control routine at each control time-step. The vector, \tilde{r} , in Figure 3.2 represents a modified net output control effort allocation, given by the dot product of ρ and r ,

$$\tilde{r} = \rho \cdot r \quad (3.1)$$

The state space of feasible values for \tilde{r} will be dictated by the limitations of the system’s individual control effectors, and the acceptable range for the nominal performance modifier, ρ . Sections 3.3.1 and 3.3.2 of this chapter, introduce the formulation of performance and prognostic constraints to be enforced over the domain of feasible component load allocations.

The variable, \tilde{r}_i , in Figure 3.2 represents the load allocated to component i by the ‘Parallel Control Distribution’ functional block. Because the optimization of the system output control effort vector, \tilde{r} , requires knowledge of the component failure risk that would be induced by a particular system output load demand, the mapping $\tilde{r} \rightarrow \{\tilde{r}_1, \tilde{r}_2, \dots, \tilde{r}_N\}$, should be used to optimize the ‘Series Control Modification’ policy.

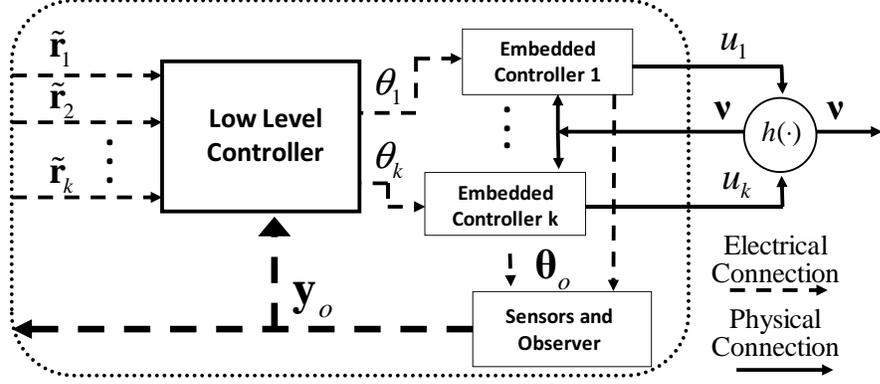


Figure 3.3: A block diagram of component load assignment and inversion of embedded control logic

3.2.3 Low Level Component Load Allocation

Figure 3.3 shows the insertion of a low level controller that is assumed to work as an outer-loop controller on top of other embedded component control logic. The main problem the implementation of such controllers lies in determining how to make clever use of the sensing and control hardware already on the system to enable accurate load allocation. Model predictive control (MPC) and dynamic inversion are commonly used to implement a desired component load allocation in a damaged system, as described in the references [56, 54, 17] and [47, 61] respectively.

Figure 3.4 depicts a generic formulation of low level control dynamics, in which the error in attaining allocated component loads is represented by the dynamic variable δ_i . A constraint on the accuracy for achieving a specified component load allocation is given as:

$$|\delta| \leq \varepsilon \quad (3.2)$$

where ε denotes the error bound on a component load controller.

Acceptable ranges for δ_i and ε would need to be developed on an application by application basis to ensure system stability and model fidelity. The discussion of the higher level control effort allocation routines will proceed in this document with the assumption that

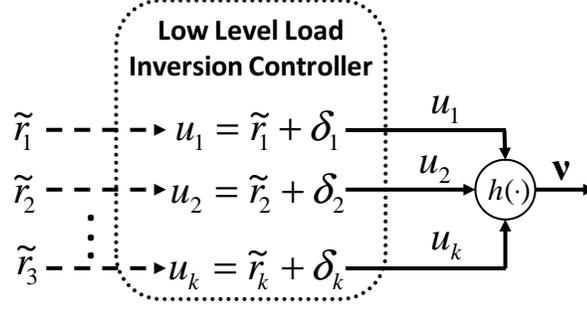


Figure 3.4: A reduced order system model obtained using the assumptions: $\mathbf{u}_i = \mathbf{r}_i + \delta_i$ and $|\mathbf{u}_i - \tilde{\mathbf{r}}_i| \leq \varepsilon$

δ_i and ε are small enough to be ignored.

3.3 Stochastic Modeling of Fault Growth in terms of Component Load and Net Output Control Effort Allocations

Consider a state space of feasible component load allocations that may be enacted by a supervisory controller at current and future decision making epochs. The general stochastic component degradation process given in Eqn. 2.1 is formulated here in terms of a discrete Markov process,

$$\begin{aligned}
 p_{i,j}^l(u_l) &= p(\gamma_l(k+1) = s_j | \gamma_l(k) = s_i, u_l) \\
 &= p(\xi_l(k) = \xi_j) \text{ where } f(s_i, u_l, \xi_l) = s_j, \quad s_i, s_j \in S_l, u_l(k) \in U_l(k), k \in \mathbb{N} \quad (3.3)
 \end{aligned}$$

$$\sum_{j=0}^m p_{i,j}^l = 1, \quad \forall i \in \{0, 1, \dots, m\} \quad (3.4)$$

Here, S_l and U_l represent a uniformly quantized state space of potential fault magnitudes and potential load allocations for component l respectively, and $p_{i,j}^l(u_l)$ represents the probability of transitioning from damage state s_i to damage state s_j , given a particular component load allocation, u_l . The function, $f(s_i, u_l, \xi_l)$, represents a model of fault mode dynamics,

$$\gamma_l(k+1) = f(\gamma_l(k), u_l(k), \xi_l(k)) \quad (3.5)$$

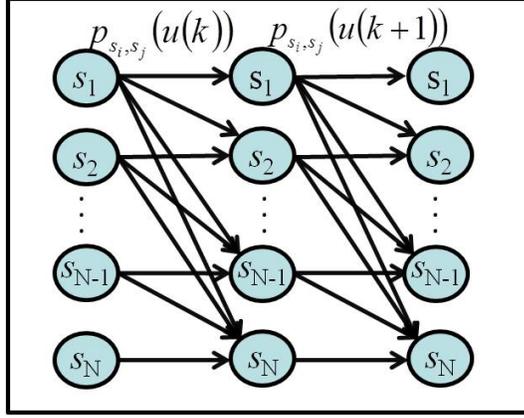


Figure 3.5: A discrete Markov process representation of fault growth modeling

where γ_l represents the damage state of component l , and ξ_l is a random variable that is used to represent uncertainty in the fault dynamics model.

Equation 3.4 specifies that the sum of all transition probabilities defined at each system state must always be equal to one. Figure 3.5 shows an illustration of a discrete Markov process representation of a fault growth model. Such discrete Markov process models capture the entirety of modeled dynamics using a quantized set of fault states to be observed at discrete time-steps, and a fixed set of state transition probabilities.

A mandate of monotonically increasing component fault modes is incorporated into the Markov process notation given in Eqn. 3.3 as:

$$p_{i,j}^l = 0, \text{ if } j < i, \quad \sum_{j=0}^m p_{i,j}^l = 1 \quad \forall i \in \{0, 1, \dots, m\} \quad (3.6)$$

This constraint will be problematic for other fault growth modeling techniques that represent process uncertainty with an analytical distribution that lacks an explicit lower bound. For example, in the case of Kalman filtering or Gaussian process models of fault growth, an assumption of Gaussian uncertainty would introduce some probability that the fault mode will be smaller in the future than it was known to be in the past. It would be necessary, in such cases, to assure that the probability attributed to non-realizable outcomes,

$P(\gamma(\tau) < \gamma(t))$ for $\tau > t$, will be acceptably small.

The state transition probabilities defined in Eqn. 3.3, are directly derivable from the fault mode dynamics model represented in Eqn. 3.5, given a model for the statistics of ξ , and kinematic modeling sufficient to determine the component loads that result future control actions and environmental inputs.

As described in Section 3.2, the net output control effort, v , exerted by a system at each control time-step is expressed in terms of a nominal control effort output response r , and a performance metric, ρ , as:

$$v = \rho \cdot r \quad (3.7)$$

The nominal control effort output response of a given system is defined here using the following functional mapping:

$$r = G(w, x, x_c) \quad (3.8)$$

where $G(w, x, x_c)$ maps a system state, x , a system state command, x_c , and an exogenous system load, w , to an expected nominal control effort output response, r .

The minimum damage mapping between the net output loads exerted by a system on its environment, and the loads on constituent components is denoted by the mapping

$$u(k) = H(v(k)) \quad (3.9)$$

where,

$$H(v(k)) = \min_u E \{f(\gamma, u, \xi)\}, \quad \text{s.t. } v = B_u(x) \cdot u \quad (3.10)$$

Derivation of the minimum damage mapping is described in Appendix A.

Equation 3.3 can now be rewritten in terms of the performance metric, ρ , as:

$$p_{i,j}^l(\rho(k)) = p(\gamma_l(k+1) = s_j | \gamma_l(k) = s_i, \rho = \rho(k)) = \sum_{w \in W} \sum_{\xi \in \Xi} p(w(k) = w) \cdot p(\xi_l(k) = \xi),$$

$$\text{where } f(s_i, H(\rho \cdot G(w, x, x_c))_l, \xi_l) = s_j, s_i, s_j \in S \quad (3.11)$$

The process uncertainty terms, ξ_l , are considered to belong to a stationary distribution, and a potentially non-stationary distribution is considered for the exogenous inputs to the system, w . This formulation of fault growth dynamics provides a convenient means for analyzing the prognostics-based control problem, because it directly relates modifications of system output performance to predictions of component degradations.

3.3.1 Performance Constraints

Operational constraints on minimum allowable output performance and maximum allowable fault growth risk are considered here. An operational constraint on minimum allowable system performance is considered to be defined in terms of a maximum allowable deviation from some commanded system state:

$$|x_c(k) - x(k)|_i \leq \Delta_i(k), \quad i \in \{1, 2, \dots, n\} \quad (3.12)$$

where Δ_i specifies a maximum acceptable error between the i^{th} dimensions of x and x_c .

Equation 3.12 is approximated in terms of a maximum allowable deviation from a given nominal control effort output profile as:

$$|r_i(k) - v_i(k)| \leq \tilde{\Delta}_i(k), \quad i \in \{1, 2, \dots, p\} \quad (3.13)$$

where $\tilde{\Delta}_i(k)$ specifies a maximum acceptable error between the i^{th} dimensions of v and r .

The constraint is expressed in terms of the performance metric ρ by substituting Eqn. 3.7

into Eqn. 3.13;

$$|1 - \rho_i(k)| \leq \frac{\tilde{\Delta}_i(k)}{r_i(k)}, \quad i \in \{1, 2, \dots, p\} \quad (3.14)$$

$$\bar{\Delta}_i(k) = \frac{\tilde{\Delta}_i(k)}{r_i(k)} \quad (3.15)$$

where $\bar{\Delta}_i$ now represents a constraint on the maximum allowable deviation from a system's nominal control effort output in dimension i at time-index k .

3.3.2 Prognostic Constraints

Value at risk (VaR) metrics are commonly used to judge the financial risk of investments [30]. The VaR of a random variable X at a confidence level ψ is defined as:

$$\text{VaR}_\psi(X) = \inf\{x \in \mathbb{R} : P(X < x) > \psi\} \quad (3.16)$$

Constraints on component longevity may be enforced at a fixed prognostic horizon through the specification of a lower bound on acceptable VaR assessments of system health over the range $t = [t_p, t_M]$. This metric has been applied to a wide variety of risk management problems in areas such as industrial systems management [81] and agricultural resource planning [95].

A finite horizon constraint on future fault growth risk is written as:

$$\text{VaR}_\alpha(\gamma_l(t_p + \tau)) > \zeta_l(t_p + \tau) \quad (3.17)$$

where γ_l represents a state of health (SOH) index for component l , t_p represents the time at which a fault growth prediction is made, τ specifies a time horizon at which prognostic constraints will be evaluated, ζ_l specifies a lower bound on the VaR estimates of component health over the range $t = [t_p, t_M]$, and α specifies the maximum acceptable probability that the health of component l is less than ζ_l at time $t_p + \tau$.

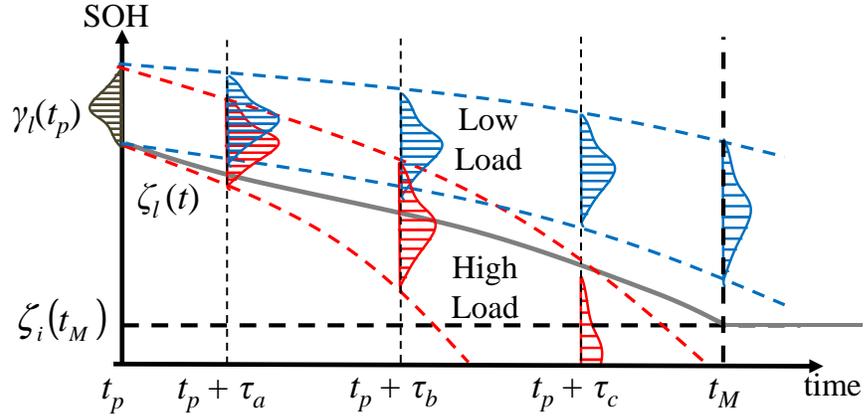


Figure 3.6: Evaluation of the prognostic constraint at several finite horizons for two sample component loading profiles denoted ‘High’ and ‘Low’

Constraints on the minimum acceptable component failure time may be specified in terms of a maximum acceptable probability that the component will reach a given SOH index before a given time:

$$\text{VaR}_\alpha(\gamma_l(T)) > \zeta_l \quad (3.18)$$

where α represents an upper bound on the probability that component l will be lower than the SOH index, ζ_l , at time, T .

Figure 3.6 illustrates the spread of uncertainty for future SOH predictions, depicted at finite horizon prognostic horizons, $(\tau_a, \tau_b, \tau_c, t_M)$, for two sample component loading profiles denoted ‘High’ and ‘Low’ . The growth of uncertainty in future SOH predictions over large time horizons will tend to diminish the utility of those estimates, which motivates the use of shorter horizon predictions in risk management control loops. A sample definition for the prognostic constraint, ζ_l , is shown in the figure.

The specification of an appropriate profile for the hazard threshold, $\zeta_l(t)$, will be left as a design choice. Some general considerations for the specification of an appropriate profile include:

- If component health is expected to be monotonically decreasing and it is ever the

case that $VaR_{\beta_l}(\gamma_l(t_p)) < \zeta_l(t_p + \tau)$, then the prognostic constraint is unsatisfiable.

- The greater the difference between $VaR_{\beta_l}(\gamma_l(t_p))$ and $\zeta_l(t_p + \tau)$, the greater the control freedom allowed under the prognostic constraint.

Many publications on the topic of prognostics-based control elect to simplify the prognostics problem by assuming that component loadings will be unvarying over the prediction horizon:

$$u(t) = u(t_p) \quad \forall t \in [t_p, \infty] \quad (3.19)$$

However, in most cases, time-varying environmental loading conditions and time-varying component health estimates are expected to result in time-varying loadings on a system's components. The production of accurate prognostic estimates in such cases, requires future component loadings to be modeled as a stochastic process that accounts for the statistics of all parameters affecting component load allocations within the controlled system.

3.3.3 Verifying the Feasibility of the Constrained Load Allocation Problem

The existence of feasible solutions to the constrained load allocation problem is verified in simulation studies by demonstrating the existence of at least one feasible solution. The candidate load allocation used to demonstrate the existence of controls that satisfy both performance and prognostic constraints is selected to be one that supplies the minimum output force allowed under the performance constraint.

The minimum allowable output force profile is defined here as:

$$\tilde{\mathbf{v}} = \min \{ \mathbf{v} \}, \quad \text{s.t. } |\mathbf{v} - \mathbf{r}| = \tilde{\Delta} \quad (3.20)$$

where $\tilde{\mathbf{v}}$ represents the minimum output control effort that is allowed under the performance constraint.

Feasible solutions to the component load allocation problem then exist if there is a distri-

bution of component loads that result in $\tilde{\mathbf{v}}$ and do not violate the prognostic constraint at the end of the mission. This condition is written as follows:

$$\Pr(d_i(t_M) > \gamma_i | \mathbf{v}_i(t)) \leq \alpha, \text{ s.t. } \tilde{\mathbf{v}} = h(\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_N), \quad t \in [t_p, t_M] \quad (3.21)$$

where $h(\cdot)$ represents the kinematic mapping of component loads to the net output forces exerted by a system on its environment, and t_M represents the end of mission (EOM) time.

The dimensionality of space of feasible component load distributions corresponding to $\tilde{\mathbf{v}} = h(\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_N)$ is equal to the degree of active redundancy in the system. The search for feasible control solutions can be terminated as soon as one feasible solution is identified.

3.3.4 Strategic Definition of Performance Objectives and Risk Aversion Metrics

The proposed fault risk assessment and risk management paradigm will indirectly address the task of designating risk assessment metrics through comparative studies of the efficacy of various formulations for the bounded load allocation problem on selected application examples.

It is proposed that system designers could start from the definition of an analytical expression to evaluate the desirability of future control outcomes. For example, the ultimate objective for operations and maintenance of industrial robotics will likely be to maximize future profits, and for a NASA rover, the objective may be to maximize the vehicle's useful lifetime and landmass explored. The definition of such outcome evaluation expressions provides a means of assessing the relative desirability of control policies through iterative simulation studies. Iterative simulation studies are examined here to tune risk assessment metrics based on the relative desirability of control outcomes corresponding to the online optimization of a given set of risk metrics. The application of this type of iterative empirical approach to the definition of risk quantification metrics is more of an academic exercise than something that would actually be performed on a real system. On real systems such

iterative run-to-failure experiments would be very costly and impractical. However, once this paradigm is understood it can be an instructive conceptual tool.

A general form of the risk metric to be optimized by supervisory load allocation controllers is:

$$J(u, \rho) = J_p(\rho) + \lambda \cdot J_d(p(\gamma(t_p + \tau) | u(t))) \quad (3.22)$$

where $p(\gamma(t_p + \tau) | u(t))$ represents component health predictions evaluated at a finite horizon, using the component loading vector, $u(t)$. $J_p(\rho)$ represents a penalty function that assigns a cost to induced deviations from a system's nominal output performance, $J_d(p(\gamma(t_p + \tau) | u(t)))$ penalizes aversion to prognostic pdf estimates, and λ is a weight that captures the relative importance of performance and reliability.

Chapters 4 and 5 present empirical simulation studies to demonstrate the effects that various risk metric formulations have on resulting prognostics-based supervisory control policies for several simulated application examples.

CHAPTER IV

AN EXPLORATION OF LOAD ALLOCATION AMONG SIMILAR DEGRADING COMPONENTS

4.1 Introduction

This chapter describes the load allocation for risk management problem in the context of two example systems: a triplex redundant electromechanical actuator (EMA), and a skid-steered unmanned ground vehicle (UGV). The two example systems are used to demonstrate the allocation of load in systems with multiple similar degrading effectors. An empirical model for winding insulation degradation in DC motor systems is adapted to simulate the degradation of the motors in both example systems over simulated missions.

The EMA and UGV systems are used to illustrate the challenge of deciding how load is best split among similar cooperating components at different states of health. In both cases substantial conceptual and computational challenges are shown to stand in the way of optimizing system output performance and prognostics-based risk metrics. Analysis of the prognostics-based component load allocation problem on the EMA and UGV systems was previously published in [8, 17] and [12] respectively. In addition to both examples being informative to the general theoretical discussion of prognostics-based control allocation problems, the analysis of prognostics-based decision making on these example systems is also motivated by the fact that both systems are seeing increasing use in the safety minded domains of aerospace and exploration vehicles respectively.

EMAs are commonly used in commercial aircraft, military air/land vehicles, robotics, and industrial process [87], in which system failures would be extremely costly and possibly dangerous. A triplex redundant EMA of the type used on aircraft is shown in Figure 4.1. Diagnostic and prognostic studies for similar EMA systems are found in [17, 16, 99, 37].

An illustration of a skid-steered vehicle is shown in Figure 4.2. Each of the four wheels of the skid-steered vehicle are fixed to the vehicle frame, and are pointed straight forward. Each of the wheels is powered by an independently driven DC motor. The four wheel motors will interact with each other and the vehicle through their mutual contact with the ground.

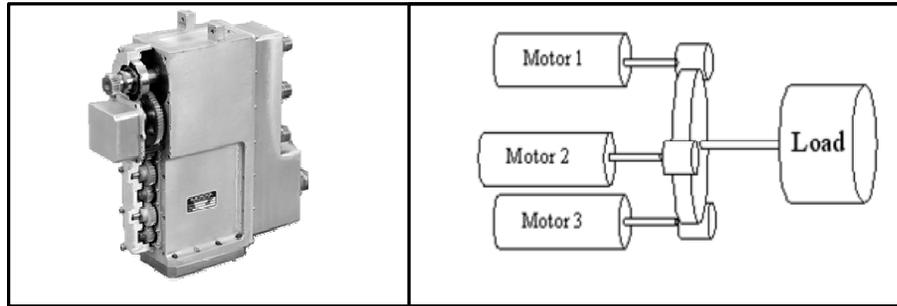


Figure 4.1: Photograph and illustration of mechanical coupling in a triplex redundant electromechanical actuator

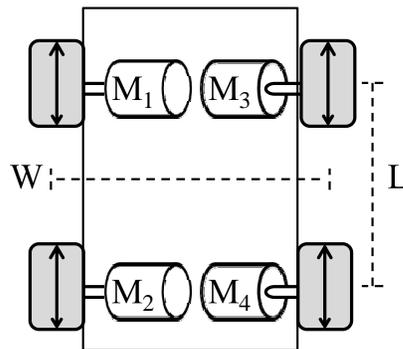


Figure 4.2: An illustration of a skid-steered vehicle

4.2 A Fault Growth Model

Winding insulation breakdown is considered to be the primary failure mechanism for the example systems considered in this chapter. The rate of motor winding insulation breakdown is assumed to be driven primarily by the thermal loads on motor windings. A standard heuristic model for this relationship is the ten-degree rule, introduced in 1930 by Montsinger [57]. The useful lifetime of motor winding insulation is observed to be roughly halved for each increase of 10°C in the exposure temperature. This exponential relationship

between winding temperature and expected winding life is expressed as:

$$L_N(T_w(t)) = \alpha e^{-\beta T_w(t)} \quad (4.1)$$

where the α and β coefficients parameterize the linear part and exponential part of the anticipated lifetime model respectively, L_N represents the anticipated lifetime of new insulation in seconds if it is held at a given temperature, and $T_w(t)$ represents the winding temperature at time t . A halving of insulation life for each 10°C rise in temperature corresponds to $\beta = 0.069$ in this equation.

The empirical model of expected winding lifetime as a function of temperature will be shown to play an illustrative role in the analyses presented here. Note that, if the reader is interested in implementing motor winding insulation diagnostics and prognosis, more sophisticated fault detection and fault progression models can and should be explored.

Each of the parameters in the fault model defined in Eqn. 4.1 will be viewed here as random variables with priors defined based on available empirical and/or heuristic knowledge of the winding degradation process. The value of α in the model will depend on the grade and thickness of the materials used to insulate the motor windings. The α value for a particular winding could be estimated experimentally by running several windings at a fixed temperature and observing the distribution of realized motor winding lifetime.

Motor health is estimated at a given time, t_p , by accumulating the portion motor health that was expended during past operation.

$$\gamma(t_p) = \gamma(0) - 100 \cdot \int_0^{t_p} \frac{1}{L_N(T_w(\tau))} d\tau \quad (4.2)$$

where $\gamma \in [0, 100]$ represents winding SOH percentage, T_w represents the average winding temperature, and $L_N(T_w)$ is given by the exponential model defined in Eqn. 4.1. Here, $\gamma = 100$ indicates perfect health, and $\gamma = 0$ indicates failure.

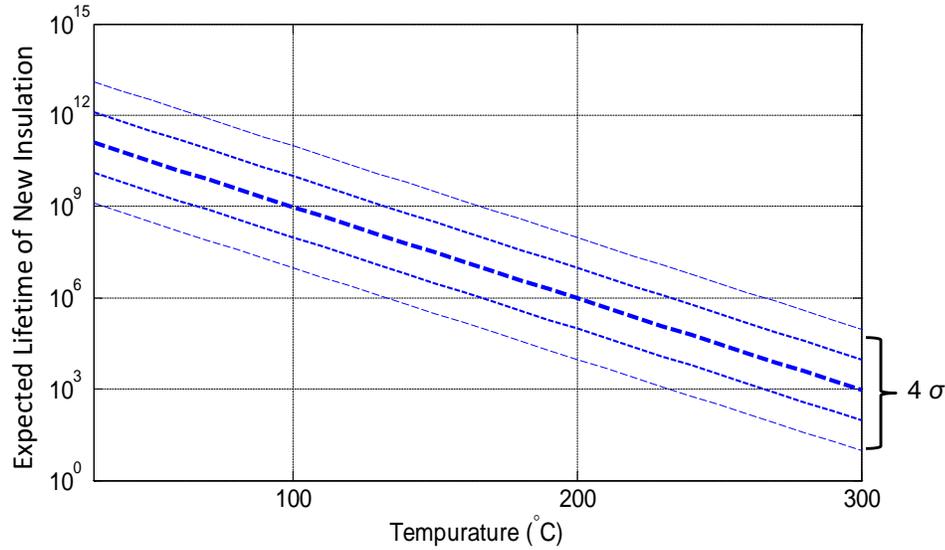


Figure 4.3: Addition of uncertainty to the insulation breakdown model

Figure 4.3 shows a probabilistic insulation life versus temperature estimate resulting from adding a probability distribution to the α coefficient in Eqn. 4.1. The value of the β coefficient in Eqn. 4.1 is set to $\beta = 0.069$. The probability distribution assumed for α in this example has a mean of $\alpha = 10^{11}$, and two standard deviations above and below the mean are given by $\alpha = (10^{12}, 10^{13})$ and $\alpha = (10^{10}, 10^9)$ respectively. Uncertain beliefs about the true value for β in the winding degradation model could similarly be incorporated into the winding lifetime model shown in Figure 4.3.

4.3 A Thermal Model

The thermal dynamics modeling for the DC motors used here is split into two parts. First, winding temperature is expressed as a function of the Ohmic power loss in motor windings. Second, the Ohmic power loss in the windings is expressed as a function of torque measured at the motor output shaft. Although the modeling presented here represents a fairly simplistic approach to representing winding thermo-electrical dynamics, the realism is deemed sufficient for the higher level control problem analysis that is the focus of this work.

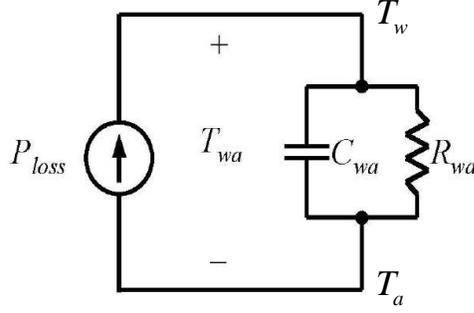


Figure 4.4: An equivalent circuit model for motor winding temperature dynamics

A first order equivalent circuit model [59], shown in Figure 4.4, is used to represent the thermo-electrical dynamics of the motor windings. Winding-to-ambient temperature is represented by T_{WA} in the model, C_{WA} and R_{WA} represent the thermal capacitance and thermal resistance between the motor winding temperature and the ambient temperature, and P_{loss} represents Ohmic power loss in a motor's windings.

Appendix B describes the derivation of the following ODE expression relating motor torque output to winding temperature dynamics

$$\dot{T}_W(t) = -\frac{T_W(t) - T_A}{R_{WA}C_{WA}} + \left(\frac{T_M(t)}{k_T}\right)^2 \frac{R_t}{C_{WA}} \quad (4.3)$$

T_W represents the temperature of motor windings. T_A represents the ambient temperature of the motor housing. R_{WA} and C_{WA} represent the thermal resistance and thermal capacitance respectively of the motor winding and motor housing. R_t represents the motor winding resistance. T_M represents the torque at a motor's output shaft, and k_T represents a coefficient of proportionality that relates winding current and motor torque.

4.4 Application Example - A Triplex Redundant Electromechanical Actuator

Performance and prognostic metrics are discussed and demonstrated here for a relatively easy to understand EMA example. The primary failure mechanism of the EMA is assumed

to be the degradation of motor windings. The EMA system considered here consists of three independently operated DC motors that are each geared to the output shaft of a linear actuator. The results and analysis generated here for a multiple input single output system will provide insight into the more general and more difficult to analyze case of multiple input multiple output (MIMO) systems.

4.4.1 Mechanical Dynamics Model

The torque allocated to each motor shaft in the mechanically coupled EMA motor arrangement is designated to be the ‘load’ variable that drives component fault dynamics and system output effector dynamics.

The EMA simulation model is described and its coefficients are defined in Appendix C. The mechanical dynamics of torque transmission in a three motor system are expressed by the following linear model:

$$\dot{x} = Ax + B_1u + B_2w \quad (4.4)$$

$$A = \begin{bmatrix} 1 & 0 \\ \frac{k_L}{J_L+3\lambda J_M} & \frac{-b_L-3b_M\lambda}{J_L+3\lambda J_M} \end{bmatrix}$$

$$B_1 = \frac{\lambda}{J_L+3\lambda J_M} \begin{bmatrix} 0 & 0 & 0 \\ 1 & 1 & 1 \end{bmatrix}, \quad B_2 = \frac{1}{J_L+3\lambda J_M} \begin{bmatrix} 0 \\ 1 \end{bmatrix} \quad (4.5)$$

$$x = \begin{bmatrix} \theta_L & \omega_L \end{bmatrix}^T, \quad u_c = \begin{bmatrix} T_1 & T_2 & T_3 \end{bmatrix}^T,$$

The state vector, x , in this model consist of the position, θ_L , and the rotational velocity, ω_L , of the EMA output shaft. Two control input vectors are considered in the model; u represents the three independently assignable motor torques, (T_1, T_2, T_3) , and w represents a non-controllable and time-varying environmental loading on the output shaft. The model uses only Coulomb damping to model friction. Damping forces measured at the individual motor output shafts, and at the actuator output shaft are represented by b_M and b_L . The k_L term represents an elastic force that drives the output shaft to an equilibrium position of

$\theta_L = 0$. The inertias of the three constituent EMA motors, and the inertia of the load driven by output shaft are represented by the terms J_M and J_L respectively. The λ term represents a gearing ratio between the motor shafts and the output shaft.

