Building Condition Monitoring

By

Stephen Samouhos

B.S. Mechanical Engineering
Massachusetts Institute of Technology, 2004

M.S. Mechanical Engineering
Massachusetts Institute of Technology, 2006

SUBMITTED TO THE DEPARTMENT OF MECHANICAL ENGINEERING IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY IN MECHANICAL ENGINEERING
AT THE
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

JUNE 2010

Signature of Author:_____
Department of Mechanical Engineering
June 2010

Certified by:_____
Leon Glicksman
Professor of Mechanical Engineering
Thesis Supervisor

Accepted by:_____
David E. Hardt
Professor of Mechanical Engineering
Chairman, Committee for Graduate Students

Building Condition Monitoring

By

Stephen Samouhos

Submitted to the Department of Mechanical Engineering on June 1st, 2010, in Partial Fulfillment of the Requirements for the Degree of Doctor of Philosophy in Mechanical Engineering

Abstract:

The building sector of the United States currently consumes over 40% of the United States primary energy supply. Estimates suggest that between 5 and 30% of any building's annual energy consumption is unknowingly wasted due to pathologically malfunctioning lighting and comfort conditioning systems. This thesis is focused on developing analytical methods embodied within useful software tools to quickly identify and evaluate those building system faults that cause large building energy inefficiencies.

The technical contributions of this work include expert rules that adapt to HVAC equipment scale and operation, a general framework for applying probabilistic inference to HVAC fault detection and evaluation, and methods for sorting fault signals according to user-defined interests such as annual cost of energy inefficiencies. These contributions are particularly unique in their treatment of model and measurement uncertainty within the fault inference, and the careful consideration of user interests in fault evaluation.

As a first step to developing this general framework for fault detection, I targeted first order faults such as simultaneous heating and cooling and imbalanced air flows within several large air-handling units in three buildings on the MIT campus. Experiments included the purposeful implementation of mechanical and software control programming faults on otherwise fault-free equipment. Between the five pieces of equipment, the software system successfully identified all previously known and experimentally implemented faults, as well as additional faults that had not been previously identified or imposed during the experiment. User testing and experiments show that embracing uncertainty within HVAC fault detection and evaluation is not only paramount to judicious fault inference but it is also central to gaining the trust and buy-in of system users who ultimately can apply fault detection information to actually fix and improve building operations.

Thesis Supervisor: Leon Glicksman
Title: Professor of Mechanical Engineering

Foreword and Acknowledgements

I have lived the dream that is MIT for almost 10 full years. I came here in September of 2000, not knowing where I would go, what could happen, or much of anything at all. Despite the unknown, and my initial fear of that unknown, I knew that it was the start of the most amazing adventure. Indeed that adventure turned out to be much longer than I could have ever imagined at the time; as a 17 year old, getting into MIT was certainly a dream come true, but thinking that I would stay here for ten years and earn a doctorate was well beyond a dream; it was reaching for the stars themselves.

And now, just days before leaving MIT for good, and hours away from completing a document that has taken years to forge, I feel that the real voyage is just about to begin. Like the calm before the storm, today the air is still, and everything moves just a bit slower, allowing me to soak in these final sweet moments of my home for the past decade. I don't know what life will be like in the fall, or next year, or even several years from now, but I stand ready for whatever the seas may bring. I have learned to not fear what I don't know, but instead to seek it out, to understand, to learn, to grow, and to conquer that which I do not yet know.

The experiences that I had here at MIT, especially in the period of history when I was here, are overwhelming and life altering. I am so very proud to say that I have absolutely no regrets; I really did try to do everything, and I couldn't have asked for more. From being a fraternity brother at Delta Tau Delta, to a resident at East Campus on Pi-West, to working with a team to win the IDEAS contest, and the ISN Soldier Design Contest, all the way to "borrowing" MIT's golf carts as a freshman, serenading sororities only to steal their silverware, living in the heart of Kenmore Square when the Red Sox won their first world series in several decades, travelling to the Middle East with my father and missing a month of classes, taking and passing my qualifiers after a weekend ski trip party, earning a Hertz fellowship that supported me for 5 wonderful years, and oh so many more things; the list just keeps going. I cannot be more thankful for the fullest cup of life that I have had at MIT.

I have, over the past 10 years, also had the pleasure and honor of meeting some of the most amazing friends that a man could ever have. Corey Fucetola was the very first person that I met at MIT, and to this day he has been my best friend, confidant, brother, and favorite critic; bubbles and S'mooch, washing dishes was never more fun. Deanelle Symonds was the big sister that I never had, wished that I did have, but will be lucky enough to always keep in my life; batman and robin will set flight again sometime. Susan Brown, who I've known since year two, was the first person to help me see the big picture at MIT and beyond, and she'll always help me keep that in mind. Bryan Schmid and Greg Mark, both my big brothers from Delts, taught me to look outside MIT for answers; when there is doubt, there is no doubt. Andrew Mcinnes, my friend, teacher, mentor and colleague, without whom I wouldn't have reached this point, and who always told me how it was; go big or go home. John Anastasio, my partner in crime from Northeastern, and colleague in various "adventures". Kaitlin Thaney, who's seen all of this unfold for 8 years from the sidelines and still has the patience to occasionally knock some sense into me. Tim Ferris, who has been my little brother and one of my most trusted friends since Pi-West. Uncle Michael Fishelson and Grant Duers, who were the first to teach me about "resources" and built my confidence that we could make something great happen. Nicholas Gayeski and Sian Kleindienst, the real dojo-masters of the BT ninjas academy, my best friends,

my inspiration and anchor, my partners in the grand experiment of KGS Buildings, and my "balance point" for ideas on how to fix buildings everywhere. And Danielle Brown, who's love and support got me through in the end; "persistence" never meant so much.

I had always wanted to find my teacher at MIT; someone who could help me see further into the science and calculus than I could ever on my own; I am very blessed to have had several teachers while at MIT. Ioannis Yannas, who was my first teacher, gave me the initial confidence to believe in myself, my ideas, and the power of my heritage. John Brisson, from whom every MechE should be lucky enough to learn thermal-fluids engineering, drove me to think deeper about thermodynamics and physics than I had ever before. Markus Zahn taught me the magic of continuum electro-mechanics, and the art of electro-magnetism. Gareth McKinley, my Master's thesis advisor, and the first professor that I ever called by a first name, taught me the rules and art of academic fencing and one of the most important lessons of my education; MIT gives you the tools to solve any problem, but not a detailed list of the problems to solve.

My committee was composed of my last four teachers at MIT. Sanjay Sarma gave me the confidence to be what I am today, and helped me understand that true knowledge is far more respected and useful than a big equation. Les Norford taught me many things, especially the virtue of patience and humility; we all want to improve the world around us, but change requires patience, and the humility to listen to those who have come before us. Neil Gershenfeld helped me stretch my thinking farther than I ever would have on my own; I'm incredibly grateful to have had someone in my life that challenged me so much to think so differently and in ways that I would never have considered. Finally, Leon Glicksman, my chair advisor, and the man who four years ago willingly took a chance to advise an opinionated, difficult, stubborn, mouthy (yet very convincing), self-funded (dangerous), and extremely ambitious graduate student. I'm not sure if an advisor knows how things will turn out when they first accept a new student, but I can't imagine how Leon felt after our first year or so together! I was lucky though; Leon is patient, challenging, honest and devoted to making a difference in people's lives and the world around him. I learned from Leon that simplicity usually leads to success; big problems are only tough while you keep them as big problems. Energy in buildings is a tremendous problem, perhaps one of the biggest possible problems, but that doesn't mean it can't be broken down, made manageable and thereby resolved.

Finally, we reach the end, at 2:41 a.m. on June 1st, the day that I hand in the final piece of my education at the Massachusetts Institute of Technology. I would not be who I am today, or where I am today without the unceasing love and support of my family. The seeds that eventually brought forth the ideas in this thesis came from a man who more than 30 years ago arrived in this country from Greece with nothing more than his ambition, strength of heart, and dream of reaching the stars. My father is the greatest man that I will ever know, and I am blessed beyond any measure to have such a man as a teacher, father, and friend; he has reached the stars, and I hope that I may be sufficiently strong to one day follow him there. Neither of us, however, would be what we are without my mother; she is our rock, our foundation, our un-ending support. I would never have even applied to MIT 10 years ago without her staunch encouragement and I wouldn't have made it this far without her strength; like any fantastic mother, she has the magic ability to just always "know" everything. To the Samouhos family, Vinnie, Leah, Christina, Mom, Baba, all I can say is: quantum leap accomplished... found to be improbable, but not impossible.

I know that Για–Για and Παποο Σαμουχο can see us from where they are; I dedicate this thesis to their memory, and to that of the entire Samouhos family.

Table of Contents

1 Introduction 14

1.1 Thesis Motivations and Direction 14

1.2 Building Faults and Energy Efficiency 14

1.3 Fault Definition and Identification 16

1.4 Contemporary "No-tech" Fault Detection in Buildings 16

1.5 Thesis goals 18

1.6 Thesis Chapter Outline 19

2 Literature Review 20

2.1 Introduction 20

2.2 Predicting Building Energy Consumption 21

2.2.1 Introduction 22

2.2.2 Physical Modeling Techniques for Building Energy Models 22

2.2.3 Black and Grey Box Modeling Techniques for Building Energy Models 23

2.2.4 Summary for Predicting Building Energy Consumption 29

2.3 Fault Detection and Diagnosis of Building Systems 30

2.3.1 Introduction 30

2.3.2 Overview of Fault Detection and Diagnosis Methods 30

2.3.3 Quantitative Methods of FDD 32

2.3.4 Qualitative Methods of FDD 40

2.3.5 FDD Based on Process History 50

2.3.6 Summary of FDD 54

3 Modeling of HVAC Equipment and Faults 55

3.1 Introduction 55

3.2 Air handler Description 56

3.2.1 Mixing Box Air Flow Modeling 62
3.2.2 Mixing Box Recirculation Damper Fault Modeling 65
3.3 Mixing Box Simulations 67
3.3.1 Assumptions 67
3.3.2 Results for Boston, Massachusetts 69
3.4 Measured performance of a mixing box 74
3.4.1 Fault-free operation 74
3.4.2 Fault-laden operation 82
3.4.3 Real energy impact of faults 86
3.5 Conclusions 95
4 Building Data and Data Uncertainty 96
4.1 Introduction 96
4.1.1 The Role of Data in Buildings Today 97
4.1.2 Leveraging the Experience of Building Management 97
4.1.3 Making decisions with incomplete knowledge 98
4.2 Collecting data from buildings 98
4.2.1 Building control systems for data collection 99
4.2.2 Cost of Data Collection 100
4.2.3 What to Measure 101
4.2.4 Instrumentation Options 102
4.2.5 Sensor Measurement Uncertainty 103
4.3 Propagation of Uncertainty in HVAC analysis 105
4.4 Modes of uncertainty 105
4.4.1 Random Measurement Error 105
4.4.2 Systematic Measurement Error 105

4.4.3 Methods for accommodating uncertainty in analysis 108
4.5 Probabilistic Inference 117
4.5.1 Example of cost prioritization 118
4.6 Taking action under uncertainty 118
5 Experimental work on MIT equipment 119
5.1 Introduction 119
5.2 AHUs 9 and 10 in MIT Building 46 120
5.2.1 Physical Description of AHUs 9 and 10 121
5.2.2 Relevant Controls Programming 124
5.2.3 Instrumentation and Data Acquisition 126
5.2.4 Commissioning AHUs 9 and 10 130
5.2.5 AHU Models 133
5.3 FDD System 149
5.3.1 FDD Expert Rules 149
5.3.2 FDD Financial Estimates 151
5.3.3 Fault Experiment Program 152
5.4 AHUs 2 and 3 in MIT Building 56, and AHU 2 in MIT Building 16 156
5.4.1 Physical Characteristics 156
5.4.2 Relevant Controls Programming 159
5.4.3 Instrumentation and Data Acquisition 160
5.4.4 Commissioning 160
5.4.5 AHU Models 161
5.4.6 FDD System 161
5.5 Summary of Experimental Setup 162
6 Results 162

6.1 Introduction: 163
6.2 FD&E Software 163
6.2.1 Software architecture and setup at MIT 163
6.2.2 Software interface 165
6.3 Experiments in Building 46: AHUs 9 and 10 168
6.3.1 Results from MIT Building 46 168
6.3.2 Summary of Results for MIT Building 46 179
6.4 Analysis of Buildings 16 and 56 180
6.4.1 Findings on 2009 and 2010 performance data 180
6.4.2 Summary of results for MIT buildings 16 and 56 186
6.5 Results of Monte Carlo Simulations 186
6.5.1 Introduction 186
6.5.2 MC Simulations 186
6.5.3 Summary of results from MC simulations 189
6.6 Feedback from users on system design features 190
7 Conclusions 191
7.1 Limitations of current approach 192
7.2 Further Development 193
7.3 Final Remarks 199
8 Bibliography 201

Table of Figures

Figure 1 Katipamula and Brambley's approach to structuring the field of research in building fault detection, diagnostics and prognostics 31
Figure 2 A plan diagram of the test facility at IEC 34
Figure 3 A schematic diagram of the AHUs at IEC 35
Figure 4 Table of faults tested in ASHRAE 1020, and their method of implementation 36
Figure 5 Diagram of the "innovations" based fault detection and diagnosis algorithm used in ASHRAE 1020 38
Figure 6 The decision tree used by the passive OAE fault detection module in Katipamula's WBD 43
Figure 7 From Katipamula et al, 2003, the diagram shows a test sequence for proactive FDD of erroneous temperature measurement 44
Figure 8 From House et al (51), list of expert rules per category of AHU operation. Note that House only relied on temperature and humidity measurements about the system, as well as occupancy schedule and valve control signals 46
Figure 9 From House et al (50), schematic of a typical economizing AHU for which House et al developed APAR. 47
Figure 10 From House et al (51), diagram of transitions between classifications of AHU operational modes. 49
Figure 11 Schematic system model of the variable-air-volume air-handler studied in this project 57
Figure 12 Families of curves with varying damper authority for percent of maximum recirculated flow through opposed (left image) and parallel (right image) dampers, as a function of percent opening of the damper. ASHRAE Guideline 16-2003 63
Figure 13 Baseline operation of Air-Handler under fault-free operating conditions for Boston, Massachusetts 69
Figure 14 Energy loss index simulated for a recirculation damper stuck at various positions 70
Figure 15 Energy loss index simulated over three damper leakage levels 71
Figure 16 Energy impact of various mixing box faults 72
Figure 17 Outdoor air fraction versus recirculation damper position for several fault cases, simulated over a year's worth of Boston weather data 73

Figure 18 Measured fault-free operation of Air-Handler B at the IEC 75
Figure 19 Measured fault-free operation of Air-Handler A at the IEC 76
Figure 20 Time-series of rad and g measurements on a spring day within ASHRAE 1020-RP for air-handler B; the time series begins at 9 am 77
Figure 21 Time series measurements for rad and g, collected from air-handler A over 6 different days spanning summer, winter, and spring seasons 79
Figure 22 Time-series for rad, g, and supply air flow rate scaled by its maximum possible value (denoted by ssaf) 80
Figure 23 Measurements of faults in a real mixing box: stuck and leaking dampers 83
Figure 24 Data from three different leak levels imposed on air-handler A at IEC during the winter testing season of ASHRAE 1020 85
Figure 25 Distribution of supply air flow rates for AHSRAE 1020 air-handlers; the data is drawn from over 30 days of operating history, measured at 1 minute intervals during the regular working hours of 9 am to 5 pm. 88
Figure 26 Distribution of humidity for return and supply air, across 30 days from the ASHAE 1020 data set 89
Figure 27 Simulation results for Des Moines, Iowa: annual outdoor air temperature and damper positions per month 92
Figure 28 Simulation results for Des Moines, Iowa: scatter of outside air fraction versus return air damper position 93
Figure 29 Simulation results for Des Moines, Iowa: monthly energy consumption per fault status 94
Figure 30 Modes of uncertainty for measuring a variable, X: a) two discrete readings b) infinite number of readings (adapted from Coleman and Steele, 1995). In both cases, the bias error, $\boldsymbol{\beta k}$ represents a systematic difference between the averaged measured value of X and its true value. The precisions error, $\boldsymbol{\epsilon k}$ is a measure of the data scatter between the mean measurement of X, and its k^{th} sample 110
Figure 31 Propagation of measurement bias and precision errors into experimental results 111
Figure 32 Screenshot from the building control system for MIT building 46, detailing a schematic of AHU 09; the schematic for AHU 10 is identical. The design of these AHUs is

reminiscent of the schematic VAV AHU design that was included as figure 1 in chapter 3 of this thesis 122
Figure 33 Picture of the mixing box interior showing the use of turbulators to improve mixing of return and outdoor air; the averaging mixed air temperature string is also shown traversing the turbulator surface. As a length scale, the length of a horizontal turbulator fin is 4 feet 124
Figure 34 Picture of the i2-920 system controller for AHU 9 in MIT building 46 128
Figure 35 Screenshot of the setup screen for implementing extended logging on an Infinity point; the data logging feature is supported on MIT's building control system for all types of control points 129
Figure 36 Control volume analysis of CO2 species balance around air-handler 141
Figure 37 control volume analysis of CO2 species balance around set of conference rooms 141
Figure 38 control volume analysis of CO2 species balance around atrium 142
Figure 39 control volume around mixing box plenum for atrium and conference room returns 143
Figure 40 (upper graph) CO_2 concentration contours per supply air flow and outdoor air fraction, and (lower graph) critical curve between supply and outdoor air flow rate for a CO_2 concentration of 800 PPM. In both graphs a red-line has been included to show the minimum outdoor air flow that is possible for the AHU 145
Figure 41 Picture of an AHU mixing box that is suffering from broken damper linkages, rusted open dampers, and disconnected pneumatic actuators. The picture was taken during a tour of 14 AHUs across the MIT campus; we observed similar pathologies in a number of AHUs across the campus. 152
Figure 42 Interior picture of the mixing box for AHU 9, showing the outdoor and recirculation air dampers, and their corresponding actuators, and actuator linkages. 153
Figure 43 Cascade of mixing box images and close-up pictures of fault implementations; the lower left picture shows the disconection and capping of the pneumatic supply to the outdoor and recirculation dampers at the pneumatic supply manifold for the AHU. The lower right picture shows the removal of a linkage between the outdoor air damper and its actuator. 154
Figure 44 Schematic diagram of AHU 2 in MIT building 16; the image is a screenshot from the building control system interface to the AHU 157
Figure 45 Schematic diagram of AHU 2 in MIT building 56; the image is a screenshot from the building control system interface to the AHU 159

Figure 46 Schematic of the software architecture that underpins the FD&E system developed in this thesis 164
Figure 47 Screenshot of the software interface, showing features that link the underlying analysis to user defined settings 166
Figure 48 Close up view of the software pop-up screen that facilitates closer inspection of analyzed equipment performance; this particular view is for AHU 9 in MIT building 46, for a period when the damper pneumatic pressure supply was removed and the damper control programming was altered 169
Figure 49 Results screenshot for analysis on air-flows across AHU 9 in MIT Building 46 172
Figure 50 Pictures of faults diagnosed due to inspection of results from software analysis; the right hand image verifies suspicions of flow reversal in the mixing box, while the left hand image shows evidence of a broken pressure transducer for the outdoor air dampers 174
Figure 51 Results of experimentation on AHU 9 for January 29th, 2010 through February 5th, 2010 177
Figure 52 Condition monitoring on AHU 10 in building 46, from January 22nd, 2010 through January 29th, 2010 179
Figure 53 Analysis results on AHU 2 in MIT building 16 181
Figure 54 Results of FD&E analysis on AHU 2 in Building 56 184
Figure 55 Results of FD&E analysis for AHU 2 in building 16, with current 2010 data 185
Figure 56 Q-Q plot of empirical CDF on MC simulation results; a Lilliefors test for normality was applied to the distribution of energy consumption and found the null hypothesis valid at the 5% level 188
Figure 57 Two-sample Kolmogorov-Smirnov test of homogeneity between MC and single sample distributions; the KS test rejeceted the null hypothesis at the 5% level for almost all test samples 189

1 Introduction

1.1 Thesis Motivations and Direction

Buildings use 40% of the primary energy supplied in the United States (US), and over 70% of all generated electricity (1), primarily for heating, cooling and lighting. In commercial buildings, about 20% of their energy consumption can potentially be saved by correcting faults, including malfunctions and unnecessary operation of equipment (2). And research and initial deployments of advanced control systems suggest that they can save an additional 10-20% (3). The energy efficiency resource recoverable through such improved building controls and fault detection correspond to the output from hundreds of power plants, equivalent to more than a third of the coal-fired power production in the United States (4). Realizing these substantial savings will require introducing intelligence into the infrastructure of buildings (4).

This thesis is intended to contribute to the foundation of intelligent building infrastructure by providing a framework and methods for the automated detection of building energy in-efficiency faults. The anticipated extensions of this research include application of the fault detection framework towards the construction of next generation phenomenological building control systems, and enhanced intelligent infrastructure for energy efficiency.

1.2 Building Faults and Energy Efficiency

According to a recent report from TIAX (2), commercial buildings can typically reduce their energy consumption by 20-30% through continuous commissioning practices[1] and the implementation of a handful of energy efficiency strategies. These statistics are especially pertinent to large commercial buildings whose conditioned surface area exceeds 50,000 square

[1] Commissioning is the practice of reviewing and testing the operation of installed equipment in buildings to ensure their proper installation, controls programming and maintenance; continuous commissioning typically uses real time data feeds to automate and continuously carry out this process.

feet. Of the 5 million commercial buildings in the United States, less than 20% of them can be considered as large, however that sub-set of the stock represents over 80% of the energy used in the commercial building sector, or equivalently 25% of the energy consumed by all buildings (1), commercial and residential. Furthermore, over 80% of the energy inefficiency, or energy wasted by those large commercial buildings can be attributed to five common pathologies: simultaneous heating and cooling, extraneous lighting, extraneous heating, cooling or ventilation, imbalanced ventilation systems, and leaky ventilation ductwork (2).

A small sampling of energy audits and inspections for a variety of buildings on the eastern seaboard of the US, between New Jersey and Massachusetts, further supports the hypothesis that there is no shortage of efficiency opportunities within the US building stock. Amongst a dozen residential and commercial buildings that were inspected as part of this thesis, all of them were found to have efficiency opportunities whose value matched or exceeded 15% of their annual utility bills. Even with a small sample of buildings, I discovered systemic in-efficiencies ranging from thermostat programming errors that caused equipment short cycling, to hot and cold water loops whose connections had been reversed and incurred year-round simultaneous heating and cooling. Less conspicuous efficiency opportunities were revealed through sub-metering programs in large, mixed-use buildings; for example a 60 story mixed-use tower in New York City, NY, experienced a decrease in common-area-maintenance (CAM) charges by 8% following a re-allocation of utility bills according to tenant actual consumption rather than the more commonly used allocation metric, tenant occupancy square footage.

Still, other buildings were found with oversized heating, ventilation and air-conditioning (HVAC) equipment that would short cycle and waste energy simply because their control programming was intended for equipment of a smaller size. While there is no publication or report that has explicitly evaluated or even estimated the total energy efficiency opportunity within the entire US building stock, historical precedence and practical experience suggests that most buildings possess an efficiency opportunity equivalent to or exceeding 15% of their annual utility costs.

1.3 Fault Definition and Identification

Interviews with a variety of building stakeholders in conjunction with results found in the literature and through practical experience suggest that a building pathology, or fault, that wastes money or causes discomfort can be due to a very wide range of issues: erroneous controls programming, inherent design flaws, mechanical degradation, occupant disposition and many others. Likewise, for each fault that exists, different stakeholders may interpret those faults to have different values or priorities for remediation; financial officers may prioritize energy waste whereas facility managers may prioritize occupant discomfort. As a consequence of this rich fault environment, the definition and implications of a building fault can vary widely depending on the perspective of a stakeholder. The design of a useful fault detection system must therefore consider the technical aspects of faults that the system will identify, as well as the information content that is desired by the intended audience of the fault detection system. The focus of my research is to identify HVAC equipment that consumes more energy than expected, and to communicate those results to an audience comprised of building management, operators, and engineers who control the funds needed to resolve building energy in-efficiencies.

1.4 Contemporary "No-tech" Fault Detection in Buildings

Despite the apparent opportunity for fixing building faults and recovering energy efficiency resources, very few of today's buildings contain a dedicated system for the detection of energy in-efficiency faults. The past twenty years of research and development into sophisticated and practical building fault detection has yielded numerous potential solutions, as will be discussed later in this thesis, however market proliferation of fault detection products remains very small (2).