4.4.2 Performance and Prognostics Based Load Allocation Constraints

The performance constraint used in simulation studies of the EMA system is:

$$0.8 \cdot T_c \leq T_o \leq T_c \quad (4.6)$$

where T_o represents the sum of the shaft torques supplied by the each of the three constituent motors, and T_c represents the summed output torque that is required at the current time instant to satisfy the nominal performance expectation for the system.

The net output torque allocated at each control time-step is expressed in terms of the performance metric ρ , introduced in Chapter 3, Section 3.2.2 as:

$$T_o = \rho \cdot T_c \quad (4.7)$$

The performance constraint given in Eqn. 4.6 is expressed in terms of the performance metric as:

$$0.8 \leq \rho \leq 1 \quad (4.8)$$

As shown in the following equation, making a selection for ρ and for the torque loads on two of the three motors, determines the torque load on the third EMA motor;

$$T_3 = \rho \cdot T_c - T_1 - T_2 \quad (4.9)$$

A uniformly quantized component load allocation space satisfying the performance constraint is illustrated in Figure 4.5.

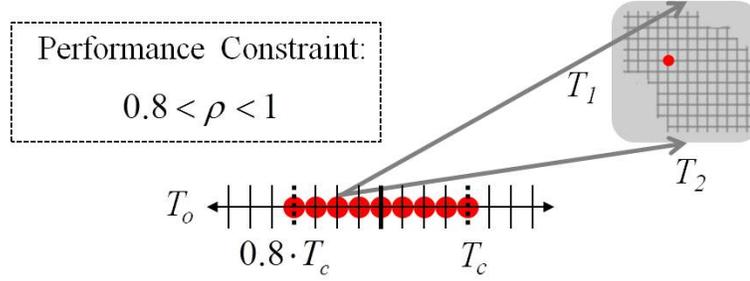


Figure 4.5: Illustration of the feasible load allocation domain for EMA motors under the performance constraint $0.8 \cdot T_c \leq T_o \leq T_c$

The prognostic constraint used in simulation studies of the EMA system is:

$$VaR_{.98}(\gamma_i(t_M)) > 10\% \quad (4.10)$$

where t_M represents the end of mission (EOM) time and $\gamma_i(t)$ represents the state of health (SOH) of motor i at time-index t . The SOH at EOM constraint mandates that there be a 98% confidence that each of the system's three motors have a health of greater than 10% at t_M .

As discussed in Chapter 3, Section 3.3.2, the growth of uncertainty over long prognostic horizons can lead to SOH at EOM predictions with very large uncertainties. It may be more useful in such cases to perform fault growth prediction only up to a finite prognostic horizon, at which the estimation uncertainty is more manageable.

A finite horizon prognostic value at risk constraint is expressed as:

$$VaR_{\beta_i}(\gamma_i(t_p + \tau)) > \zeta_i(t_p + \tau) \quad (4.11)$$

Two sample definitions for the finite horizon prognostic threshold, ζ_i , are shown in Figure 4.6. The illustration in Figure 4.6a shows the use of linear interpolation to derive a finite horizon prognostic constraint from an SOH at EOM constraint. The illustration shown in Figure 4.6b illustrates the use of a similar piecewise linear interpolation process. A 98%

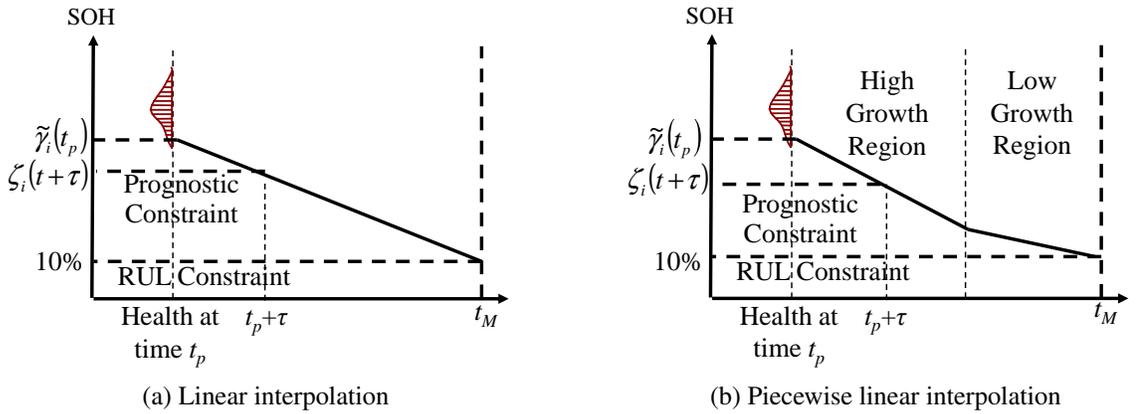


Figure 4.6: Illustrations of finite horizon prognostic constraint evaluation using linear (a) and piecewise linear (b) interpolation from a state of health at end of mission constraint

confidence bound is chosen for the VaR measure in both cases in order to match that chosen for the SOH at EOM constraint defined in Eqn. 4.10. The 98% confidence VaR measure at the time of prediction, t_p , is denoted by $\tilde{\gamma}_i$ in both figures.

The use of a piecewise linear prognostic constraint potentially allows components to degrade more rapidly during high demand regions with knowledge that the degradation will be mandated to be much slower during low demand regions. The simulation results presented in Section 4.4.4 demonstrate this flexibility.

4.4.3 Quantification of Performance and Risk Metrics

The performance and risk metrics to be optimized online by supervisory component load allocation routines will penalize the predicted deviation from a commanded EMA output profile, and the predicted damage to components over a defined prognostic horizon. A cost function of the following general form is considered here:

$$J = \int_{t_p}^{t_p+\tau} \rho(t) dt + \lambda \sum_{i=1}^3 p(\gamma_i(t_p + \tau) < \psi_i(t_p + \tau) | T_i(t_p)) \quad (4.12)$$

where the term, $\int_{t_p}^{t_p+\tau} \rho(t) dt$, penalizes the summation of induced output control effort deviations relative to a given nominal control policy. The term $p(\gamma_i(t + \tau) < \psi_i(t + \tau) | T_i(t))$

is used to assign a cost to the probability that a torque, T_i , on motor i will result in that motor's SOH dropping below the threshold $\psi_i(t_p + \tau)$ at time $t_p + \tau$. The variable λ is used to weight the performance and prognostic penalty terms in the expression.

The future loads driving fault growth over a defined prognostic horizon are assumed, in this case, to remain constant over the interval $(t_p, t_p + \tau)$. However, motor torque allocations will in fact be subject to change at each control time-step, as the output demands on the system change.

4.4.4 Simulation Results

An exhaustive search is used to optimize the performance evaluation and prognostic risk metrics defined in Section 4.4.3 over a quantized domain of allowable load allocations. The simulation results presented in this section are used to explore the effect of various cost function definitions on the resulting optimal component load allocation policies. The EMA output speed profile given in Figure 4.7 is used as the sample mission to be optimized over. The coefficients that define the thermal model for EMA motor windings are given in Table 4.1.

Each of the three EMA motors will be initialized at different states of health in the simulation studies presented here in order to illustrate the challenge of controlling several similar

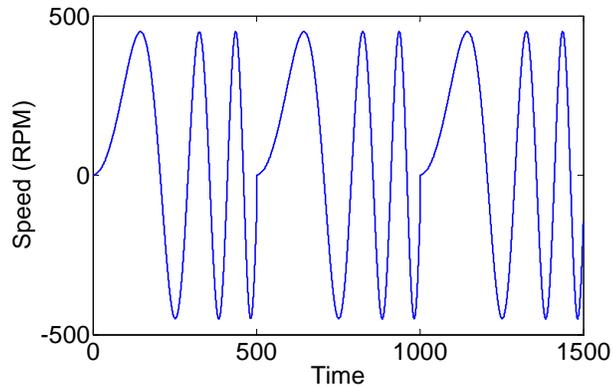


Figure 4.7: A sample EMA head speed profile command

Table 4.1: Parameter definitions for the EMA motor winding temperature dynamics model

Sym	Description	Units	Value
R_{wa}	Thermal resistance	Ω	1.1
C_{wa}	Thermal capacitance	farad	100
T_a	Ambient Motor Temperature	$^{\circ}C$	30

components at different states of health. The mean SOH estimates for the three EMA motors are set to begin simulated missions at 95%, 65%, and 45% SOH for motors 1, 2, and 3 respectively. Uncertainty on initial SOH estimates is represented by placing the first two standard deviations of an SOH prior distribution at a distance of $\pm 2.5\%$ and $\pm 5\%$ above and below the mean SOH estimate. The degradation of motor winding health is simulated for the mean and the ± 2 standard deviation points in the motor winding SOH estimate pdf using the winding lifetime model described in Section 4.2. The trajectories of the motor winding health estimate mean and ± 2 standard deviations are simulated using $\beta = 0.025$ and $\alpha = (10^5 \cdot 2, 10^5 \cdot 1.5, 10^5, 10^5/1.5, 10^5/2)$.

The simulation example presented is designed such that the risk of motor failure will be unacceptably high for a system under nominal control, while some solutions that lie within the allowed output performance modification domain are shown to be sufficiently safe. The optimization of load allocation policies is examined for the EMA example, with special attention paid to the influence of the prognostic horizon length and the formulation of finite horizon prognostic constraints.

4.4.4.1 A Nominal Control Law

The EMA nominal output torque command is assumed to be updated continuously using the following proportional control law:

$$T_c(t) = \begin{cases} k \cdot (\omega_o(t) - \omega_c(t)) & \text{if } |k \cdot (\omega_o(t) - \omega_c(t))| \leq 170 \\ -170 & \text{if } k \cdot (\omega_o(t) - \omega_c(t)) < -170 \\ 170 & \text{if } k \cdot (\omega_o(t) - \omega_c(t)) > 170 \end{cases} \quad (4.13)$$

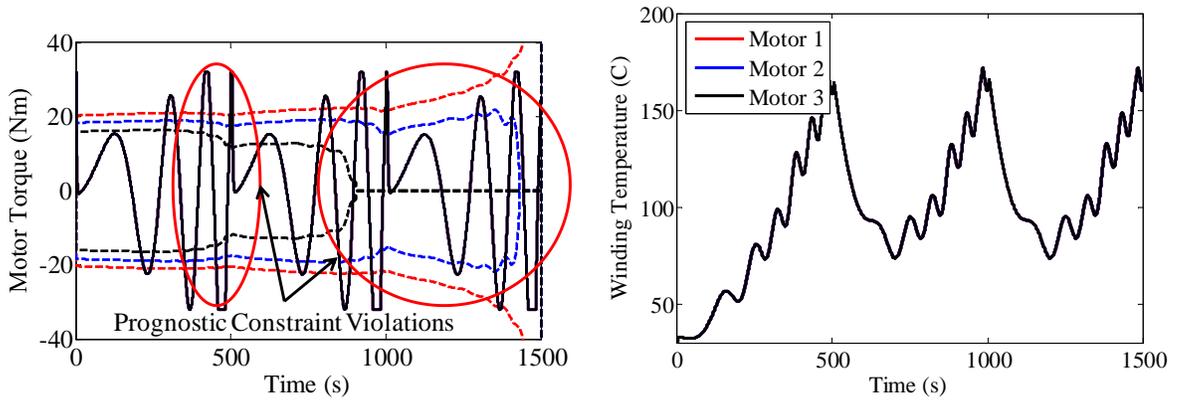
where $\omega_o(t)$ is the EMA output speed, ω_c is a commanded speed, and k is a proportional gain ($k = 15$ in simulation studies). The maximum allowable actuator torque command is shown to be clipped at a range of ± 170 Nm to avoid over stressing the system

The nominal motor load allocation law is assumed to split the commanded output load equally among all three motors:

$$\begin{bmatrix} T_1 \\ T_2 \\ T_3 \end{bmatrix} = T_c \cdot \begin{bmatrix} \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \end{bmatrix} \quad (4.14)$$

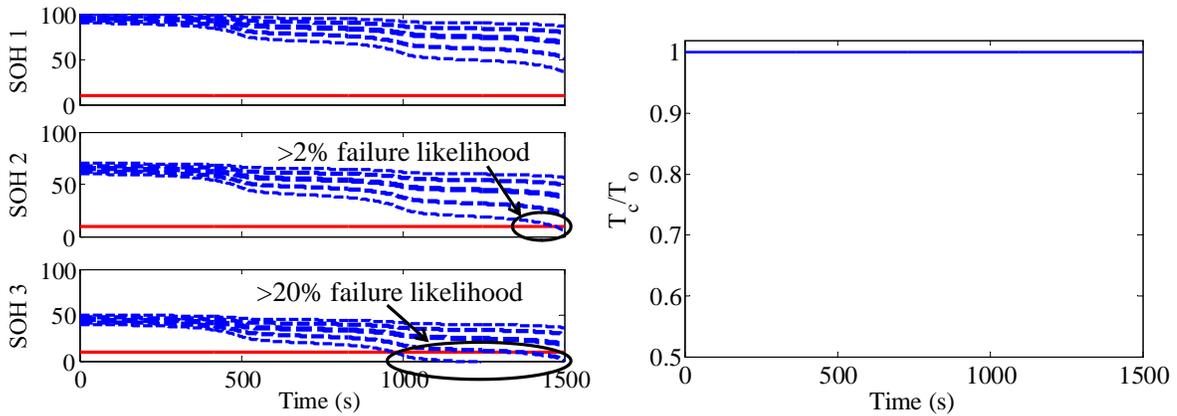
Figure 4.8 shows simulation results obtained using the nominal control law given in Eqns. 4.13 and 4.14 on the EMA thermal model and winding degradation models defined earlier.

The plot in Figure 4.8a shows the load allocated to each of the three EMA motors by the nominal control law, along with dashed lines that shown the maximum motor torque that can be sustained for the remainder of the mission without violating the SOH at EOM constraint. The plot in Figure 4.8b shows the winding temperatures of the three EMA motors. The simulated evolution of probabilistic winding insulation SOH estimates as a function of motor temperature is given Figure 4.8c. The supervisory output control effort modifier, ρ , is plotted in Figure 4.8d, ($\rho = 1$ for nominal control).



(a) Motor torques and estimated bounds imposed by the prognostic constraint

(b) Motor winding temperatures



(c) Estimated winding insulation SOH pdfs (the lines represent the mean and ± 2 standard deviations)

(d) Profile for the supervisory output control effort modifier, $\rho = T_c/T_o$

Figure 4.8: Simulation results for an EMA under nominal control

The simulation results show that using the nominal control law to follow the sample EMA output speed profile given in Figure 4.7 is predicted to result in the violation of the SOH at EOM constraint given in Eqn. 4.10 for motors 2 and 3. The regions of constraint violation are circled in the predicted motor winding insulation degradation plots given in Figure 4.8c.

4.4.4.2 Verification of Performance and SOH at EOM Constraint Feasibility

As discussed in Section 3.3.3, the feasibility of the commanded speed profile is checked by verifying that at least one load allocation policy satisfies the given performance and SOH at EOM constraints. The existence of at least one feasible control provides proof that

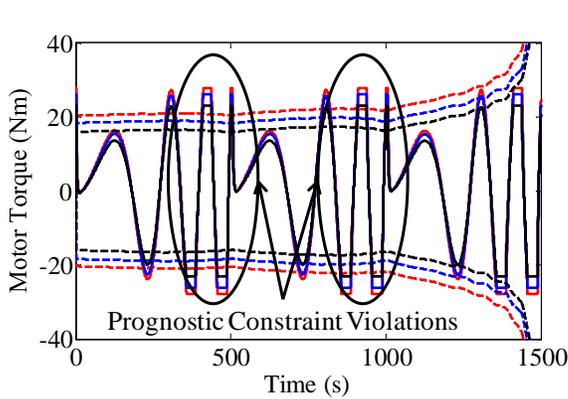
the domain to be optimized over by risk management controllers is non-empty. Figure 4.9 shows simulation results obtained using the minimum allowable value of the performance metric, $\rho = 0.8$, as defined by the performance constraint given in Eqn. 4.8.

The ratio of the EMA net output load demand, T_o , allocated to each of the three EMA motors is set to the following constants, identified experimentally:

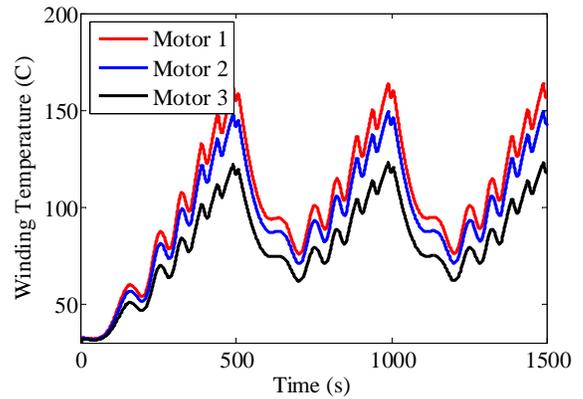
$$\begin{bmatrix} T_1 \\ T_2 \\ T_3 \end{bmatrix} = T_o \cdot \begin{bmatrix} \frac{9}{25} \\ \frac{17}{50} \\ \frac{3}{10} \end{bmatrix} \quad (4.15)$$

The plot in Figure 4.9a shows the load allocated to each of the three EMA motors by the sample minimum output performance control law. Dashed lines in Figure 4.9a show the maximum motor torque that can be sustained for the remainder of the mission without violating the SOH at EOM constraint. The plot in Figure 4.9b shows the winding temperatures of the three EMA motors. The simulated evolution of probabilistic winding insulation SOH estimates as a function of motor temperature is given Figure 4.9c. The supervisory output control effort modifier, ρ , is plotted in Figure 4.9d, ($\rho = 0.8$ for minimum allowable output performance).

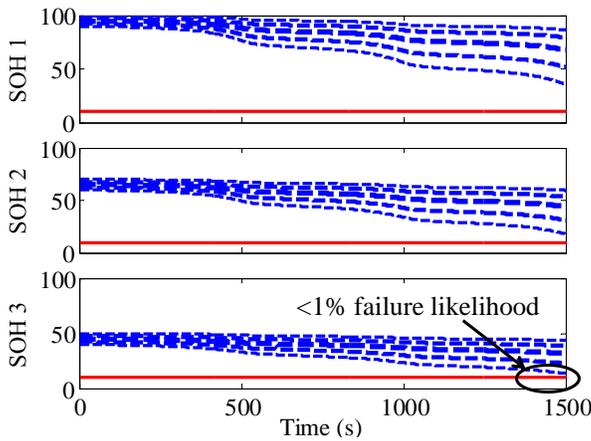
The plot in Figure 4.9c shows that there is greater than 99% confidence that no motor will be below 10% SOH at the EOM. The SOH at EOM constraint is thus shown to be satisfiable for some set of component load allocation policies that also satisfy the net output performance constraint. Note however that the estimated maximum allowable sustained motor torque allocations, shown with the dashed lines in Figure 4.9a, are repeatedly violated using the sample policy. This indicates that the maximum allowable sustained motor torque constraint is overly conservative and more clever prognostic constraints will need to be devised to account for the future load variability observed in the sample mission.



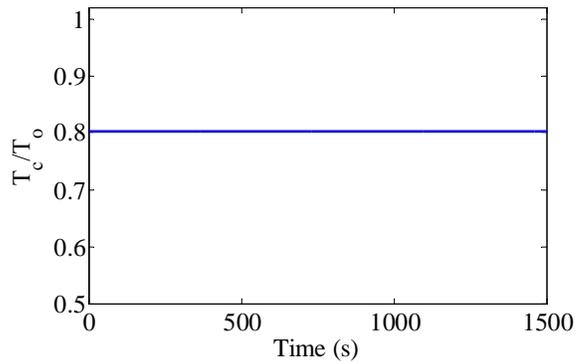
(a) Motor torques and estimated bounds imposed by the prognostic constraint



(b) Motor winding temperatures



(c) Estimated winding insulation SOH pdfs (the lines represent the mean and ± 2 standard deviations)



(d) Profile for the supervisory output control effort modifier, $\rho = T_c/T_o$

Figure 4.9: Simulation results for an EMA following the minimum allowable performance path

4.4.4.3 Optimization of Performance and Risk Metrics within a Bounded Load Allocation

Domain

As described in Section 4.4.2, the EMA output regulation task has one degree of freedom, the choice of ρ , and values for the remaining two degrees of freedom in the problem are specified by the component load allocation task. Given a choice of ρ at each control time-step, the fault risk optimization problem to be solved by the component load allocation

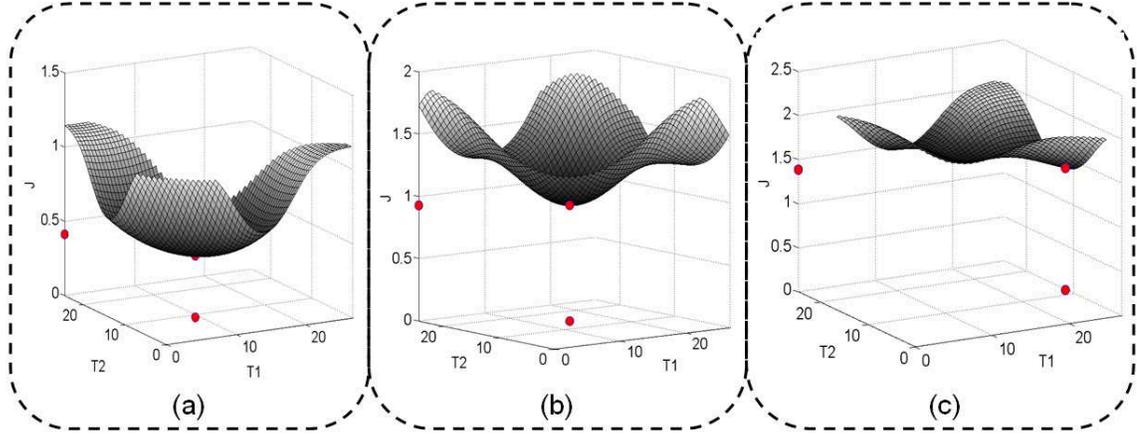


Figure 4.10: Plots of the parallel distribution cost function using $\rho \cdot T_c = 0.4\tilde{T}$ (a), $\rho \cdot T_c = 0.6\tilde{T}$ (b), and $\rho \cdot T_c = 0.7\tilde{T}$ (c)

routine is:

$$\min_{T_1, T_2} \sum_{i=1}^3 p(\gamma_i(t_p + \tau) < \psi_i(t_p + \tau) | T_i), \text{ where } T_3 = \rho \cdot T_c - T_1 - T_2 \quad (4.16)$$

Figure 4.10 shows plots of the cost function given in Eqn. 4.16 for low, moderate, and high load levels. Here, low, moderate, and high load levels are defined as a proportion of the maximum output torque allowed under the prognostic constraint, denoted \tilde{T} . Low, moderate, and high load levels are defined as $0.4\tilde{T}$, $0.6\tilde{T}$, and $0.7\tilde{T}$ respectively in this example.

As shown in Figure 4.10, the optimal prognostic cost is monotonically increasing with system load. In this section, both the output regulation and component load allocation tasks are performed using exhaustive searching on a uniform grid over the space of all feasible motor load allocations satisfying the performance constraints. This somewhat naive approach to global optimization was successful because of the linearity of the motor's thermal and degradation models, and the efficiency of arithmetic on large matrices in MATLAB.

Simulation results are shown in Figure 4.11 for the control policy that is found to optimize the receding horizon risk management problem defined in Eqn. 4.12 at 0.1 second control

time-steps. The prognostic horizon is set to $\tau = 50$ seconds, and the weighting term in Eqn. 4.12 is set to $\lambda = 25000$. Finite horizon prognostic constraints are defined using linear interpolation, as described in Section 4.4.2. The risk metric $\psi_i(t_p + \tau)$ is defined to be the average of the current 98% confidence VaR estimate of SOH, and the prognostic constraint at the end of the prognostic horizon.

The plot in Figure 4.11a shows the component load allocations that optimize the given risk metric at each control time-step over a simulated mission. The dashed lines in Figure 4.11a represent the maximum motor torque allocation allowable under the prognostic constraint evaluated at the prognostic horizon. The plot in Figure 4.11b shows the winding temperatures of the three EMA motors. The simulated evolution of probabilistic winding insulation SOH estimates as a function of motor temperature is given Figure 4.11c. The supervisory output control effort modifier, ρ , is plotted in Figure 4.11d.

Note that the finite horizon prognostic constraint shown with the dashed lines in Figure 4.11a is seen to tighten and loosen as the motor windings warm up and cool down. This is an effect of the prognostic horizon ($\tau = 50$ sec) being small with respect the time constant in the thermal dynamics of the motor winding. If a longer prognostic horizon were used then the increased uncertainty of prognostic predictions and the decreased significance of the thermal capacitance's effect on winding temperatures would lead to more conservative prognostic bounds with less variation over the mission.

The plots in Figure 4.11c show that the control policy used is estimated to result in a sufficiently low risk of any motor windings having less than 10% SOH at EOM. However, the plot of the output performance modifier, ρ , given in Figure 4.11d, shows that enforcing the finite horizon prognostic constraint at each control time-step forces the performance constraint to be violated during the more aggressive regions in the commanded load speed profile.

It was proven in Section 4.4.4.2 that a set of motor load allocation policies exists such

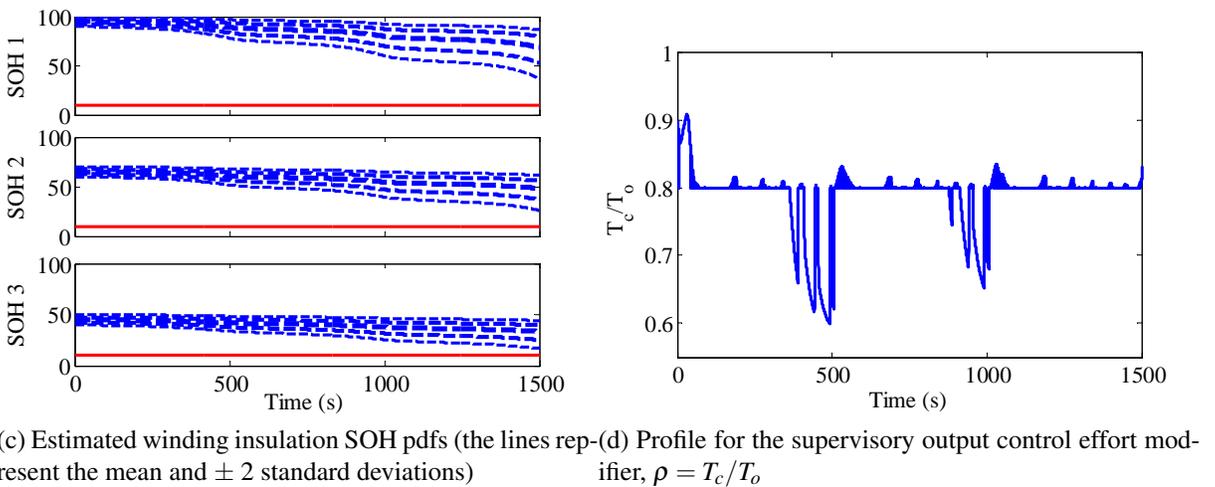
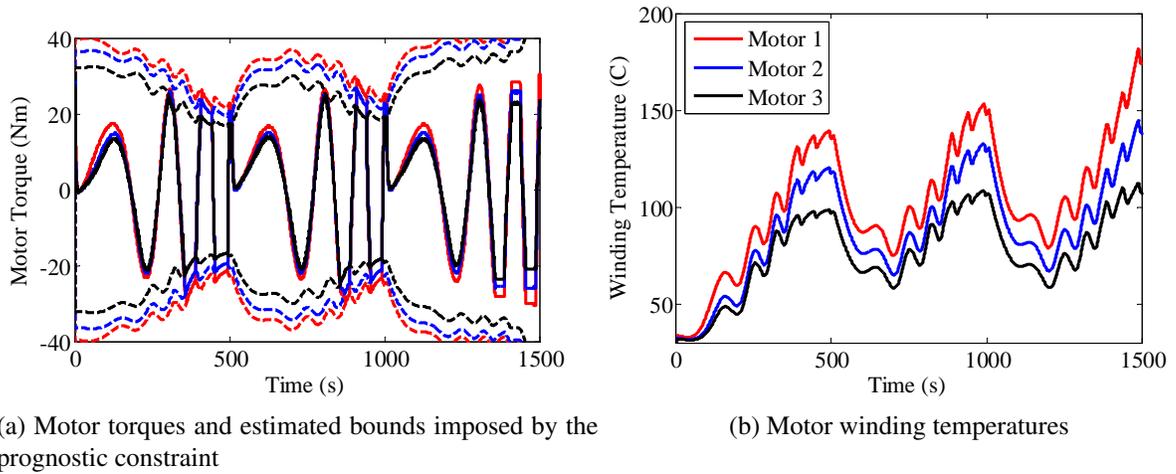


Figure 4.11: Simulation results using $\tau = 50$ sec, $\lambda = 25000$

that there is $> 98\%$ confidence that all motors have $> 10\%$ SOH at the EOM, and the performance constraint is never violated. Therefore there must exist some more clever definition of the finite horizon prognostic constraint that will result in feasible component load allocation policies.