One large hurdle for 3rd party or dedicated building fault detection systems is the belief amongst building managers that their traditional means of fault detection are already sufficient; the traditional approach to building fault detection by building managers is through alarms on the building's HVAC control system, occupant feedback and complaints, and routine or un-

scheduled maintenance. Most of the facility and technical stakeholders interviewed in this research stated that they often find equipment inefficiencies while tracking down the causes for occupant discomfort; inefficiencies are often "discovered" by accident during these un-scheduled equipment inspections. While occupant complaints are certain to raise the fastest response amongst maintenance personnel, it is not a guaranteed solution to finding the most wasteful of in-efficiencies. For example, simultaneous heating and cooling within an HVAC system is an extremely wasteful and prevalent fault, however it typically does not yield uncomfortable space conditions. Consequently, simultaneous heating and cooling often persists for long periods of time without any conspicuous indications of the fault's existence.

Routine preventative maintenance of equipment is another form of building equipment fault detection and remediation. Because the method involves a periodic physical inspection of equipment by an experienced technician, it may appear to be the most reliable amongst fault detection schemes. The technical skill of preventative maintenance personnel, however, is not guaranteed to be sufficient for identifying or fixing all buildings faults. At the same time, preventative maintenance itself is a process that focuses on specific equipment components; if the preventative maintenance schedule does not actually cover the section of the equipment that is defective, then the preventative maintenance technician will not observe the defect. Moreover, many equipment in-efficiencies are artifacts of control system programming errors or operator overrides; these types of software-derived faults are difficult to identify during a routine or un-scheduled equipment inspection.

The contemporary standard for software-based building fault detection is single-variable alarms from the building's HVAC control system. All building control systems support alarm functionality whereby control system installers and designers can implement conventional threshold and duration alarms on measured or computed variables. Such alarms are standard features for making building operators aware of equipment operation that may be un-safe or uncomfortable for building occupants, or deleterious to the useful life of the equipment. While control system software alarms seem like the most obvious solution for continuous detection of HVAC energy in-efficiencies, they have not yet found extensive use in that capacity. Interviews with building stakeholders suggest that the lack of exploitation of control alarms in detecting

energy in-efficiencies is partially due to the miss-match between the intended audience of control system alarms and the audience that wants to identify wasteful equipment operation. Likewise, the technical personnel who are responsible for programming building control systems typically do not have specific training on how to write building control alarms that explicitly identify energy inefficiency opportunities. While building control systems seem to be the most obvious platform for supporting software-based HVAC fault detection, control system companies have not yet delivered explicit software solutions to do so.

As buildings have become more complex and managers find the classical fault detection paradigm to be insufficient, other 3rd party solutions for identifying energy in-efficiency faults have begun to emerge. The most noticeable amongst new solutions in the market are explicit monitoring services where control system alarms are used to filter and condition building operation data for expert fault-detection analysts. Those expert analysts consume the data that is filtered by the building control alarms to create monthly reports that inform building management stakeholders about efficiency opportunities that exist in their buildings. While successful at identifying and communicating efficiency opportunities, interviews with building stakeholders suggest that such analyst services are still far too expensive for wide-scale or even continuous use over a building's entire lifetime. The high cost of analyst services may be an artifact of the immature market for 3rd party building fault detection; as time goes on, better technology and extension of the market beyond early-adopters will ultimately drive down the cost for 3rd party fault detection.

1.5 Thesis goals

This thesis is motivated by the goal of providing a 3rd party software system that can use building control data to automatically identify, communicate and track building energy efficiency opportunities for building management. The software automation of what is currently an expert-analyst service could potentially yield a continuous commissioning program that is sufficiently low-cost for building stakeholders to use it across their entire building stock and for the full lifetime of their buildings. Especially within the context of a growing building stock,

rising energy costs, and deeper political and environmental implications of energy waste, low-cost systems for identifying building energy in-efficiencies may be critical to our energy future.

This thesis has three primary technical goals that should culminate in a prototype automated HVAC fault detection system

i. Create a fault detection system architecture that could potentially scale to include any type of equipment, in any building, for any fault analysis

ii. Create a fault detection framework that supports probabilistic inference, and expresses results where possible in financial terms that characterize some of the risk as well as the reward of diagnosing or resolving a potential fault

iii. Create fault classifiers that adapt to the equipment or systems that they analyze

The ensuing chapters of this thesis will explore the origins of these goals, their development, and their demonstration within a proto-type automated fault detection system. Additional thesis chapters will discuss the initial performance of that prototype fault detection system and detail the future research needed to further cultivate the underlying ideas.

1.6 Thesis Chapter Outline

The following chapter, chapter 2, will enumerate some of the standard prior art for HVAC fault detection systems, as well as the apparent research gaps in that literature and how this thesis intends to fill some of those gaps. Chapter 3 will focus on energy-modeling of Air-Handling Units (AHUs), and the simulation of several equipment and control faults that are typically found in that type of equipment. I will continue to use AHUs throughout the thesis as a test subject for exploring novel fault detection ideas, and chapter 3 provides the physical basis for measuring and predicting AHU energy consumption under fault-free and fault-laden

operating conditions. In addition to defining equipment energy and fault models, chapter 3 also explores the magnitude of impact that various faults can have on AHU energy consumption, as well as the effects of seasonal weather on discriminating between fault-free and fault-laden equipment operation. Chapter 4 of the thesis concerns the more practical aspects of data and uncertainty within the built environment, and how data and uncertainty could be used to detect faults in HVAC equipment. Chapter 4 will explore how model and measurement uncertainty can propagate through the fault detection analysis in order to yield probabilistic inference. In that same chapter I will discuss the methods by which data is typically extracted from a building, and how that process might be improved to further reduce the cost of implementing automated building fault detection systems. Chapter 5 details the application of the core thesis ideas into a tangible prototype system for detecting HVAC faults on the campus of the Massachusetts Institute of Technology (MIT). Chapter 6 includes and discusses the results of applying that prototype HVAC fault detection system to several AHUs across MIT's campus. Finally chapter 7 includes concluding remarks about the ideas and experiments that were developed and explored in this thesis, as well as a roadmap for continued cultivation of these ideas towards the next generation of building control and fault detection systems.

2 Literature Review

2.1 Introduction

Building modeling, simulation and fault detection research has a rich, interwoven history that spans over five decades. Research into software-based methods for detecting faulty building operation extends back to the early 1980's when direct digital control equipment was first deployed in buildings. The government sponsorship for developing building energy modeling tools began in the 1960's with programs for the DOE-2 and BLAST software packages. Today, building modeling for energy prediction or fault detection has evolved into numerous software tools and its own distinct profession; numerous engineering firms exist whose sole service is building energy modeling. Furthermore, the earlier distinctions between fault detection and

energy modeling have begun to fade; building modeling today is used to identify when buildings are not working as we might want them too as well as to quantify the value of correcting that divergent behavior.

There exist three primary approaches to the modeling of buildings:

- **Physical Modeling:** thermodynamics, mass transport and heat transfer are used to derive equations of state that describe the performance of the building

- **Black Box Modeling:** purely empirical state machines, trained by historical building data, that generate an output for a given set of data inputs

- **Grey Box Modeling**: empirical models that are enhanced by physical models, or vice-versa. This approach typically combines physics with building data in order to provide a more accurate, semi-empirical perspective on performance.

Most contemporary design-build construction projects now require a physical energy model of the building in order to achieve certain construction certifications and help the design team make educated choices between various building design options (5). Black box models have been extensively developed for applications in both the detection and diagnosis of building system faults (6) (7) and the prediction of building energy consumption (8). Likewise, physical and grey box models have also been developed to diagnose buildings faults (9) and predict building energy consumption (10). While this thesis is focused on fault detection and diagnostics of building systems, the interaction between fault detection and energy modeling warrants the inclusion of a literature review on both subjects. The following sections will briefly review some of the more recent developments in building energy modeling, and then focus more deeply on the prior art of fault detection.

2.2 Predicting Building Energy Consumption

2.2.1 Introduction

Holistic building energy modeling techniques were originally motivated by the commercial need for tools that could predict and verify the energy-cost implications building design and renovation options. To that end, modeling techniques can now utilize historical building utility records, building design documents or both to create a baseline and forecast for the energy consumption of a building. Several energy models can be generated and compared for a building, each one representing a different design or renovation case, and thereby provide the basis for judicious selection of design and renovation options. On the other hand, baseline building energy models can also be applied retro-actively after a renovation or design change in order to verify the impact of pursuing those options; this is a technique often used in the measurement and verification of energy savings after a renovation.

2.2.2 Physical Modeling Techniques for Building Energy Models

Numerous software tools have been commercialized and developed for the physical modeling of building energy consumption (11). Consequently, physical modeling techniques for buildings are very mature and the majority of contemporary work on such techniques has focused on streamlining user interfaces, improving model accuracy, and expanding model libraries (12) (13). Using today's tools, however, the cost and labor associated with forming an accurate physics-based building energy model is still prohibitively high for all but large and luxury type buildings to have an energy model. The primary cause of that high hurdle is the effort associated with collecting and entering all of the relevant building information into the modeling software. In some cases the high level of effort is unavoidable simply because many buildings do not have as-built drawings or equipment schedules to instruct the modeling agent on how the building is put together or what is in it. In all cases, however, the modeling agent must requisition building information from the building management or designer, and take some action to verify that their building information is accurate.

In order to get around the *physical* limitations of physical models, industry and academia have focused on the development of empirical modeling techniques that can *learn* from utility records and other data how a building consumes energy and reacts to its environment. By eliminating the need for physical information about the building, a black or grey box model may be easier, faster, and less expensive to develop for a building than a physical model.

2.2.3 Black and Grey Box Modeling Techniques for Building Energy Models

ASHRAE's "Great Building Energy Shoot-Out" I and II in the mid-1990's were high-profile events in the history of the major international efforts devoted to developing black box techniques for predicting building energy consumption. The contests were intended to encourage the development of models that could be used to estimate building energy savings due to HVAC and lighting equipment retrofits; this goal was motivated by the commercial need for tools that could help companies sell and execute building retrofits on the basis of improved energy efficiency. The second contest was particularly successful and focused on predicting the hourly building energy data for the fourth week of every month within a year's worth of building energy data from two buildings at Texas A&M university (10).

Twelve whole months of building energy data were used in the second shoot out contest. The first three weeks of each month were provided as model training data in the contest. The 44 international contestants were judged based on the coefficient of variation (CV) and the mean bias error (MBE) between their predictions and the actual values for the energy data of the missing fourth week of each month. These performance metrics were chosen by the contest organizers and had been used in the first shoot-out as well (14).

$$CV = 100 \times \frac{\sqrt{\frac{\sum_{i=1}^{N}\left(y_{pred,i} - y_{data,i}\right)^2}{N}}}{\bar{y}_{data}}$$

Equation 2-1

$$MBE = 100 \times \frac{\frac{\sum_{i=1}^{N}(y_{pred,i} - y_{data,i})}{N}}{\bar{y}_{data}}$$

Equation 2-2

The training data set was composed of local weather data and five thermal and electrical energy measurements for each building. In both shoot-outs, the best results were produced by an artificial neural network (ANN) model that also included statistical tests for relevance of input data and refinement of the ANN structure. The winners of the second contest, Dodier and Henze from the University of Colorado and Nebraska, respectively, created an ANN that was perfected through the Wald test, and resulted in an average root mean square error (RMSE) of 17% (15).

Shootout contestants utilized a wide variety of input relevance tests, such as Bayesian relevance estimation (16), statistical tests for nonlinear correlation (17), autocovariance and the Wald test (15), and principal component analysis (18). Contestants consistently identified input selection and model structural design as the principal challenges in model synthesis. Initial input variable selection was made by the organizers based on their broad knowledge of building parameters and weather characteristics that are known to influence building energy usage. These initial variables included outdoor air temperature and humidity, wind speed, insolation, and the energy variables themselves. Contestants independently created additional time variables to capture the influence of annual period, day-type, and other temporal features that are known in the field to impact energy consumption.

In Dodier and Henze's winning entry, a separate ANN was developed for each of the five energy variables in each test building. The model inputs were chosen from a consistent set of inputs, but their network weights were pruned according to the Wald test, which was applied to each separate model. The architecture of the ANN (feed-forward, single hidden layer with 25 nodes) was chosen for ease of implementation and balance between modeling accuracy and ability to generalize. The Wald statistic was computed for each input in each model by training the model with and without the input variable; if the Wald statistic was found to be larger than a given threshold, then the input was deemed relevant.

Dodier and Henze report that the selection of this threshold is another critical parameter that must be empirically determined during model synthesis. A full discussion of the Wald statistic can be found in Dodier and Henze (15), but in general they report that basic engineering knowledge of input relevance was reflected in the results of the Wald statistical analysis. Moreover, Dodier and Henze found that their temporal input variables generated greater values of the Wald statistic than did their environmental input variables; they interpreted this result as a reflection of the overwhelming influence of human occupancy on building energy consumption. On the other hand, outdoor air temperature and insolation also follow strict time variations thus further strengthen the role of time in predicting building energy consumption. Finally, Dodier and Henze also suggested that model performance may be improved by taking a Bayesian estimation approach, or constructing multiple models for each prediction, and taking an average output from all plausible models as the final prediction.

In addition to reporting the statistical results of their models, the energy shoot-out contestants were also asked to predict the dollar savings from retrofits made at each test building. Contestants applied their pre-retrofit models to post-retrofit weather data in order to simulate the building energy consumption that might have occurred in the absence of any retrofits. The simulated energy data was compared to actual post-retrofit energy costs in order to create a margin of savings incurred by the retrofits. Energy savings predictions made by the top five contestants for the larger test building spanned from $83,399 to $189,655, with an average savings of $163,058.

On the other hand, predictions made by those contestants for the smaller test building spanned from -$23,333 to $22,822 with an average savings of $4,761. The reduced accuracy for predictions made on the smaller building were suspected to be due to un-accountable occupancy schedules of the building (10). Overall, however, this wide range of predictions on the economic performance of retrofits is partially responsible for the limited proliferation of forecasting techniques into the mainstream business of selling energy efficiency retrofits.

Continuing from the ASHRAE energy shoot-out, Karatsou et al (19) reduced the problem of ANN driven building energy prediction to three main challenges:

- Input selection
- Gross model structure selection
- Model refinement

Various statistical tests have also been developed to help design the architecture of ANN models. Anders and Korn (20) highlight tools such as hypothesis testing, information criteria and cross-validation. Rivals and Personnaz (21) examine methods based on least squares estimation and statistical tests applied to a two step process of building up a preliminary model followed by refinement through subtracting away irrelevant input and middle layer nodes. Karatsou applied the principles found in Rivals and Personnaz to build ANNs that predicted annual building energy data for buildings in Athens, Greece, and Austin, Texas. The latter building data set was the same data set used in ASHRAE's great building energy shoot-out I. In their experiments as well those performed by most other researchers in the field, Karatsou et al used a conventional single hidden-layer, feed-forward ANN model, with hyperbolic tangent hidden layer nodes.

With a constant set of inputs, Karatsou et al incrementally increased the number of hidden layer nodes according to the growth conditions set in Rival and Personnaz (21). Once a satisfactory ANN was built, the network was subjected to a pruning process that removed irrelevant inputs and nodes. Once again, time of day and an occupancy indicator appeared as the most influential input variables, followed closely by environmental variables such as temperature and insolation, and additional time-lag variables for temperature and insolation. Synthetic, time-shifted environmental input variables have been used across the literature to account for the thermal inertia of buildings. In general, the two-step statistical approach used by Karatsou et al resulted in ANN models that outperformed the contestants of ASHRAE's building energy shoot-out II, with CVs ranging between 8 and 13%.

Seem (22) in 2006 presented an alternative approach to identifying anomalous energy consumption in commercial buildings via outlier detection algorithms. The novel method employed outlier detection to identify occasions when building electricity consumption was significantly different from it's historical performance. Seem's method evolved from the conventional human-operator approach to anomaly detection, where building operators had to manually pick out anomalous consumption from graphical presentations of utility data. The

intelligent data analysis approach presented by Seem is intended to eliminate the tedious practice of *hunting* for anomalies, and provide the building operator with data sets populated only by likely anomalous performance data.

Seem's method begins by grouping utility data into day-type categories, with historically similar consumption profiles, and then organizes statistically significant outliers in each category according to their modified z-scores. An outlier in this case is judged to be statistically significant if the associated studentized deviation of the outlier exceeds a critical value. The severity of the outlier is then ranked according to its modified z-score, which is equal to the number of robust standard deviations between the outlier value and the robust mean (robust statistics are computed from the data set that excludes all outliers). The method was applied to 97 buildings, and over the course of several months the method successfully identified three buildings with anomalous periods of energy consumption. The causes of those anomalies were investigated and attributed to chiller failures, poor ventilation design and poor controls operation. Seem reports that those errors were corrected soon after their identification, yielding significant energy and operational cost savings for the buildings. A year earlier, Seem had also published a similar pattern recognition algorithm for determining days of the week with similar daily energy consumption profiles (23).

Another more recent addition to the spectrum of building energy intelligence is the application of support vector machines (SVMs) to energy forecasting and fault detection (24), (25), (26). A popular rival to ANNs, SVMs are typically regarded as having a more rigorous mathematical foundation than ANNs, and over the past decade SVMs have in fact displaced ANNs in most machine-learning applications (27), (28), (29). SVMs have been particularly successful in learning problems composed of sparse, yet high-dimensionality training sets, such as those data sets found in DNA micro-array analysis (30). Success in that application suggests that SVMs might also be well suited to similar energy learning problems composed of monthly or daily energy and environmental data. While the building energy literature is well stocked with publications concerning ANNs and other popular intelligence technologies, there is a noticeable dearth of attention paid to SVMs applied to building energy problems. It appears that SVMs are

not even mentioned in the building energy literature until 2005, despite their origins three decades prior.

Recently in 2008, Lai et al reported their results of training SVMs to forecast the building energy consumption of a 2,000 sq.ft. residential building in Tohoku, Japan. Using 15 months of hourly electrical energy consumption and indoor environmental conditions, and daily climate data recordings, Lai et al explored daily and monthly building energy forecasting. In addition, they also tested the ability of their model to identify artificially induced anomalous building energy consumption. The performance of their models was measured by the Pearson and robustness coefficients

$$P = \frac{(x_i - \bar{x})(y_i - \bar{y})}{(n-1)S_x S_y}$$

Equation 2-3

$$R = \frac{(x_i - \bar{x})^2}{(y_i - \bar{y})^2}$$

Equation 2-4

where x and y denote the measured and predicted data sets, respectively, S is the standard deviation with subscripts denoting the data sets described by the parameter, and n is the size of the measured data set. The Pearson coefficient, P given by Equation 2-3, is used to measure the correlation between the measured and predicted data, and the robustness coefficient, R defined by Equation 2-4, is the ratio of the variances for the measured and predicted data sets. Lai's month-ahead and day-ahead forecasting was found to produce Pearson and robustness coefficients roughly equal to 0.9. Although it is not clear how the conclusion was reached, Lai et al report that outdoor air temperature was the most significant contributing factor to their forecasts. By identifying significant differences between predicted and actual energy consumption, their SVM model was also capable of identifying an artificially induced period of anomalous energy consumption. In their conclusion, Lai et al suggest that since SVMs are sufficiently fast and easy to train on building data, they should be applied as *model factories* to

continuously learn from daily building operations and produce accurate daily or monthly energy forecasts.

In an earlier paper, Dong et al (25) explore SVMs to create forecasting models for the purpose of whole building energy baselining. Their model was applied to four commercial office buildings in Singapore, to predict whole building energy consumption based on monthly averaged outdoor air temperature, humidity, and solar insolation. Three years of monthly electric utility bills and weather data were used for the training and testing of their SVM models. Dong et al used a Gaussian radial-basis function kernel for their SVM, as well as a step-wise search algorithm to estimate the best model parameters for the learning problem. Consequently, Dong's models for each building produced CV values, defined by Equation 2-1, under 3% using a single year of whole building energy (WBE) training data. The performance of Dong's model exceeded the performance of all other WBE ANN models produced through ASHRAE's great energy shoot-out, which achieved a minimum WBE modeling CV of 10.36%. In their conclusions, Dong et al also claim that their publication was the first to ever explore the application of SVMs to the prediction of whole building energy consumption.

2.2.4 Summary for Predicting Building Energy Consumption

While physical methods for building energy modeling are mature and manifest in commercial products, black and grey box approaches to building energy modeling represent an active area of academic and industrial research. The literature suggests that the algorithms used for black and grey model input selection and model structure definition in building energy forecasting are of greater importance than the algorithms that are used to train the models themselves. Cross-over between energy modeling and fault detection is apparent in the literature, and indeed other reviews of the literature typically treat both in the same context of modeling techniques (8).

2.3 Fault Detection and Diagnosis of Building Systems

2.3.1 Introduction

The following examination of prior art for building fault detection and diagnostics explores those applications at all levels of the building (31), ranging from a high-level whole-building perspective (32), (33), (34) down through specific components and systems within the building (35), (36), (37), (38), (39), (40). Like building energy modeling, the research in fault detection and diagnostics can be segmented according to physical, black, and grey box modeling techniques.

2.3.2 Overview of Fault Detection and Diagnosis Methods

Fault detection and diagnostics (FDD) is implicitly based on forming and evaluating comparisons between measured and predicted building performance. Consequently, FDD is a multi-part system comprised of data collection, analysis and inference engines. Because of that multi-part nature, FDD systems can have multiple points of failure; data collection can suffer from broken sensors, analysis can suffer from inaccurate models, and inference can suffer from poor tuning. The complexity in FDD systems has given rise to a rich field of academic and industrial research and development, which will be partially examined in the following sections. Despite the volume of research, however, there does not exist today any set of FDD products that have found as wide application or use in practice as have building energy modeling tools.

The most recent authoritative review of the state of the art of building FDD was prepared by Katipamula and Brambley in 2005 (6), (7) and we will adopt and expand their organization of the literature in this review.

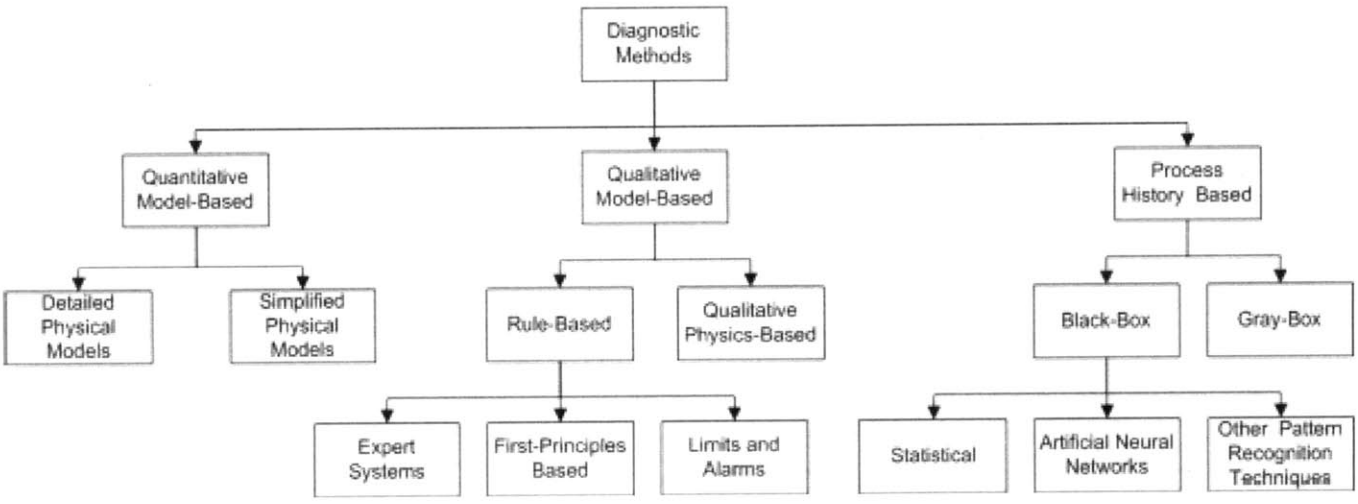

Figure 1 Katipamula and Brambley's approach to structuring the field of research in building fault detection, diagnostics and prognostics

Katipamula and Brambley organized the field of FDD research according to their interpretation of the modeling and inference techniques that are being used by FDD researchers. In what they termed quantitative models, Katipamula and Brambley included research where system models were formulated from physical or *first principle* engineering models, and faults were detected by analysis of residuals between model predictions and measurements. On the opposite end of their research classification scheme, Katipamula and Brambley included FDD methods that were primarily based on *learning* faults from historical process data. The approach in that category of research could otherwise be categorized as purely black box, or even semi-empirical, where the fundamental process of FDD is based on pattern recognition analysis of current data against a historical database. In between these extremes are what Katipamula and Brambley have termed qualitative methods; research in that category includes FDD systems comprised of expert inference rules that are derived from engineering first principles or practical experience.

Overall, the categories introduced by Katipamula and Brambley reflect the manner in which system knowledge is stored in a computer, and how that system knowledge is used to infer the existence of a fault. The three primary method of storing system knowledge include first-principle or physical model derivation, practical engineering models, or historical records of

performance. Likewise, the three corresponding methods of fault inference include residual comparison between models and measurements, violation of expert rules, and detection of statistical anomalies. FDD researchers are not restricted to operate in any one of these research categories, and in fact much of the recent research activity has focused on blending the best attributes of each category.

2.3.3 Quantitative Methods of FDD

2.3.3.1 Overview of Quantitative FDD Methods

Quantitative model-based approaches to FDD rely on detailed physical modeling of system thermodynamics, heat transfer and mass transfer. While the fundamental equations and concepts of physics are universal, the algorithms for using such models to discern and diagnose a specific fault have remained a very active area of research. The quantitative modeling approach to fault detection and diagnosis typically employs the continuous computation of residuals between predicted and measured system performance. Typically, thresholds and simple discrimination functions are used to convert those residuals into fault signals of varying intensity. A statistical test may also be applied to these residuals in order to compute the likelihood of a particular fault (39), (41). In general, quantitative, first-principle models are prone to difficulties incurred by un-modeled behaviors, complex models of transient interactions, loss of generality by over-development towards a specific system, and requirements for additional system sensors. Although many physical HVAC&R system models have been developed (42), these limitations have resulted in very few quantitative models that have actually been reduced to practice and tested in realistic situations (9), (43), (44), (45).

2.3.3.2 Detailed Physical Models for FDD

The ASHRAE 1020 research project (9) was a significant demonstration of empirical and physical model based FDD methods. The research was performed at the Iowa Energy Center's Energy Resource Station (IEC), in their building test facility. The IEC research space is particularly unique in that it is a fully operational, yet unoccupied office-style building equipped with three AHUs that serve multiple test rooms. The complete variable-air-volume (VAV) HVAC system is controlled by a commercial BCS, which was also serves as the data collection tool for HVAC research data.

Over the course of several weeks in multiple seasons, two different FDD methods were tested at the IEC for their ability to identify faults about the three AHUs installed at the building. One FDD method was based on detailed physical models that identified faults by comparing real time operational data against engineering-based performance models of the equipment. The second FDD method was based on anomaly detection within semi-empirical correlations between sub-metered electrical power measurements about the mechanical equipment and nominal HVAC data from the BCS. The details of the semi-empirical FDD approach will be discussed in a later section of this review that deals explicitly with grey and black box FDD techniques.

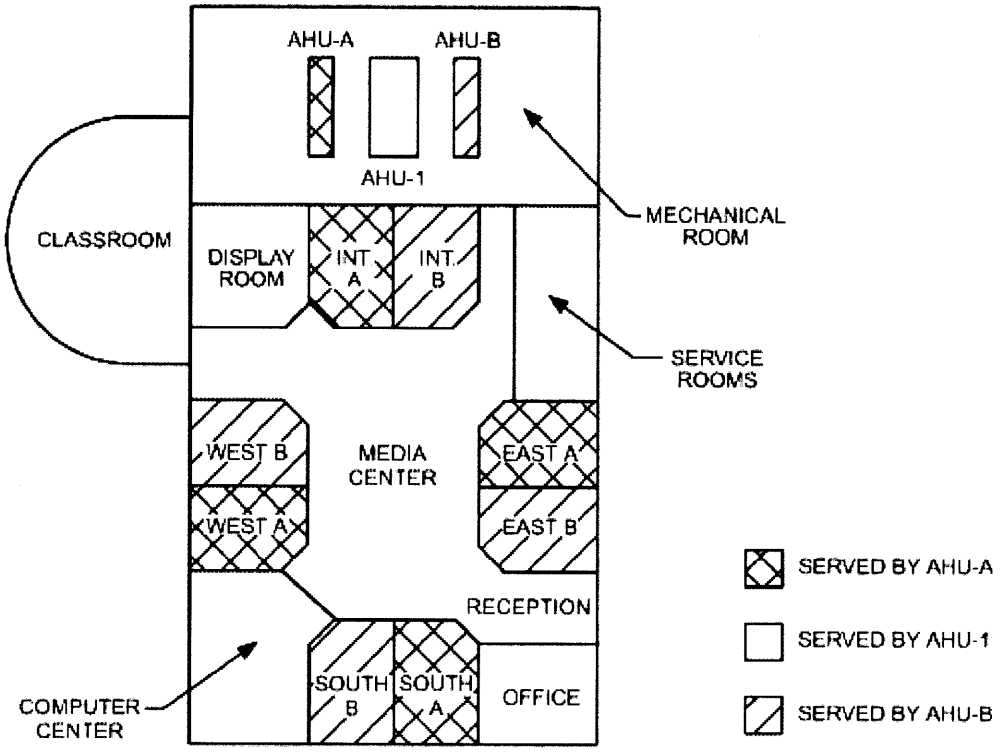

Figure 2 A plan diagram of the test facility at IEC

The test facility at IEC has a total floor space of 9,272 square feet that includes several teaching classrooms, office space, reception, and commons areas. The matched pair of AHUs (denoted AHU-A and AHU-B in Figure 2) serves four classrooms each, where the eight classrooms were paired on each face of the building. The building itself has a true north-south alignment which provides each test room pair with nearly identical exposure to external loads. The test rooms were unoccupied but equipped with electric heaters to simulate normal thermal loads as well as electric lighting. The third AHU (denoted as AHU-1 in Figure 2) serves the remaining areas of the facility including offices, commons space and a reception area.

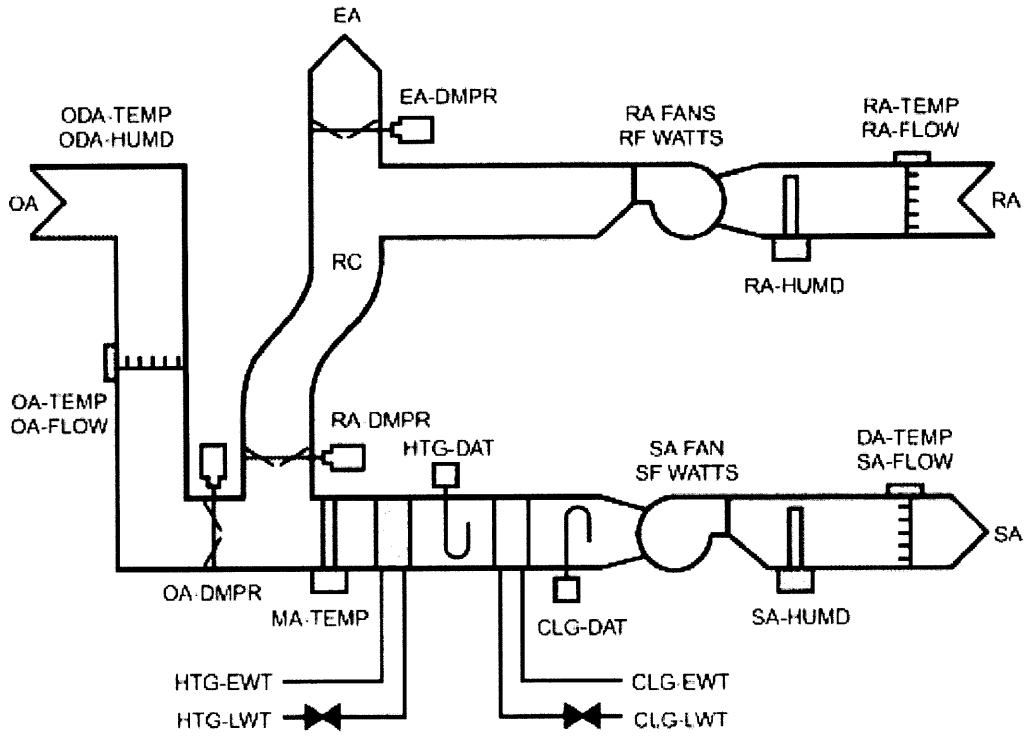

Figure 3 A schematic diagram of the AHUs at IEC

All three AHUs had a similar variable-air-volume (VAV) configuration shown in Figure 3. The cooling plant for the system included a 10 ton, two-stage, reciprocating air-cooled chiller. The major components of the AHUs include recirculated air, exhaust air, and outdoor air dampers; cooling and heating coils with control valves; and supply and return fans with variable frequency belt drives (VFDs). The heating coils denoted in Figure 3 were not used during any of the test periods. Local space heating was performed at the level of terminal VAV units that were equipped with electric or hydronic re-heat coils. The supply fan speed in all AHUs was regulated to maintain a constant supply duct static pressure. The return fans on AHU-A and AHU-B were controlled to maintain a constant percentage of the supply airflow; in AHU-1, the return fan control signal was a constant percentage of the supply fan control signal. The chilled water flow rate through the cooling coils in AHU-A and AHU-B was controlled by a three-port mixing valve in a diverting application. A two-port valve was used to control chilled water flow through the cooling coil in AHU-1.

Several common AHU faults were investigated in this experiment, two of which also represented the impact of long-term component degradation. Faults were implemented via software and mechanical component alterations, including system modification with special bypass piping and valves. Faults included stuck and leaking recirculation dampers, leaking coil valves, internally fouled cooling coil, drifting pressure sensor, unstable fan control, and slipping fan belt. Each fault was implemented in at least two of the three week test periods held during the summer, winter and spring seasons. Each test included a two week measurement period, one for calibrating a performance baseline for the system, and the second for a blind test on fault detection and diagnosis. For tests with AHUs A and B, the list of possible faults for a particular test period was known ahead of time to the investigators; they did not, however, know when the faults would be implemented during the test week. On the other hand, tests with AHU-1 did not permit the investigators to know the nature or timing of faults. Consequently, tests with AHU-1 were more representative of situations that exist in real buildings.

Fault	Type	Implementation
Air-Mixing Section		
Stuck-closed recirculation damper	Abrupt	Application of a control voltage from an independent source to maintain the damper in the closed position.
Leaking recirculation damper	Degradation	Removal of the recirculation damper seals, with one seal removed for the first fault stage, two for the second, and all seals for the third stage.
Filter-Coil Section		
Leaking cooling coil valve	Degradation	Manual opening of a coil bypass valve.
Reduced coil capacity (water-side)	Degradation	Manual throttling of the cooling coil balancing valve, to 70%, 42%, and 27% of the maximum coil flow of 1.7 L/s (27.5 gpm) for the three fault stages.
Fan		
Drifting pressure sensor	Degradation	Introduction of a controlled leak in the pneumatic signal tube from the supply duct static pressure sensor to the transducer, to a maximum reduction of 225 Pa (0.9 in. of water).
Unstable supply fan controller	Abrupt	Introduction of alternative gains for the PID controller that adjusts fan speed to regulate static pressure.
Slipping supply fan belt	Degradation	Adjustment of fan belt tension to reduce maximum fan speed by 15% at 100% control signal for the first stage and 20% for the second stage. The third stage had an extremely loose belt with variable fan speed.

Figure 4 Table of faults tested in ASHRAE 1020, and their method of implementation

The first principles approach to fault detection relied on thermo-mechanical sub-system models for the entire AHU. In particular, sub-system models for the fan, economizer and cooling coil were used to predict performance from a given set of input variables and compare those predictions against measured system output values. The fan sub-system model was based on the quadratic relationship between air flow rate, system resistance and static pressure change. The economizer model was based on a thermodynamic representation of the mixed air conditions as a function of outdoor air conditions, mixing damper position and return air conditions. The cooling coil model was based on the NTU method for modeling heat exchangers. In all three cases, the thermal models are complimented with actuator and motor models that account for latency and hysteresis of response, as well as thermal gains due to various modes of electro-mechanical energy dissipation. All data intended for the first principles model was first treated with a low pass filter in order to extract the steady state data needed by the model. In fact, a high rejection rate at the filter was used as a fault detection method for identifying inappropriate equipment performance oscillations. To account for measurement errors and inaccuracy of modeling, discrepancies between output predictions and measurements were compared against a statistically significant threshold to reduce the rate of false fault alarms.

The magnitude by which the predictions exceeded the thresholds were termed ``innovations", and were split into three bins based on a three way organization of the AHU operating space. The bins stored average innovation magnitudes, weighted exponentially by their age. Expert rules were then applied to the average values of the bins in order to diagnose a fault. As an alternative to this method (which is illustrated schematically in Figure 4), fault detection and diagnosis was also explored in ASHRAE 1020 through recursive parameter estimation within first-principle system models. In this latter FDD approach, the inputs and outputs of the system are used to iteratively update the values of certain model parameters. The FDD algorithm in this case tracks the change in those specific parameters until they surpass a certain threshold or exhibit a time series pattern indicative of a certain fault.

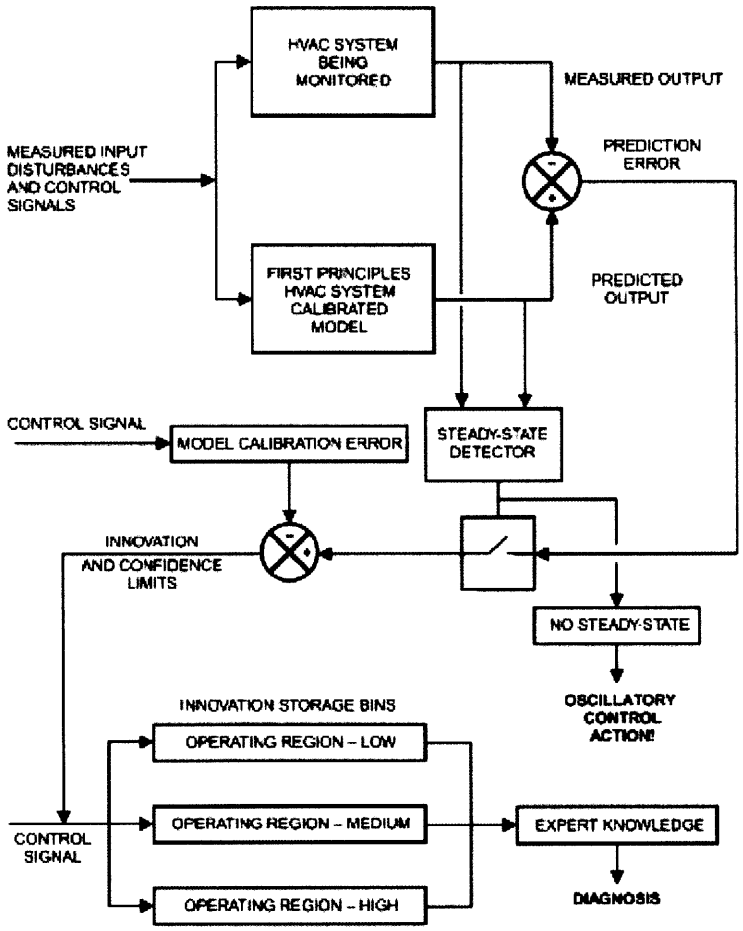

Figure 5 Diagram of the "innovations" based fault detection and diagnosis algorithm used in ASHRAE 1020

ASHRAE 1020 exposed a series of practical difficulties in using physical models to detect and diagnose AHU faults. Various components of the FDD method were affected by

- Measurements from sensors placed in positions not conducive for collecting observations needed by the model
- Un-modeled effects in the AHU stemming from un-balanced air flows and flow reversals
- Model training data that did not encompass the full range of possible observations
- Insufficient modeling at the device level to accommodate for small-signal fault data or unknown disturbances
- Preponderance of transient operation, precluding the capabilities of the steady state thermo-mechanical model

- Dependent model parameter estimation

Overall, these deficiencies in the approach were mitigated throughout the test periods by adjusting parameters that governed the sensitivity and efficacy of the FDD algorithm. The results of the study showed that fault diagnosis was typically harder to do successfully than fault detection. This deficiency is partially to blame on the practical need for building an FDD system from sensor data that is traditionally collected on commercial VAV AHUs. Additional AHU sensors would have enabled more detailed physical models or expert classifier rules, either of which may have supported more robust and extensive capture and classification of AHU faults. On the other hand, this study also relied on instrumentation grade HVAC sensors that were calibrated and well tended as part of the research center; it remains unclear what additional technical difficulties would arise from commercial grade HVAC sensors that are not typically re-calibrated during their operational lifetime or even at installation. Beyond sensor data, the physical models themselves also required an extensive training period, typically consisting of a whole day's worth of measurements from a series of step input signals to the AHU.

ASHRAE 1020 acknowledged that the first-principles FDD method relied on a model that only represents what is normally considered as "ideal" operation of the HVAC system. Design faults inherent to a system (of which many commercial installations are prone) present a considerable challenge to the robustness and sensitivity of this FDD method. In the ASHRAE 1020 research, non-ideal behavior was present as both unexpected relationships between outside air flow rates and supply fan speed, as well as the mechanical integrity of the damper linkages. Due to constant mechanical alterations for various experiments, the tuning of equipment like the damper linkages was not consistent and eventually played a role in determining the efficacy of the FDD method. This phenomenon can be extrapolated to consider situations where the FDD method is rendered ineffective due to training or operational data from equipment that is not tuned properly or in a state of disrepair. ASHRAE 1020 pointed out that the FDD method should recognize the long-term degradation of equipment once FDD has been initiated, but then the calibration process itself requires not only data but also consummate commissioning of the equipment.

To mitigate the influence of non-ideal starting conditions, statistical fault detection thresholds were typically reset and re-calibrated for each test period. This threshold adjustment was performed in a subjective manner in order to reduce the false alarm rate during periods where the system was known to operate "normally". Testing with AHU-1 was particularly troublesome in this respect and required two days of effort just to properly adjust the fault detection threshold. Even with those adjustments of AHU-1, however, non-ideal system behavior was predominant and created several false alarms during the related test period. In other cases, the false alarm rates were curtailed by setting wider thresholds for fault alarms. The eventual result of this strategy was to reduce the sensitivity of the system below that needed to capture small-signal faults that were induced in the AHU.

2.3.3.3 Summary for Physical Methods of FDD

Real world issues such as sensor placement, fault magnitude and significance, and imperfectly understood equipment stand between laboratory successes in FDD and commercial solutions for improved building operations. The ASHRAE 1020 project was instrumental in exposing these key issues that surround continued development of commercially viable FDD solutions, especially with respect to thermo-mechanical system models and non-intrusive load monitoring.

2.3.4 Qualitative Methods of FDD

2.3.4.1 Overview of Qualitative FDD

Qualitative models employ knowledge based relationships to infer the state of systems from their measured data. This modeling approach can be further subdivided into rule-based mechanisms and those that rely on qualitative physical models. The latter subdivision uses characteristics of physical equations that govern machine operation to create heuristics for

interpreting data from that machinery. For instance, the order and components of differential equations that govern thermo-fluid processes can be used to create dimensionless groups and order of magnitude relationships that describe measurable data. In this fashion, incomplete or imprecise models about systems may yield basic relationships for qualitatively assessing the measured performance of the system. Despite the apparent advantage of not requiring detailed physical models or complete data sets, the qualitative-physical modeling approach has not found wide application to fault detection and diagnosis about building HVAC&R systems (6).

2.3.4.2 Rule Based FDD

Models based on expert rules for FDD, unlike qualitative-physical models, are well represented in the relevant HVAC&R literature. Such models rely on expert or engineering knowledge about physical systems in order to create a decision tree for classifying system data. Rules based on expert knowledge are typically built from a collection of anecdotal and best practice guidelines regarding the variables in the measured system data set. Due to their computational simplicity, expert rule-based FDD tools have been researched for use on packaged or stand-alone HVAC&R equipment as well as for larger systems that are under the control of an energy management control system (EMCS). Despite the long history of expert-rule research (46) (47) (48) no tool based on this modeling approach has been widely commercialized for FDD of building systems. Expert systems are desirable due to their ease of development, ability to deal with uncertainty and transparent reasoning towards a conclusion, but these advantages are also the source of the model's weaknesses; expert systems tend to be very specialized towards a particular system and they are prone to failure when they encounter situations that extend beyond the boundaries of their programmed knowledge.

A close cousin to the expert rule method, FDD based on engineering first principle rules utilizes physical first principles to derive basic quantitative relationships between measurable system variables. Such engineering rule systems have been developed into full scale, tested diagnostic tools for air handlers by Katipamula and Brambley (48), (49) and House et al (50), (51). Both tools operate under similar principles where data from a building control system is

used to navigate through a decision tree composed of physical rules in order to reach a conclusion about the operation of an air handling unit.

Katipamula et al described their FDD efforts for automated AHU commissioning in 2003 (52). Their motivation for the technology stemmed from the need for reducing the cost and time of implementing human commissioning of complex building systems such as a large AHU. The design of Katipamula's automated FDD system included four fundamental modules:

- fault detection and diagnosis based on passive measurements
- proactive fault detection and diagnosis
- fault evaluation
- decision engine for suggesting solutions to faults, if any exist

The detection of economizer faults through passive monitoring of AHU data was accomplished by Katipamula et al in the Outdoor Air Economizer (OAE) software module of their Whole Building Diagnostician (WBD) package. The OAE module is capable of identifying 20 different ventilation-related faults, however it does not operate to find any faults related to water-side AHU components such as hydronic heating and cooling coils.

Figure 6 below illustrates the implementation of physical rules within decision tree architecture to facilitate passive FDD about AHU economizer operation. The diamonds and squares show the action of rules upon HVAC measurements, and the logical linkage between different operations as the process seeks a conclusive reaction to the input data. Beyond passive FDD means, Katipamula et al also explore proactive FDD that utilizes brief, but real-time test control sequences applied to an AHU in order to actively probe for AHU faults.

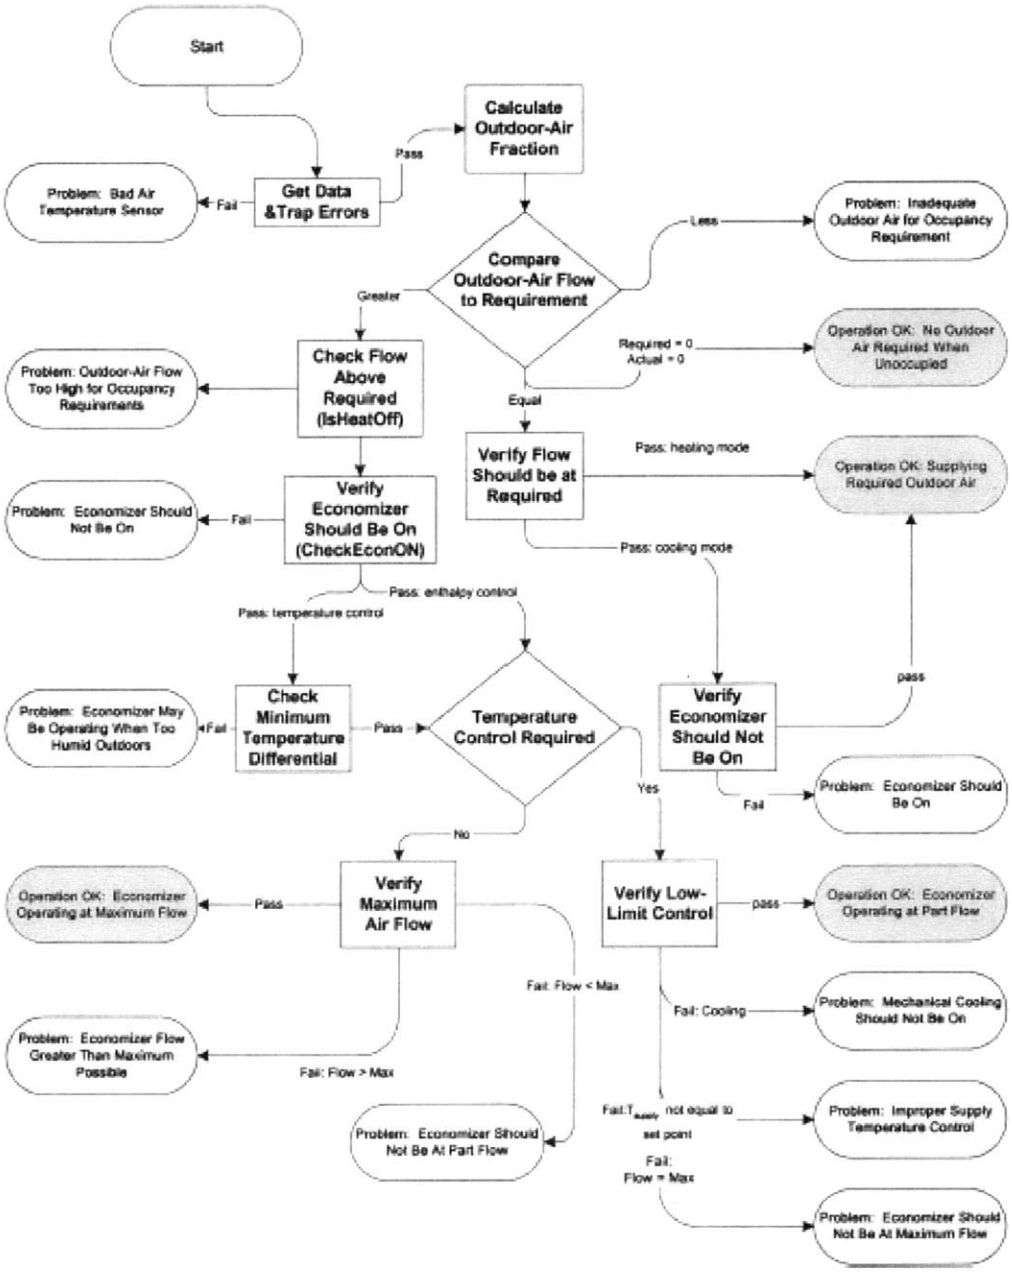

Figure 6 The decision tree used by the passive OAE fault detection module in Katipamula's WBD

In their 2003 paper, Katipamula et al describe an example of proactive FDD for testing of erroneous temperature sensor measurement. Figure 7 below illustrates that proactive FDD test sequence imposed on an AHU in order to create operating conditions to satisfy a physical rule that probes for faulty mixed air temperature measurement. As with all other FDD decision rules in Katipamula et al (both passive and proactive), physical rules are used to create analytical

redundancy about a physical measurement. Katimapula et al identify faults by measuring the difference between analytical predictions about a measured physical variable, and the real measured value of that variable. In essence, each FDD decision block in their approach is actually a streamlined first principles model about certain aspects of AHU behavior. This approach is repeated in their work for a wide variety of typical AHU faults, ranging from malfunctioning dampers to drifting pressure sensors.