Consideration is given here to how finite horizon prognostic constraint may be tuned based on foreknowledge of the alternating periods of high and low load in the sample mission. The maximum allowable fault growth curve was made steeper during the more aggressive parts of a mission, and less steep during the less aggressive parts of a mission.

Figure 4.12 shows simulation results for a control policy that was obtained using the same τ and λ values as in Figure 4.11, with a piecewise linear definition of the prognostic constraint. Here, the slope of the prognostic constraint, $\zeta_i(t)$, was defined to be twice as steep during the final 200 seconds than during the first 300 seconds of the repeating actuator head speed command profile, shown in Figure 4.7. The plots shown in Figure 4.12c and 4.12d prove that the use of this type of piecewise linear definition of $\zeta_i(t)$ allows controls to be found that do not violate the finite horizon performance and prognostic constraints at any time over the sample mission.

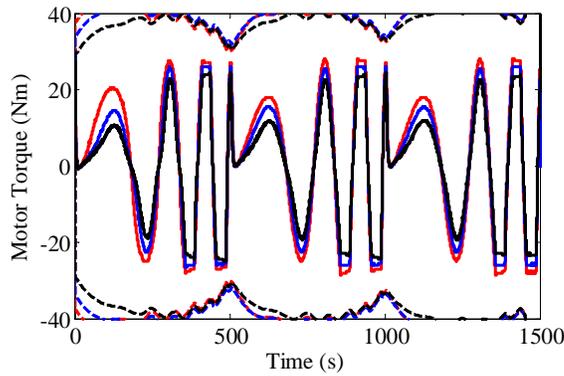
The simulation studies presented in the next section will present a simulation based exploration of how the τ and λ terms in the risk quantification function may be tuned based on repeated simulation trials.

4.5 Application Example - A Skid-Steered Vehicle

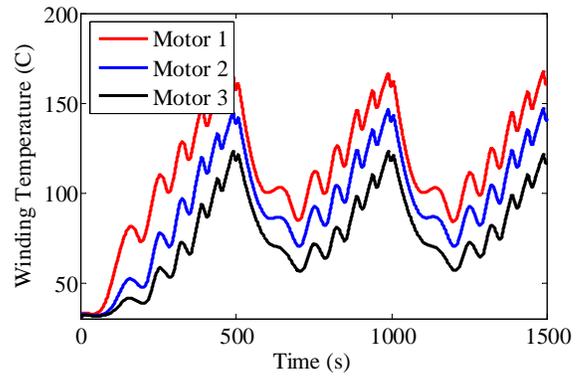
This application example explores the allocation of load among four independent motors that each power one of the four wheels on a skid-steered unmanned ground vehicle (UGV). The overactuation present in this application example is more complex than the simple active redundant configuration explored in the EMA example. Here, rather than having the motors directly geared together, as was the case in the EMA, the four motors in the UGV are linked through their mutual contact with the ground.

4.5.1 Mechanical Dynamics Model

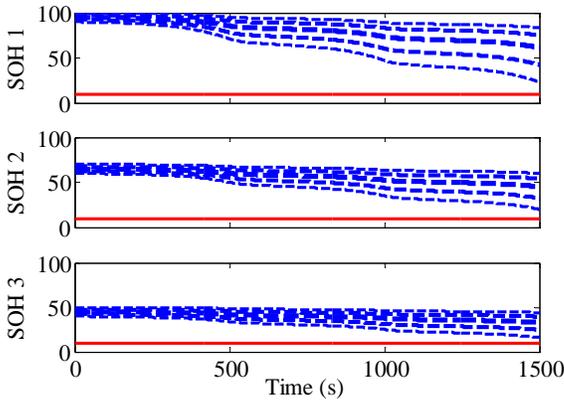
An illustration of the motor torque contributions to vehicle locomotion is given in Figure 4.13. The skid-steered vehicle model and coefficients are defined in Appendix D. Assuming that all of the robot's wheels are getting approximately the same traction, then the skid-steered wheeled vehicle will behave much like a treaded vehicle [96]. In the simulation studies presented here, the skid-steered vehicle's modeling is simplified by treating it as a treaded vehicle and making other simplifying assumptions.



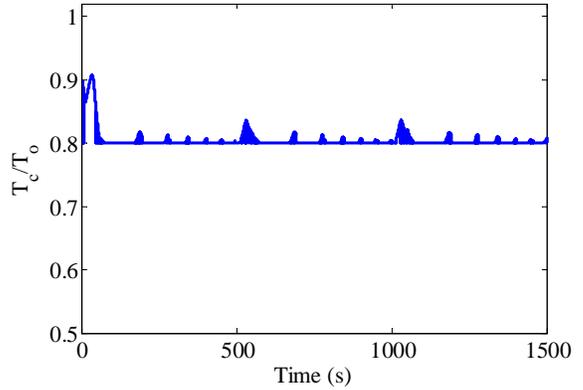
(a) Motor torques and estimated bounds imposed by the prognostic constraint



(b) Motor winding temperatures



(c) Estimated winding insulation degradation pdfs (the lines represent the mean and ± 2 standard deviations)



(d) Profile for the supervisory output control effort modifier, $\rho = T_c/T_o$

Figure 4.12: Simulation results using a piecewise linear constraint

The simple model used here is adequate to serve our purposes of exploring some of the conceptual and notation complexities associated with control of MIMO systems. However, the model used in the simulation studies presented here may be overly simplistic for some applications, especially for applications that rely on dead reckoning for self location.

The skid-steered vehicle model used here expresses vehicle dynamics in terms of wheel

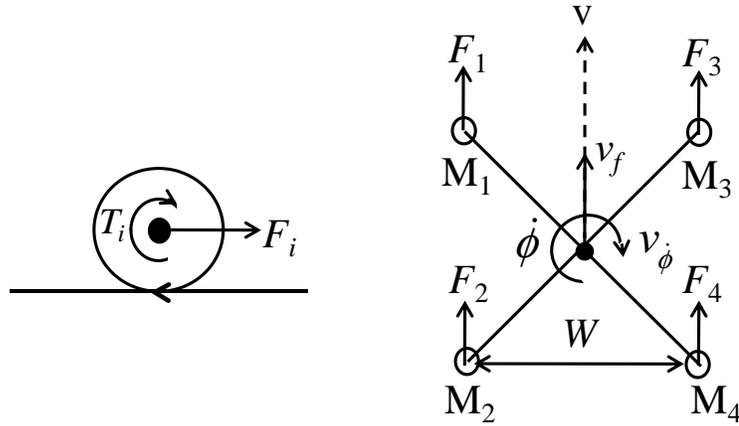


Figure 4.13: Visualization of motor torque contributions to system locomotion on a four wheeled skid-steered vehicle

speeds and motor torque outputs with the following ODE expression:

$$\begin{bmatrix} \dot{\omega}_L \\ \dot{\omega}_R \end{bmatrix} = \begin{bmatrix} \frac{mr^2}{4} + \frac{r^2 I}{\alpha W^2} & \frac{mr^2}{4} - \frac{r^2 I}{\alpha W^2} \\ \frac{mr^2}{4} - \frac{r^2 I}{\alpha W^2} & \frac{mr^2}{4} + \frac{r^2 I}{\alpha W^2} \end{bmatrix}^{-1} \left(-\frac{k_r}{2} \begin{bmatrix} \omega_L \\ \omega_R \end{bmatrix} + \begin{bmatrix} T_1 + T_2 \\ T_3 + T_4 \end{bmatrix} \right) \quad (4.17)$$

$$\begin{bmatrix} v \\ \dot{\phi} \end{bmatrix} = \begin{bmatrix} \frac{r}{2} & \frac{r}{2} \\ \frac{-r}{\alpha W} & \frac{r}{\alpha W} \end{bmatrix} \begin{bmatrix} \dot{\omega}_L \\ \dot{\omega}_R \end{bmatrix}$$

The coefficients used in this model are defined in Table 4.2.

The model has four input degrees of freedom (the torques assigned to the four motors), and two output degrees of freedom (the vehicle's linear speed and angular velocity).

4.5.2 Performance and Prognostics Based Load Allocation Constraints

The performance constraint used in simulation studies of the skid-steered vehicle system is:

$$\begin{aligned} 0.8 \cdot r_L &\leq T_L \leq r_L \\ 0.8 \cdot r_R &\leq T_R \leq r_R \end{aligned} \quad (4.18)$$

Table 4.2: Parameter definitions for the skid-steered vehicle model

Symbol	Description	Units	Value
r	Wheel radius	m	0.1
W	Vehicle width	m	0.5
L	Vehicle length	m	0.5
I	Wheel rotational inertia	kg·m ²	0.1
m	Vehicle mass	kg	1
k_r	Rolling resistance term	-	.19
α	Longitudinal slip during turn	-	2
v	Linear vehicle speed	m/s	-
ϕ	Vehicle angular velocity	rad/s	-
ω_L	Left side wheel speeds	rad/s	-
ω_R	Right side wheel speeds	rad/s	-
T_i	Torque allocated to motor i	N/m	-

where r_L and r_R represent a desired net control effort output from the left-hand and right-hand motors respectively. T_L and T_R represent the net torque output of the left-hand and right-hand motors respectively,

$$T_L = T_1 + T_2, \quad T_R = T_3 + T_4 \quad (4.19)$$

The desired net effector output will be determined at each control time-step by a nominal control law that will be defined in Section 4.5.4. The net motor output torque is expressed in terms of the performance metric ρ , introduced in Chapter 3, Section 3.2.2, as:

$$\begin{bmatrix} T_L \\ T_R \end{bmatrix} = \begin{bmatrix} \rho_1 & \rho_2 \end{bmatrix}^T \begin{bmatrix} r_L \\ r_R \end{bmatrix} \quad (4.20)$$

The performance constraint given in Eqn. 4.18 is expressed in terms of the performance metric as:

$$0.8 \leq \rho_i \leq 1 \quad (4.21)$$

The EOM damage constraint used in simulation studies of the skid-steered vehicle is the same as was used for the EMA example:

$$VaR_{98}(\gamma_i(t_M)) > 10\% \quad (4.22)$$

where t_M represents EOM time and $\gamma_i(t)$ represents the SOH of motor i at time-index t . This SOH at EOM constraint mandates that there be a 98% confidence that each of the system's motors have a health of greater than %10 at t_M .

The finite horizon prognostic value at risk constraint is the same as was used in the EMA example

$$VaR_{98}(\gamma_i(t_p + \tau)) > \zeta_i(t_p + \tau) \quad (4.23)$$

Linear interpolation is used here to enforce SOH at EOM constraints at the prognostic horizon, just as was done in the EMA example.

4.5.3 Quantification of Performance and Risk Metrics

An expression used to evaluate preference over a space of various possible system output profiles and component damage estimates at EOM is defined for the skid-steered vehicle simulation study as:

$$J_M(\phi_c(t) - \phi(t)) + J_d(VaR_{.98}\gamma_i(t_M)) \quad (4.24)$$

where $\phi_c(t)$ and $\phi(t)$ represent points on the desired and the actual path followed by the skid-steered vehicle respectively, and $VaR_{.98}\gamma_i(t_M)$ represents the 98% confidence VaR estimate of the SOH for motor i at the EOM. J_M and J_d represent functions that penalize vehicle deviations from a desired path and ending motor damage estimates respectively.

The path deviation penalty is defined for this example as:

$$J_M(\phi_c(t) - \phi(t)) = \frac{1}{t_M} \int_0^{t_M} \exp|\phi_c(t) - \phi(t)| \quad (4.25)$$

The ending motor damage penalty is defined as:

$$J_d(\tilde{d}_i(t_M)) = \frac{4}{3} \cdot \max_i [\exp(100 - \text{VaR}_{.98} \gamma_i(t_M))] \quad (4.26)$$

Implemented component load allocation routines will assign risk to uncertain predictions of future system dynamics using performance and prognostic penalties;

$$J_P(\rho(t_p, \tau)) + \lambda \cdot J_f(\text{VaR}_{.98} \gamma_i(t_p + \tau)) \quad (4.27)$$

where J_P represents a performance penalty, J_f represents a penalty on future damage estimates, and λ is a weighting coefficient, representing the relative value of maximizing performance and minimizing component degradations

The suggested formulation for the performance penalty is:

$$J_P(\rho(t_p, \tau)) = \int_{t_p}^{t_p + \tau} \exp \left(\left\| \begin{bmatrix} r_f - \rho \cdot r_f \\ r_\phi - \rho \cdot r_\phi \end{bmatrix} \right\| \right) dz \quad (4.28)$$

Here, r_f and r_ϕ represent the net forward torque and the net turning torque requested by a nominal control law;

$$r_f = T_L + T_R \quad (4.29)$$

$$r_\phi = T_L - T_R \quad (4.30)$$

The suggested formulation for the prognostic penalty is:

$$J_f(\text{VaR}_{.98} \gamma_i(t_p + \tau)) = \sum_{i=1}^4 [\exp(\text{VaR}_{.98} \gamma_i(t_p + \tau) - \zeta_i(t_p + \tau))] \quad (4.31)$$

The simulation studies presented in the following section explore the effect that varying the τ and λ terms in the risk assessment metric defined in Eqns. 4.27 - 4.31, has on the outcome evaluation function defined in Eqns. 4.24 - 4.26.

Table 4.3: Parameter definitions for the skid-steered vehicle motor winding temperature dynamics model

Sym	Description	Units	Value
R_{wa}	Thermal resistance	Ω	0.75
C_{wa}	Thermal capacitance	farad	20
T_a	Ambient Motor Temperature	$^{\circ}C$	30

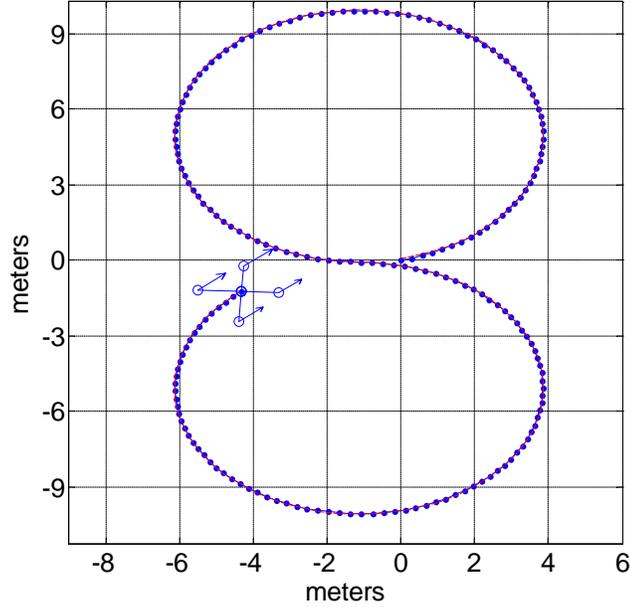


Figure 4.14: Sample skid-steered vehicle path command

4.5.4 Simulation Results

The simulation results presented in this section show the optimal load allocation among the four motors on a skid-steered vehicle following the figure-eight path shown in Figure 4.14. The coefficients that define the thermal model for the vehicle motors are given in Table 4.3.

Component load allocations for best risk management are found at each time-step by evaluating the objective function defined in Eqn. 4.27 on a sufficiently dense uniform grid over the space of all component load allocations satisfying the performance constraints.

The two motors on each side of the skid-steered vehicle are initialized to different SOH in the simulation studies presented here. The mean SOH estimate at the beginning of

simulated missions is set to 95% for vehicle motors 1 and 3, and it is set to 65% for motors 2 and 4. Uncertainty on initial SOH estimates is represented here by placing the first two standard deviations of an SOH prior distribution at a distance of $\pm 2.5\%$ and $\pm 5\%$ above and below the mean SOH estimate.

The degradation of motor winding health is simulated for the mean and the ± 2 standard deviation points in the motor winding SOH estimate pdf using the winding lifetime model described in Section 4.2. The trajectories of the motor winding health estimate mean and ± 2 standard deviations are simulated using $\beta = 0.035$ and $\alpha = (10^4 \cdot 7, 10^4 \cdot 8.5, 10^5, 10^5 \cdot 1.25, 10^5 \cdot 1.5)$.

4.5.4.1 A Nominal Control Law

The net output control effort commanded of the skid-steered vehicle is updated continuously over the mission using the following proportional control law:

$$\begin{bmatrix} r_f \\ r_\phi \end{bmatrix} = r_{\text{ref}}(t) + \begin{bmatrix} p_1 \cdot \cos(|\phi_e|) \cdot e_d \\ p_2 \cdot \sin(|\phi_e|) \end{bmatrix} \quad (4.32)$$

where $r_{\text{ref}}(t)$ is output control effort that would be used at time t if the vehicle followed the reference path exactly, p_i are the proportional control coefficients, ϕ_e is the vehicle's heading error with respect to the reference path, and e_d is the vehicle's position error with respect to the reference path.

The nominal control law defined in Eqn. 4.32 is assumed to split motor load equally among the two motors on each side of the vehicle:

$$\begin{bmatrix} T_1 \\ T_2 \\ T_3 \\ T_4 \end{bmatrix} = \begin{bmatrix} \frac{1}{2} & 0 \\ \frac{1}{2} & 0 \\ 0 & \frac{1}{2} \\ 0 & \frac{1}{2} \end{bmatrix} \cdot \begin{bmatrix} r_L \\ r_R \end{bmatrix} \quad (4.33)$$

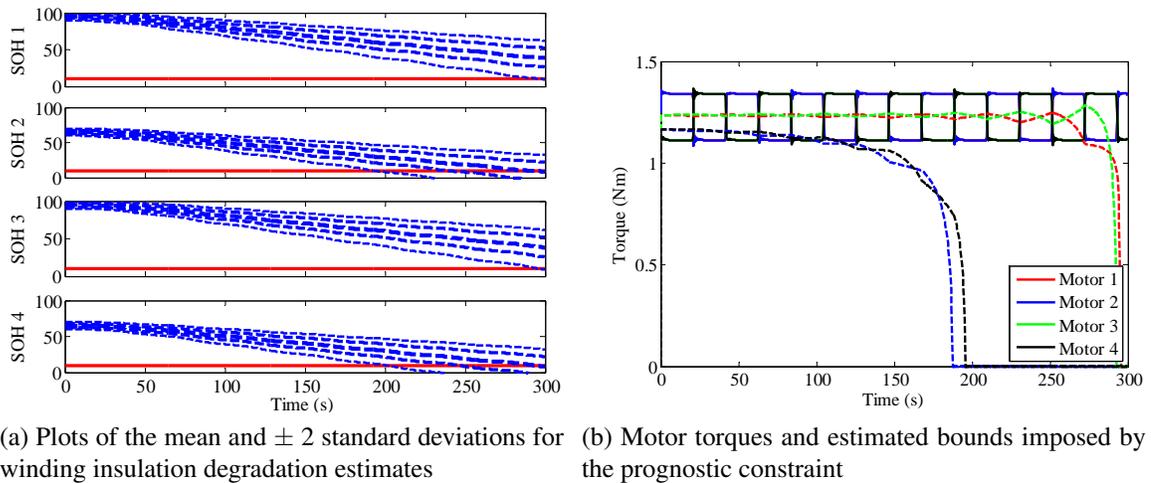


Figure 4.15: Simulation results for a skid-steered vehicle under nominal control

Simulation results for the nominal control policy are shown in Figure 4.15. The plot given in 4.15a shows the simulated evolution of probabilistic winding insulation SOH estimates as a function of winding temperature. Motor winding temperature is modeled as a function of motor torque loads using the thermal model described in Section 4.3. Figure 4.15b shows the load allocated to each vehicle motor, along with dashed lines showing estimates of the maximum motor torque that can be sustained over the mission without violating the SOH at EOM constraint. It is seen in Figure 4.15a that all motors are estimated to accrue an unacceptable probability of ending the mission with $< 10\%$ SOH if the nominal control law was used over the given figure-8 mission.

4.5.4.2 Verification of Performance and SOH at EOM Constraint Feasibility

The feasibility of given performance and prognostic constraints is verified by checking that the minimum allowable vehicle performance over the mission will allow all of the motors to end the mission with adequate health. Figure 4.16 shows simulation results obtained using the minimum allowable value of the performance metric, $\rho = [0.8, 0.8]$, as defined by the performance constraint given in Eqn. 4.21.

The load split between the two motors on each side of the vehicle is set to the following

constants, identified experimentally:

$$\begin{bmatrix} T_1 \\ T_2 \end{bmatrix} = T_L \cdot \begin{bmatrix} \frac{11}{20} & \frac{9}{20} \end{bmatrix}, \quad \begin{bmatrix} T_3 \\ T_4 \end{bmatrix} = T_R \cdot \begin{bmatrix} \frac{8}{15} & \frac{7}{15} \end{bmatrix} \quad (4.34)$$

The plot given in 4.16a shows the simulated evolution of probabilistic winding insulation SOH estimates as a function of winding temperature. Figure 4.15b shows the load allocated to each vehicle motor, along with dashed lines showing estimates of the maximum motor torque that can be sustained over the sample mission without violating the SOH at EOM constraint.

The plot in Figure 4.16a shows that there is greater than 98% confidence that no motor will be below 10% SOH at the EOM. The SOH at EOM constraint is thus shown to be satisfiable for some set of component load allocation policies that also satisfy the net output performance constraint. Note however that, as was also the case with the EMA example, the estimated maximum allowable sustained motor torque allocations, shown with the dashed lines in Figure 4.16a, are repeatedly violated using the sample policy. This indicates that the maximum allowable sustained motor torque constraint is overly conservative and more clever prognostic constraints will need to be devised to account for the future load variability observed in the sample mission.

4.5.4.3 Control with Foreknowledge of the Mission and the Fault Growth Model

Because the desired path for the skid-steered vehicle to follow and the fault growth model are known in advance in simulation studies, the optimal load allocations over the given mission can be found using future knowledge rather than prognostic information. Analysis of the direct optimization of the outcome evaluation function, formulated in Eqn.4.25, will provide substantial insight into the development of prognostics-based risk management controllers. The optimized control results obtained using future knowledge also provide a useful means of benchmarking control policies that use prognostic information.

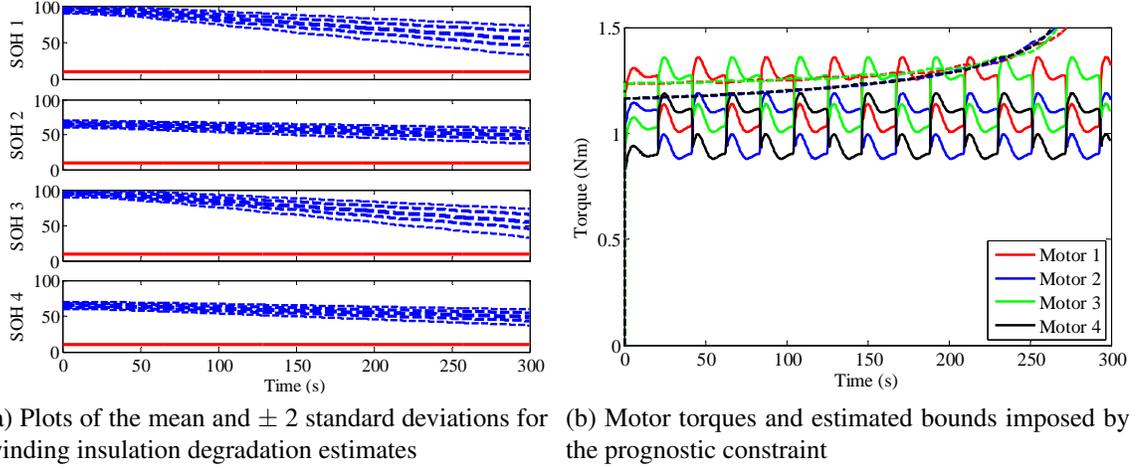


Figure 4.16: Simulation results for a skid-steered vehicle following the minimum allowable performance path

The search for optimizing control policy is formulated here as a nested process. First, a candidate path to be followed by skid-steered vehicle is defined by a set of waypoints and a third order spline that interpolates between those points. Then, an exhaustive search is used to identify the motor load split (over a quantized space) that minimizes the motor degradation cost assessed over the candidate path.

The space of allowable adjustments to a given set of waypoints is bounded by the performance constraint given in Eqn. 4.18. The net output control effort required to follow a given path is found by inverting the modeled skid-steered vehicle dynamics given in Eqn. 4.17,

$$\begin{bmatrix} T_L(t) \\ T_R(t) \end{bmatrix} = f^{-1}(\phi_p(t)), \quad \forall t \in [0, \dots, T] \quad (4.35)$$

where $\phi_p(t)$ is the (x,y) position of the skid-steered vehicle at time t . Gradient descent is used to search the space of feasible waypoint modifications.

Individual motor load allocations over the sample mission are derived using the following expression for splitting load proportionately among the two motors on each side of the

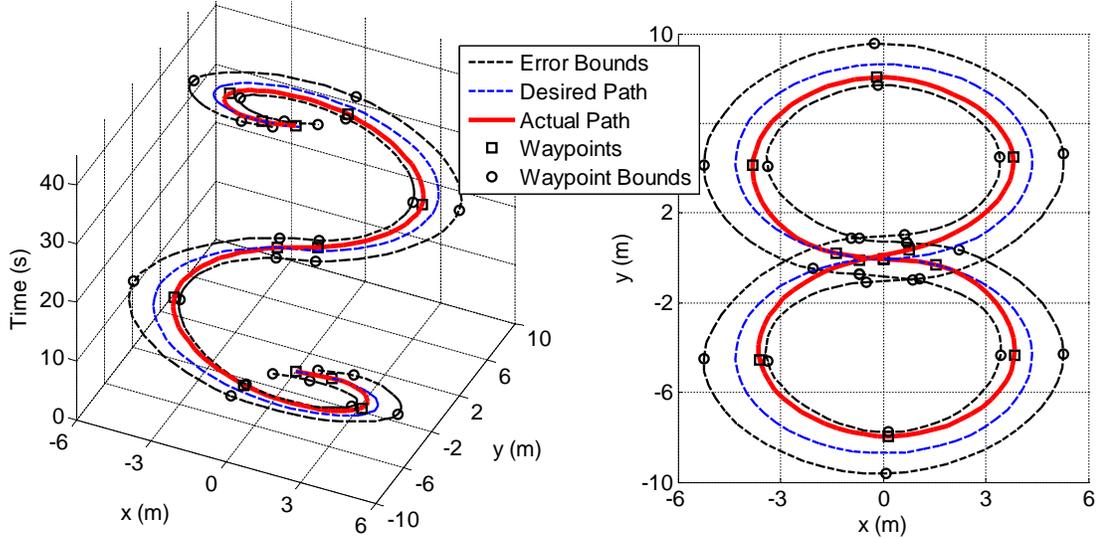


Figure 4.17: Results of a search for an optimal skid-steered vehicle load allocation policy over the repeated figure-8 path

vehicle:

$$\begin{aligned}
 T_1(t) \cdot k_1 &= T_2(t), & T_1(t) + T_2(t) &= T_L(t) \\
 T_3(t) \cdot k_2 &= T_4(t), & T_3(t) + T_4(t) &= T_R(t)
 \end{aligned}
 \tag{4.36}$$

Optimal motor load allocations for a given path are derived by evaluating Eqn. 4.25 over sufficiently dense uniform grid on $k_1, k_2 \in [0, 2]$, and selecting the value resulting in minimum cost.

Figure 4.17 shows plots of the desired skid-steered vehicle path, bounds on allowable path error, and an approximation of the optimal skid-steered vehicle path over one cycle of the commanded figure-8 maneuver. The vehicle path is optimized using a gradient descent search over the space of allowable adjustments to a set of eleven waypoints shown in the figure. The space of allowable adjustments to each waypoint is shown in Figure 4.17 as the linear region between the black circles.

The simulated evolution of probabilistic winding SOH estimates and the motor torque profiles applied over the path shown in Figure 4.17 are plotted in Figures 4.18a and 4.18b respectively. Because the control cost penalizes only the lowest 98% confidence motor SOH assessment at EOM, the optimal control over the mission is expected to allocate load

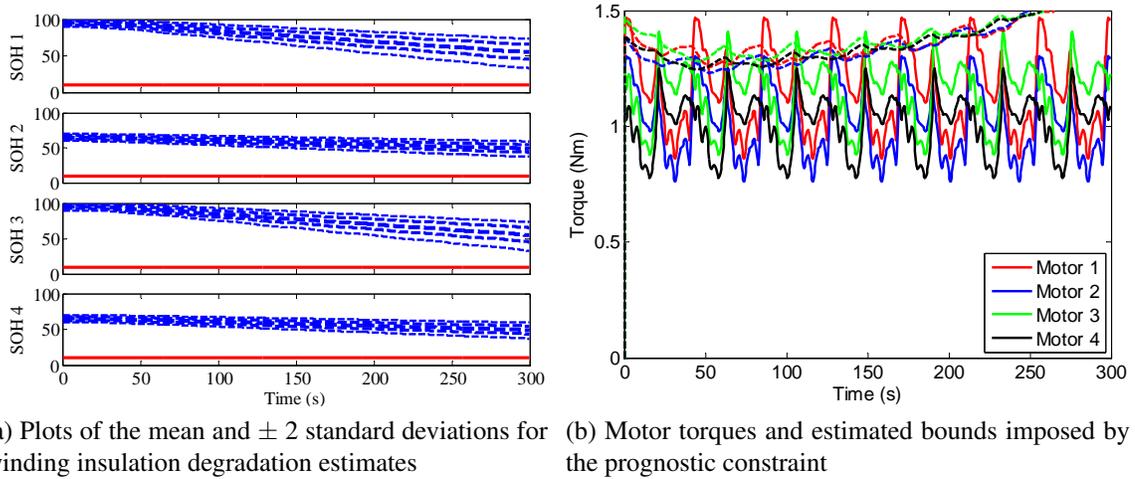


Figure 4.18: Simulation results for a skid-steered vehicle following the minimum allowable performance path

to healthier motors, such that all of the motors will end the simulated mission with nearly equal 98% confidence SOH estimates. This behavior is in fact seen in Figures 4.18a and 4.18b. Also, note that the error between the commanded and the estimated optimal path, plotted in Figure 4.17, is greatest in the extreme upper and lower regions of the figure-8 path because introducing an error in those regions results in the greatest reduction in the total distance traveled by the skid-steered vehicle. Although it will be very difficult to design a controller that can match these optimal behaviors using uncertain predictions of future environmental demands, it remains instructive to be able to identify those behaviors in non-causal control optimization studies

4.5.4.4 Prognostics-Based Control

At each control time-step, a prognostics-based controller will allocate motor loads to best manage the risk posed by estimates of future system performance and fault growth models. Motor load allocations for best risk management are derived in simulation studies by evaluating Eqn. 4.27 on a sufficiently dense uniform grid over the space of all motor loads satisfying the performance constraint.