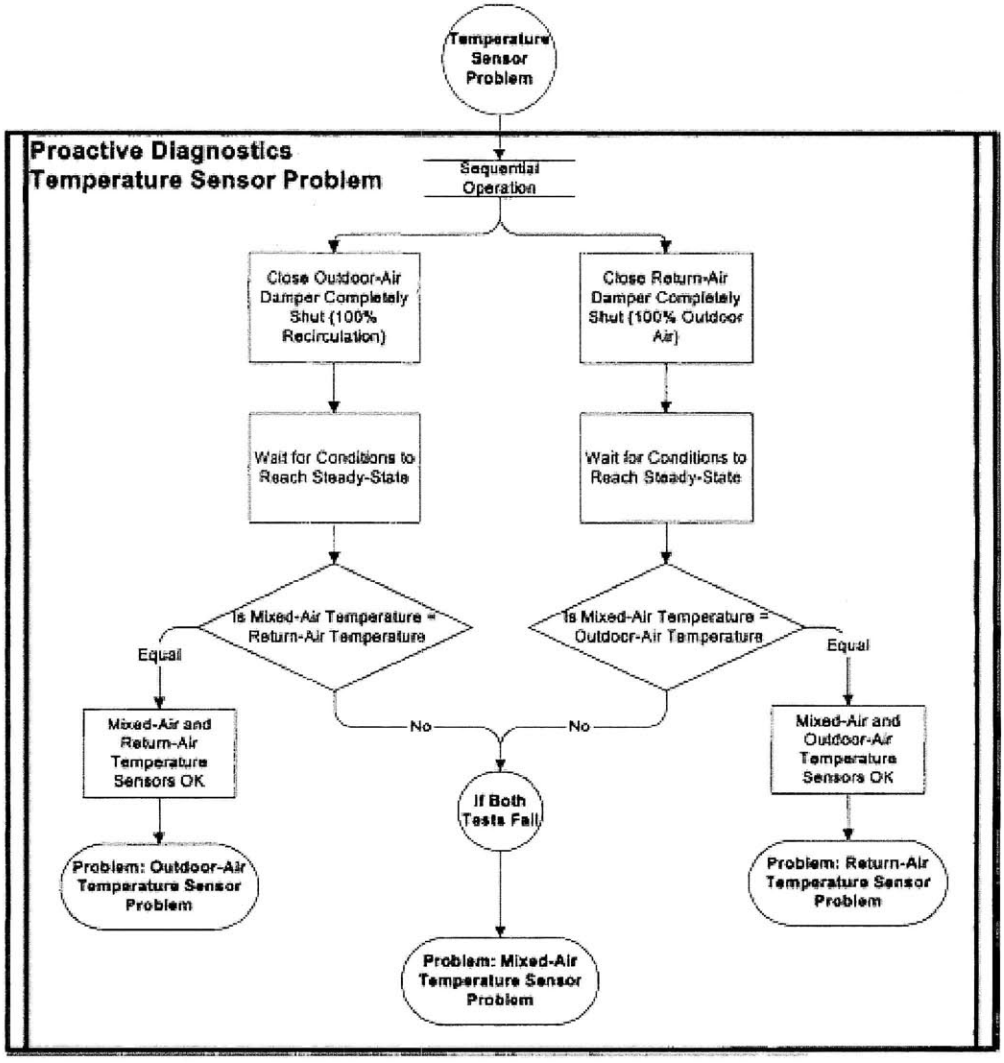

Figure 7 From Katipamula et al, 2003, the diagram shows a test sequence for proactive FDD of erroneous temperature measurement

While the work done by Katimapula et al has been successfully applied in several real building test beds, the modeling approach suffers from limitations that are common amongst all FDD that are based on engineering first principles (17):

- Lack of uniform AHU design across building stock
- Existence of un-modeled AHU behavior
- Transient behaviors in measurement data
- Poor tuning of thresholds and tolerances for fault classification

The last two items in the above list are complimentary in their efforts of causing false fault alarms. Thresholds and tolerances for classification rules must be set according to the dynamics of real time data measurement; first principle FDD systems must be sufficiently flexible to handle noise in the data as well as natural variations due to weather and occupancy. Such variations may be periodic, or even at times apparently random, but in either case they frequently cause a well tuned first-principles-based FDD to miss-fire a fault alarm. Un-modeled effects in AHUs have a similar impact on first-principle FDD because they create data that extends beyond the physical knowledge base of the FDD classifier. Finally, even if a first-principles-based FDD could adjust to noise and be resilient against un-modeled behavior, it is not likely that the FDD tool will be applicable across the entire building stock. This is attributed to the customized nature of large AHUs that are tailor designed, built, and programmed to the needs of a specific building. Consequently, first principle FDD rules that apply in one building may not necessarily apply in the building next door.

House et al (50), (51) developed a technology that was similar to that created by Katipamula et al in its use of first principles to derive expert rules about AHU operation. The new technology, termed APAR (AHU performance assessment rules), was composed of roughly two dozen rules about AHU operation that were derived from mass and energy balances applied about the AHU. House et al also relied on their unique method for classifying the operation of the AHU into five distinct categories to which operational rules were then assigned and used to infer corresponding faulty AHU operation.

APAR rule set

Mode	Rule no.	Rule expression (true implies existence of a fault)
Heating (mode 1)	1	$T_{sa} < T_{ma} + \Delta T_{sf} - \varepsilon_t$
	2	For $\lvert T_{ra} - T_{oa}\rvert \geq \Delta T_{min}$: $\lvert Q_{oa}/Q_{sa} - (Q_{oa}/Q_{sa})_{min}\rvert > \varepsilon_f$
	3	$\lvert u_{hc} - 1\rvert \leq \varepsilon_{hc}$ and $T_{sa,s} - T_{sa} \geq \varepsilon_t$
	4	$\lvert u_{hc} - 1\rvert \leq \varepsilon_{hc}$
Cooling with outdoor air (mode 2)	5	$T_{oa} > T_{sa,s} - \Delta T_{sf} + \varepsilon_t$
	6	$T_{sa} > T_{ra} - \Delta T_{rf} + \varepsilon_t$
	7	$\lvert T_{sa} - \Delta T_{sf} - T_{ma}\rvert > \varepsilon_t$
Mechanical cooling with 100% outdoor air (mode 3)	8	$T_{oa} < T_{sa,s} - \Delta T_{sf} - \varepsilon_t$
	9	$T_{oa} > T_{co} + \varepsilon_t$
	10	$\lvert T_{oa} - T_{ma}\rvert > \varepsilon_t$
	11	$T_{sa} > T_{ma} + \Delta T_{sf} + \varepsilon_t$
	12	$T_{sa} > T_{ra} - \Delta T_{rf} + \varepsilon_t$
	13	$\lvert u_{cc} - 1\rvert \leq \varepsilon_{cc}$ and $T_{sa} - T_{sa,s} \geq \varepsilon_t$
	14	$\lvert u_{cc} - 1\rvert \leq \varepsilon_{cc}$
Mechanical cooling with minimum outdoor air (mode 4)	15	$T_{oa} < T_{co} - \varepsilon_t$
	16	$T_{sa} > T_{ma} + \Delta T_{sf} + \varepsilon_t$
	17	$T_{sa} > T_{ra} - \Delta T_{rf} + \varepsilon_t$
	18	For $\lvert T_{ra} - T_{oa}\rvert \geq \Delta T_{min}$: $\lvert Q_{oa}/Q_{sa} - (Q_{oa}/Q_{sa})_{min}\rvert > \varepsilon_f$
	19	$\lvert u_{cc} - 1\rvert \leq \varepsilon_{cc}$ and $T_{sa} - T_{sa,s} \geq \varepsilon_t$
	20	$\lvert u_{cc} - 1\rvert \leq \varepsilon_{cc}$
Unknown occupied modes (mode 5)	21	$u_{cc} > \varepsilon_{cc}$ and $u_{hc} > \varepsilon_{hc}$ and $\varepsilon_d < u_d < 1 - \varepsilon_d$
	22	$u_{hc} > \varepsilon_{hc}$ and $u_{cc} > \varepsilon_{cc}$
	23	$u_{hc} > \varepsilon_{hc}$ and $u_d > \varepsilon_d$
	24	$\varepsilon_d < u_d < 1 - \varepsilon_d$ and $u_{cc} > \varepsilon_{cc}$
All occupied modes (mode 1, 2, 3, 4, or 5)	25	$\lvert T_{sa} - T_{sa,s}\rvert > \varepsilon_t$
	26	$T_{ma} < \min(T_{ra}, T_{oa}) - \varepsilon_t$
	27	$T_{ma} > \max(T_{ra}, T_{oa}) + \varepsilon_t$
	28	Number of mode transitions per hour $> MT_{max}$

Figure 8 From House et al (51), list of expert rules per category of AHU operation. Note that House only relied on temperature and humidity measurements about the system, as well as occupancy schedule and valve control signals.

In contrast to Katipamula et al and their extensive use of system model parameters, House et al reduced the set of variables to only a few temperature, humidity, occupancy and control signal measurements with corresponding input parameters for measurement tolerance and fault threshold. With fewer model parameters than Katipamula et al, House's FDD technology was designed to be simple and streamlined for rapid application to a typical economizing AHU.

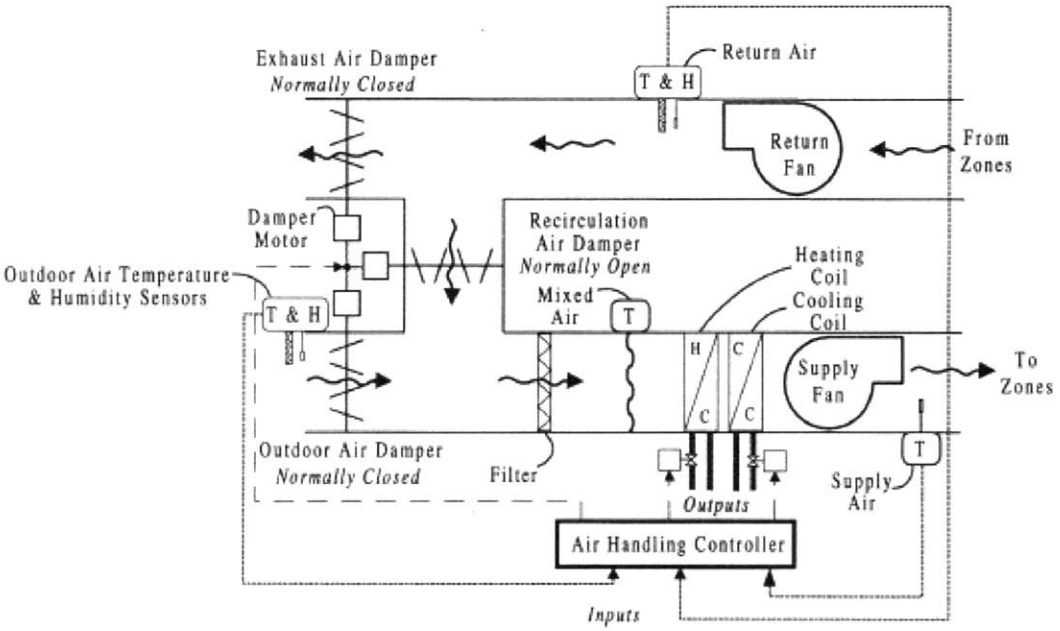

Figure 9 From House et al (50), schematic of a typical economizing AHU for which House et al developed APAR.

In fact, the design intent of APAR was to overcome the limitations imposed on FDD by the scarcity of fault-less system training data that spans the full range of AHU operating conditions. The APAR rules express engineering knowledge of how the AHU should operate and hence obviate the need for extensive training data about fault-free AHU operation. Furthermore unlike other expert HVAC FDD systems that focus purely on mechanical faults, the rules invoked by APAR seek out both mechanical and control based AHU faults.

The approach taken by House et al begins with identifying the occupancy status of the building and in the case of occupancy, further categorizing the mode of operation for the AHU. For the test AHU shown in Figure 9 House defines the following AHU operational modes

- Heating
- Cooling with outdoor air (free cooling)
- Mechanical cooling with 100% outdoor air
- Mechanical cooling with minimum outdoor air
- unknown mode

In mode 1, system heating, APAR assumes that the heating coil valve is controlled to maintain the supply air temperature at its set point value while the cooling coil valve remains closed. Furthermore, the mixing box dampers are modulated to allow the minimum outdoor air fraction mandated to satisfy building ventilation requirements.

Mode 2, cooling with outdoor air, is activated once the outdoor air temperature rises above the cooling outdoor air setpoint. In this mode, the heating and cooling coil valves are both closed and the mixing box dampers are modulated to maintain the desired discharge air temperature from the AHU. As the cooling load rises, however, the outdoor air flow rate will saturate and the building will require mechanical cooling in order to sustain indoor comfort conditions.

At this point there are two subsequent mechanical cooling modes: mode 3 supplements free economizer cooling with some mechanical cooling action, while mode 4 all together relies on mechanical cooling with the minimum required outdoor air fraction. Modes 3 and 4 represent opposite extremes of the cooling control logic, with the latter case most often found in very warm and humid ambient conditions. The control logic for modulating the AHU dampers in mode 3 can also vary at the discretion of the controls engineer; for APAR, the expert rules assume that the mixed air controls are driven by the comparison of the return and outdoor air enthalpies.

Finally, mode 5, unknown operation, is declared for any periods of AHU operation that cannot be classified as modes 1 through 4. The classification of the AHU in modes 1 through 5 depends on the measured control signals from the AHU control system.

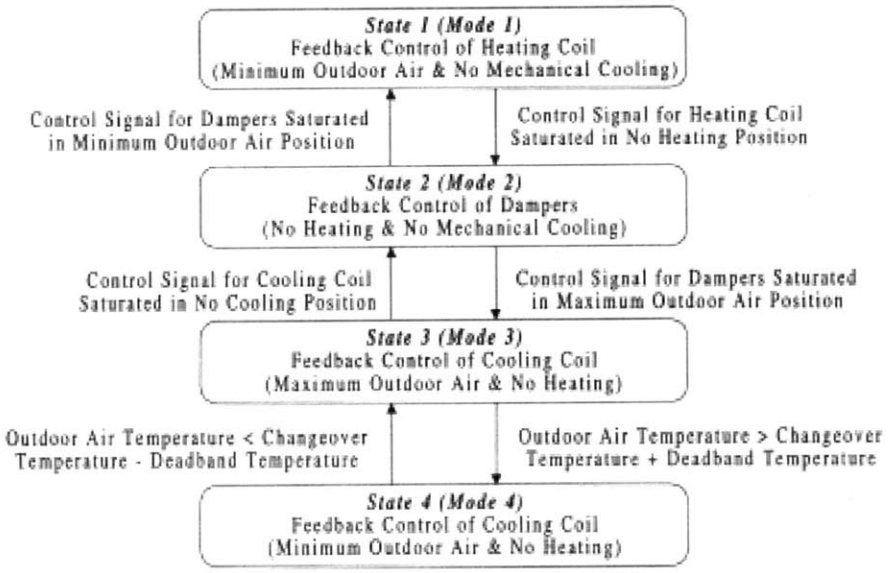

Figure 10 From House et al (51), diagram of transitions between classifications of AHU operational modes.

Once the AHU operational mode has been classified, families of expert rules are used to infer faulty AHU operation according to the schedule shown in Figure 8. The rules are written such that faults are identified when a rule is satisfied, however, APAR does not search for or guarantee the existence of a particular fault; instead, the satisfaction of a rule provides the basis for inferring that a range of faults or faulty-systems could exist. Despite the lack of specificity, the rules were designed to isolate faults related to actuator failures or degradation, sensor malfunction and drift, inherent mechanical design or installation flaws, erroneous control logic, and inappropriate operator intervention. In practice, individual APAR expert rules can infer several distinct fault mechanisms, with increasing resolution as more rules are activated. The design intent of APAR, however, was not to yield precise fault diagnosis but instead to better inform building operators of potential problems in their system.

Like most other expert rules systems, APAR exhibited a higher rate of false positives when applied to transient data instead of steady state data. To mitigate that failure mode, APAR was designed to operate only once per hour on exponentially weighted moving averages of the measured data. Furthermore, APAR rules were only applied to AHU data if the AHU remained in any one particular mode of operation for an entire hour. This constraint was obviously not

applied to rule 28 of APAR which measured the number of mode transitions executed each hour. Despite these constraints, APAR required a reasonably small set of commonly available input variables, most of which were temperature measurements and actuator control signals. In addition to the measured system data, APAR also required fundamental AHU specifications in order to setup and verify the suitability of the expert rule set for performing FDD about the AHU. Such system specifications included:

- Minimum and maximum expectation values for the actuator control signals
- Minimum outdoor air fraction
- Outdoor air enthalpy reset for economizer cycle
- Description of AHU control sequence

While the classification performance of APAR was not explicitly measured in simulation and field testing, the results of all tests have shown that APAR can successfully detect faults and often diagnose them to a high degree of specificity. Despite their success, House et al suggest that future research should focus on methods for dynamic updating of user-defined fault thresholds and measurement tolerances, field-testing and development of APAR rules for more AHU designs, and a rigorous study of APAR false alarm rates that occur during field trials.

2.3.4.3 Summary of Rule Based FDD

In general, rule based FDD falls into two categories; residual methods that infer faulty operation from the comparison of measurements and predictions about system variables, and fault rules that raise alarms when system variables exceed user-defined thresholds.

2.3.5 FDD Based on Process History

2.3.5.1 Overview of Process History FDD

In contrast to physical models that derive from first principles, or expert rules that derive from expert experience, process history FDD is derived from a mathematical regression over process data. There are a wide variety of mathematical techniques available for creating a data-driven multi-input and multi-output (MIMO) FDD model, including linear and multi-linear regression (LR and MLR), artificial neural networks (ANN), fuzzy logic (FL), support vector machines (SVM) statistical process control (SPC), Bayesian networks (BN) and combinations therein. Often times these analytical techniques are combined with first-principle derivations to yield grey box MIMO models for FDD, and in other cases the data regression is used on its own resulting in completely black box models. Complimentary to Katipamula and Brambley's treatment of the subject, an authoritative review of process history methods for equipment fault detection was published by Venkatasubramanian in 2003 (53) for chemical plant applications.

2.3.5.2 Grey Box FDD Methods

Parameter estimation in grey box FDD models has been explored for a wide variety of HVAC FDD applications by many authors. Katipamula and Brambley include several examples and a list of appropriate citations for the time frame of their publication. An important recent addition to the literature was made by Najafi et al in 2008 (54), where the authors extended the grey box modeling field to include Bayesian Network models of HVAC equipment. The novel approach captures a system model within the probabilistic framework of a Bayesian Net, and thereby supports posterior probabilistic inference on the existence of equipment faults. Najafi et al applied their method to model the mixing box of a variable-air-volume (VAV) air-handler (AHU), and detect some common faults such as stuck and leaking dampers. While the method was successfully tested and shown to work with experimental data from the Iowa Energy Center (IEC) that included known faults, the publication does not discuss how the approach can identify faults that are not a-priori expected in a data set.

Aside from the detection of purely anomalous behavior (in other words, data that cannot be explained by historical data), black box classifiers such as Bayesian Networks, Support Vector Machines, and Artificial Neural Networks cannot diagnose a specific fault without a-

priori or in-situ supervised learning of that fault. The a-priori approach to supervised learning would include a pre-existing database of fault-laden data that was labeled with diagnosed faults. In-situ supervised learning of faults would require an on-line labeling mechanism where expert analysts could label fault-laden data and re-train the machine classifier with the new classifications. In both cases, the space of faults is not fully defined or even known to be finite. Najafi et al discuss the additional, semi-supervised learning approach of grey box machine classifiers where first principle models and engineering rules can generate psuedo-data for training machines on fault-free and fault-laden equipment behavior. While this approach has the added benefit of excluding an expert analyst and labeling of fault-free and fault-laden behavior, it is still limited by the ability of physics and engineering rules to adequately model how equipment operates correctly and in-correctly. While the probability calculus of Bayesian network models has useful properties for judicious inference, the grey box approach is still ultimately limited by the model's perception of equipment physics, and the sufficiency of a historical database.

2.3.5.3 Black Box FDD Methods

Black box FDD models are developed in a similar fashion to grey box models, however the estimated model parameters have no physical significance. Katipamula and Brambley include a long list of authors who have developed statistical and non-statistical approaches to black box FDD modeling, with techniques that range from Artificial Neural Networks and fuzzy logic to various linear and non-linear regressions. Another more recent publication by Choi et al (55) showcases an expansion of the black box modeling technique to include support vector machine classifiers.

Choi et al discuss a data-driven approach to detecting faults in HVAC chiller performance through three different fault classifiers, most notable of which is a support vector machine. The classification of faults was only one component within their overall fault detection and isolation (FDI) scheme; other components included fault severity estimators, and a likelihood ratio test for residuals between model predictions and system measurements. The fault

signal for classification was generated in a similar fashion to the research in ASHRAE RP 1020; residuals between system measurements and predictions were passed through a statistical test for significance, which Choi et al termed a generalized likelihood ratio (GLR) test. Fault classification and severity were both judged by independent support vector machine classifiers in order to yield total fault detection and isolation. While the classification and severity ranking of faults was near perfect for their SVM classifiers, the authors do not mention the extent of training that was necessary to yield such accurate classifiers. Likewise, the entirety of the research project was based on noisy simulation data gathered from a detailed model of a centrifugal chiller (56), and a predefined space of chiller faults; the authors do not discuss how the system may perform in a more realistic setting. Like most other black box modeling techniques, SVM's require a learning algorithm (which may include historical data) in order to learn correct and in-correct operation of equipment. While the details of that learning algorithm are not included in their publication, Choi et al do show a promising new direction for extending black box models to include modern machine learning techniques such as the SVM.

2.3.5.4 Summary of Process History FDD Methods

Process history based FDD models are well suited to operate on equipment or systems where theoretical or engineering models of behavior are poorly developed or insufficient to describe the full range of operation. Despite their ostensible convenience for systems that are too complex to accurately model, black box models carry the additional requirements of training data sufficiency in order to learn complex models. Furthermore, as is often the case with ANNs, the over-use of training data can result in black box models that are hyper-trained towards a particular aspect of system behavior. Modern machine learning techniques such as the SVM and Bayesian Network are inherently designed to overcome the threat of over-training, and in fact minimize the need for training data all together. While such modern techniques may seem to be a universal solution for effective black box modeling, they still carry a need for some level of supervised learning of representative fault-laden behavior. As more live-building data is collected and labeled over time, it may be possible to create a master library of real-world faults

that is sufficient to train flexible machine classifiers to accurately identify a diverse variety of HVAC faults.

2.3.6 Summary of FDD

A wide variety of techniques and tools exist for performing fault detection and diagnostics on building systems. Modern FDD technologies leverage machine learning techniques and statistical analysis, expert rules derived from expert experience and engineering, and first principle physical models for system operation. Through this literature review we have identified several existing technical challenges that seem pervasive to the continued research and development of FDD technology:

- Lack of a real-world database of system and equipment operating data that is labeled with fault classifications
- The need for adaptive FDD thresholds in order to accommodate customized equipment installations
- The need for FDD systems to deliver dollar values, or hard costs, of faults that are identified
- The need for FDD systems to communicate the uncertainty of their inference

FDD research and development has attempted to answer these challenges through more accurate physical and grey box models of systems and equipment, FDD techniques that yield dollar estimates for fault diagnoses, and probabilistic inference in the form of Bayesian Networks, Fuzzy Logic, and other statistical analysis. Despite these attempts, it also appears difficult for current FDD implementations to strike a careful balance between ease of implementation without loss of utility or generalization; it seems that the majority of techniques have so far succeeded in the opposite extreme of yielding highly detailed and useful implementations with a large sacrifice of generality.

Based on this research review we have chosen to pursue the research and development of an FDD system that incorporates dynamic FDD thresholds for creative expert rules, within a

FDD framework that could scale to multiple equipment types and installations, while making straightforward probabilistic inferences that yield estimates on the dollar value of fault diagnoses.

3 Modeling of HVAC Equipment and Faults

3.1 Introduction

Fault detection is implicitly based on having a model of the system or piece of equipment that is under scrutiny; chapter two reviewed a wide variety of physical, black, and grey box modeling techniques for this application. Of the possible modeling techniques, physical and some grey box approaches are preferable by building engineers and technicians because the mathematical relationships tend to have physical interpretations and context that are more familiar to them. Following those preferences, we will also develop FDD models that are based on the physics of system and equipment operation.

In addition to choosing physical modeling techniques, we have also chosen to focus our research on the detection and valuation of faults in large air-handling units (AHUs). This is because AHUs and their associated air distribution systems support four out of the top five energy in-efficiencies that are commonly found in commercial buildings:

1. Simultaneous heating and cooling
2. Extraneous heating, cooling or ventilation
3. Imbalanced air flows
4. Leaking air distribution systems

AHUs also represent an intellectually challenging system of interacting components that exchange mass, momentum and thermal energy across multiple phases and media; the large degrees of freedom in these systems supports a rich and interesting space of possible faults. Likewise, AHUs are responsible for maintaining both the health and comfort of occupied spaces,

hence fault detection on those systems carries an impact on energy as well as occupant safety and health. In this particular chapter we will develop models of AHU equipment and some of their common faults in order to ascertain the impact of those faults, at various fault severities, and the difficulty of detecting those faults under various seasonal and occupancy conditions. Our fault space for this chapter will include:

1. Stuck AHU mixing box dampers
2. Leaking AHU mixing box dampers

Through the course of this chapter we will model and simulate the operation of AHUs under fault-free and fault-laden conditions, and qualitatively compare the results of simulation against experimental results collected in ASHRAE RP 1020 for the same faults on similar equipment.

3.2 Air handler Description

Figure 11 shows the schematic representation of a variable-air-volume (VAV) AHU; the VAV AHU is commonly found in large commercial buildings that have a single duct air distribution network.

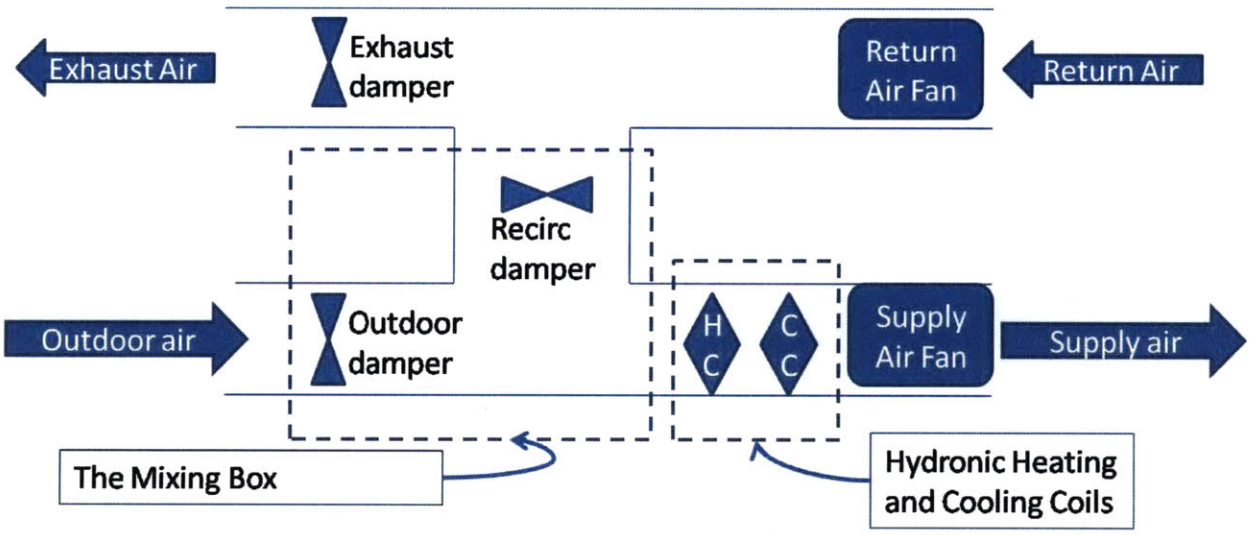

Figure 11 Schematic system model of the variable-air-volume air-handler studied in this project

The AHU design is termed variable-air-volume because the displaced air volume of the supply and return air fans can be varied by slowing down or speeding up the rotation rate of the fans. This is a modern, energy-efficient approach to ventilation since it allows the AHU to modulate ventilation rates in response to changing building occupancy conditions and thermal loads. In addition to the supply and return fans, the other primary sections of the representative air-handler include a mixing box for combining fresh outdoor air with recirculated air from the building, heating and cooling coils for adjusting air temperature and humidity, and a set of dampers that control the outdoor air content of the supply air flow. The air-handler modulates the circulation and quality of air within the building by adjusting the orientation of the exhaust, recirculation and outdoor air dampers, and the rotational speed of the return and supply air fans.

The temperature and humidity of the air that is supplied to the building are typically adjusted by heat exchange with the heating and cooling coils that are found within the air-handler. Given appropriate weather conditions, however, the desired temperature and humidity of the air supplied to the building may also be achieved simply by mixing adequate quantities of outdoor and recirculation air. This "air-mixing" approach to tempering the supply air temperature and humidity is typically termed free-cooling since it obviates the use of the heating and cooling coils in the air handler. There are other names for this approach to supply air conditioning, including air-side economizing, or just economizing.

Free-cooling is a very attractive method for providing comfortable air to a building since it requires no thermal energy consumption by the air-handler. The success of this approach, however, is predicated on the existence of appropriate outdoor weather conditions, as well as the fault-free operation of dampers within the mixing box. Free cooling is possible during at least half the year for buildings in a cool or mild climate zone, such as the Northern United States, but its usefulness decreases as the mean annual outdoor air temperature of a location increases beyond 65 ^{o}F. Since free-cooling is dependent on the appropriate control of outdoor air flow into the AHU, mixing box faults such as stuck or leaking dampers may render free cooling less effective, or even completely impossible.

Dampers within a mixing box may fail in a wide variety of ways, all of them rendering the air-handler unable to rely solely on free cooling to condition the supply air to a desired temperature and humidity. In the ensuing discussion we will focus on just two common damper faults, a stuck damper and leaking damper, at three different levels of severity for each of them. These two faults are prevalent in commercial buildings with large air-handlers and they are considered by engineers as a significant source of energy in-efficiency in buildings with large heating, ventilation and air-conditioning (HVAC) systems. Beyond the energy implications of these faults, lack of control over the intake of outdoor air by the air-handler can also result in catastrophic mechanical failures such as freezing coils or uncomfortable and unhealthy building conditions due to poor air quality.

$$\dot{V}_{SA}\rho = \dot{V}_{RC}\rho + \dot{V}_{OA}\rho$$

Equation 3-1

$$\dot{V}_{SA}c\rho T_{MA} = \dot{V}_{RC}c\rho T_{RA} + \dot{V}_{OA}c\rho T_{OA}$$

Equation 3-2

The subscripts OA and RC indicate variables for the outdoor air and recirculated air, respectively. The variables $\dot{V}, \rho, c, T$ correspond to volumetric flow rate, density, thermal capacitance and temperature of the air, respectively. The subscript MA refers to the mixed air

conditions that exist at the boundary between the outlet of the mixing box and inlet to the hydronic coil control volume. A species balance for air across the hydronic coil control volume shown in Figure 11 reveals that the mass flow rate of air at the mixed air boundary must equal the mass flow rate of air supplied to the building; this result is included in Equation 3-1 and Equation 3-2 with the subscript SA on the volumetric flow terms. Contributions from humidity are also important in an enthalpy balance and mass balance, but they have been neglected here in order to simplify the analysis; those contributions will be included later on in the final formulation of our results.

The heat transfer between air and the air-handler's heating and cooling coils is intended to adjust the mixed air thermal conditions to achieve the supply air conditions, which are the desired temperature and humidity of air that is distributed within the building. The energy required to raise the mixed air conditions to the supply air conditions is found through an enthalpy balance about the hydronic coil control volume shown in Figure 11.

$$E_{H\&C} = \dot{V}_{SA}\rho c(T_{MA} - T_{SA})$$

Equation 3-3

The new term in equation three, $E_{H\&C}$, is the enthalpy exchange rate between the air flow and hydronic heating and cooling coils; this is the thermal energy cost for operating the air-handler. Combing enthalpy balances across the mixing box and hydronic coil sections yields an overall energy cost for comfort conditioning that is based on supply, return and outdoor air flows and temperatures.

$$E_{H\&C} = \dot{V}_{RC}c\rho T_{RA} + \dot{V}_{OA}c\rho T_{OA} - \dot{V}_{SA}c\rho T_{SA}$$

Equation 3-4

Outdoor air is drawn into the air-handler in an attempt to yield mixed-air conditions that require the least amount of hydronic heating or cooling enthalpy exchange in order to achieve supply air conditions. A minimum outdoor air flow rate is also required by building code in order to satisfy mandates on indoor air quality. Depending on the ambient conditions outside the building,

mixtures of recirculated and outdoor air flow may yield the required supply air conditions and altogether eliminate the need for mechanical cooling or heating. Conversely, mandates on minimum outdoor air flow rate may incur extraneous cooling or heating energy expenditure when the outdoor air cannot mix with the recirculated air to yield the supply air conditions.

The role of outdoor air in heating and cooling energy expenditure suggests that the enthalpy exchanged defined in equation four may be scaled by the energy that is needed to create supply air conditions from the recirculated air flow alone.

$$E_s = \dot{V}_{SA}\rho c(T_{RA} - T_{SA})$$

Equation 3-5

Dividing the heating and cooling energy cost function in Equation 3-4 by the above energy scale yields a dimensionless energy term that isolates the impact of outdoor air flow on the thermal energy cost of the air-handler. Eliminating the recirculation flow rate through the mass balance around the mixing box section, and dividing through by the supply air flow rate yields a dimensionless energy loss index that is a function of the outdoor air fraction and the supply, return, and outdoor air temperatures

$$E^* = \frac{(1-g)T_{RA} + gT_{OA} - T_{SA}}{T_{RA} - T_{SA}}$$

Equation 3-6

$$g = {\dot{V}_{OA}}/{\dot{V}_{SA}}$$

Equation 3-7

The energy loss index show in Equation 3-6 is a measure of the impact of outdoor air flow on the thermal energy consumption of the air-handler. The total energy consumption of the air-handler would require additional terms that represent electricity consumption of the supply and return fans. Those additional fan energy terms would require fan models that yield fan energy as a function of pressure and flow, and ultimately ventilation in the building. Models for ventilation in buildings are typically complicated and time-consuming to develop, and the purpose of Equation 3-6 is to create an energy metric that can be evaluated without a model for real-time building ventilation. The thermal energy analysis represented by equation 6 is a simple alternative to total energy calculations that isolates the impact of faulty dampers on air-handler energy consumption.

The outdoor air fraction, denoted as g and defined in Equation 3-7, is the ratio of the outdoor air flow rate to the supply air flow rate; it is typically controlled by varying the recirculation damper position. When the outdoor air fraction is equal to zero, the energy loss index, shown in equation six, is equal to one. For non-zero values of outdoor air fraction, the energy loss index is a function of outdoor, return and supply air temperatures. The energy loss index will be negative or positive, respectively, for cases where heat must be added or removed from the mixed air flow in order to meet supply air conditions. Outdoor air fractions that minimize energy consumption for any given set of supply, return, and outside air temperatures can be derived from setting the numerator of equation six equal to zero.

In this simple model the outdoor air fraction might be controlled through the following set of rules that satisfy requirements on indoor air quality and guarantee the smallest possible energy loss index across all thermal conditions:

$$T_{OA} \geq T_{RA}, \quad g = g_{min}$$

Equation 3-8

$$T_{SA} \leq T_{OA} < T_{RA}, \quad g = 1$$

Equation 3-9

$$T_{OA} < T_{SA}, \quad g = \frac{T_{RA} - T_{SA}}{T_{RA} - T_{OA}}$$

Equation 3-10

The control rules for the outdoor fraction are such that the outdoor air fraction is a minimum whenever the outdoor air temperature exceeds the return air temperature. Conversely, the outside air fraction is at a maximum value whenever the outdoor air temperature is between the return and supply air temperatures. For days that are sufficiently cold, the outdoor air fraction can be modulated between its minimum and maximum values to yield an energy loss index equal to zero. This set of control rules could be modified to include humidity contributions by replacing the temperature terms with their respective total air enthalpies.

3.2.1 Mixing Box Air Flow Modeling

The flow of outdoor air into an air-handler, and commensurate outdoor air fraction, is controlled by the pressure differential that exists between the core of the mixing box and the air pressure at the inlet to the building's outdoor-air intake duct. This pressure differential is physically dependent on several aspects of air-handler design and operation, but the proper adherence to ASHRAE design guidelines should yield a primary dependence on recirculation damper position. In particular, ASHRAE guideline 16-2003 details the selection of dampers in order to achieve certain levels of recirculation damper authority,A, that reduces the influence of other variables on outside air flow.

$$A = \frac{\Delta P_{damper,100\%\ open}}{\Delta P_{damper\ duct\ total}}$$

Equation 3-11

The damper authority is defined as the ratio of the pressure drop across the recirculation dampers when they are 100% open to the total pressure drop in the duct section that contains the recirculation dampers, both measured under maximum flow conditions. Other equipment such as outside air louvers, exhaust air dampers, and variations in supply air fan speed can also influence the mixing box pressure, but we will assume for now that those effects are muted by proper damper selection.

Damper performance is reported in ASHRAE Guideline 16-2003 as the percent of maximum possible flow through the damper versus percentage opening of the damper.

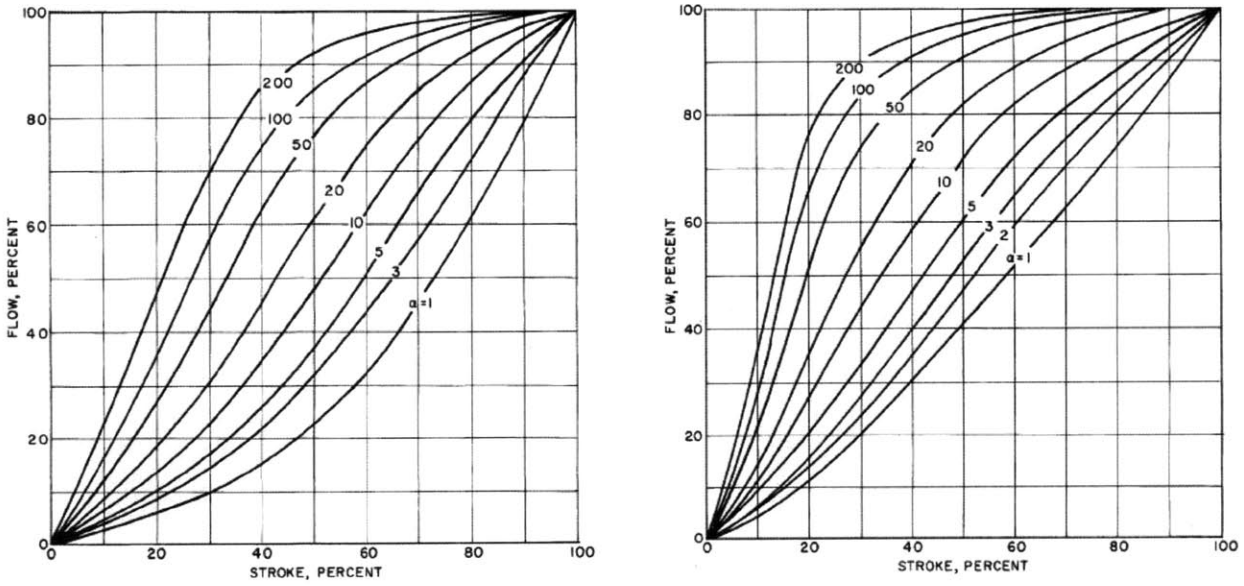

Figure 12 Families of curves with varying damper authority for percent of maximum recirculated flow through opposed (left image) and parallel (right image) dampers, as a function of percent opening of the damper. ASHRAE Guideline 16-2003

The maximum possible recirculation flow through the recirculation damper is a function of the actual supply air flow rate and the minimum required outdoor air flow rate.

$$\dot{V}_{RC,max} = \dot{V}_{SA} - \dot{V}_{OA,min}$$

Equation 3-12

ASHRAE guideline 16-2003 is intended to help engineers design recirculation system with a damper authority that results in the most linear relationship between the recirculation damper position and maximum recirculated air flow. This linear relationship is convenient and desirable for the programming of building controls, because it simplifies the code needed to modulate recirculation air flow based on recirculation damper position. Non-linear relationships between recirculation damper position and air flow would require more complex controls programming to accommodate for that non-linearity, or the purchase of additional sensor equipment to support a local feedback control loop. For parallel blade recirculation dampers, the right hand image in Figure 12 shows that a design damper authority between 2 and 5 results in an approximately linear relationship between damper opening and percentage of maximum recirculation flow.

$$\frac{\dot{V}_{RC}}{\dot{V}_{RC,max}} \approx rad$$

Equation 3-13

Equation 3-13 is a linear constitutive model for the recirculation damper that includes the percent opening of the recirculation damper,rad, the recirculation flow rate and maximum possible recirculation flow rate. Combining the linear constitutive model in Equation 3-13 with the definition of maximum recirculation flow in Equation 3-12 yields a new relationship between outdoor air flow, supply air flow, recirculation air flow and recirculation damper position.

$$\dot{V}_{RC} = rad \cdot \dot{V}_{SA} + rad \cdot \dot{V}_{OA,min}$$

Equation 3-14

The recirculation air flow can be eliminated in favor of the outdoor air flow and supply air flow by using Equation 3-1 for the mass balance around the mixing box section. Dividing through by the supply air flow yields the following dimensionless relationship between outside air flow and recirculation damper position for air-handlers with parallel blade recirculation dampers that have a damper authority between 2 and 5.

$$g = 1 - rad \cdot (1 - g_{min})$$

Equation 3-15

When the return air dampers are fully closed, the outdoor air fraction is equal to one. Conversely, when the return air dampers are fully open, the outdoor air fraction is reduced to its minimum value. Equipped with a constitutive model that relates the recirculation damper position to outdoor air fraction, we can now use Equation 3-7 through Equation 3-10 to create control laws that minimize the energy loss index by modulating the recirculation damper position. We can also leverage Equation 3-14 to generate additional constitutive models that represent a variety of faulty damper conditions.

3.2.2 Mixing Box Recirculation Damper Fault Modeling

A stuck damper is the simplest generic fault to model because the damper remains in one position and cannot move. This is an important fault to model because it represents a wide variety of physical situations that waste energy, such as binding of damper linkages, loss of air-pressure in pneumatic actuators, or loosened connections between dampers and their actuators. Many dampers also come pre-loaded with springs that drive the damper blades fully open or closed when the damper linkages fail. The fault-free constitutive model in Equation 3-14 can be modified to represent all of these stuck-damper situations by replacing the variable rad with a constant, α.

$$g = 1 - \alpha \cdot (1 - g_{min})$$

Equation 3-16

The fault-free model in Equation 3-14 can also be modified to represent leaking dampers, which is another class of damper faults. In this case the damper is still able to move, but the damper blades cannot provide an air-tight seal when they are completely closed. In contrast to a stuck damper, the effects of a leaking damper will also vanish as the damper continues to open.

The effect disappears because the contribution of the leakage air-flow to the total air-flow becomes very small as the damper opens wider. The subtraction of an exponential decay from the original linear model in Equation 3-14 is one of many flexible representations for a leaking damper,

$$g = 1 - rad \cdot \delta - \beta e^{-(\gamma \cdot rad)}$$

Equation 3-17

The parameter β in this equation represents the damper leakage that occurs when the damper is fully closed, and the parameter γ measures how fast the effects of the leak vanish. The parameter δ is a constant that represents the minimum required outdoor air flow rate. In order to qualify as a model for the damper, the parameters β, γ and δ must satisfy two conditions: the outdoor air fraction is always a decreasing function of return air damper position, and the minimum outdoor air fraction is achieved when the return air damper is fully open. These boundary conditions lead to a set of constraints that can be solved to simulate the effects of various damper leakage levels.

$$\beta = \frac{1 - g_m}{\gamma + e^{-\gamma}}$$

Equation 3-18

$$\delta = 1 - g_m - \beta e^{-\gamma}$$

Equation 3-19

The mathematical properties of the exponential decay provide a convenient representation of damper leakage because the response is fully specified by only two boundary conditions and yields a decay that asymptotically approaches the fault-free behavior. Polynomials or trigonometric functions can also be used to represent damper leakage, however they may require additional boundary conditions to fully specify the response; the exponential decay appears to be a simple choice amongst these options.

Beyond these fault-models, there are numerous ways in which the mixing box may actually malfunction and alter the relationship between outdoor air flow and recirculation damper position. The relationships developed in Equation 3-12 through Equation 3-17 allow us to simulate mixing box performance under a variety of these circumstances, and estimate their impact on the air-handler's energy consumption. Our expectation is that we can determine a-priori what faults and fault severities create the greatest opportunities for energy efficiency improvements, and thereby triage fault signals to only those that are most important to energy efficiency.

3.3 Mixing Box Simulations

3.3.1 Assumptions

In conjunction with typical meteorological year (TMY) data for a given geographic location, we can use the models derived in Equation 3-14 through Equation 3-20 to explore the impacts of various mixing box faults on the relative annual energy consumption of an air-handler. To simplify the analysis, we will assume that the return air damper control signal, rad, follows ideal control laws that are derived from Equation 3-7 through Equation 3-9 and reproduced here:

$$T_{OA} \geq T_{RA} \quad rad = 1$$

$$T_{SA} \leq T_{OA} < T_{RA} \quad rad = 0$$

$$T_{OA} < T_{SA} \quad rad = \left(1 - \frac{T_{RA} - T_{SA}}{T_{RA} - T_{OA}}\right) \cdot \frac{1}{1 - g_{min}}$$

Under fault-free conditions, the control signal and physical position for the damper are assumed to be identical. In that case, Equation 3-14 applies to compute the outdoor air fraction

from the damper control signal. The results from Equation 3-14 are subsequently consumed by equation 6 to yield the energy loss index for the thermal conditions that gave rise to that damper position. The same computation process applies to fault-laden simulations, with the exception that equation 14 is replaced with the appropriate constitutive model for the fault under investigation.

In order to easily evaluate Equation 3-20 without an additional model for the supply air flow rate, we will assume that the minimum outdoor air fraction is a constant equal to 0.2. This assumption is reasonable for certain types of buildings, such as commercial office space, where we hope to deploy fault detection systems. The minimum outdoor air flow rate for an office building with large cooling loads could be on the order of 10% of the maximum supply air flow rate for the air-handler. If the building is equipped with a typical single-duct, variable-air-volume distribution system, then the supply air flow from the air-handler will be modulated to maintain the static duct pressure of the building. A typical variable-frequency drive can modulate the supply air flow between 50 and 100% of its maximum value, so for this type of building the minimum outdoor air fraction should only vary between 0.1 and 0.2. By assuming the larger value for the minimum outdoor air fraction, we can compute an upper bound on the energy loss index for periods when the air-handler would ideally draw the least amount of outside air. To completely specify the year round operation of the building, we must also make assumptions about the supply and return air conditions.

The air-handler in a commercial office building is typically controlled to maintain supply air conditions at 55 ^{o}F. The return air conditions are a function of the building loads, which will vary over the course of the day and season. For simplicity and the purposes of this sensitivity analysis, however, we will assume that the building loads vary such that the return air temperature is always 72 ^{o}F. This assumption is also reasonable during the day-time operating hours of a commercial office building where the primary cooling load of the building is due to the people and equipment in the building. More sophisticated simulations may relax these assumptions and leverage existing building energy simulation programs such as energy-plus in order to create more realistic operating conditions. At this point, however, only a simple building

simulation is needed in order to identify the energy impact of various mixing-box faults relative to the base-case of fault-free mixing box operation.