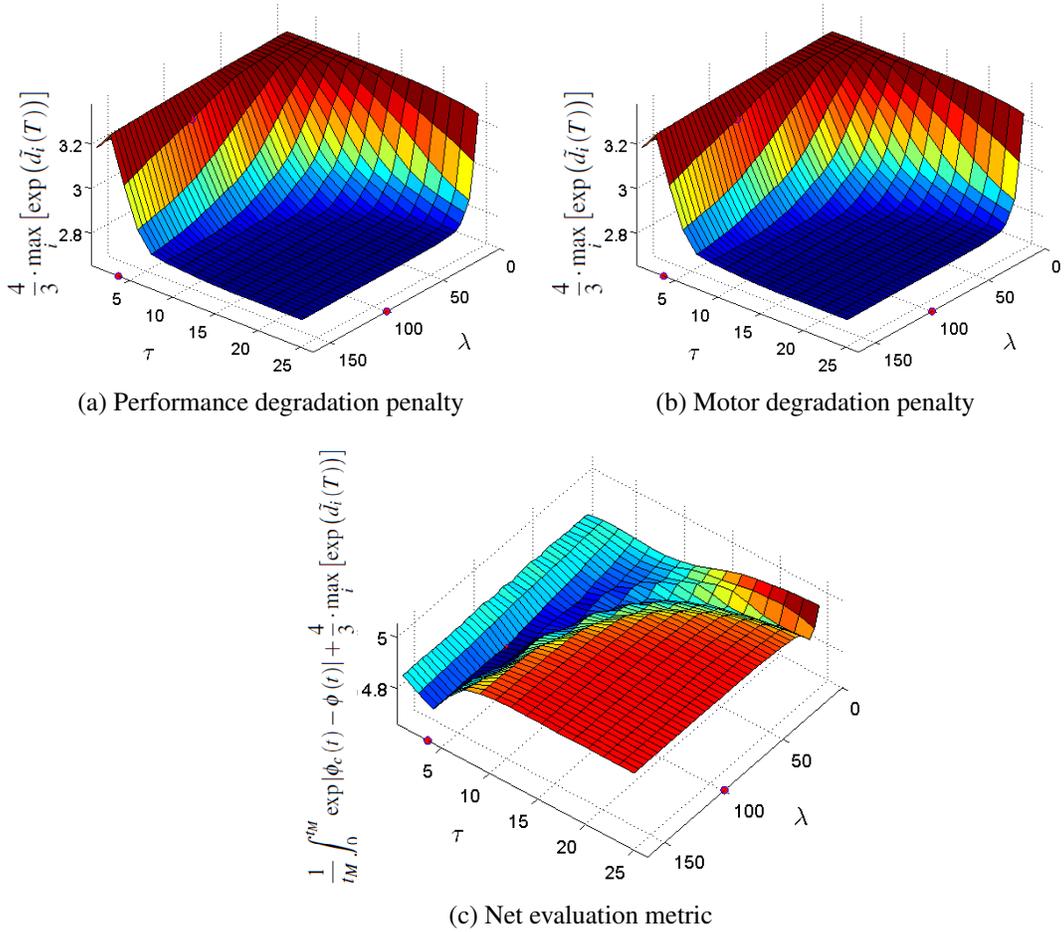


Figure 4.19: Plots of the performance degradation penalty (a), the motor degradation penalty (b), and the net outcome evaluation metric (c) for control policies identified using $\lambda = [0, 3, 6, \dots, 160]$ and $\tau = [.4, 1.6, 2.8, \dots, 24.4]$ (optimal value found at $\lambda = 100$ and $\tau = 3.6$ s)

Two degrees of freedom in the formation of performance and risk metrics are considered here. The effects of varying both the prognostic horizon length, τ , and the risk-reward weighting factor, λ will be explored.

Plots of J_d and J_M (defined in Eqns. 4.25 - 4.26) versus τ and λ are shown in Figures 4.19a and 4.19b respectively. Figure 4.19c shows plots of the net outcome evaluation cost ($J_M + J_d$), versus values of λ and τ used to define the risk metric to be optimized at each control time-step.

A general property apparent from the figure is an increase in the conservatism of control

Table 4.4: Outcome evaluation costs and 98% confidence VaR estimates for motor state of health at end of mission, using three sample control routines

	$\tilde{d}_1(t_M)$	$\tilde{d}_2(t_M)$	$\tilde{d}_3(t_M)$	$\tilde{d}_4(t_M)$	$J_d _0^{t_M}$	$J_M _0^{t_M}$	$(J_M + J_d) _0^{t_M}$
Min performance path	51%	27%	51%	27%	2.76	2.16	4.93
With future knowledge	36%	36%	36%	36%	2.52	1.69	4.21
Without future knowledge	30%	14%	30%	14%	3.17	1.58	4.76

policies as the prognostic horizon, and the damage weighting term, λ , are increased. Control conservatism is observed in Figures 4.19a and 4.19b as an increase in the performance evaluation cost and a decrease in the SOH at EOM cost. Conversely, the control policies are seen to become more greedy as the prognostic and the damage weighting term are decreased. Greedy behaviors are observed in Figures 4.19a and 4.19b as a decrease in the performance evaluation cost and an increase in the SOH at EOM cost. The trough seen in Figure 4.19c indicates a domain of τ and λ values corresponding to controls that are neither overly conservative nor overly aggressive. The optimal outcome evaluation cost is found to result from setting $\lambda = 100$ and $\tau = 3.6$.

Table 4.4 shows the outcome evaluation costs and 98% confidence VaR estimates for motor SOH at EOM using three sample control routines. The table shows control outcomes for the sample feasible control solution on the minimum performance path (defined in Section 4.5.4.2), the vehicle load allocations optimized using future knowledge (defined in Section 4.5.4.3), and the prognostics-based policy identified using $\lambda = 100$ and $\tau = 3.6$ seconds. The 98% confidence VaR estimate for motor i is denoted by \tilde{d}_i in the table. The ending motor damage term, J_d , and the path deviation penalty, J_M , were defined in Eqns. 4.24 - 4.26. The tabulated results show that the control policy optimized using future knowledge results in a slightly higher performance degradation penalty than is observed for the controller that was optimized for uncertain prognostic modeling information. However, the control policy optimized using future knowledge is seen to trade off this slightly higher performance degradation penalty for significantly less motor damage at EOM than is observed for the controller that was optimized for uncertain prognostic modeling.

CHAPTER V

A GENERALIZED MDP FORMULATION OF PROGNOSTIC MODELING AND RISK BASED DECISION MAKING

5.1 Introduction

Markov process notation provides a conceptually convenient and extremely flexible means of encoding uncertainty into the representation of component SOH estimates and future fault growth estimates. The science of encoding uncertainty into Markov process representations of system dynamics based on frequentist theory or Bayesian inference is well developed and widely applied in academia and in practice. The down side of using Markov processes to represent system dynamics however is the number of states and state transition probabilities necessary to describe a given process quickly becomes intractable as the dimensionality of the represented dynamics is expanded.

A generalized Markov process representation of fault dynamics is presented in this chapter for the case that uncertainty in component degradation models and uncertainty in predictions of the future exogenous stresses to be placed on degrading components will both be incorporated into the computation of prognostic estimates. Both sources of uncertainty are represented here as explicit stochastic processes. The action space of the supervisory component load allocation process is defined here, as in the previous chapter, in terms of a metric representing the relative deviation between the system's nominal output response and the net system output that is actually enacted by an implemented prognostics-based control routine at each control time-step.

The risk based decision making problem may be generally expressed as a Markov decision process (MDP) [86, 66], defined over the space of allowable supervisory control actions available at fixed decision making epochs. MDPs have been widely applied to the representation of problems involving sequential decision making in the presence of uncertain or

stochastic modeling information in the areas of economics [41], supply chain management [65], scheduled maintenance [86], and health care [88], in addition to being a widely used tool for describing fault-adaptive and robust control problems [103]. A formal description of fault growth modeling and RUL estimation in terms of Markov process models, as well as a survey of similar stochastic modeling techniques are given in [6].

This chapter uses the generalized Markov process formulation of component fault growth dynamics originally described in [10]. A multi-variable stochastic system example, originally introduced in [11], is used to illustrate the compounding of modeling uncertainties and the application of dynamic programming to identify optimal control policies. Section 5.2 introduces a generalized representation of component degradation dynamics in terms of a multi-variable stochastic Markov process. Section 5.3 describes the problem of identifying control policies that optimize the expectation of a given fault risk evaluation function over a finite time window via dynamic programming. Section 5.4 introduces an illustrative and representatively computationally challenging example of a system in which uncertainty in fault growth physics models is represented by a uniformly distributed random process, and uncertainty in future demand modeling is represented by a discrete random walk. Finally, concluding remarks are given in Section 5.5.

5.2 Markov Process Modeling of Fault Growth Dynamics

Chapter 3, Section 3.3 introduced the following Markov process description of component SOH evolution in terms of an outer-loop performance modifier, ρ :

$$\begin{aligned}
 p_{i,j}^l(\rho(k)) &= p(\gamma(k+1) = s_j | \gamma(k) = s_i, \rho = \rho(k)) \\
 &= \sum_{w \in W} \sum_{\xi \in \Xi} p(w(k) = w) \cdot p(\xi_l(k) = \xi) \\
 &\quad \text{where } f(s_i, H(\rho \cdot G(w, x, x_c))_l, \xi_l) = s_j, \quad s_i, s_j \in S \quad (5.1)
 \end{aligned}$$

Here, the process uncertainty terms, ξ_t , are considered to belong to a stationary distribution, and a potentially non-stationary distribution is considered for the exogenous inputs to the system, w . The component state of health (SOH) variable, γ_t , represents a percent health index, ranging from 100% to 0% (failed).

The incorporation of uncertain beliefs about the present state of a system at fixed control decision epochs can be found in publications on partially observable Markov decision processes; see the survey paper by Lovejoy for more information [53]. The additional notation necessary to include state estimation uncertainty in the system health management problem is omitted from this chapter in order to promote clarity in the presentation of the system health prognostics-based risk management problem and the case-study example.

The Markov process notation given here may be used to describe all probabilistic fault growth process models in which the following assumptions are satisfied:

Assumption 5.1. The fault growth dynamics are taken to be memoryless; i.e., the conditional probability distribution for future states depends only on the present state of the process, and not the past. This assumption is referred to as the Markov assumption in stochastic systems.

Should it be the case that a fault growth process of interest is not completely memoryless, but future states only depend on a finite number, m , of previous states, then the Markov process notation given in this document could be extended to satisfy the Markov assumption by defining the state space of the process to be the ordered m -tuple of the current state and the m previously visited states, as described in [94]

Assumption 5.2. State transition probabilities are considered to be time invariant; although, it may be the case that fault growth models are not precisely known a priori and must be adapted online using techniques such as particle filtering [64] or Bayesian learning [73].

Assumption 5.3. At all discrete time-steps, the state space, the action space, and the space of environmental and other exogenous inputs to the system are adequately represented by a finite set of states, which is bounded from above by the availability of computational resources.

In the event that fault growth must be modeled as a continuous time process, a representation of fault growth modeling similar to that given here may be expressed in terms of a continuous time Markov process [83] or a semi-Markov process [27].

These required assumptions are mild enough to allow a wide array of fault growth processes to be described in terms of the Markov process notation given in Eqn. 5.1 [38, 91]

5.2.1 Incorporating a Stochastic Process Model for Environmental Loading Demand

A generic process model for the environmental demand term, w , is:

$$p(w(k) = a) = p(w(k) = a | w(k-1) = b) \cdot p(w(k-1) = b),$$

$$a, b \in W, \quad k \in [t_p + 1, \infty) \quad (5.2)$$

where an observation of w at time-index t_p is used to initialize the random process. Stochastic modeling of environmental dynamics is represented by $p(w(k) = a | w(k-1) = b)$.

The Markovian state transition model of the degrading system can now be expressed as a four dimensional matrix that incorporates the process model for w .

$$p_{(i,j),(l,m)}^l(\rho) = p(w(k) = z_l | w(k-1) = z_m) \cdot p(\xi_l(k) = \xi)$$

$$\text{where } f(s_i, H(\rho(k)) \cdot G(w, x, x_c)), \xi) = s_j, \quad s_i, s_j \in S, \quad z_l, z_m \in W \quad (5.3)$$

Here, $p_{(i,j),(l,m)}^l(\rho)$ represents the probability of component l transitioning from fault state s_i and exogenous loading demand z_l to fault state s_j and exogenous loading demand z_m .

This notation enables available knowledge of system kinematics, fault growth modeling, exogenous demand modeling, measurement uncertainties, and modeling uncertainties to be represented by a finite set of state transition probabilities.

5.3 Encoding Risk Aversion into a Markov Decision Process

Over the past several decades much has been published on the theory of encoding various forms of risk aversion into the specification of MDP state transition costs [42, 68]. An MDP formulation of the stochastic optimization problem requires control costs to be assessed using the accumulation of state transition costs,

$$\text{Total Cost} = E \left\{ g_N(\gamma(N)) + \sum_{k=0}^{T_M-1} g_k(\gamma', w', \gamma, w, \rho) \right\} \quad (5.4)$$

where $g_k(\gamma', w', \gamma, w, \rho)$ denotes a state transition cost assigned to the possibility of transitioning from one system state, (γ, w) , to another, (γ', w') , at time-index k , given a control action, ρ . A terminating cost that penalizes the total component degradation over a simulated time window is denoted by $g_N(\gamma(N))$. Cost discounting and average cost formulations are used in the formulation of infinite horizon MDPs, not considered here.

The sequence of control actions taken over the domain $k = [0, N - 1]$ in response to online observations of system and environmental states is referred to as a control policy, denoted $\pi = \{\mu_0, \dots, \mu_{N-1}\}$, where μ_k represents a control mapping. Here, μ_k maps observations of γ and w obtained at time-index k into a control action;

$$\rho(k) = \mu_k(\gamma(k), w(k)) \quad (5.5)$$

The expected cost for enacting a particular control policy when starting at given initial

values for γ and w at time-index t_p is denoted:

$$J_{\pi}(\gamma(t_p), w(t_p)) = E \left\{ g_N(\gamma(N)) + \sum_{k=t_p}^{N-1} g_k(\gamma', w', \gamma, w, \rho) \right\} \quad (5.6)$$

An optimal control policy is defined as a one that minimizes J_{π} ,

$$J_{\pi^*}(\gamma(0), w(0)) = \min_{\pi \in \Pi} J_{\pi}(\gamma(0), w(0)) \quad (5.7)$$

Here π^* represents an optimal control policy.

5.3.1 Identification of Optimal Finite Horizon Control Policies with Dynamic Programming

Optimizing control policies may be identified using well studied MDP optimization techniques such as backwards induction for finite horizon optimization problems, and linear programming, value iteration, and policy iteration for discounted and average-reward infinite horizon optimization problems. The well known dynamic programming algorithm uses backwards induction to identify an optimal control policy over the time window $k \in \{N-2, N-1\}$, and then for $k \in \{N-3, N-2, N-1\}$ and so on, until the optimal policy is found over the entire time-window of interest. The control costs used in the backwards induction approach are defined for the multi-variable stochastic system discussed here as:

$$\begin{aligned} J_N(\gamma(N), w(N)) &= g_N(\gamma(N)) \\ \mu_k^*(\gamma, w) &= \min_{\rho} E \{ g_k(\gamma', w', \gamma, w, \rho) + J_{k+1}(\gamma', w') \} \\ J_k(\gamma, w) &= E \{ g_k(\gamma', w', \gamma, w, \mu_k^*(\gamma, w)) + J_{k+1}(\gamma', w') | \rho = \mu_k^*(\gamma, w) \} \\ \gamma \in \mathcal{S}, w \in \mathcal{W}, \pi^* &= \{ \mu_0^*, \dots, \mu_{N-1}^* \}, k = \{0, 1, \dots, N-1\} \end{aligned} \quad (5.8)$$

The computational burden of this solution method is:

$$O(mn^2N) \tag{5.9}$$

where N is the time horizon to be optimized over, and m and n are the cardinalities of the discrete spaces defined for the action vector, ρ , and state space of (γ, w) respectively.

While this is a great improvement over the computational burden of an exhaustive search, which is $O(m^{m^N})$, the reader should note that the cardinalities of the state space and action space used in the MDP will grow exponentially with the dimensionalities of γ , w , and ρ . Therefore, the dynamic programming method quickly becomes computationally infeasible for higher dimensional problems.

5.4 Consideration of a Multi-Variable Stochastic System Example

This section considers an example multi-variable stochastic system that is intended to illustrate some of the fundamental difficulties in assessing and managing the prognostic uncertainty introduced when stochastic models for fault growth physics and environmental loading demands are incorporated into prognostic estimates.

Consider a health depletion model of the form:

$$\gamma(k+1) = \gamma(k) - \lambda \cdot |u(k)| \cdot \xi(k) \tag{5.10}$$

where the rate of component fault growth is defined to be proportional to the magnitude of component load u multiplied by a process noise variable ξ , and λ represents a constant of proportionality in this mapping.

The net control effort output is considered to be equal to the component load, u ,

$$v = u \tag{5.11}$$

Environmental loading demands on the system are taken to be represented by the sum of consecutive draws from the set $\{-1,0,1\}$, constituting a discrete random walk process. The probability mass function for environmental demands at time-index k is:

$$\Pr(w(k) = a|w(k-1) = b) = \begin{cases} \frac{1}{3} & a - b \in \{-1,0,1\} \\ 0 & \text{else} \end{cases}, \quad k \in [t_p + 1, \infty) \quad (5.12)$$

where the variable t_p represents the time-index at which prognostic predictions are made.

A nominal control law is considered to command a net system output control effort that exactly matches the environmental loading demand at each control time epoch,

$$r(k) = w(k) \quad (5.13)$$

Substitution of Eqns. 3.7, 5.11, and 5.13 into Eqn. 5.10 yields:

$$\gamma(k+1) = \gamma(k) - \lambda \cdot \rho(k) \cdot |w(k)| \cdot \xi(k) \quad (5.14)$$

Process uncertainty in the example fault growth model is taken to be represented by independent random draws from the set $\{.7, .8, .9, 1.1, 1.2, 1.3\}$. The resulting probability mass function for ξ is:

$$\Pr(\xi(k) = \xi) = \begin{cases} \frac{1}{6} & \xi \in \{.7, .8, .9, 1.1, 1.2, 1.3\} \\ 0 & \text{else} \end{cases}, \quad k \in [t_p, \infty) \quad (5.15)$$

State transition costs for the example system were designated to penalize the proportional deviation from nominal output loading commands at each control time-index,

$$g_k(\gamma', w', \gamma, w, \rho) = 1 - \rho \quad (5.16)$$

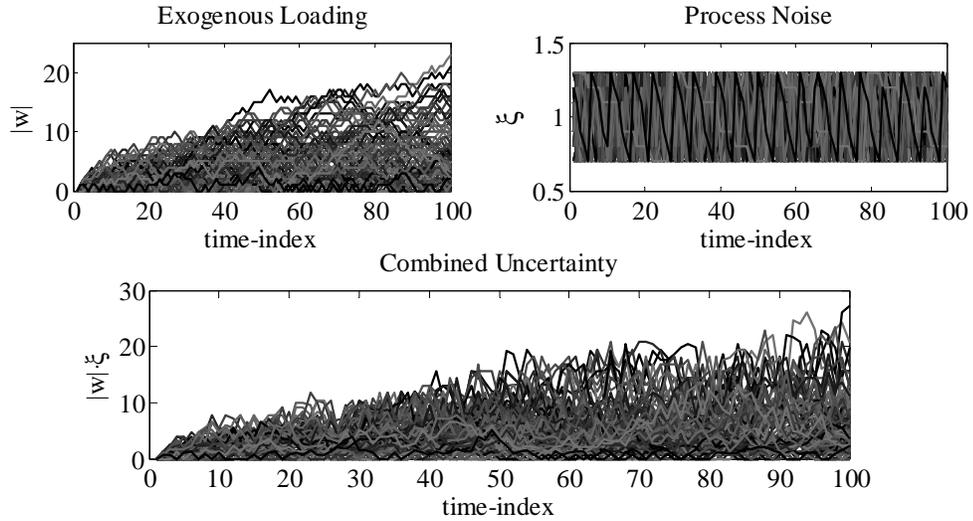


Figure 5.1: Time-series plots for 100 simulations of the random variables, w and ξ (top), and the product of w and ξ (bottom)

where, $\rho = 1$ corresponds to no deviation from the system's nominal control.

The terminating cost for this example is designated to be inversely proportional to the square of component health at EOM.

$$g_N(\gamma(N)) = (100 - \gamma(N))^2 \quad (5.17)$$

5.4.1 Simulation Results

Both w and γ are assumed to be observable in this example. Simulation studies use $\lambda = 0.33$ in the fault growth process model, and $\rho \in \{.2, .3, \dots, 1\}$ to define the space of supervisory control actions that may be enacted at each simulated time-index. The time-series behavior of the example system was simulated using a pseudo random number generator initialized with a unique 'seed' value provided for each simulation run. Figure 5.1 shows time-series profiles for $|w|$, ξ , and $|w| \cdot \xi$ using 100 repeated randomized simulations of the example system, each using a unique seed value.

Figure 5.2 shows simulation results for two sample control policies; $\rho = .2$ (top) and $\rho = 1$

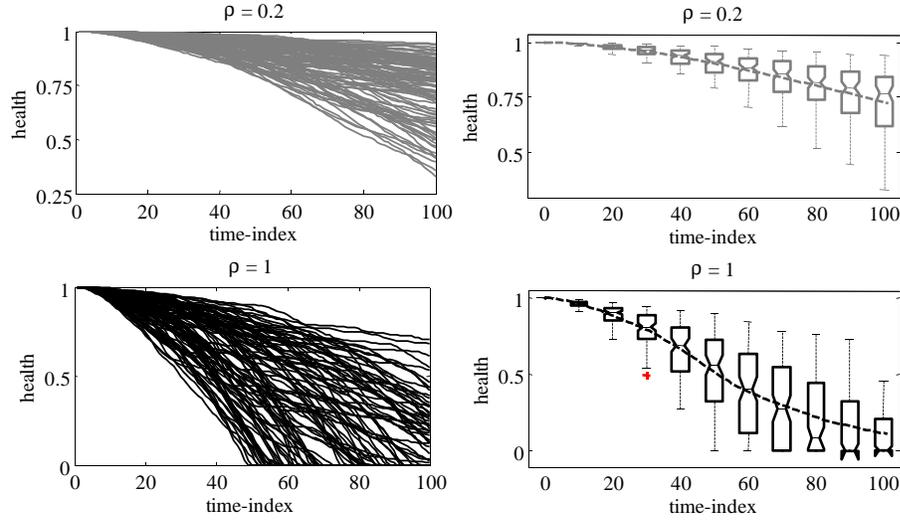


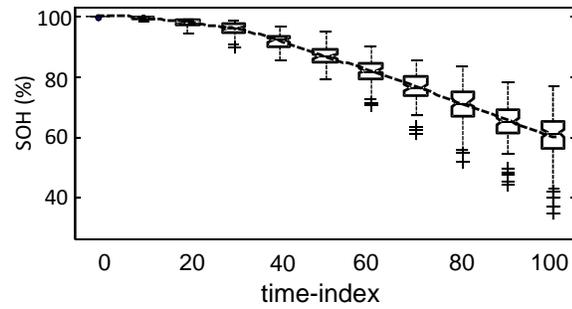
Figure 5.2: Time-series data (left) and box plots (right) for 100 simulations of the example fault growth process with 80% deviation ($\rho = 0.2$) from nominal output demands (top), and no deviation ($\rho = 1$) from nominal output demands (bottom)

(bottom). Setting $\rho = .2$ for all time-indexes results in the maximum allowable degradation of the system's nominal load allocations at all time-indexes. Setting $\rho = 1$ for all time-indexes enacts no changes to the system's specified nominal load allocation policy over any of the sample simulations of environmental loading and fault growth dynamics.

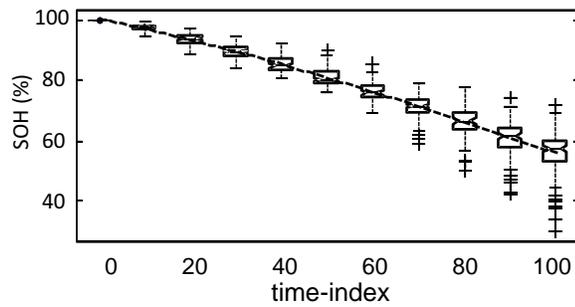
The box plots shown on the right hand side of Figure 5.2 provide a convenient means of representing the simulated statistics of the fault growth process. The top and bottom of the boxes plotted in Figure 5.2 represent the 25th and 75th percentiles of the data range at a given time-index. The notch in each box represents the median of the data points, the dashed line represents the mean value, and the whiskers in the box plots extend to the most extreme points falling within the range,

$$q_1 - 1.5 \cdot (q_3 - q_1) < d_i < q_3 + 1.5 \cdot (q_3 - q_1) \quad (5.18)$$

where q_1 and q_3 are the 25th and 75th quantiles of the data respectively, and d_i represents a datapoint.



(a) Optimal control with stochastic modeling information



(b) Optimal control with future knowledge

Figure 5.3: Plots of the stochastic health degradation process over repeated simulations of the optimal control policies computed with (a) and without (b) model uncertainty

Note that the spread of possible system trajectories observed in these simulations become relatively large over the 100 increment time-window used in this example. This is typical for prognostics-based control applications. It can be observed from the sample results shown that always enacting the minimum allowable system performance would be very ‘safe’, but likely overly conservative in many cases. On the other hand, always enacting nominal system performance would likely be unacceptably ‘risky’, and result in failure in many cases.

Figure 5.3a shows the distribution of component health as it evolves over 100 repeated simulations of a control policy that minimizes expected cost over a 100 time-increment window. Figure 5.3b shows the distribution of component health profiles observed over repeated simulations of an optimal control policy computed using future knowledge of the fault growth process, as was described in Chapter 4, Section 4.5.4.3. The dynamic programming algorithm given in Eqn. 5.8 was used to identify an optimal control policy for

Table 5.1: Statistics of the accumulated state transition costs computed over repeated simulations of four sample control policies

	$\mu(\sum_k g_k)$	$\sigma(\sum_k g_k)$	$\mu(g_N)$	$\sigma(g_N)$	$\mu(J_\pi)$	$\sigma(J_\pi)$
Sample control policy: $\rho = 1$	0	0	81.6	26.7	81.6	26.7
Sample control policy: $\rho = .2$	78.4	0	9.9	10.5	88.3	10.5
Optimal with stochastic modeling	55.1	18.2	16.8	7	71.8	23
Optimal with future knowledge	40.7	20.1	20.4	7.3	61.1	25.7

both cases. In the case that future knowledge is available, the dynamic programming algorithm is run using a deterministic Markov process model. In the causal case where future knowledge is not available, the dynamic programming algorithm is run using a stochastic Markov process model.

Note that while the discovery of an optimizing control policy through finite horizon dynamic programming is a computationally challenging problem, as described in Section 5.3.1, the optimal policy is computed offline, and requires no online optimization as long as the Markov process model used to generate the policy is still applicable. If online updates to the Markov process modeling of environmental loading and fault growth dynamics were considered, then the optimizing policy would need to be recomputed online.

The mean and standard deviation of the control costs evaluated over repeated simulations of the two optimized control policies and the two example control policies are given in Table 5.1. Comparison of the health deterioration plots given in Figure 5.3, show a somewhat more conservative behavior early in the mission from the control policy lacking perfect future knowledge, which is to be expected. The more conservative nature of the optimal policy computed using the stochastic system model results in an increased performance penalty and a decreased ending health penalty, compared to the optimal policy computed with future knowledge. Both control policies are seen to show substantial variation in evolution of SOH over repeated trials. This complicates the V&V of control policies, as discussed in Chapter 2. Many more quantitative and qualitative observations may be explored using this illustrative and fairly computationally challenging example. The reader

is recommended to program this simple example and observe the aggregate efficacy of various risk management control policy formulations.

5.5 Summarizing Remarks

A generalized Markov modeling representation of fault dynamics is shown for the case that available modeling of fault growth physics and available modeling of future environmental stresses may be represented by two independent Markov process models. A notationally convenient MDP formulation of the prognostics-based control problem is provided for a system with multiple degrading effectors. Formulation of the component fault growth process in terms of a metric representing the magnitude of system output performance degradation induced by a given control policy was used to directly relate supervisory health management control actions to their effects on system performance and component degradations. Finite horizon dynamic programming is described to solve for the optimal risk mitigating control policy over a finite time window for the general case that stochastic models representing physics of failure and future environmental stresses are known, and the states of both stochastic processes are observable by an implemented supervisory health management controller.

The effects of compounding uncertainties in physics of failure and environmental demand modeling were illustrated via a simulation study presented for a multi-variable stochastic process model. The formulation and implementation of prognostics-based control on the multi-variable stochastic system example is intended to be structurally analogous to the prognostics-based control problem on a wide variety of real-world applications. Causal and non-causal implementations of dynamic programming were used to illustrate the dependence of the optimal health management control policy on uncertainties in fault growth and environmental dynamics models.

CHAPTER VI

SIL/HIL REPLICATION OF POWERTRAIN LOAD DYNAMICS AND ONLINE PROGNOSTICS FOR A UAV APPLICATION

6.1 Introduction

Previous chapters presented analysis of the prognostics-based decision making problem using simulation examples. This chapter considers the development and testing of prognostics and supervisory decision making on a real system.

As previously described, the development and tuning of prognostics-based decision making routines is heavily reliant on the availability of trustworthy models for fault growth physics, future exogenous loading dynamics, and system physics. A testbed intended for offline software-in-the-loop (SIL) and hardware-in-the-loop (HIL) testing of battery charge prognostics and battery charge management routines is described in this chapter. SIL testing refers to tests conducted using only software representations of system physics and embedded control routines. HIL testing refers to tests that include some hardware components from the target system.

A framework is described for the offline recreation of dynamic loads on simulated or physical aircraft powertrain components. The dynamic loads on aircraft powertrain components are identified using a real-time simulation of airframe dynamics running on a flight simulator, an inner-loop flight control policy executed by either an autopilot routine or a human pilot, and a supervisory fault management control policy that can interact with a human pilot or autopilot. The analysis presented here is mostly drawn from the following two recent publications: [13, 22]. SIL development of battery SOC management for a similar application, a hybrid electric vehicle, was presented in [9].

An early investment of resources into the development of an offline V&V testing infrastruc-

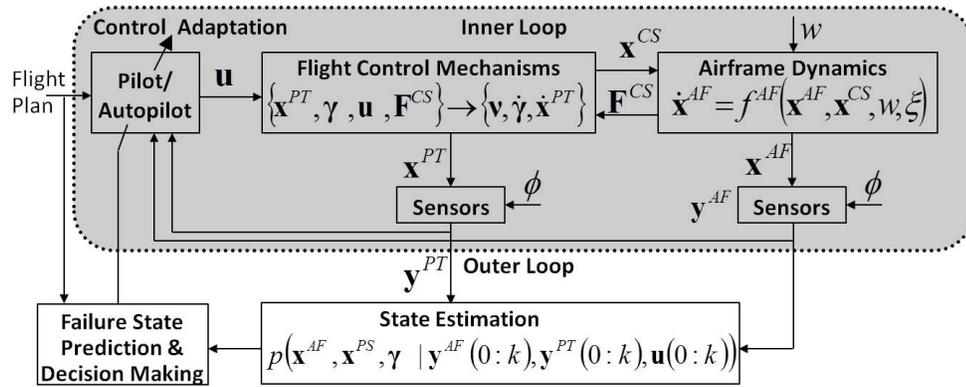
ture for prognostics and supervisory health management algorithms is easily justified for complex systems in which online testing is substantially more time consuming and costly than offline testing. The SIL and HIL testing of failure prognostics and decision making tools for aircraft systems will facilitate much more comprehensive and cost-effective testing than what is practical to conduct with flight tests. It is often not feasible to inject faults and run to failure during full-scale tests without compromising the vehicle or operator safety, thus it is valuable to implement a framework for the offline verification and validation (V&V) of algorithm performance during failure scenarios. An offline V&V framework will also facilitate testing over a wide range of potential environmental conditions, including extreme conditions that are rarely encountered in practice.

Offline testing of supervisory health management algorithms in a laboratory setting will not only improve safety, but, as many issues can be resolved during offline tests. Offline testing also reduces the number of full-scale deployment tests required for algorithm V&V, reducing testing costs and development time. That said however, offline V&V testing is limited by the accuracy of SIL and HIL replications of nominal and off-nominal system dynamics. Full-scale deployment testing is thus still a necessary part of the algorithm development and V&V process.

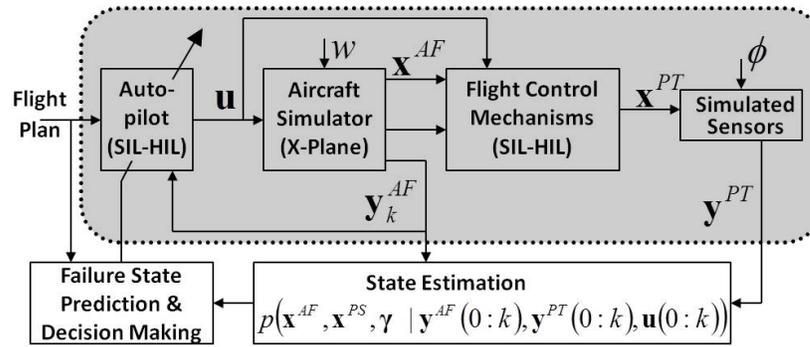
The SIL/HIL testing framework described in this chapter uses the X-Plane¹ flight simulator package and an X-Plane Toolbox for MATLAB² to facilitate prognostic based control algorithm V&V over a range of potential operating conditions. Examples of other offline testbeds making use of X-Plane for aerodynamics simulation and MATLAB/Simulink softwares for simulation of control routines are found in [44, 14, 70].

¹www.x-plane.com/

²Developed at NASA Ames and scheduled for open source distribution in 2014



(a) Inner and outer control loops for online flight testing



(b) Inner and outer control loops for offline flight testing

Figure 6.1: Closed-loop control diagrams for online and offline flight testing

6.2 A General Software-in-the-Loop / Hardware-in-the-Loop Framework

Offline V&V tests of supervisory failure prognosis and decision making routines will allow supervisory health management algorithms to interact with system controllers and measured data exactly as they would during full-scale tests. The offline testing of health management algorithms may be conducted using strictly software models of flight systems or a combination of simulated and physical components.

Block diagrams illustrating the structure of control loops used for online and offline flight testing are shown in Figures 6.1a and 6.1b respectively. The notation used in the figures is defined in Table 6.1. An inner-loop controller is assumed to update flight control inputs based on a known flight plan and observations of the system state. The inner-loop controller

Table 6.1: Symbol definitions for the software-in-the-loop / hardware-in-the-loop testing framework

x^{AF}	airframe state vector
y^{AF}	observation of airframe state vector
x^{PT}	electrical power dist. system state vector
y^{PT}	observation of x^{PT} states
u	pilot or autopilot control output vector
v	mechanical loads on electromechanical components
F^{CS}	net mechanical loads exerted on airframe control surfaces
w	environmental state parameter vector
γ	Component SOH vector
ξ	captures uncertainties in physics of failure models
ϕ	captures noise in sensor measurements

is denoted by the ‘Pilot/Autopilot’ blocks in Figures 6.1a and 6.1b. Failure prognostics and supervisory decision making operations are performed by an outer-loop process, denoted by the ‘State Estimation’, ‘Failure State Prediction’, and ‘Decision Making’ blocks in Figure 6.1. Both the inner-loop and outer-loop controllers would be unchanged in either online or offline flight testing.

6.2.1 Description of Inner-Loop Control Dynamics

Inner-loop control of the aircraft is assumed to be provided by either a human pilot or a pre-programmed autopilot. Both human pilot and autopilot will henceforth be referred to as just ‘the pilot’, for convenience. The pilot will update the control vector, u , based on the observed states of the aircraft and a desired system state directed by a given flight plan.

The ‘Flight Control Mechanisms’ block shown in Figures 6.1a and 6.1b represents the internal electrical and mechanical dynamics of the vehicle’s powertrain. The inner-loop control signals sent by the pilot, and the aerodynamic forces exerted on the vehicle’s control surfaces by the surrounding environment, F^{CS} , are inputs to this block. The F^{CS} vector consists of forces like the drag on the aircraft propeller, and the torque on control flaps. These inputs result in the loading of powertrain components, represented by the vector v , which in turn determine the dynamics of powertrain component states, x^{PT} , and the

dynamics of component health, represented by the vector γ .

The loads exerted by the vehicle's active components at a given time index, k , are expressed as a function of the pilot control action vector, the current states of powertrain components, and the states of component fault modes that may reduce component effectiveness.

$$v(k) = f^{PT}(u(k), x^{PT}(k), \gamma(k), \xi(k)) \quad (6.1)$$

where the ξ term is used in this expression to represent a vector of unknown or uncertain model parameters.

The 'Airframe Dynamics' block shown in Figure 6.1a represents the aerodynamic interactions between the vehicle airframe, vehicle control surfaces, and the operating environment. The inputs to this block are the current states of aircraft control surfaces, x^{CS} , and the current state of the operating environment, w . Environmental states represented by w may include atmospheric pressure, air temperature, wind speed, and turbulence. The current state of the airframe is represented by, x^{AF} ; it includes the position, heading, linear and rotational speed, and linear and rotational accelerations of the airframe in a given coordinate system.

Vehicle control surfaces are mechanically connected to powertrain components, so they should be a known function of x^{PT} ,

$$x^{CS} = f^{CS}(x^{PT}) \quad (6.2)$$

The forces exerted on the vehicle's control surfaces due to their motion through surrounding air is represented here as a generic non-linear function of the airframe state, the states of vehicle control surfaces, and current environmental states,

$$F^{CS}(k) = f^{FCS}(x^{AF}(k), x^{CS}(k), w(k), \xi(k)) \quad (6.3)$$

where the additional ξ term is used to represent a vector of unknown or uncertain model parameters.

Powertrain state dynamics and airframe dynamics are generically expressed in terms of the loading vectors \mathbf{v} and F^{CS} as:

$$\dot{x}^{PT}(k) = f^{PT} \left(x^{PT}(k), \mathbf{v}(k), F^{CS}(k), \xi(k) \right) \quad (6.4)$$

$$y^{PT}(k) = h^{PT} \left(x^{PT}(k), \phi(k) \right) \quad (6.5)$$

$$\dot{x}^{AF}(k) = f^{AF} \left(x^{AF}(k), x^{CS}(k), \mathbf{w}(k), \xi(k) \right) \quad (6.6)$$

$$y^{AF}(k) = h^{AF} \left(x^{AF}(k), \phi(k) \right) \quad (6.7)$$

The progression of component health degradation is represented as:

$$\dot{\gamma}(k) = f^{\gamma} \left(x^{PT}(k), \gamma(k), \mathbf{v}(k), \xi(k) \right) \quad (6.8)$$

The deterioration of control surfaces and electromechanical components in aircraft powertrains as a function mechanical loading forces has been a topic of study for some time; examples include: electromechanical actuators [5] and composite wing structures [35], to name a few. The degradation and failure of electrical components as a function of electrical power loading has also been examined for aircraft components such as batteries [73] and power electronics [21].

Measurements from simulated vehicle powertrain components, y^{PT} , are generated in offline testing using models for both the underlying component dynamics, and sensor response dynamics. Data from actual system hardware is obtained in offline testing by applying mechanical loads to hardware components in real-time, in accordance with the environmental loads, \mathbf{F}^{CS} , reported by the X-Plane simulator.

A hardware-only recreation of the ‘Flight Control Mechanisms’ portion of the inner-loop

vehicle dynamics, illustrated in Figure 6.1b, could be accomplished in a laboratory setting using an aircraft battery pack, power electronic motor/actuator drivers, electromechanical components, and associated interconnection cabling. Pilot controls could be sent directly to an electrical power distribution system assembled in the laboratory and additional loading hardware could be used to apply mechanical loads to the electromechanical components of the powertrain in order to recreate the environmental loads estimated by an aircraft simulator. This approach is similar in nature to dynamometer testing commonly performed in the testing of automotive systems [50, 92]. Software models may be switched in for some or all of the hardware components in this setup; however, small errors in modeling the behavior of a given component may have outsized effects in observed system behavior over long time periods.

Measurements of the airframe states are represented by the vector, y^{AF} . In offline testing, aircraft dynamics are simulated using X-Plane. The X-Plane simulator includes sensor models that are used to generate y^{AF} . Autopilot routines are connected to the X-Plane simulator in offline testing using the open-source program APM Mission Planner³. Human pilots may interface with the X-Plane simulator using a simulator driven interface, and any desired I/O mechanisms that may be connected to computer.

Adequate control of aircraft does not in most cases require a pilot to understand environmental dynamics or the internal dynamics of the flight vehicle in great detail. Pilots are autopilots are considered here to make decisions based on an internal decision making policy that maps observations of $y^{AF}(k)$ and $y^{PT}(k)$ at time-index k to appropriate control outputs, $u(k)$. An autopilot will use an embedded control policy to map $(y^{AF}(k), y^{PT}(k))$ onto u . For human pilots, the mapping $(y^{AF}(k), y^{PT}(k))$ onto u will be determined by the pilot's situational awareness and judgment. The mechanism for interaction between an autopilot and supervisory prognostics-based decision making routines can be for the decision making routines to directly update the autopilot's control policy. Policy updates for human

³<http://code.google.com/p/ardupilot-mega/wiki/Mission>

pilots could be prompted indirectly by presenting the pilot with system health information and suggested risk mitigating actions, as described in [18].

6.2.2 Description of Outer-loop Failure Prognostics and Decision Making

Supervisory outer-loop control routines make use of sensor measurements to estimate current and future system states given models of system state dynamics and physics of failure models. Probability distributions for belief in the current states of \mathbf{x}^{PT} , \mathbf{x}^{AF} , and γ , based on a history of observations of \mathbf{y}^{PT} , \mathbf{y}^{AF} , and \mathbf{u}^{AF} are generically represented in Figure 6.1 as:

$$p\left(\mathbf{x}^{PT}, \mathbf{x}^{AF}, \gamma | \mathbf{y}^{PT}(0:k), \mathbf{y}^{AF}(0:k), \mathbf{u}(0:k)\right) \quad (6.9)$$

Many Bayesian and machine learning methods have been published for the estimation of such probability distributions in the aviation domain [52, 58].

In offline simulations, stochastic beliefs about the manner in which the operating environment or internal system dynamics will evolve over time may be validated against repeated randomized simulations of flight scenarios. Although it will not be demonstrated in this chapter, the proposed SIL/HIL testing framework could be used to conduct repeated randomized trials to validate prognostics-based control routines and online constraints, as was demonstrated in Chapters 4 and 5.

6.3 Application Example - Battery Depletion Modeling on an Edge 540T UAV

The aircraft platform used in this application example is a commercial-off-the-shelf (COTS) 33% scale model of the Zivko Edge 540T airplane, pictured in Figure 6.2. The wingspan of the Edge 540T is 87 inches. The mass of the instrumented vehicle is approximately 44lbs. The propeller of the UAV is driven by two tandem mounted outrunner brushless DC motors that are each powered by a series connection of two lithium polymer battery packs. Each of



Figure 6.2: Picture of an Edge 540T on runway

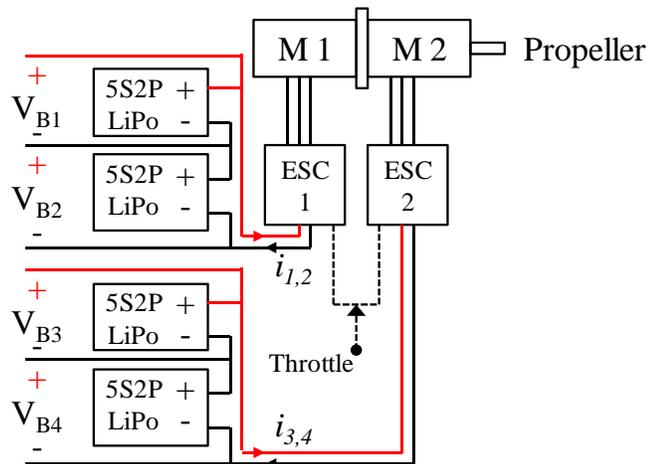


Figure 6.3: Electrical and mechanical connections of the Edge 540T UAV powertrain

the battery packs consist of five series connections of two 4.2V 3900mAh lithium polymer pouch cells wired in parallel.

The electrical and mechanical connections in the UAV powertrain are illustrated in Figure 6.3. Power flow from the battery packs to the driving motors is controlled by a Jeti 90 Pro Opto electric speed controller (ESC). The ESC sends synchronized voltages to the propeller motors at a duty cycle determined by a throttle input signal. The throttle input is either sent by remote control from a pilot, or by an onboard autopilot.

During both remote control and autonomous flight, a human pilot will maintain line of sight with the aircraft, and stand ready to execute a landing maneuver when the command

is given by other operators on the ground. The ground operations assisting the pilot monitor the battery end-of-discharge prognostic estimates and decision making outputs generated in real-time by outer-loop supervisory routines.

Charge estimation and end of charge prediction for UAV powertrain batteries has previously been examined in several publications by Bhaskar Saha at NASA ARC, Quach Chong Chi at NASA LaRC, and others [71, 74]. A separate battery system is used to power the data acquisition and other flight communications and control hardware. The two battery systems are sized such that it is very likely that the batteries powering the propeller motors will be the first to be depleted. For that reason, onboard battery discharge prognostic algorithms and supervisory decision making actions are considered to only be concerned with the propeller driving batteries.

6.3.1 Edge 540T Inner-Loop Controls

Vehicle flight plans are considered to be given in terms of an ordered set of 3D coordinates to be visited by the UAV, and a desired airspeed for making the translation from one way-point to the next. Autonomous control of the Edge 540T is performed using an ArduPilot board. The ArduPilot sends control commands to the aircraft ESCs and flight control surfaces based on a set of proportional integral derivative (PID) control parameters that are tuned prior to flight, and periodic measurements of vehicle airspeed, heading, and GPS position.

As was described in Section 6.2.1, X-Plane is used to simulate vehicle aerodynamics in offline simulations. Plane Maker, a design tool within the X-Plane package, was used to specify the aircraft mass, balance, and geometry used in X-Plane aerodynamic simulations. There is some unavoidable error between the actual geometry, drag, and mass distribution of the aircraft and that used in the X-Plane aerodynamics models; however, because the control system is closed-loop, small errors in simulating aircraft aerodynamics will not typically accumulate into large errors.

X-Plane can simulate various weather conditions and hardware configurations, and the ArduPilot can be tested with various flight plans. This configuration allows for thorough testing of algorithm performance and safety before conducting flight tests.

6.3.2 Battery Demand Modeling

The SIL/HIL framework described in Section 6.2 separates the simulation of aerodynamics and powertrain dynamics into two functional blocks. Connecting these two blocks requires that the airframe loads reported by the aerodynamics simulation be translated into loads on the system's powertrain components. It is difficult to collect direct measurements of airframe loads such as component forces and torques in flight. This measurement difficulty makes validating the load mapping used in offline simulations a complex proposition.

The tuning and validation of a propeller load mapping function is separated into two steps in this section. First, a series of characterization experiments are performed in X-Plane to identify a nonlinear mapping between propeller output power and aircraft angle of climb, speed, and acceleration. Second, the modeled propeller power is mapped to a required battery power using a fixed power conversion efficiency coefficient and a proportional drag correction coefficient.

The nonlinear relationship between propeller output power and aircraft angle of climb, speed, and acceleration is observed for a flight simulator representation of an aircraft, by simulating a series of climbing and descending maneuvers at various angle of climb and throttle setpoints. Unlike actual flight tests, there is no difficulty in observing the precise loads on aircraft components in simulated flight tests. The results of repeated experiments at different throttle and angle of climb setpoints are used to fit a general set of aircraft aerodynamics and energy conservation equations, presented below.

The equations developed in this section make use of the following assumptions:

Assumption 6.1. The propeller is mounted on the aircraft nose.

Assumption 6.2. The angle between the thrust vector generated by the propeller and the velocity vector of the aircraft is small.

Assumption 6.3. Aircraft turning forces are small in comparison to the thrust and drag forces on the aircraft in its direction of travel.

The sum of the forces acting in the aircraft direction of travel is:

$$T_{x_w} = D(v) + m \cdot g \cdot \sin(\alpha) + m \cdot \dot{v} \quad (6.10)$$

where T_{x_w} represents the thrust produced by the aircraft in the direction of travel, D represents the drag force acting in the opposite direction of aircraft motion, v represents the aircraft speed, \dot{v} represents acceleration, α represents angle of climb, m represents the vehicle mass, and g represents the earth's gravity.

The drag force on the airframe is represented by the following polynomial function of airspeed and angle of climb.

$$D(v, \gamma) = c_1 + c_2 \cdot v + c_3 \cdot v^2 + c_4 \cdot \alpha \quad (6.11)$$

Figure 6.4a shows a fit of the modeled drag to the averaged drag force reported by the X-Plane simulator over several steady speed climbing and descending maneuvers. The fitted parameter values are: $c_1 = 13.47$, $c_2 = -0.6$, $c_3 = 0.019$, $c_4 = 0.14$. During take-off and landing maneuvers when the aircraft speed is less than $15m/s$ the drag force is approximated as $D = 3 \cdot v$.

A plot of the measured and estimated propeller thrust versus airspeed is shown in Figure 6.4b. The plot was generated using the thrust model given in Eqn. 6.10, and the drag model given in Eqn. 6.11.

The aircraft ESC throttle command is fit as a linear function of thrust and airspeed,

$$\text{Throttle} = a_1 + a_2 \cdot T_{x_w}(v, \dot{v}, \gamma) + a_3 \cdot v \quad (6.12)$$

where the fitted parameter values are: $a_1 = -19.64$, $a_2 = 0.95$, $a_3 = 1$.

A plot of the measured and estimated steady state ESC throttle commands observed in X-Plane simulations of various steady state climbing and descending maneuvers is shown in Figure 6.4c.

The product of thrust and airspeed gives the motive power exerted by the aircraft,

$$P_p = \frac{1}{\eta_p} \cdot T_{x_w} \cdot v \quad (6.13)$$

where P_p represents propeller output power and η_p represents the approximate propeller output power conversion efficiency.

Figure 6.4d shows the modeled propeller power and the averaged values reported by the X-Plane simulator over several steady speed climbing and descending maneuvers. The η_p parameter for the modeled aircraft was fitted to $\eta_p = 0.7652$.

A fixed battery power conversion efficiency is assumed here for the aircraft motors and power electronics. Conversion between the net propeller output power and the net battery output power required to maintain a particular airspeed and angle of climb setpoint is achieved by applying a fixed power factor,

$$P_b = \eta_e \cdot P_p \quad (6.14)$$

where η_e represents power conversion efficiency and P_b represents net battery output power.

A proportional factor is also introduced here to correct for discrepancies between the drag

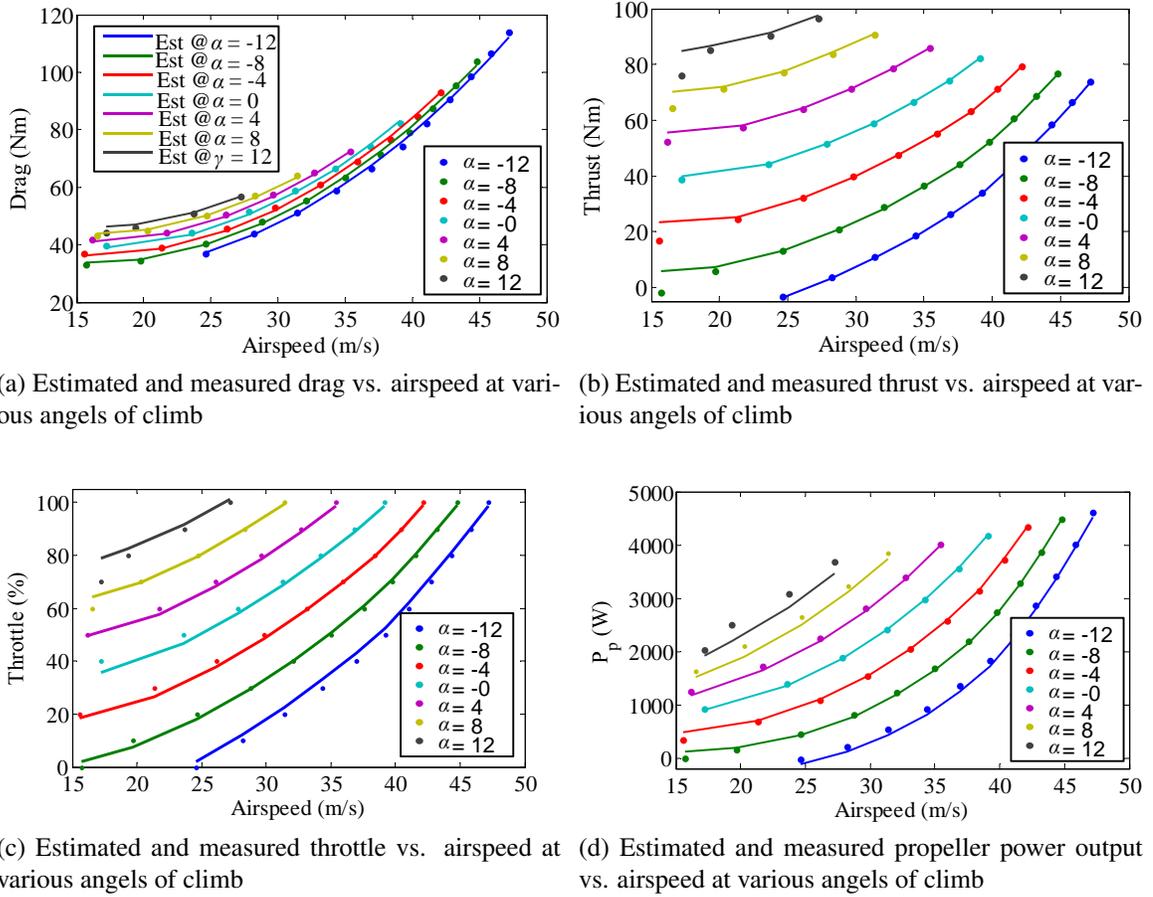


Figure 6.4: Model fitting results from X-Plane flight load characterization tests

models given in Eqn. 6.11, and the drag forces estimated for the actual aircraft. The corrective factor is expressed as:

$$D_A(v, \gamma) = \lambda_D \cdot D_M(v, \alpha) \quad (6.15)$$

where D_A and D_M represent the drag force estimated for the actual aircraft and drag force estimated for the X-Plane model respectively. λ_D represents a constant corrective factor that may be fitted by comparing modeled and actual aircraft powertrain load dynamics over sample flights.

A roughly proportional deviation between the modeled and actual aircraft drag force is attributed to slight errors in modeling the aircraft geometry and surface aberrations. Small errors in modeling the true aircraft drag will cause only small effects on the aircraft han-

dling from the perspective of a pilot or an autopilot, thus the drag correction need not necessarily be made for the SIL testing of inner-loop controllers. However, small errors in approximating the loads on onboard energy storage devices will accumulate into large errors over a simulated flight.

The substitution of Eqns. 6.14 and 6.15 into Eqns. 6.10 - 6.13 yields a model for the net battery power required to fly at a particular airspeed and angle of climb.

$$\begin{aligned}
 P_B &= \frac{1}{\eta_e \eta_p} \cdot T_{x_w} \cdot v \\
 P_B &= \frac{v}{\eta_e \eta_p} \cdot (D_A(v, \alpha) + mg \cdot \sin(\alpha) + m\dot{v}) \\
 P_B &= \frac{v}{\eta_e \eta_p} \cdot (\lambda_D D_M(v, \alpha) + mg \cdot \sin(\alpha) + m\dot{v})
 \end{aligned} \tag{6.16}$$

The net battery power, P_B , input to aircraft ESCs is given by the summation of battery power input to ESC1 and ESC2. The proportion of the net battery output powers that goes to each ESC is represented by:

$$\lambda_{ESC} = \frac{P_1 + P_2}{P_3 + P_4} \tag{6.17}$$

where λ_{ESC} represents the ratio of battery power drawn by each of the onboard ESCs.

The power output from the two strings of series connected battery packs is equal to the product of current and voltage,

$$P_B = I_{1,2} \cdot (V_{B1} + V_{B2}) + I_{3,4} \cdot (V_{B3} + V_{B4}) \tag{6.18}$$

where V_{Bi} represents the terminal voltage of battery i , $I_{1,2}$ and $I_{3,4}$ represent the current flowing through the two sets of series connected batteries.

Substitution of Eqn. 6.18 into Eqn. 6.17, and solving for I gives:

$$\begin{aligned}
 I_{1,2} &= \frac{\lambda_{ESC} P_b}{(\lambda_{ESC} + 1) \cdot (V_{B1} + V_{B2})} \\
 I_{3,4} &= \frac{P_b}{(\lambda_{ESC} + 1) \cdot (V_{B3} + V_{B4})}
 \end{aligned} \tag{6.19}$$

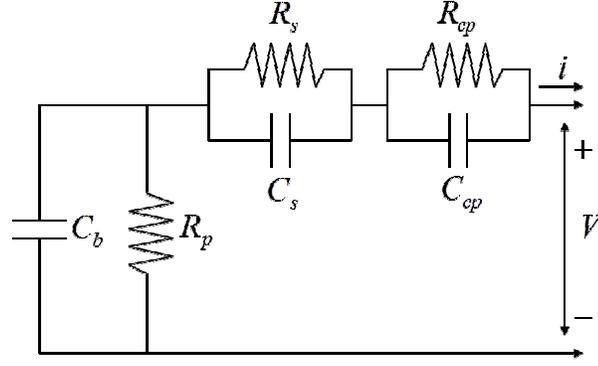


Figure 6.5: An equivalent circuit battery model

This represents an estimate of the current loads on each of the series connected battery packs, given an estimate of the battery power output required to fly a particular maneuver and knowledge of the division of power between the two propeller motors.