3.3.2 Results for Boston, Massachusetts

The stuck damper was simulated for a recirculation damper stuck at 10%, 50% and 90% open positions; these are the corresponding values of α for use in equation 15, the stuck damper constitutive model. The leaking damper was tested for a small, medium and large leak, characterized respectively by values of β in equation 15 that include 0.1, 0.3, and 0.5.

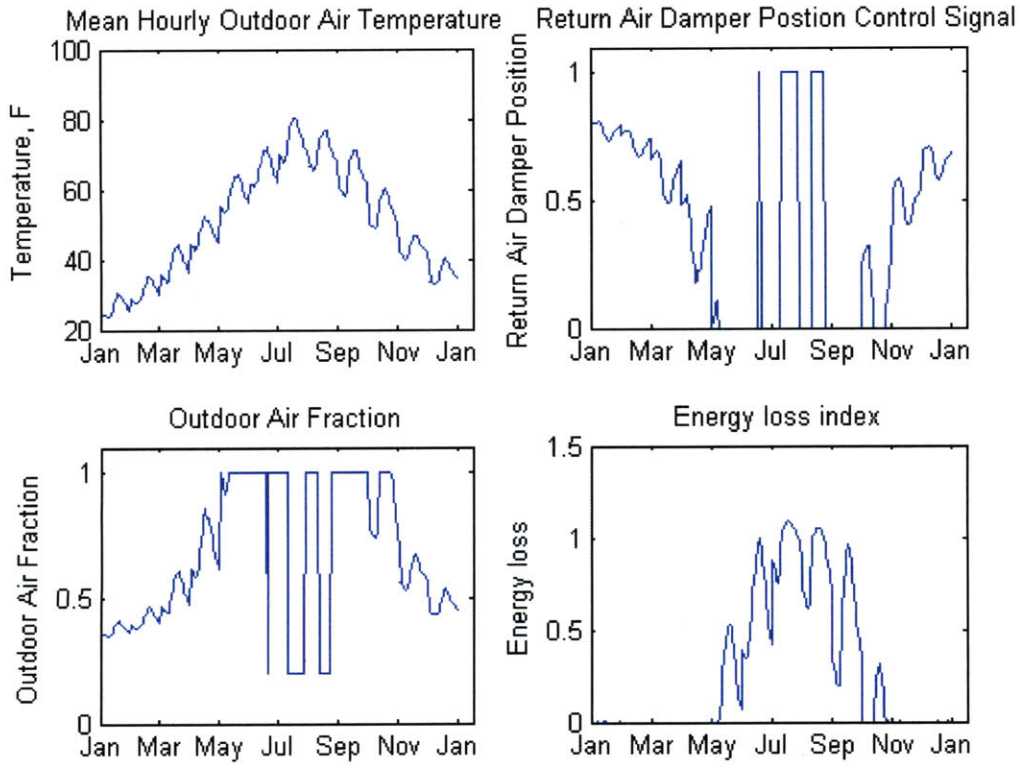

Figure 13 Baseline operation of Air-Handler under fault-free operating conditions for Boston, Massachusetts

The top left panel of Figure 13 depicts the mean hourly outdoor air temperature for twenty four hours in each month for the city of Boston, Massachusetts; the time series is in 24

hour intervals for each month. The annual variations in recirculation air damper position are shown in the top right panel of Figure 13. The majority of the year includes modulating the recirculation dampers to take advantage of free cooling with outdoor air; the damper spends more than 60% of the year at less than 60% open, with the most likely position between 0 and 10% open. Subsequent variations in outdoor air fraction and energy loss index are shown in the bottom panels of Figure 13. The lower right-hand panel of Figure 13 shows that the damper control rules are sufficient to drive the fault-free energy loss index to zero during the winter months, and keep the loss index less than one for the vast majority of the year. The same graph shows that most of the thermal energy consumption of an air-handler in Boston occurs during the summer months when the mixed air must be cooled to meet supply air conditions.

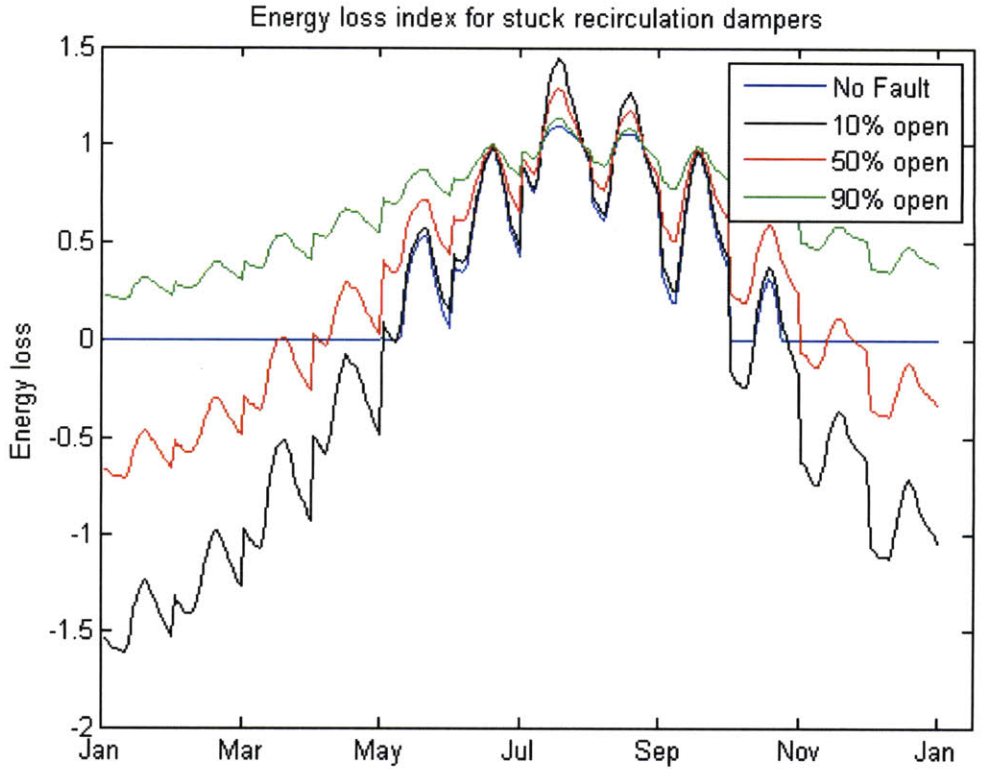

Figure 14 Energy loss index simulated for a recirculation damper stuck at various positions

Figure 14 depicts the energy loss index simulated for a year of mixing box operation with recirculation dampers stuck at various positions, and under normal operation. The net overall percent increase in energy use over the fault-free case for 10%, 50% and 90% stuck open

dampers is 169%, 82% and 132%, respectively. The largest losses relative to the fault-free case are during the winter period when the recirculation damper is stuck at 10% open; those results are shown in black within Figure 14. These losses are a result of drawing large quantities of cold, outdoor air into the air-handler that must be heated to the supply air conditions. Beyond energy implications, the excessive intake of very cold air is dangerous for the air-handler since the mixed air temperature could drop sufficiently below the freezing point of water and cause hydronic coils in the air-handler to burst due to water freezing and expansion. The opposite extreme is a damper that is stuck at 90% open and cannot draw sufficient outdoor air to take advantage of free cooling. The air-handler must constantly cool the mixed air conditions to reach the supply air temperature, which results in the green line in Figure 14 that is constantly greater than zero. A damper stuck at 50% open exhibits characteristics that are derived from both the 10% and 90% stuck open cases, but a net energy consumption that is less than the either of those two faults.

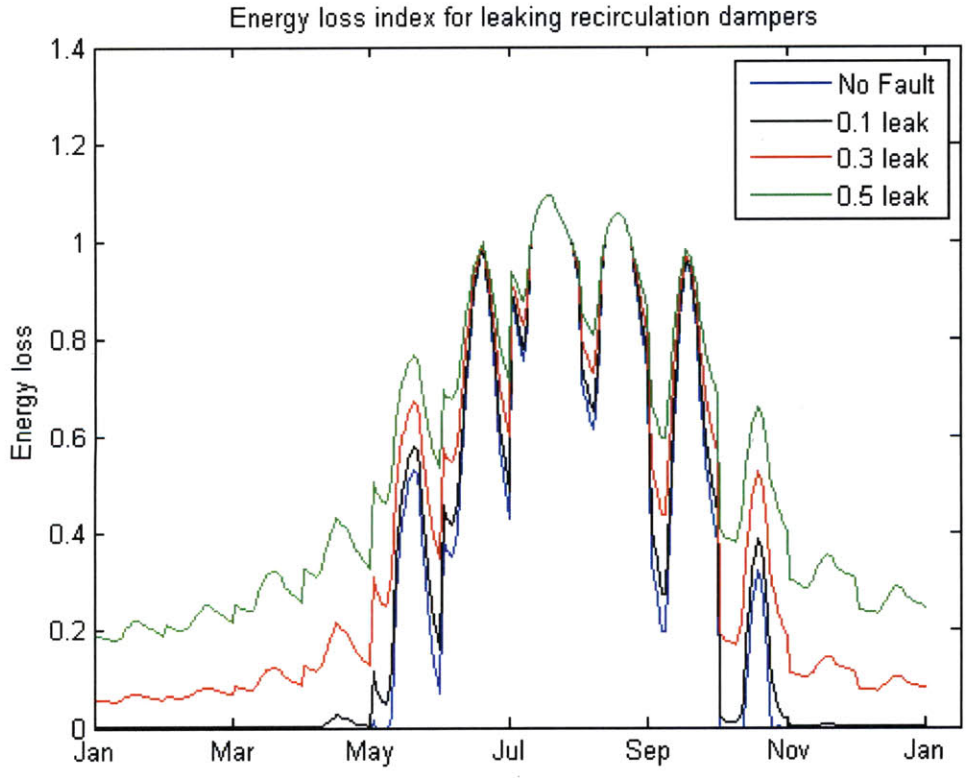

Figure 15 Energy loss index simulated over three damper leakage levels

Figure 15 depicts the energy loss index simulated for a year of mixing box operation with three different damper leak severities, and under normal operation. The net overall percent increase in energy use over the baseline case for 0.1, 0.3 and 0.5 values for β is 6.8%, 39% and 86%, respectively. Clearly, the energy loss due to leaking dampers grows with the severity of the damper leak. In contrast to the stuck damper faults, however, none of the leaking damper faults increase the amount of energy consumed for heating the mixed air to supply air conditions. In fact, all three of the simulated leaking dampers appear to increase the cooling energy consumption of an air-handler much in the same way as a damper that is stuck at 90% open.

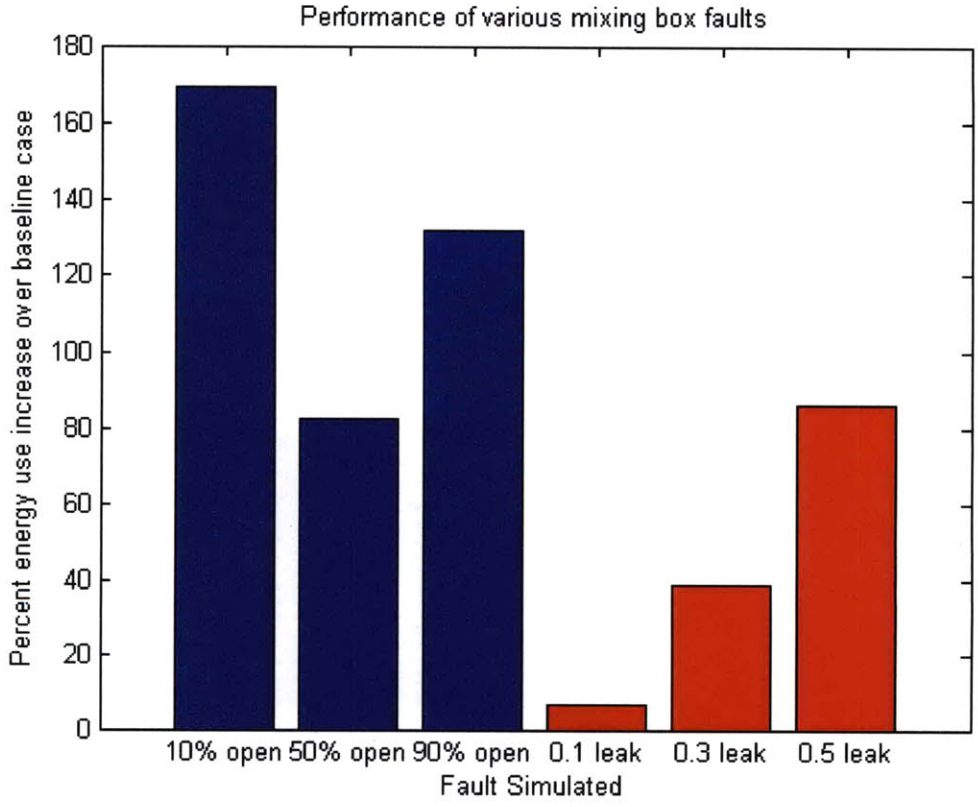

Figure 16 Energy impact of various mixing box faults

The energy impact of the different faults relative to the fault-free case is shown in the bar chart in Figure 16. The most significant increase in energy consumption follows from a recirculation damper that is stuck at 10% open, resulting in an increase of energy consumption by 170%. The smallest energy impact is only a 7% rise in energy consumption for a leaking damper with β equal to 0.1.

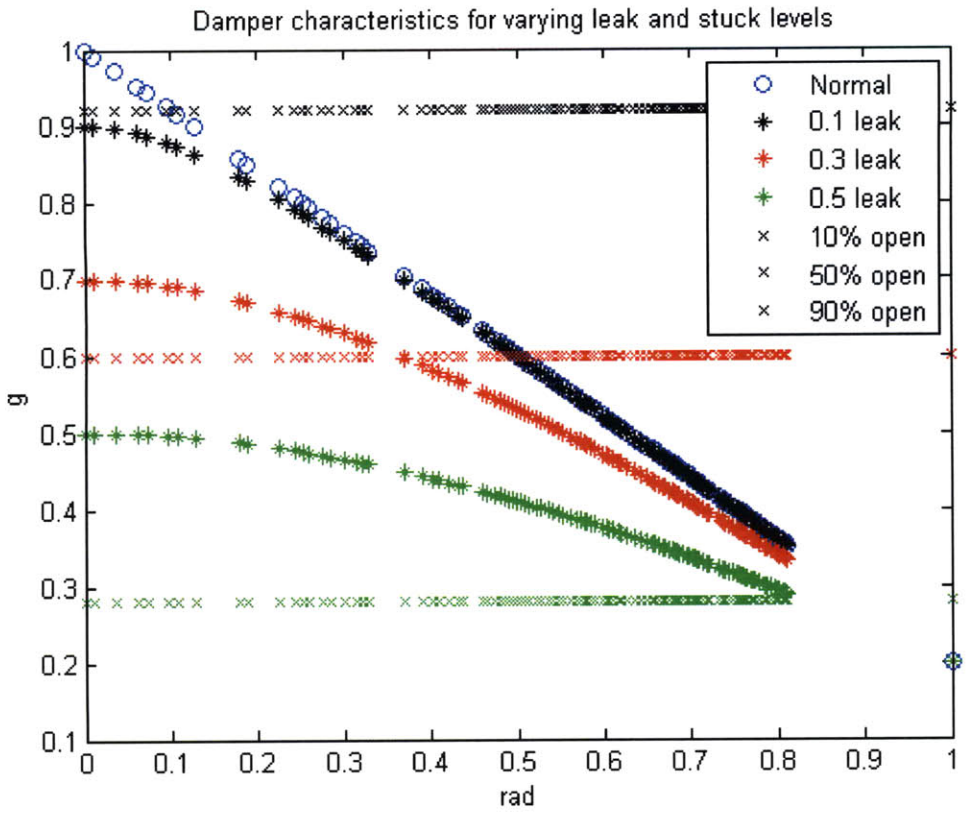

Figure 17 Outdoor air fraction versus recirculation damper position for several fault cases, simulated over a year's worth of Boston weather data

In addition to the energy loss index time-series in Figure 14 and Figure 15, the simulations also provide scatter plots of damper position versus outdoor air fraction for all fault-free and fault-laden cases. The stuck dampers are immediately identifiable in Figure 17 by their characteristic horizontal lines across the entire range of recirculation damper positions. The most severe leak with a β value of 0.5 has a sufficiently small slope for damper positions less than 40% that it resembles a stuck damper for that same range of damper positions. On the other hand, the curve for a leaking damper with β equal to 0.1 is almost identical to the fault-free curve for damper positions that exceed 20% open.

One potential metric that could be used to distinguish between these curves, and eventually support fault detection is the square root of the mean squared residuals on outdoor air fraction between the normal and faulty operating curves

$$d = \sqrt[2]{\frac{1}{N}\sum_{i}(g_i - \hat{g}_i)^2}$$

Equation 3-20

This error metric, d is regarded by engineers as a traditional measure of the difference between two sets of data; in this case it measures the difference between the fault-free outdoor air fraction, denoted by g_i, and fault-laden outdoor air fraction, denoted by $\hat{g}_i$, measured over the full range of recirculation damper positions. The error metric is computed for the simulation data where N is the total number of points, and the subscript i refers to pairs of faulty and fault-free data for each position of the recirculation damper. For dampers stuck at 10%, 50% and 90% open, the error metric is 0.37, 0.29, and 0.49, respectively. Likewise, the error for leaking dampers with leakage rates of 0.1, 0.3 and 0.5 is 0.06, 0.19, and 0.33, respectively.

The largest and smallest energy impacts are due to fault-laden behaviors that yield residual error values of 0.37 and 0.06, respectively. This result is rather counter-intuitive since we might expect that the most divergent behavior, found here with an error residual of 0.49, would yield the greatest energy impact, yet it does not. The energy impact of faulty behavior defies our intuition because energy consumption is a function of both the weather conditions and the matching operations of the equipment. For Boston, Massachusetts, it is far worse for a damper to be stuck at 10% open rather than 90% open, because of the typically low air temperatures in the northeast. In order to identify faults that have the greatest impact on energy consumption, fault detection should implicitly consider the extent of divergent behavior as well as potential energy impact when classifying measurements as either fault-negative or fault-positive.

3.4 Measured performance of a mixing box

3.4.1 Fault-free operation

The simulation results shown in Figure 13 are based on an ideal linear model of the mixing box. It is unlikely that a real mixing box in any air-handler will exhibit such convenient and simple behavior. On the other hand, we expect that the outdoor air fraction in a real air-handler will be a monotically decreasing function of increasing recirculation damper position. This assumption is based on physical intuition that resistance to air flow through the recirculation damper must continuously decrease as the dampers extends from a closed to fully open position.

Several examples of real mixing box data have been prepared in order to qualitatively compare the previous simulations against real mixing box data, and provide additional insight into the operation of a real mixing box.

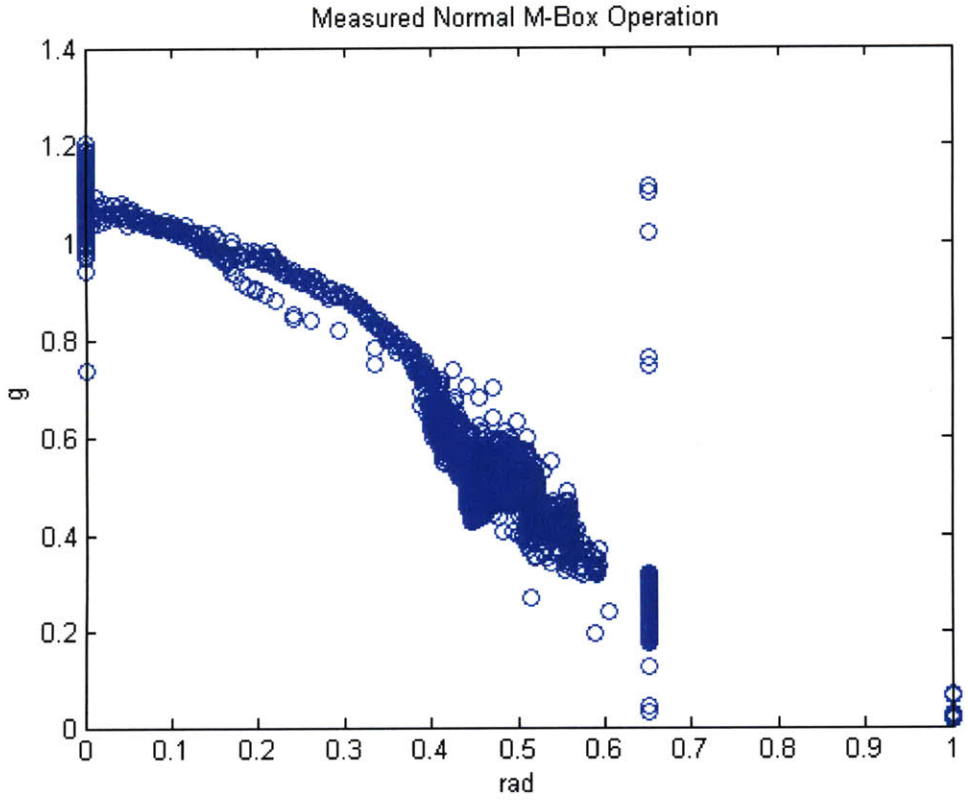

Figure 18 Measured fault-free operation of Air-Handler B at the IEC

Figure 18 shows real mixing box data measured from air-handler B at the Iowa Energy Research Center (IEC) in Des Moines, Iowa. The data for this figure has been compiled from

several days of fault-free mixing box data from the ASHRAE 1020 research project, which was conducted at the IEC. No data was available for damper positions greater than 0.65 because the recirculation dampers in the ASHRAE 1020 research project were constrained by the building control system to never exceed that maximum value unless the air-handler was off. Furthermore, we could not find any other data sources that could match even the limited data that was made available by the ASHRAE 1020 research project.

With the exception of what look to be outliers at a recirculation damper position of 0.65, the outdoor air fraction appears to be a monotically decreasing function of recirculation damper position. In contrast to the simulation results that yield an outdoor air fraction equal to 1 for closed recirculation dampers, the outdoor air fraction measured in ASHRAE 1020's air-handler B at the closed damper position yields a dense spread of values between 1.0 and 1.2.

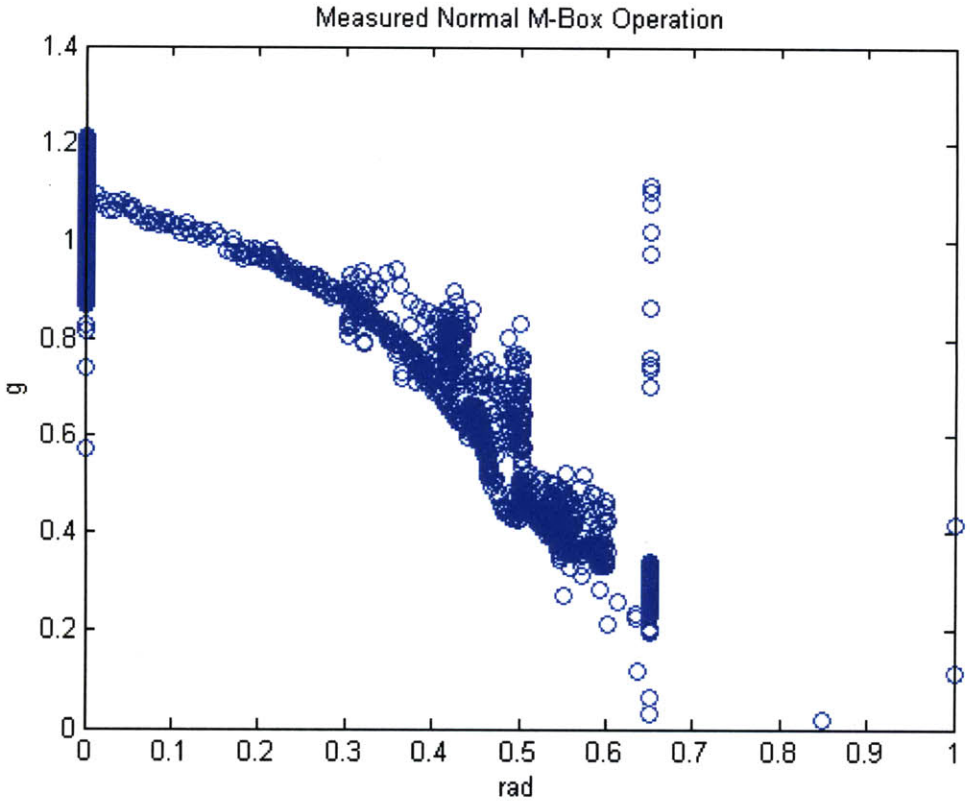

Figure 19 Measured fault-free operation of Air-Handler A at the IEC

Similar data was also acquired for air-handler A at the IEC; both air handlers have identical equipment, manufacturers, and control sequences. Furthermore, they serve identical halves of the test building that comprises the IEC. Air-handler A, like its twin, exhibits a spread of values for outdoor air fraction when the recirculation damper is fully closed, but with roughly twice the range as compared to measurements made on air-handler B.

Both air-handlers exhibit a monotically decreasing behavior that is nearly one-to-one between values of outdoor air fraction and recirculation damper position, for damper positions between 0.0 and 0.3, non-inclusive. For damper positions greater than and equal to 0.3, however, both air-handlers yield clusters of multiple values for outdoor air fraction per damper position. The most visibly extreme case of this multi-valued phenomenon appears in both air-handlers at a recirculation damper position of 0.65; at this damper position the outdoor air fraction appears to take on values across the entire range of plausible measurements, with a central tendency in the range of outdoor air fractions equal to 0.2 and 0.4.

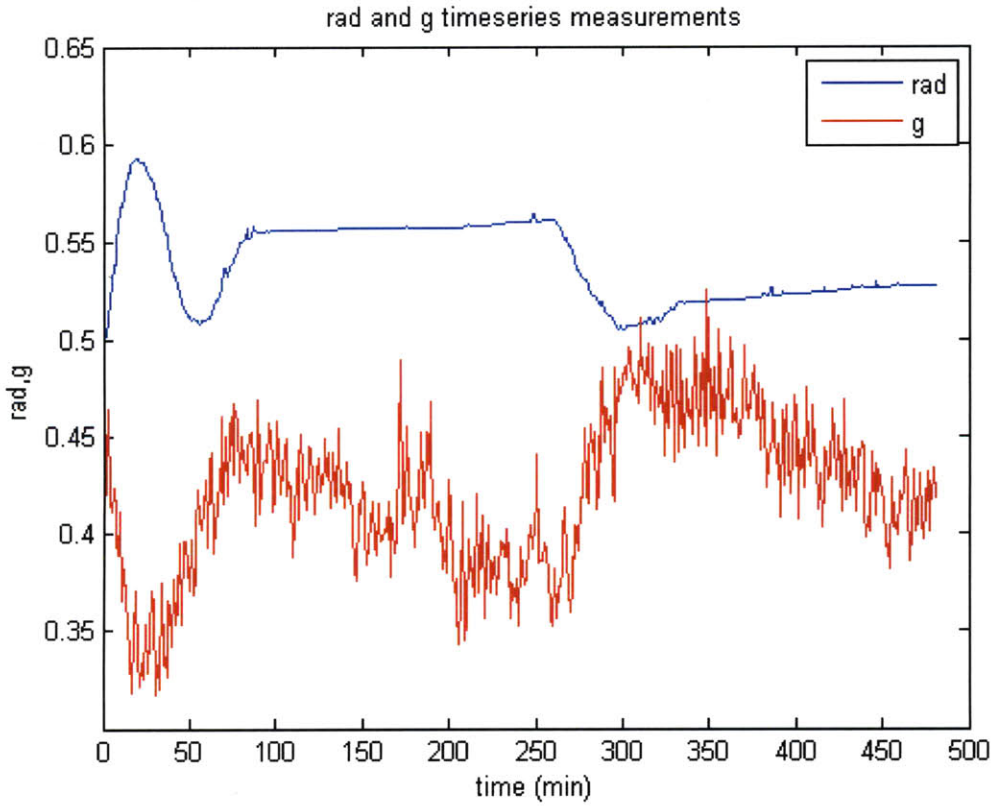

Figure 20 Time-series of rad and g measurements on a spring day within ASHRAE 1020-RP for air-handler B; the time series begins at 9 am

A day's worth of data sampled at 1-minute intervals from air-handler B is shown in Figure 20. The overlapped time-series of recirculation damper position and outdoor air fraction reveal the relative noise level of each signal, their cross-correlation, and the potential existence of any additional variables that may influence the outdoor air fraction. The maximum cross-correlation coefficient was found at 0 time-lag between the signals, indicating that the response time of the outdoor air fraction to changes in recirculation damper position is less than 1 minute. The absence of time-lag in this particular sample does not guarantee its absence for all operation of the mixing box; it remains unknown if the outdoor air fraction will exhibit a lag when the damper position changes more rapidly.

The data in Figure 20 was taken on a spring day when the recirculation air damper was modulating the outdoor air flow according to external environmental conditions. The damper control signal is visibly smoother than the outdoor air fraction, which is not surprising since the damper signal is a direct output from the building control system. The uncertainty of the outdoor air fraction was estimated by the mean standard deviation of several continuous 50-point samples drawn from visibly flat regions of the damper control signal. The mean standard deviation was found to be roughly 0.02, which is consistent with the 5% or better accuracy expected from the commercial air flow meters used to measure the supply and outdoor air flow rates of the air-handler.

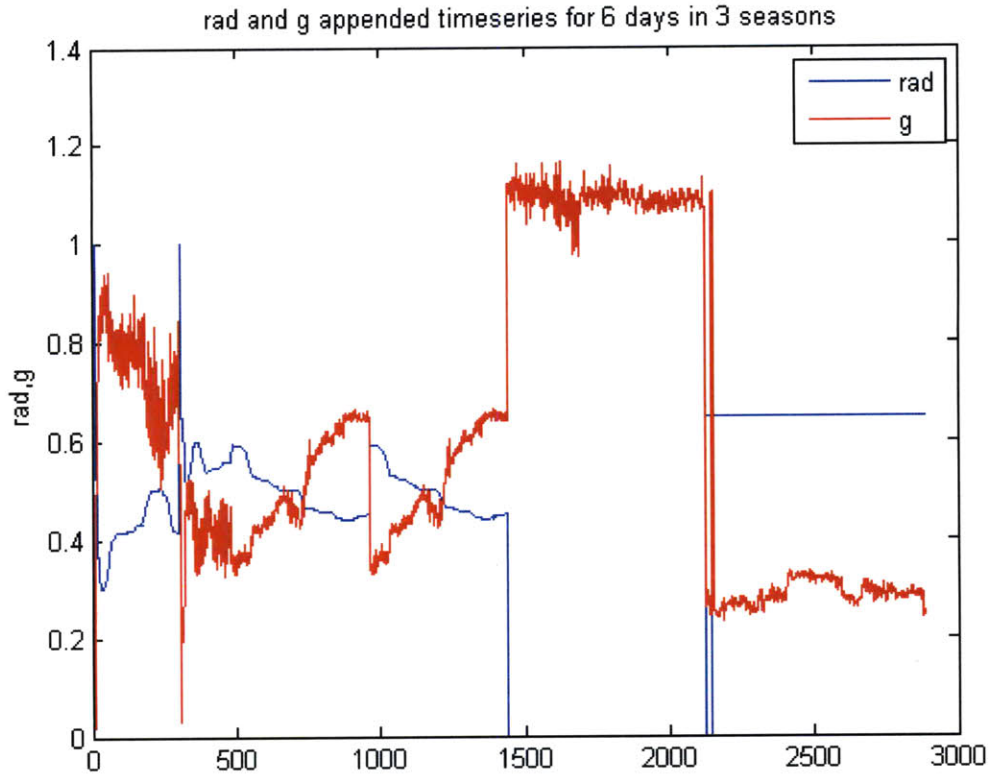

Figure 21 Time series measurements for rad and g, collected from air-handler A over 6 different days spanning summer, winter, and spring seasons

Several days of data for summer, spring, and winter seasons were sampled at 1-minute intervals from air-handler A and collected into the time-series shown in Figure 21. The winter season spans the time interval from the origin to 1450, the spring season spans the time interval 1451 to 2150, and the summer season includes the remainder of the series. The three seasonal segments of the time-series in figure 12 exhibit representative behavior for the corresponding seasonal weather conditions in Des Moines, Iowa; the outdoor air fraction modulates in the winter, is maximal in the spring, and is minimal in the summer. The uncertainty in outdoor air fraction for air-handler A was estimated from the multi-season data in Figure 21 in the same fashion as it was estimated from the data in Figure 20; the result was again roughly 0.02.

During periods when the recirculation damper position is held constant, or nearly constant, the outdoor air fraction measured for both air-handlers exhibits variations that do not appear attributable to measurement noise. The magnitudes of these variations can exceed three

times the measured or predicted measurement uncertainty of the outdoor air fraction, and persist over several minutes or more. The effect is most visible in Figure 21 for times greater than 1450. The outdoor air fraction also exhibits long-term variations with similar changes in magnitude when the damper position changes slowly over the range of 50 to 60% open; this effect is visible in Figure 20 within the time intervals of 75 to 250, and 350 to 500. These observations suggest that other variables may have a significant influence on the outdoor air fraction, especially when the recirculation damper position is held constant. The latter observations also re-enforce the non-linearity that is apparent in Figure 18 and Figure 19; the gain between damper position and outdoor air fraction appears to be greatest towards the center of the range of damper positions.

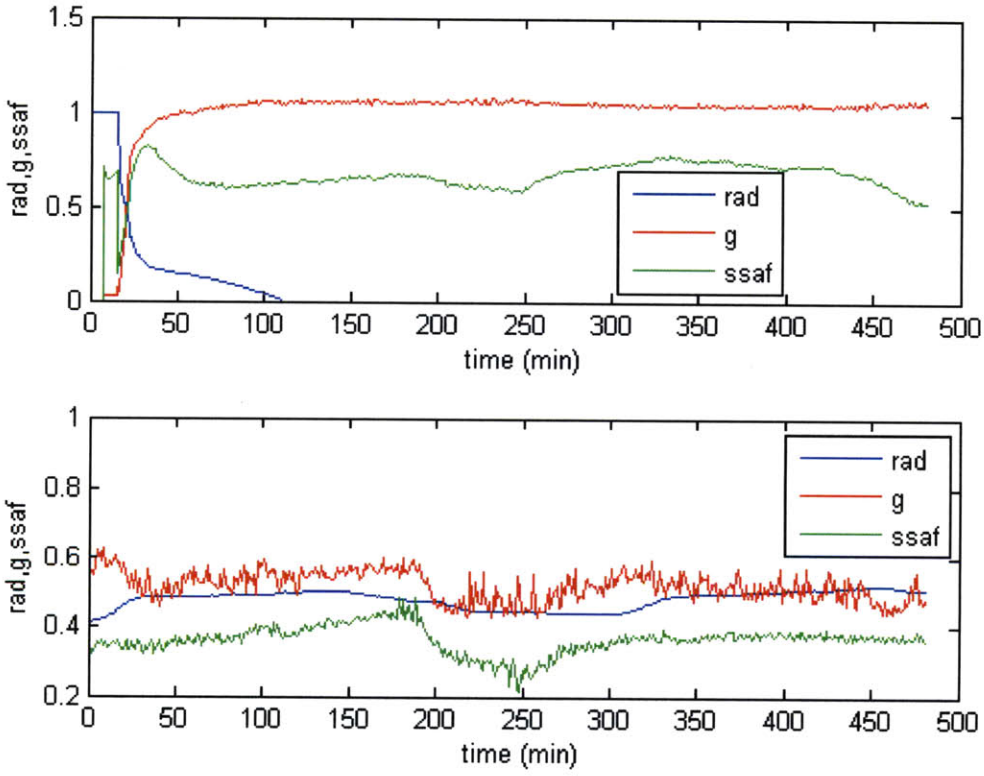

Figure 22 Time-series for rad, g, and supply air flow rate scaled by its maximum possible value (denoted by ssaf)

Figure 22 includes two time-series of recirculation damper position, outdoor air fraction, and supply air flow rate, where the supply air flow rate has been scaled by its maximum rated value of 3600 CFM; the scaled supply air flow rate is denoted by ssaf in the figure. The time-series are both sampled at 1-minute intervals for two separate days of operation during the spring

season. Literature from the ASHRAE 1020 research project suggests that the supply air flow rate can influence the outdoor air fraction, and that inter-relationship is visible for portions of the lower time-series in Figure 22. The effects of supply air flow rate on outdoor air fraction in that case are most evident in the time the interval 150 to 300, where the outdoor air fraction and supply air flow rate both appear to experience a simultaneous decline and then rise in value while the damper position remains nearly constant. On the other hand, the upper time-series in Figure 22 supports the independence of outdoor air fraction from supply air flow rate; in that case, the outdoor fraction remains nearly constant while the supply air flow rate varies over 25% of its range.

One potential explanation for this behavior is that the effects of supply air flow rate are greatest while the outdoor air fraction and damper position modulate in the middle of their ranges; as the outdoor air fraction reaches unity, or perhaps either extreme of its range, the supply air flow rate appears to exert little influence on the outdoor air fraction. This hypothesis stems from the observation that the outdoor fraction is at its maximum value in the upper time-series of figure 11, while it is in the middle of its range for the lower time-series in Figure 22. Figure 18 and Figure 19 support this saturation effect since the steepest gain in the non-linear relationship between outdoor air fraction and damper position occurs in the middle of their ranges, which is where we would expect and indeed observe additional variables to exert their influence on outdoor air fraction.

The exact relationship between supply air flow rate, outdoor air fraction and recirculation damper position remains unknown for air-handlers A and B, but the data still supports the original hypothesis that the damper position is the primary control variable for outdoor air fraction. The multiplicity of outdoor air fraction values shown in Figure 18 and Figure 19 for damper position equal 0.65 may be explained by the influence of supply air flow rate on the outdoor air fraction within that range of damper positions. The apparent multiplicity of outdoor air fractions for the closed damper position may stem from additional measurement uncertainty at that damper position, limitations of damper actuator repeatability for closed-damper position control signals, or other un-modeled physical effects such as flow reversal. In the latter case, the outdoor air fraction may exceed 1.0 because the outdoor air flow surges up through the

recirculation damper and exits via the exhaust dampers. While it is not intended to occur, flow reversal from the outdoor air duct and across the recirculation dampers is a common problem in air-handlers and may explain part of the data included here.

3.4.2 Fault-laden operation

In addition to data that is characteristic of normal mixing box operation, the ASHRAE 1020 research project also collected performance data on fault-laden mixing box operations. The prior simulation faults were in fact modeled after the leaking and stuck damper faults that were physically implemented in IEC air-handlers as part of ASHRAE 1020. Three different damper leakage levels were applied to the air-handlers at the IEC by removing in equal increments from the damper blades all of the layers of rubber sealant that are intended to yield an air-tight seal for the closed damper position. Likewise, stuck dampers were simulated on air-handlers at the IEC by disengaging the damper linkages and manually adjusting the dampers to either a fully closed or open position. The leaking and stuck damper faults were applied on several days in the winter, summer and spring seasons of the ASHRAE 1020 research project.

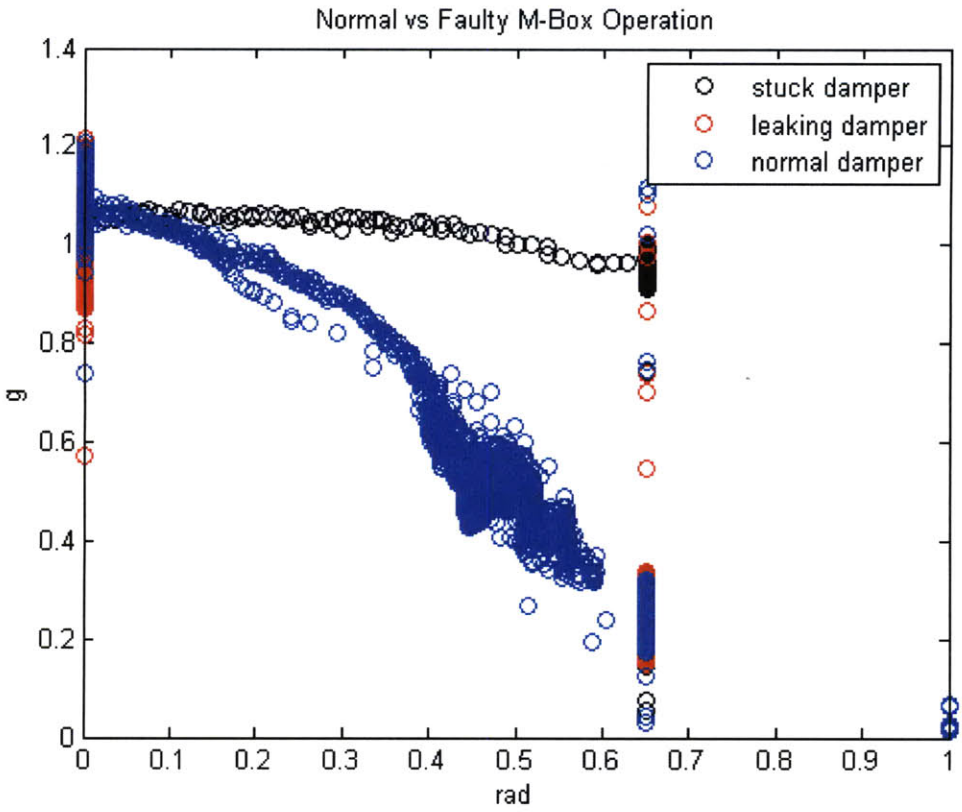

Figure 23 Measurements of faults in a real mixing box: stuck and leaking dampers

Figure 23 shows an overlay of normal mixing box data from air-handler B at the IEC with fault-laden mixing box data from the same air-handler, for stuck and leaking dampers that were tested over several winter, summer, and spring days of the ASHRAE 1020 research project. In this particular data set, the leaking faults were only applied on days where the weather was sufficient to keep the damper position held constant at either 0.65 or 0.00. Consequently, Figure 23 only shows leaking dampers as two vertical red line segments whose abscissae indicate their corresponding constant valued recirculation damper position. For the three leak levels tested at a damper position of 0.65, the red circles indicating data from the leak are barely visible at the tails of the central tendency of normal mixing box operation at that same damper position. The outdoor air fraction under leak conditions at the closed damper position appears to decline past normal mixing box operation by roughly four times the measured uncertainty of the fault-free signal. In contrast to the leaking damper data that was primarily collected at two damper positions, the stuck damper faults were imposed on days where the external weather caused the

recirculation damper to travel its full range. The results of that ideal weather condition are the black circles that horizontally traverse the graph in Figure 23, and are visually distinct from the fault free mixing box data. Comparisons between the stuck damper simulations shown in Figure 17 against the measured data about a stuck damper in Figure 23 suggest that the constitutive model used to simulate stuck dampers is in fact representative of the real fault.

While the effects of a stuck damper are immediately discernible in Figure 23 the leaking damper faults are less obvious to the casual observer. The fault detection system in ASHRAE 1020, when applied to this same fault data, provided fault identification and detection with about the same accuracy as we can deliver by visual inspection of the data shown in Figure 23. Small and medium leaks, especially at open damper positions were very difficult to isolate in ASHRAE 1020, while large leaks, especially at closed damper positions were more readily detected; stuck dampers of any nature were always detected and diagnosed. Leaking damper faults were also imposed on air-handler A at the IEC, but during the winter season where the damper position was modulated over a wider range than occurred during leaking damper tests in air-handler B.

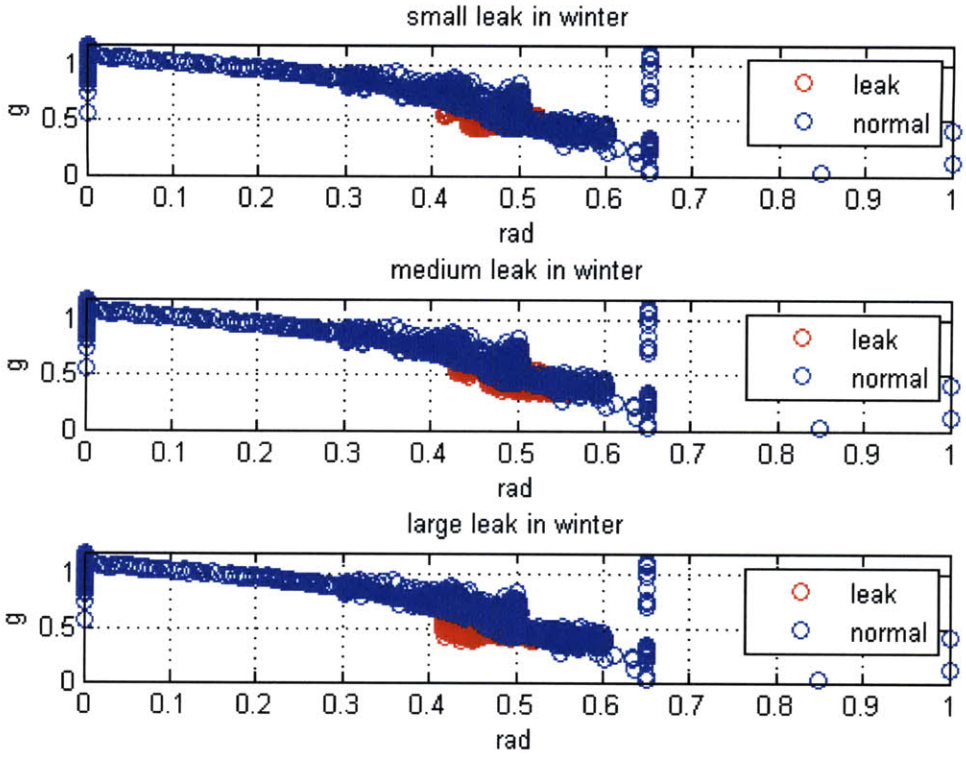

Figure 24 Data from three different leak levels imposed on air-handler A at IEC during the winter testing season of ASHRAE 1020

Figure 24 presents a juxtaposition of normal mixing box data from air-handler A against three different leakage levels imposed on the same air-handler. The data representing faulty operation is colored red in each graph of Figure 24, and is visibly discernible at each fault level from the blue data that represents the normal mixing box operation. In contrast to the leak data from air-handler B that is anchored on two damper positions, the leakage fault data from air-handler A varies over the damper position range of 0.4 to 0.55. The most visible effects of leaking dampers are observed in the bottom graph of Figure 24, corresponding to the data set from the largest leak imposed on air-handler A. The data from the large leak are seen to protrude from the bottom of the fault-free data set by roughly 4 times the measured uncertainty of the outdoor air fraction signal.

The leak data measured from air-handlers A and B does not cover a sufficient range of damper positions to support a complete comparison of the real operation of leaking dampers

against the constitutive models that were used to simulate the operation of leaking dampers. Comparing the simulation results in Figure 17 against the fault data shown in Figure 23and Figure 24 however, suggests that the constitutive model for simulating leaking dampers may approximate the measured data with Equation 3-14 for values of β between 0.1 and 0.3. In both the measured cases and simulations with β between 0.1 and 0.3, the magnitude of divergence from fault-free behavior is roughly 0.08, or 4 times the measured uncertainty of the outdoor air fraction signal.

3.4.3 Real energy impact of faults

On its own, the data available from ASHRAE 1020 is insufficient to provide a complete estimate of the impact of stuck and leaking damper faults on the annual energy consumption of an air-handler. This is because the ASHRAE 1020 data set does not include values for outdoor air fraction over the full range of damper faults and positions, or year round estimates for the ventilation rates and return air conditions experienced by the building. The annual energy consumption of the air-handler, and subsequent energy-derived cost of operations, must be estimated in order to motivate any sort of action on fault-laden operations.

We can extrapolate the impact of mixing box faults on annual energy consumption by assuming some reasonable values for year-round ventilation rates, building comfort conditions, and the relationship between outdoor air fraction and recirculation damper position. The complete enthalpy balance around the combined mixing box and hydronic coil sections of the air-handler schematic in Figure 11 reveals the year-round information that is necessary in order to estimate the true energy impact of damper faults

$$TE_{AHU} = \dot{m}_{OA}h_{OA} + \dot{m}_{RA}h_{RA} - \dot{m}_{SA}h_{SA}$$

Equation 3-21

TE_{AHU} in Equation 3-21 is the thermal power consumption of the air handler, measured in BTU/s; the new air enthalpy terms, h_{iA}, measured in BTU/lbs are defined by

$$h_{iA} = c_{A,p}T_{iA} + \chi(h_W + c_{W,p}T_{iA})$$

Equation 3-22

for each flow into and out of the air-handler.

Equation 3-22 for air enthalpy is measured in BTU/lbs wet air, and includes the variables $c_{A,p}$, the thermal capacitance of dry air measured in BTU/lbs-°F, $c_{W,p}$, the thermal capacitance of water measured in BTU/lbs-°F, h_W, the enthalpy of vaporization for water measured in BTU/lbs, and χ, the absolute humidity ratio measured in lbs H_2O/lbs dry air. The temperature dependencies of thermo-physical properties of air and water are often ignored by HVAC engineers allowing many of the terms above to be taken as constants measured at standard conditions: $c_{A,p}$ is 0.24 BTU/lbs-°F, $c_{W,p}$ is 0.444 BTU/lbs-°F, h_W and is 970 BTU/lbs.