6.3.3 An Equivalent Circuit Battery Model

The equivalent circuit model shown in Figure 6.5 is used to replicate battery current and voltage dynamics as a function of estimated battery state of charge (SOC). The equivalent circuit model used here is an extended version of the model explained in [24]. This battery model uses six electrical components that are tuned to recreate the observed current-voltage dynamics of Edge 540T powertrain batteries. Battery charge is stored in the capacitor, C_b . The R_s, C_s and R_{cp}, C_{cp} circuit element pairs capture battery internal resistance drops and concentration polarization effects, respectively. The resistor R_p accounts for the slow battery self-discharge that is seen to occur over weeks or months of storage.

The current and voltage dynamics of the equivalent circuit model are defined as:

$$x^B = \begin{bmatrix} q_b & q_{cp} & q_{Cs} \end{bmatrix}^T \quad (6.20)$$

$$\dot{x}^B = \begin{bmatrix} -\frac{1}{C_b R_p} & \frac{1}{C_{cp} R_p} & \frac{1}{C_s R_p} \\ \frac{1}{C_b R_p} & -\frac{1}{C_{cp} R_p R_{cp}} & \frac{1}{C_s R_p} \\ \frac{1}{C_b R_p} & \frac{1}{C_{cp} R_p} & \frac{1}{C_s R_p} \end{bmatrix} x + \begin{bmatrix} i \\ i \\ i \end{bmatrix} \quad (6.21)$$

$$y^B = V_p = \left[\frac{1}{C_b} \quad \frac{1}{C_{cp}} \quad \frac{1}{C_s} \right] \cdot x \quad (6.22)$$

where q_b , q_{cp} , and q_{cs} represent the charge stored in capacitors C_b , C_{cp} , and C_{cs} respectively. The total voltage drop across the battery terminals, V_p , is given by the sum of the voltage drops across the each of the three capacitors in the equivalent circuit model.

Because battery current-voltage dynamics are known to vary as a function of battery SOC, some of the resistive and capacitive (RC) components in the equivalent circuit model must be parameterized as functions of battery SOC [100]. It was decided based on qualitative observation that defining C_b , C_{cp} , and R_{cp} as parameterized functions of battery SOC gave an acceptable trade-off between the number of parameters to be identified and the resulting model error.

Battery SOC is defined as:

$$\text{SOC} = 1 - \frac{q_{max} - q_b}{C_{max}} \quad (6.23)$$

where q_b is the charge stored in capacitor C_b , q_{max} is the maximum charge of the battery, and C_{max} is the maximum charge that can be drawn from the battery. The term coulombic efficiency is used to refer to the portion of stored charge that can be withdrawn over repeated charge and discharge cycling of a battery. Resting a battery can temporarily unlock some of its lost charge storage capacity, however the overall trend is inevitably downward.

The C_b , C_{cp} and R_{cp} terms in the equivalent circuit battery model are parameterized as:

$$C_b = C_{Cb0} + C_{Cb1} \cdot \text{SOC} + C_{Cb2} \cdot \text{SOC}^2 + C_{Cb3} \cdot \text{SOC}^3 \quad (6.24)$$

$$C_{cp} = C_{cp0} + C_{cp1} \cdot \exp(C_{cp2}(1 - \text{SOC})) \quad (6.25)$$

$$R_{cp} = R_{cp0} + R_{cp1} \cdot \exp(R_{cp2}(1 - \text{SOC})) \quad (6.26)$$

Each battery pack used in Edge 540T flight tests should be characterized individually prior

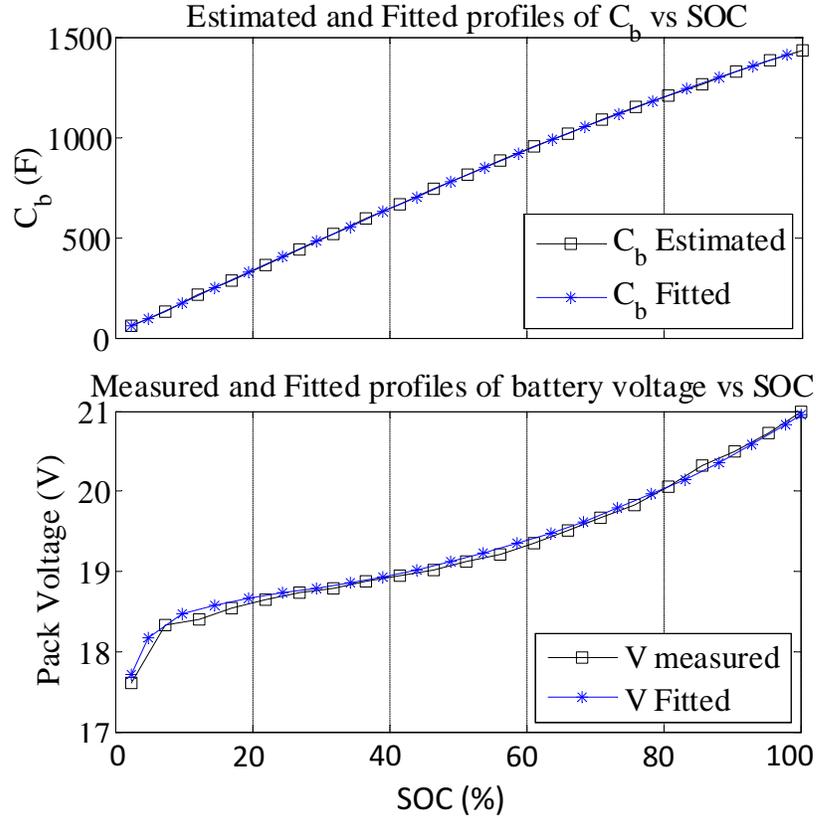


Figure 6.6: Measured and fitted profiles for C_b and battery terminal voltage

to testing, in order to account for any manufacturing and SOH variations. Two battery characterization experiments are used to identify the RC parameters in the battery equivalent circuit model.

The first experiment is a low current discharge of a battery from a fully charged state until a cutoff voltage of 17.5V is reached. This type of discharge is mostly affected by the C_b , q_b , q_{max} , and C_{max} parameters in the model. Figure 6.6 shows a polynomial fit of C_b as a function of SOC, and the battery voltage fit for the tuned parameter values: C_{Cb0} , C_{Cb1} , C_{Cb2} , C_{Cb3} , q_{max} , and C_{max} .

Next, a pulsed loading experiment is used to fit the remaining parameters in the equivalent circuit model to the observed changes in battery hysteresis behavior as a function of SOC. A gradient descent search is used to identify the remaining model coefficients using the

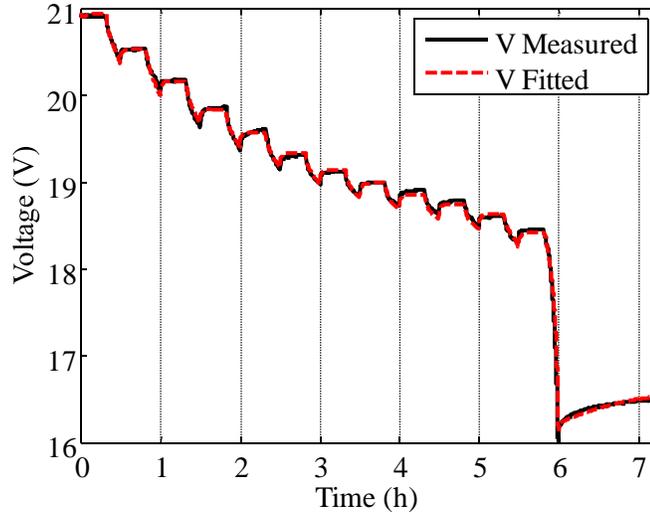


Figure 6.7: Measured and fitted profiles for terminal battery voltage over a pulsed load test

Table 6.2: Parameter definitions for the equivalent circuit battery model

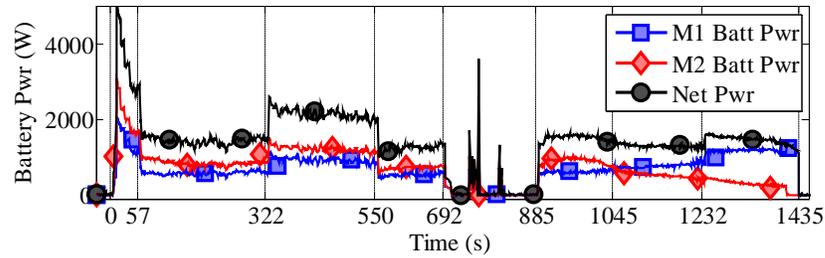
Parameter	Value	Parameter	Value
q_{max}	2.88×10^4 C	C_s	89.3 F
C_{max}	2.85×10^4 C	R_{cp0}	1.60×10^{-3} Ω
C_{Cb0}	19.4 F	R_{cp1}	8.45
C_{Cb1}	1576 F	R_{cp2}	-61.9
C_{Cb2}	41.7 F	C_{cp0}	2689 F
C_{Cb3}	-203 F	C_{cp1}	-2285 F
R_s	2.77×10^{-2}	C_{cp2}	-0.73 F

pulsed loading data. Figure 6.7 shows the battery voltage fit over a pulsed loading profile, using the tuned parameters identified in the low current experiment, and the newly tuned values of R_s , C_s , R_{cp0} , R_{cp1} , R_{cp2} , C_{cp0} , C_{cp1} , and C_{cp2} .

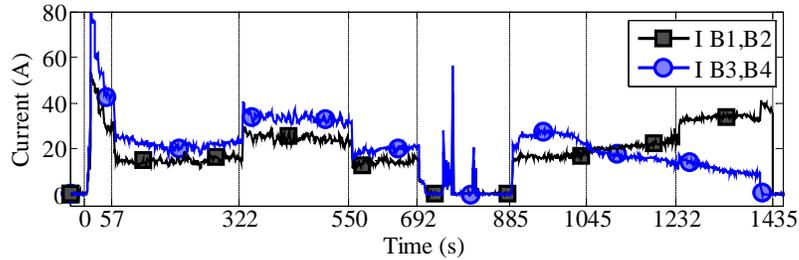
Values for all of the RC components and parameterization coefficients used in the equivalent circuit model of an Edge battery are defined in Table 6.2.

6.3.4 Battery Charge Depletion Modeling

Outer-loop supervisory health management routines are considered to be focused on the depletion of battery charge. The charge remaining in the aircraft's batteries will be estimated online using domain knowledge and periodic measurements of battery current and



(a) Individual and net battery power supplied to motor ESCs



(b) Measured current through each powertrain battery

Figure 6.8: Plots of individual and net battery power consumption (top), and battery current load (bottom) measured over a sample flight

voltage.

In offline testing, the current-voltage dynamics of aircraft batteries may be observed over simulated flights by loading the aircraft batteries with a current profile indicative of flight loads. It would also be necessary in offline tests to assure that the SOH and thermal loading of the batteries under test is similar to what will be encountered in online testing, as battery dynamics are expected vary substantially as a function battery health and temperature [49]. Similarly, offline replication of battery discharge dynamics using SIL battery models would require that the models be tuned using data from batteries at similar SOH and thermal loads to what is anticipated during online tests.

Battery power and current draw over a piloted flight of the Edge 540T are shown in Figures 6.8a and 6.8b respectively. The ticks on the x-axis denote the initiation of high level maneuvers. The sample flight consists of eight activities initiated by a pilot flying the Edge 540T in remote control mode.

1. Take-off and climb to 200 ft from 0 to 57 s

2. Circling flight in auto-mode with throttle set to 75% from 57 to 322 s
3. Throttle is increased to 85% from 322 to 550 s
4. Throttle is decreased to 75% from 550 to 692 s
5. Aircraft lands from 692 to 722 s
6. Aircraft taxis down the runway from 722 to 885 s
7. The tail of the aircraft is held on the runway and the throttle is set to 75% from 885 to 1232 s
8. The throttle is increased to 80% and held there until battery voltages drop below 17V from 885 to 1232 s

A detailed description of the sample flight and collected data is given in Appendix E.

An asymmetric loading of the two propeller motors over the sample flight is apparent from the battery power loading profiles shown in Figures 6.8a and 6.8b. Motor M2 is known to consistently draw more current than motor M1 on the Edge 540T, due to unregulated coupling of the two motor speed controls (ESCs).

Figure 6.9 shows open loop predictions of battery voltage profiles for B1 and B3 obtained using the equivalent circuit model described in Section 6.3.3 and recorded battery current profiles plotted in Figure 6.8b. The open-loop voltage predictions are seen to lie on top of the observed battery voltage profiles, providing a measure confidence in the accuracy of the software models.

Estimates of both battery terminal voltage and internal SOC will be further improved through the use of a closed-loop state estimation technique, such as Kalman filtering [24]. The output of such closed-loop state estimation techniques will be much less susceptible to initialization and measurement errors than the Coulomb counting method currently used in many battery monitoring systems.

Figure 6.10 shows the mean estimate of SOC for batteries B1 and B3 given by a Kalman

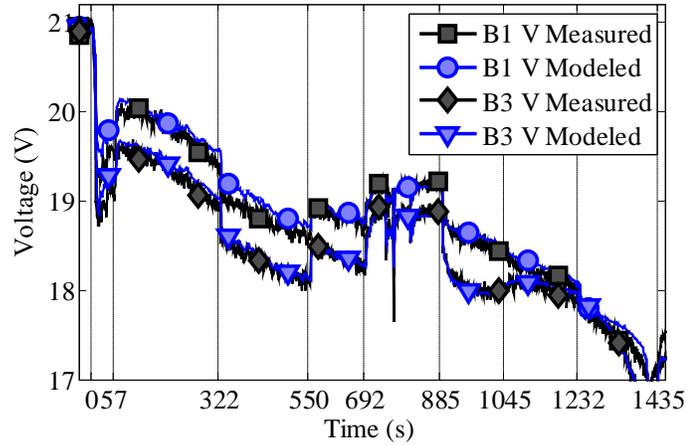


Figure 6.9: Modeled and measured voltages of batteries B1 and B3 over a sample flight

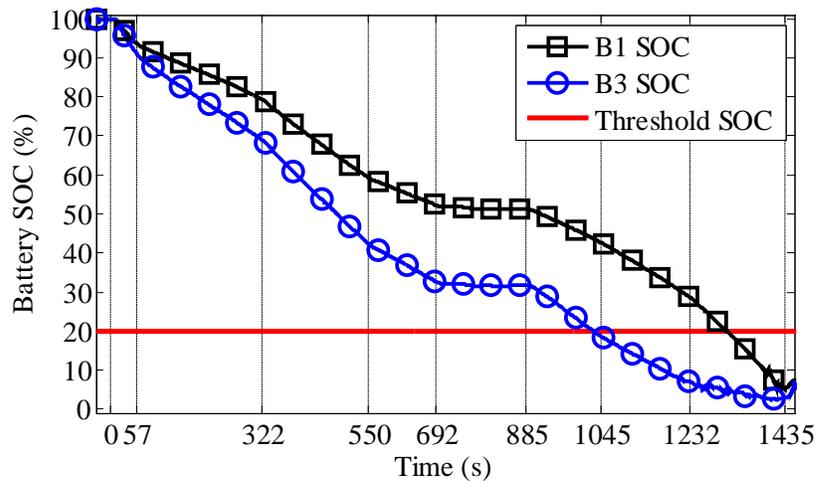


Figure 6.10: Estimated SOC for batteries B1 and B3 over a sample flight

filter with small process noise and measurement noise priors. The prior distribution for process noise is assumed to be small due to high confidence in the fitted battery model, and the measurement noise is assumed to be small due to an assumption of accurate current and voltage sensing. The low noise assumption results in negligible uncertainty around SOC estimates in this case-study. The red line in the figure shows the threshold SOC, after which the batteries will be considered no longer safe for supporting flight activities. The SOC threshold is defined to be 20% SOC for the experiment described here.

6.3.5 Prognostic Prediction

The generation of battery RUL estimates, with appropriate uncertainty bounds, is considered here. Uncertainty in state estimation and battery dynamics is assumed to be negligible for the fitted battery equivalent circuit model described in Section 6.3.3. Therefore, uncertainty in the battery discharge prognostic estimates presented here is caused solely by uncertainty in estimates of the future loads to be placed on system batteries.

Eqns. 6.10 - 6.16 provide for the estimation of battery power consumption given an angle of climb, airspeed, and acceleration dictated by each portion of a flight plan. Steps 1-5 of the sample flight are annotated below with approximated values for the angle of climb, airspeed, and acceleration, used to approximate the steady state battery power demand over each flight segment. The inversion of the throttle model given in Eqn. 6.12 is used to identify a steady state aircraft airspeed expected to correspond to given throttle and angle of climb set points.

1. takeoff and climb to ~ 200 meters (duration = 57 s) ($\alpha = 2.8^\circ$, $v_0 = 0 \frac{m}{s}$, $\dot{v} = 0.4 \frac{m}{s^2}$)
2. maintain altitude, set throttle to 75% (duration = 265 s) ($\alpha = 0^\circ$, $v = 33 \frac{m}{s}$, $\dot{v} = 0 \frac{m}{s^2}$)
3. maintain altitude, set throttle to 85% (duration = 228 s) ($\alpha = 0^\circ$, $v = 36 \frac{m}{s}$, $\dot{v} = 0 \frac{m}{s^2}$)
4. maintain altitude, set throttle to 75% (duration = 142 s) ($\alpha = 0^\circ$, $v = 33 \frac{m}{s}$, $\dot{v} = 0 \frac{m}{s^2}$)
5. land (duration = 30 s) ($\alpha = -3.2^\circ$, $v_0 = 33 \frac{m}{s}$, $\dot{v} = -1.1 \frac{m}{s^2}$)

Figure 6.11 shows measured and predicted battery power demand required over steps 1-5 in the sample flight plan, denoted S1-S5 in the figure. The dashed line in the plot denotes the steady state power demand estimated using the given values of angle of climb, airspeed, and acceleration for each step. The proportional constants η_e and λ_D used in Eqn. 6.16 were fitted to $\eta_e = 0.85$ and $\lambda_D = 0.9$ to obtain the model fit shown. The green boxes in the figure show a $\pm 30\%$ envelop around the estimated steady state power demand for each step in the flight plan. Prognostic predictions will be performed under the assumption that future

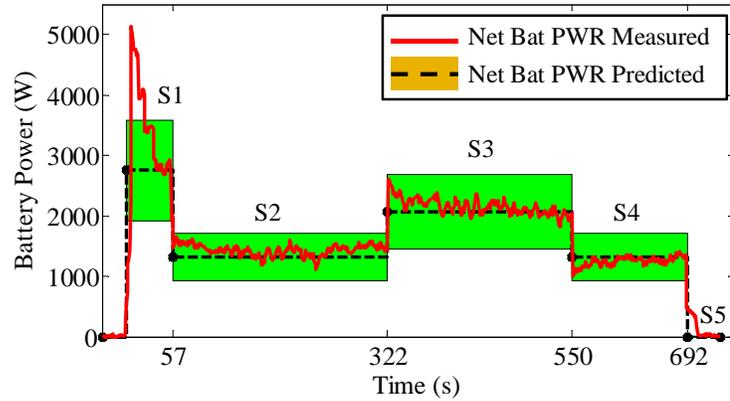


Figure 6.11: Measured and predicted net battery power consumption over steps 1-5 of the sample flight plan (annotated S1-S5 in the figure).

power demand will be selected randomly each second from a uniform distribution spanning the $\pm 30\%$ envelop shown in the figure. This assumed stochastic model for future power demand is intended to account for the unmodeled dynamics in the rather simple model for future power demand given in Eqns. 6.10 - 6.16. The measured battery power draw over the sample flight plan, shown with the red line in Figure 6.11, is seen to stay within the $\pm 30\%$ envelop to be used for prognostics, except during the peak loads occurring for the first several seconds of the flight.

Battery RUL predictions will be made by assuming that the aircraft will continue to fly at the same speed as it did in step 4 of the flight plan until one of the battery packs reaches the end of discharge (EOD) condition (min battery SOC estimate $\leq 20\%$). In reality, a landing maneuver is initiated at 692 seconds, then the aircraft taxis down the runway and is held in place while the throttle is controlled such that the batteries are placed under similar loads to those observed in step 4 until all battery voltages drop below 17 V. This procedure is described in steps 6-8 of the flight plan described in Section 6.3.4.

Figure 6.12 shows the measured and predicted net battery power consumption over the sample flight plan. The net battery power observed during the ground discharge phase of the flight plan is seen to fall within the estimated range of battery power demands over the sample mission.

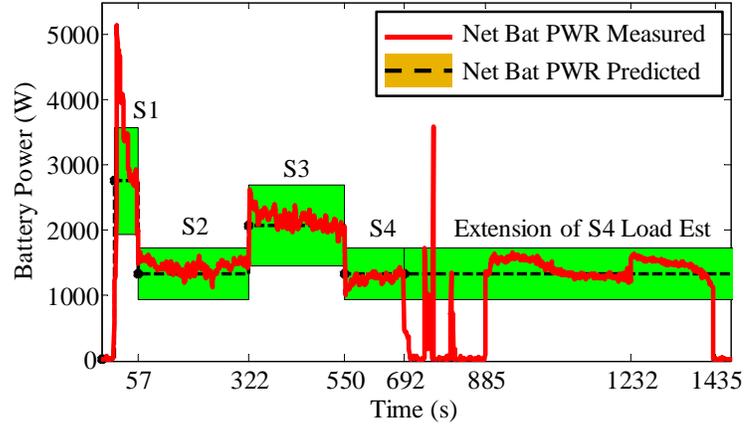


Figure 6.12: Measured and predicted net battery power consumption over the sample flight plan.

The purpose of loading the aircraft powertrain on the ground immediately after landing is to safely obtain an approximate measurement for the amount of flight time that would have been supported by the battery pack if the aircraft had continued to be flown. This measurement allows comparison between battery EOD predictions made at various points over the sample mission, and the EOD time observed experimentally. It is seen in Figure 6.10 that the EOD condition is reached at 1039 seconds after the beginning of the flight. Subtracting the time spent landing and taxiing on the runway from 1039 seconds, gives an approximate measurement for the time at which the EOD condition would have been observed if the aircraft had continued to be flown at the same speed as in step 4 of the flight plan until EOD,

$$1039 - (885 - 692) = 846 \text{ seconds} \quad (6.27)$$

A discrete Markov model, of the form described in Chapter 5, is used to propagate future battery power loading uncertainty into estimates of future SOC. A discrete Markov model is formulated over the quantized space: $q_{max} \leq q_b \leq (q_{max} - C_{max})$, where q_b was defined to be the charge stored in capacitor C_b in the equivalent circuit model described in Section 6.3.3. State transition probabilities are found using Eqn. 6.19 to map the stochastic battery power demand profile to stochastic estimates of future battery current demands.

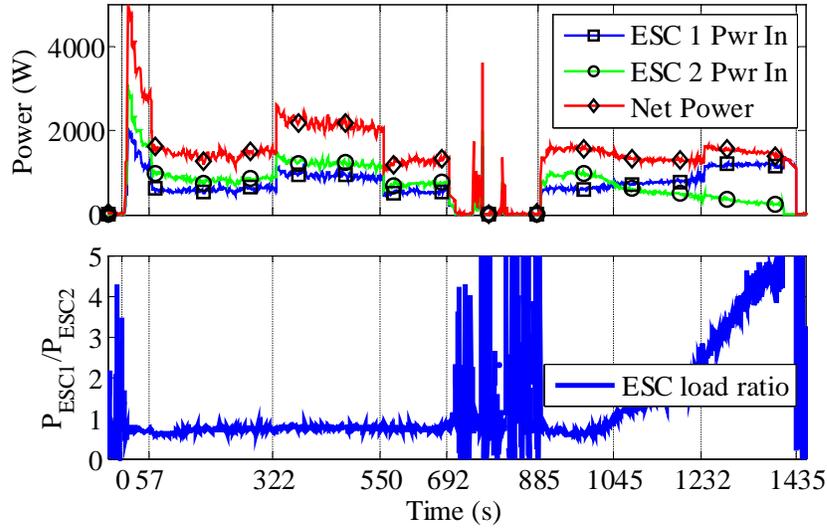


Figure 6.13: Measured battery power input to ESCs (Top) and observed ESC power ratio over a sample flight (bottom)

The dynamics of q_b were defined in terms of battery current in Eqns. 6.20 - 6.22. However, this equation is missing a fitted value for the ESC power split coefficient, λ_{ESC} . Figure 6.13a shows the observed net battery power inputs to ESC 1 and ESC 2 over a sample flight. Figure 6.13b shows the ratio of power drawn from ESC 1 and ESC 2. The split of power between the two ESCs is currently uncontrolled, but it is seen to remain close to a value of $\lambda_{ESC} \approx 0.7$ until ~ 1045 seconds into the sample flight. The EOD condition to be predicted by prognostic data occurs at 1039 seconds, which is right at the point where λ_{ESC} starts to move away from $\lambda_{ESC} = 0.7$. It remains reasonably accurate however for prognostic routines to assume a value of $\lambda_{ESC} = 0.7$ over the allowed SOC range of the batteries. The approximation for λ_{ESC} to be used in SIL and HIL testing of the vehicle power train may be improved in future work by incorporating possible dependencies on time, battery pack voltage, throttle command, and other inputs control inputs.

Figure 6.14 shows predictions for the future depletion of SOC in battery B3, and the estimated SOC in B3, using Kalman filtering, over the sample flight. Uncertainty in battery SOC depletion predictions is represented in Figures 6.14a - 6.14d using dashed lines to denote the 5% and 95% confidence VaR points in predictions of future battery SOC. The

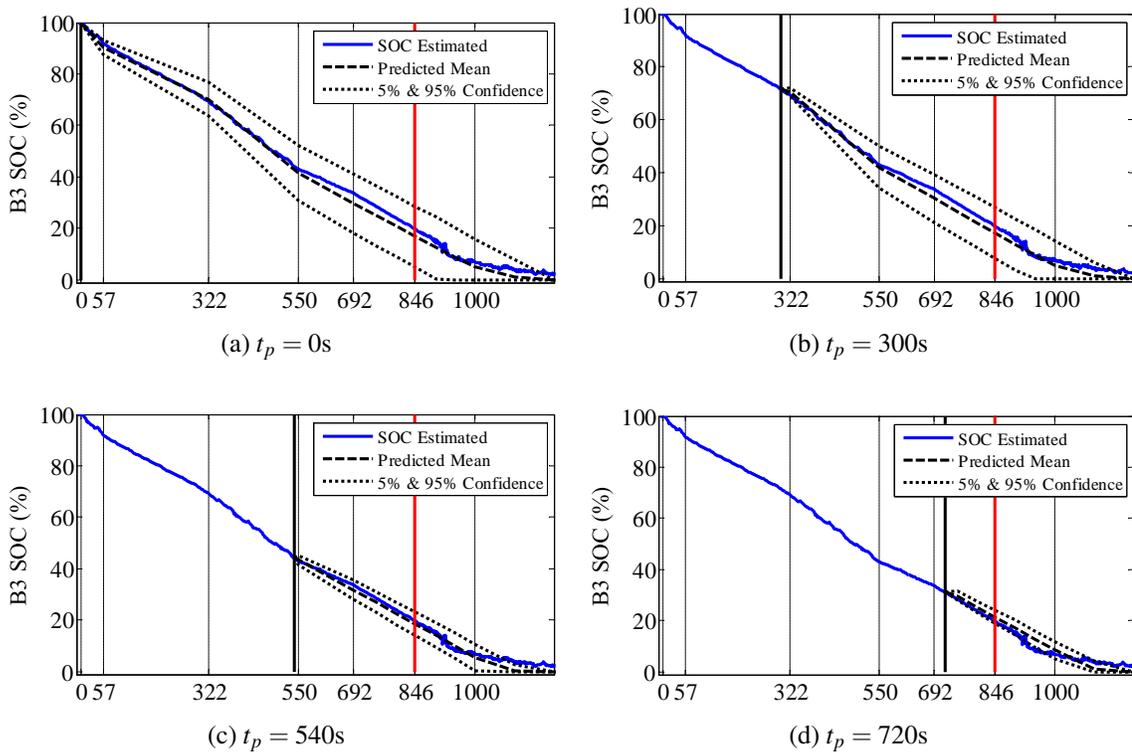


Figure 6.14: Plots of estimated state of charge for B3 versus predicted state of charge at various points over the sample flight

back line in the plots denotes the time at which SOC depletion predictions are made, and the red line denotes the time at which the SOC of battery B3 crosses the EOD threshold.

It appears from the plots shown in Figure 6.14a - 6.14d, that the estimated battery SOC always falls well within the 5% - 95% confidence estimates shown for future SOC predictions. It is also seen that predictions made later in the mission have much less uncertainty that predictions made earlier in the mission, as should be expected.

Figure 6.15 shows predictions for the SOC of battery B3 at $t = 846$ seconds, plotted at 60 second intervals over the sample flight. The red lines in the figure denote the area between the 5% and 95% VaR estimates for the predicted SOC at 846 seconds. Predictions are shown at 60 s intervals. The dashed line in the figure shows the actual battery SOC measured at 846 seconds. The actual value of B3 SOC was 20% at $t = 846$ seconds. The gray cone in the figure denotes a 30% relative accuracy cone. The relative accuracy cone is

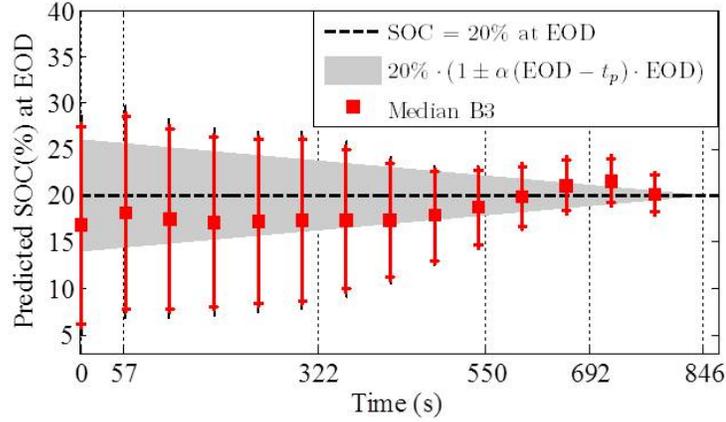


Figure 6.15: Predictions of battery B3 state of charge at $t = 846$ seconds is plotted at 60 second intervals over the sample flight.

defined in terms of a metric α as:

$$RA_{\alpha} = z_T \pm \alpha \cdot \frac{(T - t)}{T} \quad (6.28)$$

where z_T represents a measurement taken at time T and α is an accuracy modifier. In Figure 6.15, T is set to 846 seconds, z_T is set to 20% SOC, and α is set to 30%. It can be seen from the figure that the median SOC prediction stays within the 30% relative accuracy cone shown, although the relationship between uncertainty in SOC predictions and the length of the prognostic horizon appears grow faster than the relative accuracy cone.

The prognostic result shown here is a considerable improvement over previous particle filter based implementations of battery EOD prognostics, described in [74, 73, 72], which used the average of battery current over a finite window to estimate the future battery loading over a flight. Not only are the predictions more accurate, but they are also more stable.

6.3.6 Online Controls Testing

When conducting purely simulation based testing of supervisory control algorithms, control system designers may run a simulation or a batch of simulations, and assess algorithm performance upon completion. However, during testing procedures that involve real hard-

ware, the testing personnel will require some real-time interfaces to monitor a test and determine if and when it should be aborted.

A set of system status and battery prognostics displays are presented in this section to illustrate how supervisory prognostics-based decision making algorithms may be monitored in real time during a test. Because the X-Plane simulator nominally runs in real time, the development and testing of prognostics displays for electric aircraft applications is easily incorporated into the SIL/HIL testing framework described in Section 6.2 of this chapter.

The displays described here give test operators both high confidence estimates of the vehicle's current states, and uncertain predictions of future system states. Uncertainty in the future input model is represented in the displays shown here by a uniform distribution extending to $\pm 30\%$ of the predicted future battery power demand profile estimated for a predefined flight plan. Corresponding predictions of future battery state evolutions are reported in terms of mean, 5%, and 95% confidence VaR estimates. The actual confidence bounds used will depend on the conservatism of the operator; perhaps they will be interested in the 2% and 98% confidence VaR estimates used in Chapter 4.

Figure 6.16 and Figure 6.17 show a set of operator displays observed near the beginning and near the end of a sample mission to be flown by the Edge 540T UAV respectively. The sample mission considered here is plotted in Figure 6.18. Figure 6.18 shows an ordered set of 17 waypoints that specify aircraft transit speeds, acceleration, and angle of climb over the sample mission. This sample mission consists of several periods of straight and level flight, four ± 8 degree climbing and descending maneuvers, and two ± 4 degree climbing and descending maneuvers.

The right most graphic in the bottom window shown in Figures 6.16 and 6.17 plots the ordered set of waypoints to be visited by the aircraft along with past aircraft position measurements. The top left plot in the bottom window shows the measured and desired air speed of the aircraft over the sample mission. The bottom left plot in the bottom window

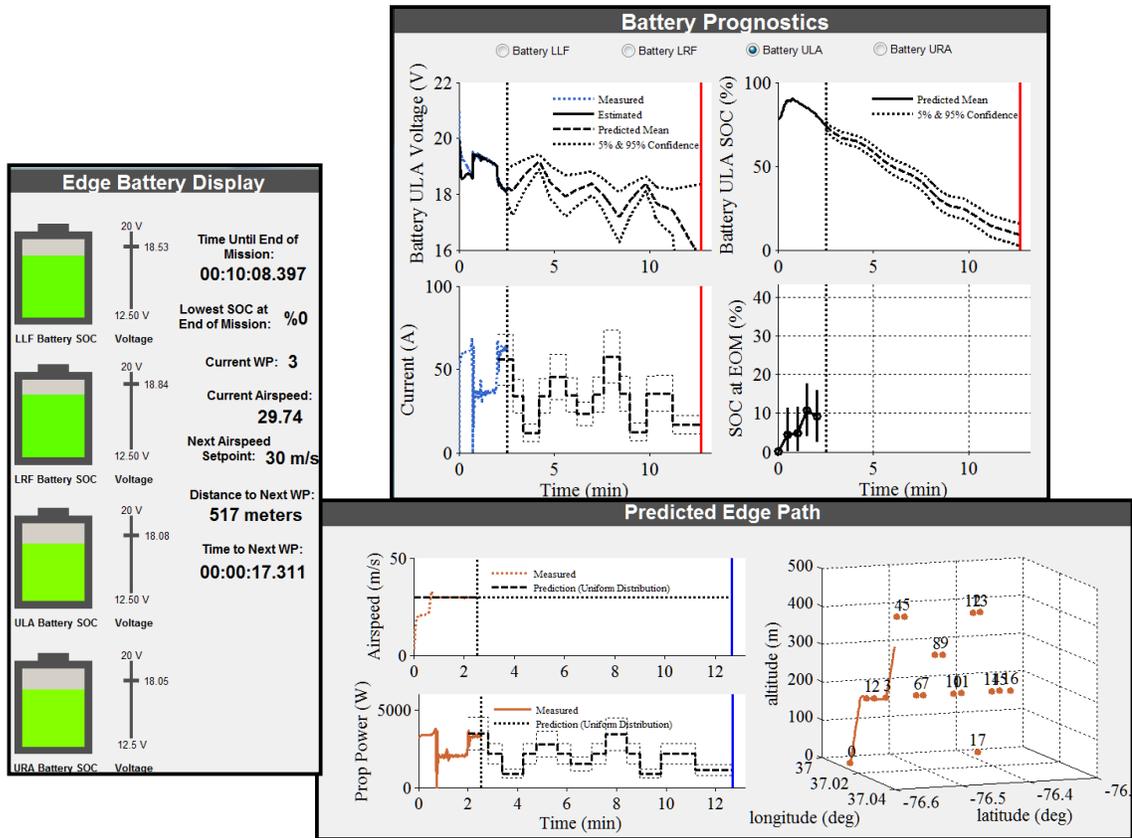


Figure 6.16: Operator displays captured near the beginning of a sample flight

shows the measured and predicted propeller power (with a 30% uncertainty factor included) over the sample mission.

The top window in Figures 6.16 and 6.17 shows detailed prediction data for a selected powertrain battery. Predictions are made using the assumption that future battery power loads are selected each second from a uniform distribution that spans $\pm 30\%$ of the mean battery power loads predicted, as described in Section 6.3.5. The top left plot in the top user display window shows the measured battery voltage, an estimate of what the battery voltage should be, using the battery model described in Section 6.3.3. The mean, 5%, and 95% VaR battery voltage predictions are also shown over the remainder of the mission.

Past battery SOC estimates and uncertain future predictions are shown in the top right plot in the top user display window. In the experiment shown here, the aircraft batteries start a mission with 100% SOC, but the battery models are initialized to 80% SOC. The

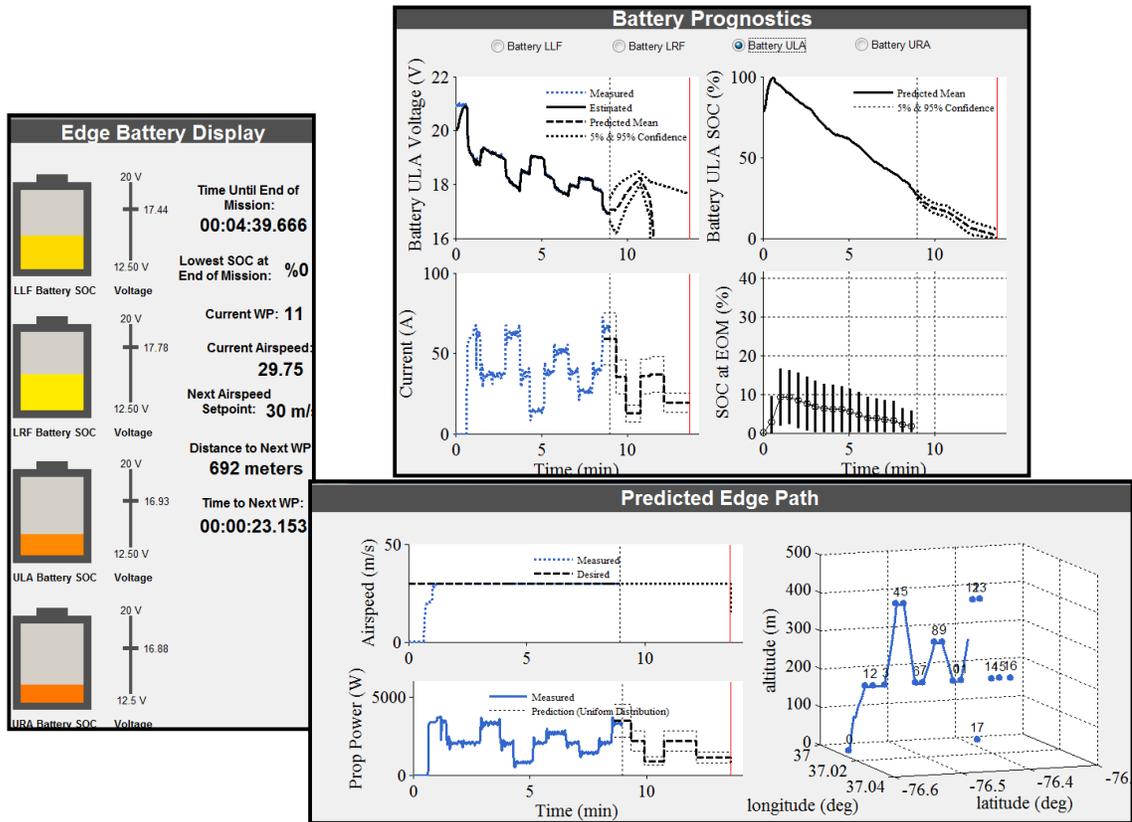


Figure 6.17: Operator displays captured near the end of a sample flight

initialization error is corrected in the first minute of the mission by the Kalman filtering routine used for state updates. The initialization error and the state corrections are seen by observing the convergence of the measured and modeled battery voltage over the first minute of the sample mission. The bottom right plot in the top display window shows past measurements and future predictions for to be drawn from the battery over the future mission. The plot in the lower right side of the top display window show the evolution of battery SOC at EOM estimates over the mission.

Comparison of the predictions made early in the sample mission (Figure 6.16) to the actual battery measurements and internal states observed later in the sample mission (Figure 6.16) shows the actual measurements to mostly fall within the assumed 5% and 95% confidence bounds. This result provides some measure of validation for the prognostic model used.

Similarly to what was shown in Figures 6.14 and 6.15 in Section 6.3.5, the SOC at EOM

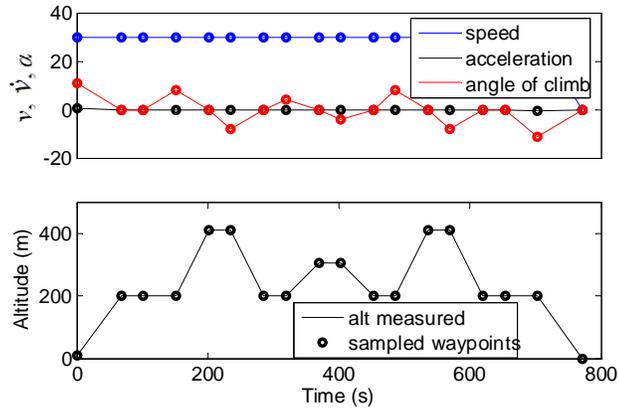


Figure 6.18: Plots of the desired aircraft speed, acceleration, angel of climb, and altitude waypoints that define a sample mission

measurements are seen to converge as the EOM draws near. In this case, it seems that the SOC at EOM estimates are converging to an estimate of 0% SOC remaining in the powertrain batteries. The setup of this example mission is similar to that of the missions analyzed in Chapters 4 and 5. The risk associated with not taking any supervisory control actions is assessed to be very high. However some relatively small deviations from the nominal system control could make the mission satisfiable.

6.3.7 Bounded Load Allocation

Assuming the prognostic model for battery charge depletion as a function of flight plan is sufficiently validated in SIL and HIL testing. The bounded load allocation problem may now be formulated. The ‘load’ variable, which drives both system dynamics and fault growth in this example, can be considered to be the instantaneous battery power draw from the powertrain batteries.

The nominal battery loading associated with a mission defined in terms of 3D waypoints and desired transit speeds is found using the load inversion approximation given in Eqn. 6.16. As was shown in Chapters 4 and 5, a range of feasible deviations from the nominal battery load demand at each time instant may be defined in terms of a relative performance

metric ρ ,

$$\tilde{P}_B = \rho \cdot P(v, \alpha) \quad (6.29)$$

where \tilde{P}_B represents a modified battery power output command, and $P(v, \alpha)$ represents the estimated net battery power required to follow a given vehicle flight plan.

The methodology demonstrated on the more simplistic case study examples discussed in Chapters 4 and 5, could also be followed at this point to quantify risk metrics, assess mission feasibility, and formulate a supervisory risk management optimization problem.

CHAPTER VII

CONCLUSIONS

The overarching analytical focus of the systems health management approach described here was on defining example scenarios that are known to present a high likelihood of failure if supervisory control actions are not taken, but which also contain a set of safe operating modes achieved by degrading nominal system performance. The challenge then is to identify prognostics-based supervisory control policies that are safe without unnecessarily degrading nominal system performance. Repeated trials were shown to facilitate quantitative performance analyses for given risk evaluation metrics and prognostics-based health management policies. The empirical analytical tools presented here represent a substantial contribution to a community that is sorely lacking in repeatable application examples.

An overview of the primary contributions to the state of the art that are claimed for this work are enumerated and explained here. Key assumptions are noted for each claim.

Claim #1 - Description of Prognostics-Based Risk Management Using a Retrofit Nominal Load Modification Architecture

The search space for supervisory health management actions was defined in terms of allowable modifications to a nominal component load allocation policy. Description of component deterioration dynamics in terms of component load allocations was shown to be notationally convenient and able to be applied across different applications. Defining health management actions in terms of modifications to a nominal system is shown to provide a clear means of benchmarking system input-output performance goals.

The utility of the search for bounded component load allocations relies on an assumption that limited deviations from the nominal input-output performance of a given system will yield sufficiently safe control outcomes. The scope presented here does not include higher level fault management actions such as changing a system's mission or scheduling a repair

operation. The formulation of the health management problem in terms of component load allocations also relies on a lower level control layer that would invert the load allocation dynamics of a given nominal control system.

Claim #2 - Analysis of Supervisory Health Management Using Finite Horizon Prognostic Estimates

The tuning finite horizon risk metrics in order to identify prognostics-based health management policies that mitigate failure risks without being overly conservative was demonstrated on two simulated systems with multiple effectors at different states of health. Analysis of prognostics-based component load optimization on the two example systems was used to illustrate the utility of risk assessments made at various prognostic horizons. Assessing prognostic risk using remaining useful life estimates, which is seen in almost all other published prognostics-based control discussions, was shown to be of limited usefulness in the two case study examples analyzed.

The empirical analyses used here rely on an assumption that an objective means of evaluating control outcomes is available. The value of the iterative empirical risk metric tuning approach presented also relies on the availability of a trustworthy means of simulating system dynamics and sample health management control policies.

Claim #3 - Demonstration of the Explicit Incorporation of Future Demand Models into Health Management Policies

The incorporation of stochastic models for future environmental loading into a risk optimizing control policy was explored using a Markov decision process formulation of the supervisory load allocation problem. Causal and non-causal implementations of dynamic programming were shown to solve for optimal component load allocation policies with and without future uncertainty respectively. Comparison of the performance of the causal and non-causal optimal policies served to illustrate the limiting effects of model uncertainty on health management controllers.

The analysis of the Markov decision process formulation of the prognostics-based load allocation problem relies on the assumption that all of the system's stochastic dynamics can be encoded as a discrete fully observable Markov process. State transition costs used in the Markov decision process formulation are assumed to encode fault risk aversion and performance degradation trade-offs. These assumptions are described to result in burdensome formulation and computational challenges. However, the ability to directly compute optimal policies from this formulation is shown to offer analytical advantages that make the approach attractive in spite of the implementation challenges described.

Claim #4 - Real-World Demonstration of Stochastic Component Health Depletion Modeling in terms of Uncertain Future Demand Estimates

Finally, the implementation and testing of a real-world prognostic system was described to illustrate model development challenges not directly addressed in the analysis of the simulated case study systems. The problem of battery charge depletion prediction for an electric aircraft was shown to offer a meaningful and repeatable failure case. Uncertainty propagation in charge depletion predictions was shown to produce a projection of future battery charge depletion with confidence bounds that in early testing appeared reasonable.

Given the prognostic model and offline simulation framework demonstrated for the real-world charge depletion failure mode example, there are still two prerequisites that must be implemented on the case study before the analytical techniques for supervisory health management that are described in this document can be applied. First, a mechanism for modifying the nominal load allocated to aircraft batteries must be identified. Second, metrics for evaluating a relative preference for control outcome distributions must be established. After these prerequisites are satisfied the supervisory health management problem can be stated in the format demonstrated in the three simulated case study systems. The amount of development effort necessary to get the real-world example to that point illustrates why the research focused on simulated examples for the initial exploration of the proposed system health management paradigm.

APPENDIX A

RESOLVING ACTIVE REDUNDANCIES IN THE LOAD ALLOCATION PROBLEM

The separation of component loading and system output regulation tasks is described in this appendix for a generic nonlinear system of the form:

$$\dot{x} = A(x) + B(x)u + C(x)w \quad (\text{A.1})$$

where $A(x) \in \mathbb{R}^n$, $B(x) \in \mathbb{R}^{n \times m}$, $C(x) \in \mathbb{R}^{n \times l}$, $x(t) \in \mathbb{R}^n$, is the state, $u(t) \in \mathbb{R}^m$ is the control effort or load on each of the m components in the system, and $w(t) \in \mathbb{R}^l$ represents the forces exerted on the system by its operating environment.

If $B(x)$ does not have full column rank, i.e., $\text{rank}\{B(x)\} = k < m \forall x$, then $B(x)$ can be factorized as:

$$B(x) = B_v(x)B_u(x) \quad (\text{A.2})$$

where $B_v(x) \in \mathbb{R}^{n \times k}$ and $B_u(x) \in \mathbb{R}^{k \times m}$ both have rank k . Now the system can be rewritten as:

$$\begin{aligned} \dot{x} &= A(x) + B_v(x)v + C(x)w \\ v &= B_u(x)u \end{aligned} \quad (\text{A.3})$$

where $v(t) \in \mathbb{R}^k$ can be interpreted as the net control effort produced by a system's m components.

Because $B_v(t)$ has full column rank, a given system state, x , desired system dynamic, \dot{x} , and environmental loading term w will uniquely determine the net control effort output, $v(t)$ required. However, since $B_u(x)$ has a nullspace of dimension $m - k$, there are $m - k$ extra degrees of freedom in assigning component loads, $u(t)$, for a given $v(t)$.

Here we could consider implementation of a nested optimization routine to identify load

allocations in the nullspace $B_u(x)$ that minimizes the expected aggregate damage of functionally redundant components. It is then only the net output control effort vector, v , and the corresponding aggregate component damage that must be considered by supervisory health management control routines.

The optimal mapping between v and u is denoted here as $u(k) = H(v(k))$.

$$H(v(k)) = \min_u E \{f(\gamma, u, \xi)\}, \quad \text{s.t. } v = B_u(x) \cdot u \quad (\text{A.4})$$

where γ is a vector representing the state of health of a system's components, ξ is a random variable representing modeling uncertainty, and $f(\gamma, u, \xi)$ represents the dynamics of component failure modes,

$$\dot{\gamma} = f(\gamma, u, \xi) \quad (\text{A.5})$$

APPENDIX B

THERMAL MODELING FOR DC MOTOR WINDING INSULATION

Winding-to-ambient temperature is defined as:

$$T_{wa}(t) = T_w(t) - T_a \quad (\text{B.1})$$

where T_w is the winding temperature and T_a is the ambient temperature of the surroundings, which are assumed to be stationary.

The first order differential equation for winding-to-ambient temperature given by this thermal model is:

$$\dot{T}_{wa} = -\frac{T_{wa}(t)}{R_{wa}C_{wa}} + \frac{P_{loss}(t)}{C_{wa}} \quad (\text{B.2})$$

The Ohmic power loss in the motor windings is:

$$P_{loss}(t) = i_M^2(t) \cdot R_t \quad (\text{B.3})$$

where i_M is the current flowing through the DC motor, and R_t is the motor winding resistance.

The motor torque roughly proportionally related to the current flowing through a DC motor, where the coefficient of proportionality in this relationship is related to the number of turns in the DC motor,

$$T_M(t) = k_T \cdot i_M(t) \quad (\text{B.4})$$

Here, T_M represents the torque at a motor's output shaft, and k_T is a coefficient of proportionality that relates winding current and motor torque.

Working backwards from Eqns. B.2 - B.4, the motor winding temperature is written in

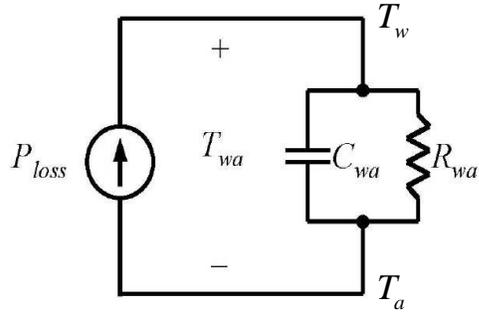


Figure B.1: An equivalent circuit model for motor winding temperature dynamics

terms of motor torque as:

$$\dot{T}_w(t) = -\frac{T_w(t) - T_a}{R_{wa}C_{wa}} + \left(\frac{T_M(t)}{k_T}\right)^2 \frac{R_t}{C_{wa}} \quad (\text{B.5})$$

APPENDIX C

TRIPLEX REDUNDANT ELECTROMECHANICAL ACTUATOR MODELING INFORMATION

The electromechanical actuator example used in Chapter 4 is discussed here.

If low level control is assumed to exist such that motor torque is assignable, then mechanical dynamics of the three motor system are much more simply expressed in terms of torque.

Torque transmission is expressed by the following linear model:

$$\begin{aligned} \dot{x} &= Ax + B_1u + B_2w \\ y &= Cx \end{aligned} \tag{C.1}$$

$$\begin{aligned} A &= \begin{bmatrix} 1 & 0 \\ \frac{k_L}{J_L+3\rho J_M} & \frac{-b_L-3b_M\rho}{J_L+3\rho J_M} \end{bmatrix} \\ B_1 &= \frac{\rho}{J_L+3\rho J_M} \begin{bmatrix} 0 & 0 & 0 \\ 1 & 1 & 1 \end{bmatrix}, \quad B_2 = \frac{1}{J_L+3\rho J_M} \begin{bmatrix} 0 \\ 1 \end{bmatrix} \\ C &= \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \\ x &= \begin{bmatrix} \theta_L & \omega_L \end{bmatrix}^T, \quad u_c = \begin{bmatrix} T_1 & T_2 & T_3 \end{bmatrix}^T, \quad w = Q \end{aligned} \tag{C.2}$$

Where model parameters are defined in Table C.1.

Compared to the consideration of imbedded speed or position controllers as described in [17], a dramatic reduction in model complexity is gained by analyzing a system's mechanical dynamics in terms of assignable component loads. This reduction in model complexity however comes at the price of needing to implement some low level control scheme that enacts component loads within acceptable tolerances. Dynamic inversion and model pre-

Table C.1: Parameter definitions for the electromechanical actuator model

Sym	Description	Units	Value
b_L	Load damping	in·lbf/rad/s	2.5×10^{-1}
b_M	Motor damping	in·lbf/rad/s	1×10^{-4}
k_L	Load stiffness	in·lbf/rad/s	2×10^{-3}
k_t	Motor torque coef.	in·lbf/A	1.01
J_L	Load inertia	in·lbf·s ²	2×10^{-3}
J_M	Motor inertia	in·lbf·s ²	2×10^{-3}
ρ	gearing ratio		8
θ_L	Load position	rad	-
ω_L	Load speed	rad/s	-

dictive control are capable tools for handling this low level control task, even in the case that system dynamics are changing due to a fault mode, as described in [47, 61].

APPENDIX D

SKID-STEERED VEHICLE MODELING INFORMATION

The skid-steered vehicle example used in Chapter 4 is discussed here.

Figure D.1 shows the structure of the four wheeled skid-steered vehicle used. The vehicle kinematics are represented here in a reference frame that is fixed to the center of the skid-steered vehicle, with an axis pointing in the direction of the velocity vector, v . The motion of the rover in a 2D inertial reference frame is given by:

$$\begin{aligned}\dot{x} &= v \cos(\theta) \\ \dot{y} &= v \sin(\theta) \\ \dot{\theta} &= \phi\end{aligned}\tag{D.1}$$

where x , y , and θ represent the 2D position and orientation of the vehicle respectively, v represents the linear velocity of the vehicle, and ϕ represents the angular velocity of the body frame.

Using the simplifying assumption that the wheels on the same side of the vehicle do not slip relative to each other, then v and ϕ are given in terms of wheel speeds as:

$$\begin{aligned}\omega_L &= \omega_1 = \omega_2 \\ \omega_R &= \omega_3 = \omega_4 \\ v &= r \frac{\omega_L + \omega_R}{2} \\ \phi &= r \frac{-\omega_L + \omega_R}{\alpha W}\end{aligned}\tag{D.2}$$

Where ω_i represents the rotational velocity of wheel i , ω_L and ω_R represent the rotational velocities of the wheels on the left-hand and right-hand sides of the vehicle respectively, r represents the wheel radius, α regulates the amount of longitudinal wheel slip that occurs during a turn, and W represents the width of the vehicle.

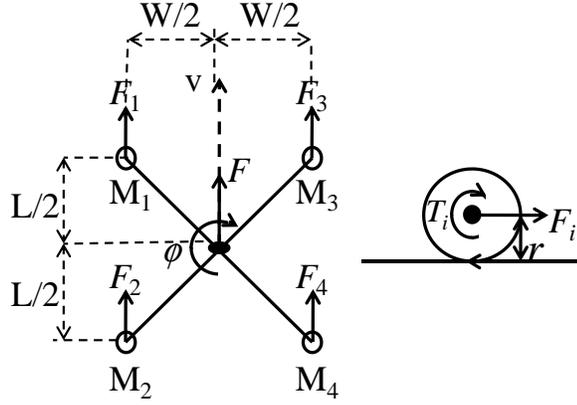


Figure D.1: Visualization of motor torque contributions to system locomotion for a four wheeled skid-steered vehicle

A linear kinetic friction term is used to model the rolling resistance force acting directly opposite to the wheel velocity,

$$f_{r_i} = \frac{m}{4} g k_r \omega_i \quad (\text{D.3})$$

where f_i represents a rolling resistance force acting on each wheel, m represents vehicle mass, g represents a gravitational force, and k_r represents a linear coefficient of kinetic friction.

Using the Euler-Lagrange principle and collection terms, as described in [1], yields:

$$\begin{bmatrix} \dot{\omega}_L \\ \dot{\omega}_R \end{bmatrix} = \begin{bmatrix} \frac{mr^2}{4} + \frac{r^2 I}{\alpha W^2} & \frac{mr^2}{4} - \frac{r^2 I}{\alpha W^2} \\ \frac{mr^2}{4} - \frac{r^2 I}{\alpha W^2} & \frac{mr^2}{4} + \frac{r^2 I}{\alpha W^2} \end{bmatrix}^{-1} \left(-\frac{k_r}{2} \begin{bmatrix} \omega_L \\ \omega_R \end{bmatrix} + \begin{bmatrix} T_1 + T_2 \\ T_3 + T_4 \end{bmatrix} \right) \quad (\text{D.4})$$

$$\begin{bmatrix} v \\ \phi \end{bmatrix} = \begin{bmatrix} \frac{r}{2} & \frac{r}{2} \\ \frac{-r}{\alpha W} & \frac{r}{\alpha W} \end{bmatrix} \begin{bmatrix} \dot{\omega}_L \\ \dot{\omega}_R \end{bmatrix}$$

The coefficients used in this model are defined in Table D.1.

This model may be considered as a placeholder for more accurate representations of kinematics and dynamics of skid-steered vehicles. Many publications have addressed improv-

Table D.1: Parameter definitions for the skid-steered vehicle model

Symbol	Description	Units	Value
r	Wheel radius	m	0.1
W	Vehicle width	m	0.5
L	Vehicle length	m	0.5
I	Wheel rotational inertia	$\text{kg}\cdot\text{m}^2$	0.1
m	Vehicle mass	kg	1
k_r	Rolling resistance term	-	.19
α	Longitudinal slip during turn	-	2
v	Linear vehicle speed	m/s	-
ϕ	Vehicle angular velocity	rad/s	-
ω_L	Left side wheel speeds	rad/s	-
ω_R	Right side wheel speeds	rad/s	-
T_i	Torque allocated to motor i	N/m	-

ing modeling accuracy for skid-steered drivetrains [98, 96].

APPENDIX E

DETAILED DESCRIPTION OF AN EDGE 540T SAMPLE FLIGHT

A sample flight of the Edge 540T electric aircraft is described in this section. The sample flight consists of eight activities:

1. Take-off and climb to 200 ft from 0 to 57 s
2. Circling flight in auto-mode with throttle set to 75% from 57 to 550 s
3. The throttle is increased to 85% from 322 to 550 s.
4. The throttle is decreased to 75% from 550 to 692 s.
5. Aircraft lands from 692 to 722 s
6. Aircraft taxis down the runway from 722 to 885 s
7. The tail of the aircraft is held on the runway and the throttle is set to 75% from 885 to 1232 s
8. The throttle is increased to 80% and held there until battery voltages drop below 17V from 885 to 1232 s

Figure E.1 shows the aircraft throttle, propeller RPM, and battery power profiles recorded over a sample flight. The blue/square and red/diamond traces shown in Figure E.1 give the net battery power input to the ESCs powering motors M1 and M2, respectively. The net battery power input to an ESC is given by the product of the current flowing through the ESC and the sum of the two series-connected battery voltages powering the ESC. The black trace in the battery power plot gives the sum of all battery power input to the two onboard ESCs.

Figure E.2 shows the current and voltage profiles for three of the four batteries used to power the two propeller motors onboard the Edge 540T. The current and voltage data for battery B4 showed anomalous readings and they are omitted from the plots. The ticks on the x-axis in Figures E.1 and E.2 denote the time of notable activities.

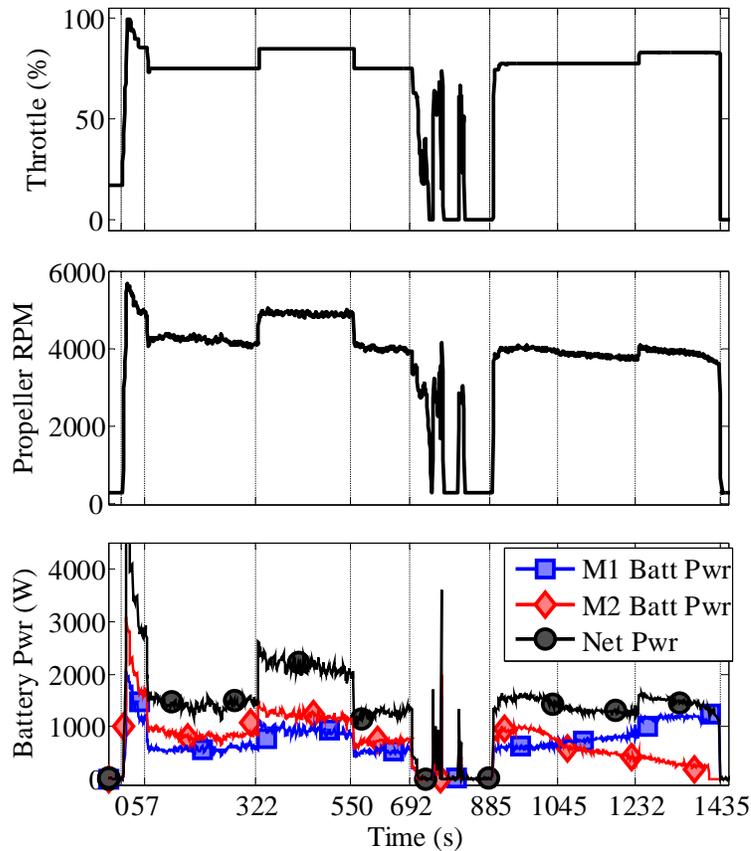


Figure E.1: Plots of throttle command, propeller RPM, and battery power draw measured over a sample flight

An interesting observation to note regarding the battery power draw over the sample flight is that motor M2 draws more power than motor M1 the majority of the time. This occurs until the batteries supplying power to motor M2 are depleted to about 18 volts each. At this point, the batteries powering motor M1 begin to take the majority load. This crossover is seen in both the battery power and battery current profiles at about the midpoint between 885 and 1232 s.

The period from 692 to 885 s in the figures shows the battery load when the pilot resumes control to land the airplane and get the airplane in position to operate the motors on the runway. Operating the motors on the runway to the end of discharge in lieu of flying until end of discharge is more desirable as it abates the risk of a "dead stick" landing or a potential crash.

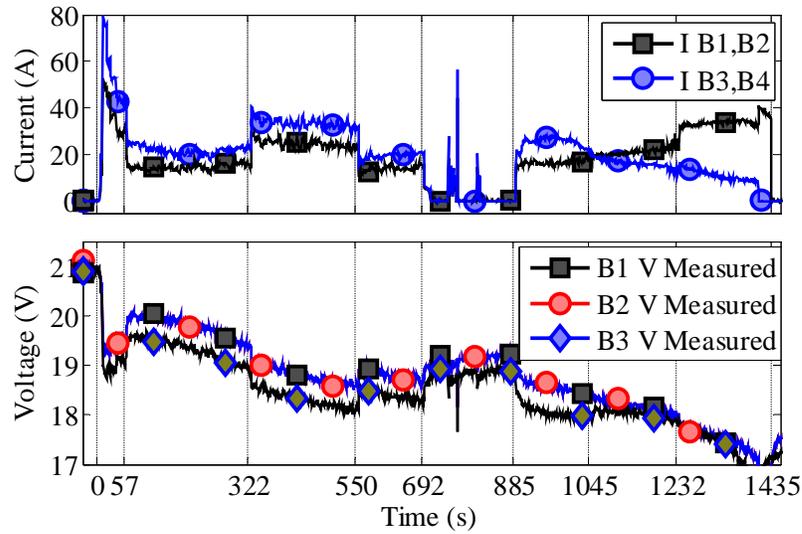


Figure E.2: Battery current and voltage profiles recorded over a sample flight

The period between 885 to 1435 s captures the ground discharge of the aircraft's batteries. This phase of the flight test is used to generate a measurement of the flight time that the batteries would have been able to support had the plane stayed in the air. For this segment, the pilot initially set the throttle to match the RPM range observed during the 75% constant throttle phase of the flight between 550 s and 692 s. At the 75% constant throttle setting in the runway discharge, the propeller RPM, shown in Figure E.1, trends downward and eventually drops below 4000 RPM. This is due to the weakened battery condition in the latter part of the discharge. Because the ground discharge is intended to improvise flight, the throttle was increased at time 1232 sec to raise the RPM above 4000 so as to meet output requirements for cruise flight.

The motors were stopped at 1435 s, after the batteries voltages were seen to fall below the 17 V threshold. Note that the net current draw is seen to increase as the net battery voltage decreases between 885 and 1232 s in order to meet the power demand dictated by the throttle set-point.

REFERENCES

- [1] *Modeling and control of a 4-wheel skid-steering mobile robot: From theory to practice*, 2004.
- [2] ABHKEN, P., “An empirical evaluation of value at risk by scenario simulation,” *Journal of Derivatives*, vol. 7, pp. 1074–1240, 2000.
- [3] ARULAMPALAM, M. S., MASKELL, S., GORDON, N., and CLAPP, T., “A tutorial on particle filters for online nonlinear/non-gaussian bayesian tracking,” *IEEE Transactions on Signal Processing*, vol. 50, no. 2, pp. 174–188, 2002.
- [4] ARULAMPALAM, M. S., MASKELL, S., GORDON, N., and CLAPP, T., “A tutorial on particle filters for online nonlinear/non-gaussian bayesian tracking,” *IEEE Transactions on Signal Processing*, vol. 50, pp. 174–188, Feb. 2002.
- [5] BALABAN, E., SAXENA, A., NARASIMHAN, S., ROYCHOUHDURY, I., GOEBEL, K., and KOOPMANS, M., “Airborne electro-mechanical actuator test stand for development of prognostic health management systems,” in *Annual conference of the prognostics and health management society*, 2010.
- [6] BANJEVIC, D. and JARDINE, A., “Calculation of reliability function and remaining useful life for a Markov failure time process,” *IMA Journal of Management Mathematics*, vol. 17, pp. 115–130, 2006.
- [7] BOLANDER, N., QIU, H., EKLUND, N., HINDLE, E., and ROSENFELD, T., “Physics-based remaining useful life prediction for aircraft engine bearing prognosis,” in *Proceedings of the Annual Conference of the Prognostics and Health Management Society 2010*, Oct. 2010.
- [8] BOLE, B., BROWN, D. W., PEI, H.-L., GOEBEL, K., TANG, L., and VACHTSEVANOS, G., “Fault adaptive control of overactuated systems using prognostic estimation,” in *Annual Conference of the Prognostics and Health Management Society*, Oct. 2010.
- [9] BOLE, B., COOGAN, S., CUBERO-PONCE, C., EDWARDS, D., MELSERT, R., and TAYLOR, D., “Energy management control of a hybrid electric vehicle with two-mode electrically variable transmission,” in *EVS26*, 2012.
- [10] BOLE, B., GOEBEL, K., and VACHTSEVANOS, G., “Markov modeling of component fault growth over a derived domain of feasible output control effort modifications,” in *Annual Conference of the Prognostics and Health Management Society*, 2012.
- [11] BOLE, B., GOEBEL, K., and VACHTSEVANOS, G., “A stochastic optimization paradigm for the development and validation of fault risk assessment and risk management techniques,” in *The American Institute of Aeronautics and Astronautics (AIAA) Infotech*, 2012.

- [12] BOLE, B., TANG, L., GOEBEL, K., and VACHTSEVANOS, G., "Adaptive load-allocation for prognosis-based risk management," in *Annual Conference of the Prognostics and Health Management Society*, 2011.
- [13] BOLE, B., TEUBERT, C., CHI, Q. C., EDWARD, H., VAZQUEZ, S., GOEBEL, K., and VACHTSEVANOS, G., "Sil/hil replication of electric aircraft powertrain dynamics and inner-loop control for v&v of system health management routines," in *Annual Conference of the Prognostics and Health Management Society*, 2013.
- [14] BROWN, A. and GARCIA, R., "Concepts and validation of a small-scale rotorcraft proportional integral derivative (pid) controller in a unique simulation environment," *Unmanned Aircraft Systems*, vol. 1, pp. 511–532, 2009.
- [15] BROWN, D., BOLE, B., and VACHTSEVANOS, G., "A prognostics enhanced reconfigurable control architecture," in *18th Mediterranean Conference on Control & Automation*, 2010.
- [16] BROWN, D., EDWARDS, D., GEORGOULAS, G., ZHANG, B., and VACHTSEVANOS, G., "Real-time fault detection and accommodation for cots resolver position sensors," in *1st International Conference on Prognostics and Health Management (PHM)*, 2008.
- [17] BROWN, D., GEORGOULAS, G., BOLE, B., PEI, H., ORCHARD, M., TANG, L., SAHA, B., SAXENA, A., GOEBEL, K., and VACHTSEVANOS, G., "'prognostics enhanced reconfigurable control of electro-mechanical actuators'," in *2nd International Conference on Prognostics and Health Management (PHM)*, Oct. 2009.
- [18] BUKOV, V., CHERNYSHOV, V., KIRK, B., and SCHAGAIEV, I., "Principle of active system safety for aviation: Challenges, supportive theory, implementation, application and future," in *ASTEC*, 2007.
- [19] CAGLAYAN, A., ALLEN, S., and WEHMULLER, K., "Evaluation of a second generation reconfiguration strategy for aircraft flight control systems subjected to actuator failure surface damage," in *Proceedings of the IEEE National Aerospace and Electronics Conference*, pp. 520–529, 1988.
- [20] CAMERER, C., *Handbook of Experimental Economics*. Princeton University Press, 1995.
- [21] CELAYA, J., KULKARNI, C., BISWAS, G., and GOEBEL, K., "A model-based prognostics methodology for electrolytic capacitors based on electrical overstress accelerated aging," in *Proceedings of the Annual Conference of the PHM Society*, 2011.
- [22] CHI, Q. C., BOLE, B., HOGGE, E., VAZQUEZ, S., DAIGLE, M., CELAYA, J., WEBER, A., and GOEBEL, K., "Battery charge depletion prediction on an electric aircraft," in *Annual Conference of the Prognostics and Health Management Society*, 2013.

- [23] DAIGLE, M., SAHA, B., and GOEBEL, K., “A comparison of filter-based approaches for model-based prognostics,” in *Proceedings of the 2012 IEEE Aerospace Conference*, Mar. 2012.
- [24] DAIGLE, M., SAXENA, A., and GOEBEL, K., “An efficient deterministic approach to model-based prediction uncertainty,” in *Annual Conference of the Prognostics and Health Management Society*, 2012.
- [25] DAVIES, A., ed., *Handbook of condition monitoring*. Chapham and Hall, 1998.
- [26] DEMETRIOU, M. and POLYCARPOU, M. M., “Incipient fault diagnosis of dynamical systems using online approximators,” *IEEE Transactions on Automatic Control*, vol. 43, no. 11, pp. 1612–1617., 1998.
- [27] DONG, M. and HE, D., “Hidden semi-Markov model-based methodology for multi-sensor equipment health diagnosis and prognosis,” *European Journal of Operational Research*, vol. 178, pp. 858–878, 2007.
- [28] DOYLE, J., GLOVER, K., KHARGONEKAR, P., and FRANCIS, B., “State-space solutions to standard h_2 and h_∞ optimal control problems,” *IEEE Transactions on Automatic control*, vol. 33, pp. 831–847, 1987.
- [29] DROZESKI, G. R., *A Fault-Tolerant Control Architecture for Unmanned Aerial Vehicles*. PhD thesis, Georgia Institute of Technology, 2005.
- [30] DUFFIE, D. and PAN, J., “An overview of value at risk,” *The Journal of derivatives*, vol. 4, no. 3, pp. 7–49, 1997.
- [31] DURHAM, W. C., “Constrained control allocation,” *Journal of Guidance, Control, and Dynamics*, vol. 16, pp. 717–725, August 1993.
- [32] FELDMAN, A., KURTOGLU, T., NARASIMHAN, S., POLL, S., GARCIA, D., DE KLEER, J., KUHN, L., and VAN GEMUND., A., “Empirical evaluation of diagnostic algorithm performance using a generic framework,” *International Journal of Prognostics and Health Management*, vol. 1, no. 2, 2010.
- [33] FOLLMER, H. and PENNER, I., “Convex risk measures and the dynamics of their penalty functions,” *Statistics and Decisions*, vol. 24, pp. 61–96, 2006.
- [34] GERGELY, E. I., SPOIALA, D. C., SPOIALA, V., SILAGHI, H. M., and NAGY, Z. T., “Design framework for risk mitigation in industrial plc control,” in *IEEE International Conference on Automation, Quality and Testing, Robotics*, vol. 2, pp. 198–202, 2008.
- [35] GOBBATO, M., CONTE, J., KOSMATKA, J., and FARRAR, C., “A reliability-based framework for fatigue damage prognosis of composite aircraft structures,” *Probabilistic Engineering Mechanics*, vol. 29, pp. 176–188, 2012.
- [36] GOEBEL, K., SAHA, B., and SAXENA, A., “A comparison of three data-driven techniques for prognostics,” in *62th meeting of the MFPT Society*, 2008.

- [37] GOKDERE, L. U., BOGDANO, A., CHIU, S. L., KELLER, K. J., and VIAN, J., “Adaptive control of actuator lifetime,” in *IEEEAC Paper 1337*, no. 1337, 2006.
- [38] GUIDAA, M. and PULCINI, G., “A continuous-state Markov model for age- and state-dependent degradation processes,” *Structural Safety*, vol. 33, no. 6, pp. 354–366, 2011.
- [39] HARKEGARD, O. and GLAD, S. T., “Resolving actuator redundancy - optimal control vs. control allocation,” *Automatica*, vol. 41, pp. 137–144, Jan. 2005.
- [40] HATTORI, Y., KOIBUCHI, K., and YOKOYAMA., T., “Force and moment control with nonlinear optimum distribution for vehicle dynamics,” in *Proc. of the 6th International Symposium on Advanced Vehicle Control*, Sept. 2002.
- [41] HAURIEA, A. and MORESINO, F., “A stochastic control model of economic growth with environmental disaster prevention,” *Automatica*, vol. 42, no. 8, pp. 1417–1428, 2006.
- [42] HERNANDEZ, D. and MARCUS, S., “Risk sensitive control of Markov processes in countable state space,” *Systems & Control Letters*, vol. 29, pp. 147–155, 1996.
- [43] HUBER, R. and MCCULLOCH, B., “Self-repairing flight control system,” Aerospace Congress and Exposition 841552,, Society of Automotive Engineers Technical Paper Series, 1984.
- [44] IBEIRO, L. and OLIVEIRA, N. M., “UAV autopilot controllers test platform using MATLAB/Simulink and X-Plane,” in *Frontiers in Education Conference (FIE)*, 2010.
- [45] INTERNATIONAL, S., “Potential failure mode and effects analysis in design (design fmea) and potential failure mode and effects analysis in manufacturing and assembly processes, reference manual technical report 51739,” tech. rep., 1994.
- [46] JOHANSEN, T., FUGLSETH, T., TØNDEL, P., and FOSSEN, T., “Optimal constrained control allocation in marine surface vessels with rudders,” *Control Engineering Practice*, vol. 16, no. 4, pp. 457–464, 2008.
- [47] JOOSTEN, D., DEN BOOM, V., J., T., and LOMBAERTS, T., “Fault-tolerant control using dynamic inversion and model-predictive control applied to an aerospace benchmark,” in *Proceedings of the 17th IFAC world congress*, pp. 12030–12035, 2008.
- [48] JORDAN, W. E., “Reliability & maintainability symposium,” in *Failure modes effects and criticality analysis*, no. pp. 30-37, 1972.
- [49] JOSSEN, A., “Fundamentals of battery dynamics,” *Journal of Power Sources*, vol. 154, pp. 530–538, 2006.

- [50] KELLY, K., MIHALIC, M., and ZOLOT, M., "Battery usage and thermal performance of the toyota prius and honda insight during chassis dynamometer testing," in *Seventeenth Annual IEEE Battery Conference on Applications and Advances*, 2002.
- [51] KRISHNAKUMAR, K., VIKEN, S., and NGUYEN, N., "Integrated resilient aircraft control - stability, maneuverability, and safe landing in the presence of adverse conditions," tech. rep., National Aeronautics and Space Administration, 2009.
- [52] LOPEZ, I. and SARIGUL-KLIJN, N., "A review of uncertainty in flight vehicle structural damage monitoring, diagnosis and control: Challenges and opportunities," *Progress in Aerospace Sciences*, vol. 46, pp. 247–273, 2010.
- [53] LOVEJOY, W. S., "A survey of algorithmic methods for partially observed markov decision processes," *Annals of Operations Research*, vol. 28, pp. 47–66, 1991.
- [54] MACIEJOWSKI, J. M. and JONES, C. N., "MPC fault tolerant control case study: Flight 1862," in *Proceedings of the 5th IFAC symposium on fault detection, supervision, and safety for technical processes*, pp. 121–126, 2003.
- [55] MAHULKAR, V., ADAMS, D. E., and DERRISO, M., "Minimization of degradation through prognosis based control for a damaged aircraft actuator," *ASME*, 2009.
- [56] MONACO, J. F., WARD, D. G., and BATEMAN, A. J., "A retrofit architecture for model-based adaptive flight control," in *AIAA 1st Intelligent Systems Technical Conference*, Sept. 2004.
- [57] MONTSINGER, V. M., "Loading transformers by temperature," *Transactions of the American Institute of Electrical Engineers*, vol. 32, 1930.
- [58] NAPOLITANO, M., WINDON, D., CASANOVA, J., INNOCENTI, M., and SILVESTRI, G., "Kalman filters and neural-network schemes for sensor validation in flight control systems," *IEEE Transactions on Control Systems Technology*, 1998.
- [59] NESTLER, H. and SATTLER, P. K., "On-line estimation of temperatures in electrical machines by an observer," *Electric Power Components and Systems*, vol. 21, pp. 39–50, Jan. 1993.
- [60] NG, K.-C. and ABRAMSON, B., "Uncertainty management in expert systems," *IEEE Expert*, vol. 5, pp. 29–48, 1990.
- [61] NGO, A. and DOMAN, D., "Dynamic inversion-based adaptive/reconfigurable control of the X-33 on ascent," in *IEEE Aerospace Conference Proceedings*, pp. 2683–2697, 2001.
- [62] OPPENHEIMER, M. W., DOMAN, D. B., and BOLENDER, M. A., "Control allocation for over-actuated systems," in *14th Mediterranean Conference on Control and Automation*, 2006.
- [63] ORCHARD, M., *A particle filtering-based framework for on-line fault diagnosis and failure prognosis*. PhD thesis, Georgia Institute of Technology, Dec. 2007.

- [64] ORCHARD, M., KACPRZYNSKI, G., GOEBEL, K., SAHA, B., and VACHTSEVANOS, G., “Advances in uncertainty representation and management for particle filtering applied to prognostics,” in *Annual Conference of the Prognostics and Health Management Society*, 2008.
- [65] PARLARA, M., WANG, Y., and GERCHAK, Y., “A periodic review inventory model with Markovian supply availability,” *International Journal of Production Economics*, vol. 42, no. 2, pp. 131–136, 1995.
- [66] PATEK, S., “On terminating markov decision processes with a risk-averse objective function,” *Automatica*, vol. 37, no. 9, pp. 1379–1386, 2001.
- [67] RAY, A. and TANGIRALA, S., “Stochastic modeling of fatigue crack dynamics for on-line failure prognostics,” *IEEE Transactions on Control Systems Technology*, vol. 4, pp. 443–451, 1996.
- [68] RUSZCZYRISKI, A., “Risk-averse dynamic programming for markov decision processes,” *Mathematical Programming*, vol. 125, pp. 235–261, 2010.
- [69] SAGLIMBENE, M., “Reliability analysis techniques: How they relate to aircraft certification,” in *Reliability and Maintainability Symposium*, pp. 218–222, 2009.
- [70] SAGOO, G., GURURAJAN, S., SEANOR, B., NAPOLITANO, M., PERHINSCHI, M., GU, Y., and CAMPA, G., “Evaluation of a fault-tolerant scheme in a six-degree-of-freedom motion flight simulator,” *Journal of Aerospace Computing, Information, and Communication*, vol. 7, pp. 47–67, 2010.
- [71] SAHA, B., QUACH, P., and GOEBEL, K., “Exploring the model design space for battery health management,” in *Annual Conference of the Prognostics and Health Management Society*, 2011.
- [72] SAHA, B. and GOEBEL, K., “Uncertainty management for diagnostics and prognostics of batteries using Bayesian techniques,” in *IEEE Aerospace Conference*, 2008.
- [73] SAHA, B., GOEBEL, K., POLL, S., and CHRISTOPHERSEN, J., “Prognostics methods for battery health monitoring using a Bayesian framework,” *IEEE Transactions on Instrumentation and Measurement*, vol. 58, no. 2, pp. 291–296, 2009.
- [74] SAHA, B., KOSHIMOTO, E., QUACH, C., HOGGE, E., STROM, T., HIL, B., VAZQUEZ, S., and GOEBEL, K., “Battery health management system for electric UAVs,” in *IEEE Aerospace Conference*, 2011.
- [75] SAMUELSON, P., “St. petersburg paradoxes: Defanged, dissected, and historically described,” *Journal of Economic Literature*, vol. 15, pp. 24–55, 1977.
- [76] SANKARARAMAN, S., DAIGLE, M., SAXENA, A., and GOEBEL, K., “Analytical algorithms to quantify the uncertainty in remaining useful life prediction,” in *IEEE Aerospace Conference*, 2013.

- [77] SANKARARAMAN, S., LING, Y., SHANTZ, C., and MAHADEVAN, S., “Uncertainty quantification in fatigue damage prognosis,” in *Annual Conference of the Prognostics and Health Management Society*, 2009.
- [78] SAXENA, A., CELAYA, J., SAHA, B., and GOEBEL, K., “Metrics for offline evaluation of prognostics performance,” *International Journal of Prognostics and Health Management (IJPHM)*, vol. 1, no. 1, p. 20, 2010.
- [79] SCHOEMAKER, P., “The expected utility model: Its variants, purposes, evidence, and limitations,” *Journal of Economic Literature*, vol. 20, no. 2, pp. 529–563, 1982.
- [80] SCHOEMAKER, P., “Multiple scenario development: Its conceptual and behavioral foundation,” *Strategic Management Journal*, vol. 14, pp. 193–213, 1993.
- [81] SCHREINER, A., BALZER, G., and PRECHT, A., “Risk analysis of distribution systems using value at risk methodology,” in *Proceedings of the 10th International Conference on Probabilistic Methods Applied to Power Systems*, 2008.
- [82] SCHWABACHER, M., “A survey of data-driven prognostics,” in *Proceedings of the AIAA Infotech@Aerospace Conference*, 2005.
- [83] SERFOZO, R. F., “An equivalence between continuous and discrete time Markov decision processes,” *Operations Research*, vol. 27, pp. 616–620, 1979.
- [84] SHEPPARD, J., BUTCHER, S., KAUFMAN, M., and MACDOUGALL, C., “Not-so-naïve bayesian networks and unique identification in developing advanced diagnostics,” in *IEEE Aerospace Conference*, pp. 1–13, 2006.
- [85] SHERTZER, R. H., ZIMPFER, D. J., and BROWN., P. D., “Control allocation for the next generation of entry vehicles,” in *AIAA Guidance, Navigation, and Control Conference*, Aug. 2002.
- [86] SMILOWITZ, K. and MADANAT, S., “Optimal inspection and repair policies for infrastructure facilities,” *Transportation science*, vol. 28, pp. 55–62, 1994.
- [87] SMITH, M., BYINGTON, C., WATSON, M., BHARADWAJ, S., SWERDON, G., and GOEBEL, K. BALABAN, E., “Experimental and analytical development of health management for electro-mechanical actuators,” in *IEEE Aerospace Conference*, 2009.
- [88] SONNENBERG, F. and BECK, R., “Markov models in medical decision making,” *Medical Decision Making*, vol. 13, no. 4, pp. 322–338, 1993.
- [89] SRIVASTAVA, A. N., M. R. W. and MEYER, C., “Integrated vehicle health management automated detection, diagnosis, prognosis to enable mitigation of adverse events during fight,” tech. rep., Version 2.02, National Aeronautics and Space Administration, Dec. 2008.

- [90] TANG, L., KACPRZYNSKI, G., GOEBEL, K., and VACHTSEVANOS, G., "Methodologies for uncertainty management in prognostics," in *IEEE Aerospace conference*, 2009.
- [91] TANG, L., KACPRZYNSKI, G. J., GOEBEL, K., and VACHTSEVANOS, G., "Methodologies for uncertainty management in prognostics," in *IEEE Aerospace Conference*, 2009.
- [92] TSANG, P., JACKO, M., and RHEE, S., "Comparison of chase and inertial brake dynamometer testing of automotive friction materials," *Wear*, vol. 103, pp. 217–232, 1985.
- [93] VERMA, V., G. G. S. R. and THRUN, S., "Particle filters for rover fault diagnosis," *IEEE Robotics & Automation Magazine*, pp. 56–64, June 2004.
- [94] WANG, H. S. and CHANG, P.-C., "On verifying the first-order Markovian assumption for a Rayleigh fading channel model," *IEEE Transactions on Vehicular Technology*, vol. 45, no. 2, pp. 353–357, 1996.
- [95] WEBBY, R., ADAMSON, P., BOLAND, J., HOWLETT, P., METCALFE, A., and PANTADOSI, J., "The mekong applications of value at risk (VaR) and conditional value at risk (CVaR) simulation to the benefits, costs and consequences of water resources development in a large river basin," *Ecological modelling*, vol. 201, pp. 89–96, 2007 201.
- [96] YI, J., SONG, D., ZHANG, J., and GOODWIN, Z., "Adaptive trajectory tracking control of skid-steered mobile robots," in *IEEE International Conference on Robotics and Automation*, pp. 2605–2610, 2007.
- [97] YOUNG, R., "Aircraft aging and durability project - technical plan summary," 2008.
- [98] YU, W., CHUY, O. Y., COLLINS, E. G., and HOLLIS, P., "Analysis and experimental verification for dynamic modeling of a skid-steered wheeled vehicle," *IEEE TRANSACTIONS ON ROBOTICS*, vol. 26, no. 2, pp. 340–353, 2010.
- [99] ZHANG, B., GEORGOULAS, G., ORCHARD, M., SAXENA, A., BROWN, D., VACHTSEVANOS, G., and LIANG, S., "Rolling element bearing feature extraction and anomaly detection based on vibration monitoring," in *16th Mediterranean Conference on Control and Automation*, 2008.
- [100] ZHANG, H. and CHOW, M.-Y., "Comprehensive dynamic battery modeling for PHEV applications," in *IEEE Power and Energy Society General Meeting*, 2010.
- [101] ZHANG, Y. M. and JIANG, J., "Active fault-tolerant control system against partial actuator failures," *IEEE proceedings on control theory and applications*, vol. 149, no. 149, p. 95, 2002.
- [102] ZHANG, Y. and JIANG, J., "Bibliographical review on reconfigurable fault-tolerant control systems," *Annual Reviews in Control*, vol. 32, pp. 229–252, 2008.

- [103] ZHANG, Y. and JIANG, J., “Bibliographical review on reconfigurable fault-tolerant control systems,” *Annual Reviews in Control*, vol. 32, no. 2, pp. 229–252, 2008.
- [104] ZHOU, K., “A new controller architecture for high performance, robust, and fault-tolerant control,” *IEEE Transactions on Automatic Control*, vol. 46, pp. 1613–1619, 2001.

VITA

Brian M. Bole graduated from the FSU-FAMU School of Engineering in 2008 with a B.S. in Electrical and Computer Engineering and a B.S. in Applied Math. Brian received a M.S. degree in Electrical Engineering from the Georgia Institute of Technology in 2011, and a Ph.D. in Electrical Engineering in 2013. Brian's research interests include: analysis of stochastic processes, risk analysis, and optimization of stochastic systems. Brian is currently investigating the use of risk management and stochastic optimization techniques for optimal adaptation of active component load allocations in robotic and aviation applications.