The mass balance in Equation 3-1 can help eliminate the recirculation mass flow rate in Equation 3-21, and upon re-arranging terms yields the thermal energy consumption of the air handler in terms of the supply and outdoor air mass flow rates

$$TE_{AHU} = \dot{m}_{OA}(h_{OA} - h_{RC}) + \dot{m}_{SA}(h_{RC} - h_{SA})$$

Equation 3-23

In order to evaluate Equation 3-23, we must estimate the supply and outdoor air mass flow rates, and the enthalpy of the recirculation, supply, and outdoor air flows. A constitutive model for the mixing box, like Equation 3-17, can be used to compute the outdoor air flow rate as a function of the supply air flow rate, and hence support the evaluation of Equation 3-23. Such constitutive equations can also model fault-laden mixing box operation and their use with Equation 3-23 facilitates the comparison of energy consumption between fault-free and fault-laden air-handler operation.

The real-time supply air flow rates in the ASHRAE 1020 building under all weather and occupancy conditions are difficult to simulate over an entire year because the fans are modulated to maintain certain conditions within the building. It is not possible to predict the year-round conditions within the buildings, or the subsequent real-time supply air flow rates, without a detailed energy and ventilation model of the entire building. While it is possible to create such detailed building energy models, a simple alternative to ascertaining the order of magnitude of energy loss incurred by various damper faults is just to assume an average ventilation rate for the building during occupied hours.

In this case, the mean supply air flow rate for both air-handlers A and B across the entire ASHRAE 1020 data set was 1,700 CFM, or roughly 50% of their maximum rated flow rate, and the standard deviation was 376 CFM. The distribution of supply air flow rates for air handlers in ASHRAE 1020 is included in Figure 25.

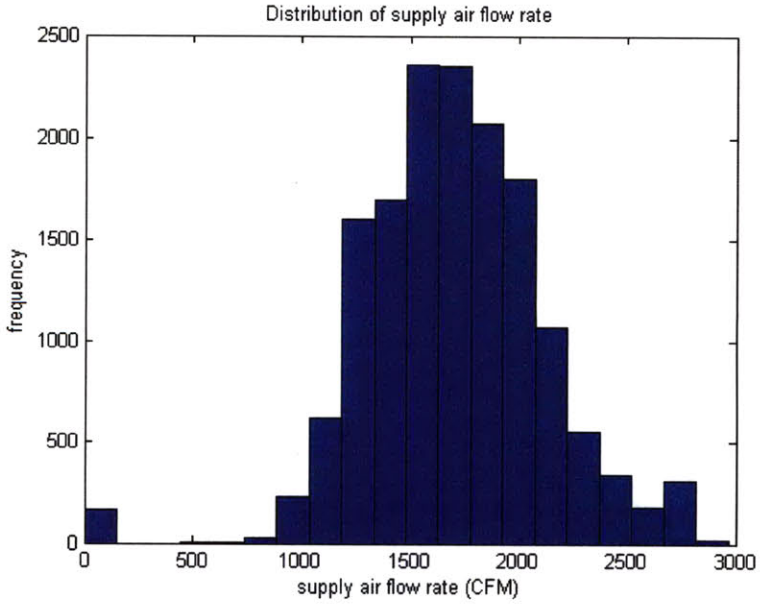

Figure 25 Distribution of supply air flow rates for AHSRAE 1020 air-handlers; the data is drawn from over 30 days of operating history, measured at 1 minute intervals during the regular working hours of 9 am to 5 pm.

Based on the distribution of supply air flow rates shown in Figure 25, we might create upper and lower bound estimates on the impact of various damper faults on air-handler energy

consumption by evaluating Equation 3-23 with constant supply air flow rates of 1,500 and 2,000 CFM.

The typical recirculation and supply air enthalpies can also be estimated from corresponding temperature and humidity data within the ASHRAE 1020 data set. Over the 30 days of data collected in ASHRAE 1020, the return air temperature distribution was broadly distributed over the range 70 °F to 74 °F with a mean of 72 °F and standard deviation of 1.5 °F. In contrast, the supply air temperature distribution had a sharp peak at 55 °F with a standard deviation of 2.5 °F. The discharge and return air humidity conditions followed a bi-modal distribution, with peaks determined by the humidity in cold and warm weather seasons.

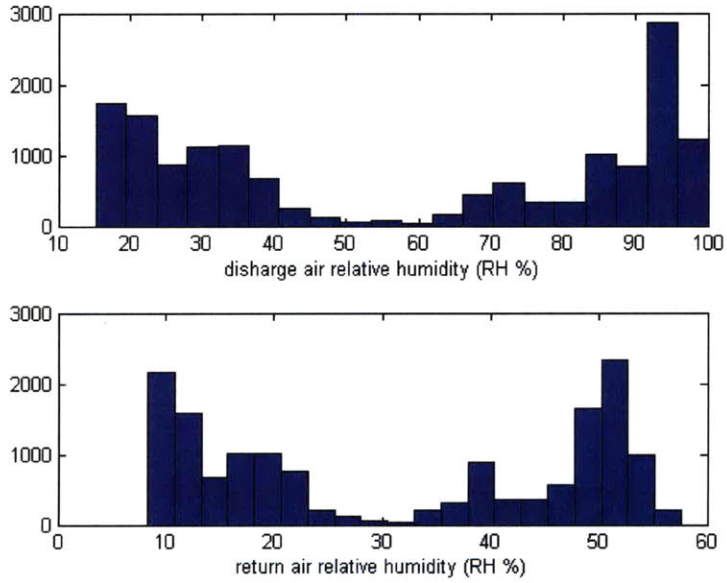

Figure 26 Distribution of humidity for return and supply air, across 30 days from the ASHAE 1020 data set

The distribution of return and supply air temperatures suggest that their respective mean values can be used to estimate the year round energy consumption of the air-handler. It is harder to assign a constant year-round value for the humidity of the recirculation and discharge air flows because those variables vary with season and building occupancy. Much like our approach to the supply air flow rate, however, we may create upper and lower bounds around the energy impact of various damper faults by assuming worst and best case values for the supply and return

air humidity. The most energy intensive operation of the air-handler results when return air arrives at the air-handler with the upper humidity value of the comfort envelope for the building, or roughly 75% relative humidity. This is a worst case scenario for air-handler energy consumption because the return air has its highest possible high water content. The best case scenario is when the return air arrives at the air-handler with a humidity that is closer to the bottom of the comfort envelope for the building, or roughly 40% relative humidity. These best and worst case values for the return air humidity do not match the two peaks that are observed in the lower histogram of Figure 26 for the distribution of return air humidity observed in ASHRAE 1020. This discrepancy is due to the lack of occupants within the ASHRAE 1020 building; human beings are a significant latent load within buildings, and their effect is noticeably lacking from the ASHRAE 1020 data set.

The supply air humidity for all simulation cases is a function of the combined water content from the outdoor and recirculation air flows, and the saturation moisture content of air at the supply air temperature. At 55 °F, which is the estimated year-round temperature for the supply air, the saturation concentration of water in air is 9.2×10^{-3} lbs H_2O/lbs dry air. If the combined water content of the recirculation and outdoor air flows exceeds that threshold, then water must condense from the supply air stream before the supply air temperature will reach 55 °F; in this case the supply air humidity will be 100%. If the combined water content of the recirculation and outdoor air flows is less than 9.2×10^{-3} lbs H_2O/lbs dry air, then no water needs to condense out of the supply air before the supply temperature reaches 55 °F; the supply air relative humidity will be a function of the water content of the recirculation and outdoor air flows.

The flow of outdoor air into the air-handler is controlled by the recirculation damper position and the supply air flow rate. The exact control rules for recirculation damper position in ASHRAE 1020 remain unknown; however the literature from ASHRAE 1020 does suggest that those rules were similar to the control rules presented previously in Equation 3-7 through Equation 3-9. Those previous control rules are shown again here for the convenience of the reader:

$$T_{OA} \geq T_{RA} \quad rad = 1$$

$$T_{SA} \leq T_{OA} < T_{RA} \quad rad = 0$$

$$T_{OA} < T_{SA} \quad rad = \left(1 - \frac{T_{RA} - T_{SA}}{T_{RA} - T_{OA}}\right) \frac{1}{1 - g_{min}}$$

The primary difference between these controls rules and the ASHRAE 1020 literature is that the control system in ASHRAE 1020 was programmed to not open the recirculation damper position beyond 65%; this feature is not represented in the above control laws, however it was included in our subsequent simulations of the ASHRAE 1020 air-handlers.

The outdoor air fraction for fault-free and fault-laden mixing box operation is given as a function of the recirculation damper position by Equation 3-14, Equation 3-15, and Equation 3-16 respectively. Subsequent simulations of this ASHRAE 1020 model used a value of β equal to 0.2 to model the worst case of leaking dampers, and a value of α equal to 0.1 to model dampers that are stuck open; these model parameters come from a visual inspection of the fault data shown in Figure 18 and Figure 19.

Simulations of the ASHRAE 1020 air-handler are intended to provide insight into the real energy costs incurred by leaking and stuck mixing box dampers. In order to compute the energy costs of heating and cooling air with the air-handler, the thermal power consumption computed in Equation 3-23 must be transformed into the primary energy consumption at the heating and cooling plant that serves the air-handler. This conversion is done by utilizing assumptions about the bulk efficiency of those heating and cooling plants; the heating plant at the IEC is a condensing boiler with fuel utilization efficiency that is likely to be approximately 92%, and the cooling plant is a 10 TON condensing chiller with a COP that is likely to be approximately 5. The condensing boiler at the IEC is fueled with natural gas and the condensing chiller is electrically powered; today's natural gas and electricity tariffs are roughly $1.50 per therm, and $0.13 per kW-hr, respectively.

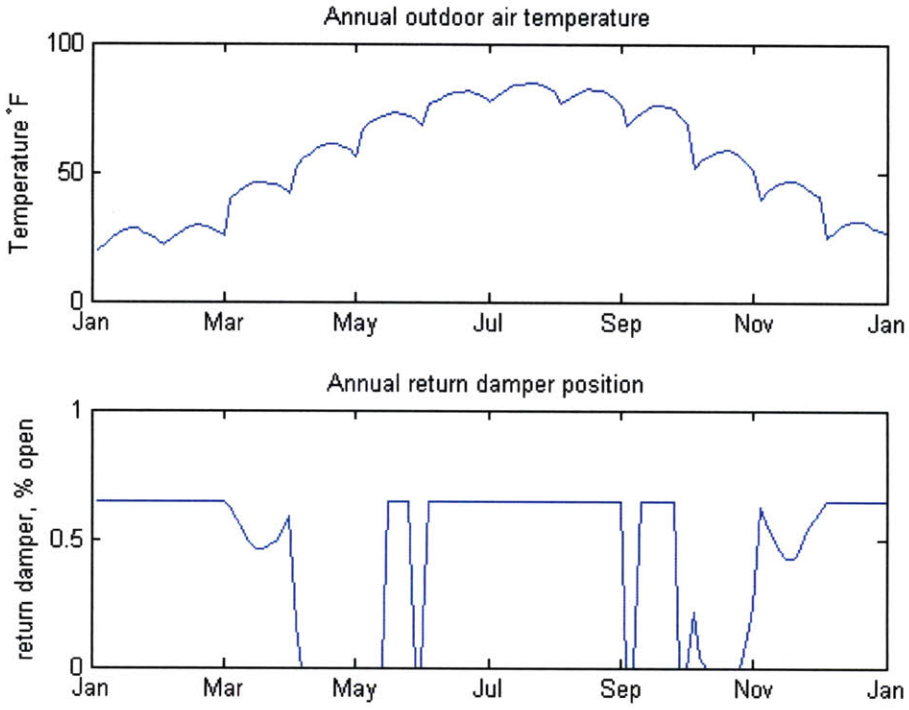

Figure 27 Simulation results for Des Moines, Iowa: annual outdoor air temperature and damper positions per month

Typical meteorological data for Des Moines, Iowa, is used to generate the graph of outdoor air temperature that is shown in Figure 27. The data is taken for 10 representative working hours in each month, from 8 am to 6 pm; the outdoor air temperature graph depicts average working day conditions for each month of the year. The corresponding control signal for the damper position is shown in the lower graph of Figure 27. The mean outdoor air temperature for Des Moines, Iowa during working hours of the IEC is 55 °F. The average summer time high for outdoor air temperature during working hours is 85 °F; the winter low is 19 °F. We have focused exclusively on data during the work week at the IEC because the air-handlers are otherwise normally off.

Fault-free, stuck, and leaking dampers were each simulated using the conditions defined above, and the outdoor air temperature and humidity data contained within the TMY file for Des Moines, Iowa. Figure 28 includes a scatter plot of outdoor air fraction versus damper position for all three fault states of the air-handler. The severities of leaking and stuck dampers included in

these simulations are representative of the same faults that were tested in the ASHRAE 1020 research project. The real fault data for stuck and leaking dampers is shown in Figure 23 and Figure 24; simulation models for those faults were chosen in order to generate data that appeared similar.

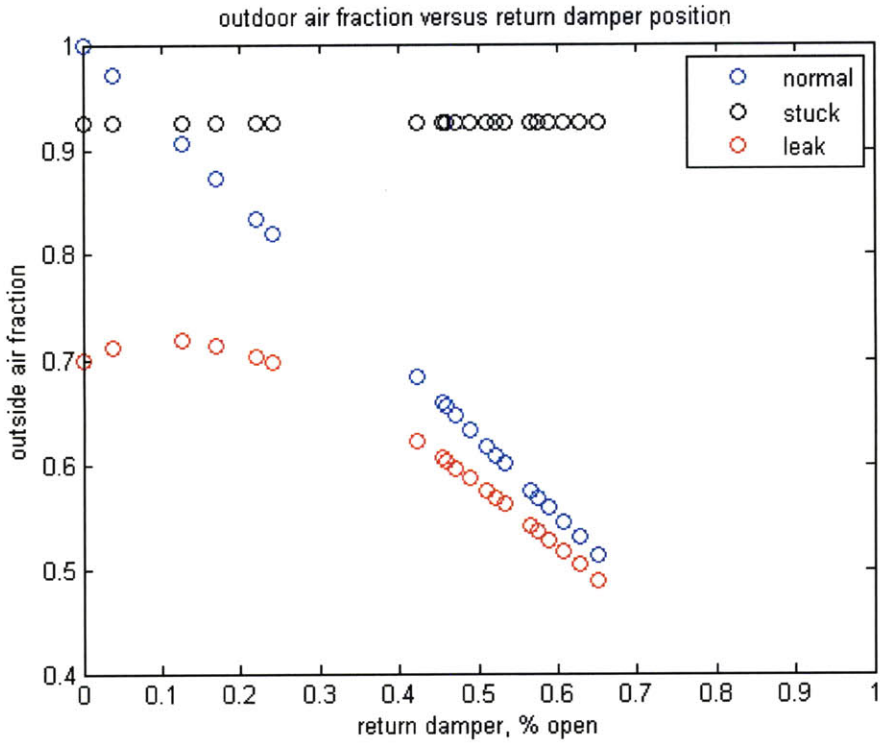

Figure 28 Simulation results for Des Moines, Iowa: scatter of outside air fraction versus return air damper position

The impact of these fault states on heating and cooling energy consumption are shown in Figure 29 as a pair of time series of monthly heating and cooling energy consumption. The energy consumption measured on the y-axis of the time series in Figure 29 reflects how much thermal energy is consumed at the air handler in order to provide desirable supply air conditions to the building.

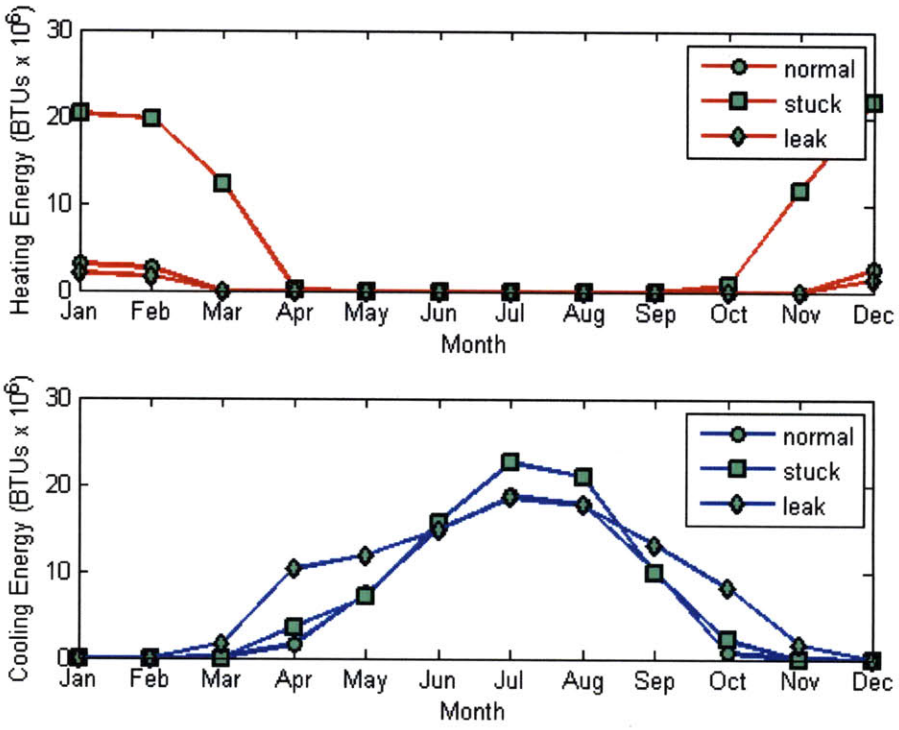

Figure 29 Simulation results for Des Moines, Iowa: monthly energy consumption per fault status

The heating energy consumption for a stuck damper is almost 10 times that of a normally operating damper; this is by far the largest increase in energy consumption amongst all cases of faulty operation simulated here. The stuck and leaking dampers increase the cooling energy consumption of the air handler by 10's of percent. Surprisingly, the leaking damper actually reduces the heating energy consumption of the air handler. The corresponding changes in energy cost per fault status are included in Table 1.

	Normal	Stuck		Leak	
Heating Cost	$136	$1,444	**962%**	$85	**-37.50%**
Cooling Cost	$544	$625	**15%**	$745	**36.95%**

Table 1Table of annual energy costs and percent difference from fault-free case, per fault status

The normal operation of the air handler incurs roughly $680 of annual thermal energy costs, with 20% of that for buying natural gas, and the remainder devoted to purchasing electricity. This does not include the cost of electricity for running the fans, pumps, or other plug loads, or the

thermal energy costs associated with the operation of the other two air handlers that are in the building. The stuck damper fault incurs the greatest waste of energy and increases the thermal energy expense of the air handler by more than a factor of three. The leaking damper has a smaller impact on the air-handler, raising the thermal energy expense by only 25%. In absolute terms, the stuck damper increases the thermal energy expense by $1,389, and the leaking damper increases thermal energy expenses by $150.

The nominal hourly rate for a mechanical service technician who could fix leaking and stuck dampers is around $100. Assuming that fixing a leaking or stuck damper requires the equivalent of three hours of labor in time and material costs, the return on investment for fixing a stuck or leaking damper is 3 months and 2 years, respectively. Even if the equivalent of 10 hours of service labor and materials is needed to fix stuck dampers, the simple pay-back period on that investment is still less than one year.

3.5 Conclusions

In this chapter we have achieved several objectives leading towards the automated identification and evaluation of air-handler faults: we have created simulation models in order to understand common damper faults that are found in air-handlers, we have compared the results of those simulations against similar fault-laden and fault-free data from two real air-handlers, and we have computed the real energy costs that are associated with stuck and leaking damper faults.

The data from simulations and real air-handlers suggest that simple faults such as stuck dampers pose a more significant threat to increase the energy consumption of an air handler than do leaking damper faults; clearly not all faults are of equal importance to either energy efficiency or occupant comfort. Review of the data available from real air-handlers shows that more data is needed on fault-laden systems over the full range of damper positions in order to quantitatively evaluate the accuracy of our simulation models. Likewise, a review of existing literature and conversations with experts in the field suggest that the research community in general does not

have a data library of fault-free and fault-laden behavior for commercial AHUs. A significant research contribution of deploying any sort of HVAC fault detection system in multiple buildings may simply be to aggregate a database of fault-free and fault-laden equipment behavior.

Estimates on the real energy costs associated with damper faults show that the dollar values of those faults can be much larger than the cost of service labor and materials needed to fix those faults. This result, in conjunction with the notion that not all faults are of equal importance, provides a basic argument for a FDD system that does not necessarily diagnose faults, but instead detects and prioritizes them in order of their financial significance. Such a fault detection and valuation system would allow system users to decide for themselves whether a fault is worth pursuing; some limited diagnostic features would also serve to help system users track the origins of a fault, instead of full diagnostics that attempt to completely specify the origins of faults.

For a comparatively small air-handler like air-handlers A and B at the IEC, the estimated annual financial loss due to stuck dampers is on the order of $ 1,000. With a simple physical assumption that AHU energy consumption scales linearly with volume, we can extrapolate the result from the IEC to larger pieces of equipment with similar control algorithms and that are located in climates similar to Des Moines, Iowa. For instance, Cambridge Ma has a climate that is similar to Des Moines, Iowa, and MIT has more than 10 VAV AHUs on its campus that exceed 30,000 CFM rated capacity; the annual financial loss for just one of those air handlers to have stuck dampers for a single year is on the order of tens of thousands of dollars.

4 Building Data and Data Uncertainty

4.1 Introduction

4.1.1 The Role of Data in Buildings Today

Data, and the possibilities for more data, abounds in buildings; ranging from records of repair and maintenance to extended logging of building control system measurements. Despite the existing and potential volume of building data, it is still not extensively archived, reviewed and converted by building managers into useful planning or execution information. Even national compendiums of construction data, such as RS Means (57) are only used as a rough starting point to estimate basic construction costs. The lack of data-driven action within buildings today is largely due to the time and effort that is required of building managers to organize, inspect and draw inference from data; they typically don't have the time or man-power to examine reams of data in order to find anomalies or wasted money. This situation is exacerbated by the fire-fighting culture of building management and services; pro-active planning and action is almost always deferred to re-active response to emergencies, complaints and discontent. Despite the contemporary proliferation of IT infrastructure, lower cost of data acquisition, databases and expert systems, our interviews with a wide variety of building stakeholders indicates that buildings are still managed and maintained much in the same way they were 20 years ago; problems come up, and if they raise enough attention they get solved.

4.1.2 Leveraging the Experience of Building Management

Data uncertainty is another key limiting agent to the use of building data for planning and execution. From a scientific perspective, engineers and scientists are comfortable with the ideas of precision and bias measurement error, but for a building manager there are many other forms of data uncertainty that are more difficult to quantify and overcome. Take for example the labor costs associated with providing mechanical, electrical or maintenance services; interviews with service managers at MIT and elsewhere clearly demonstrated to us that timesheets of labor hours don't necessarily reflect the number of hours actually worked by a technician, or in some cases they don't even reflect the true activities or location of the work.

The uncertainty of financial and execution records within an organization are something that only a member of that organization can fully appreciate and quantify. We found that in many cases, building managers don't often rely on building data because they know that it may be tainted with falsehoods or is not applicable to forecasting execution. This seems to be the underlying power of experience in buildings; cost estimates, scheduling and planning are based on the experience of the management team, and far less so on the written records of past performance. This unique nuance of building management and service suggests that FDD systems should not necessarily "solve" problems for building managers, but instead point out and prioritize decision options. The best FDD system may in fact be the one that leverages the way that buildings are actually maintained and operated; through the experience and knowledge of the actual building managers.

4.1.3 Making decisions with incomplete knowledge

Based on discussions and interviews with building stakeholders, we have found that uncertainty and risk are key considerations in the decision making process of building stakeholders. The review of prior art in chapter 2 of this thesis revealed that the majority of FDD techniques communicate their results in a deterministic fashion, and withhold from system users the inherent uncertainty of FDD results. In this chapter of the thesis we will develop techniques for managing uncertainty in FDD, and communicating the results of FDD in a framework that includes the uncertainty of inference. By embracing rather than shunning model and measurement uncertainty, we may create more robust methods of FDD that empower system users to ultimately judge for themselves whether a fault signal is important and worth investigating, based on the uncertainty of acknowledging a fault.

4.2 Collecting data from buildings

4.2.1 Building control systems for data collection

This thesis explicitly focuses on the use of physical measurements that are made about a building in order to support inference on its possible HVAC pathologies. In today's buildings, control systems are widely used to measure physical variables such as temperature, pressure and flow in order to control actuators and other devices. Because of that existing infra-structure, the most convenient approach to collect data for FDD seems to be through the data collection capabilities of installed building control systems. A survey of the major building control platforms shows that they all offer features for the extended logging and output of measurement data to third party applications. Some control systems, like that offered by Schneider Electric, are built on a standard SQL database which can facilitate the continuous electronic exchange of data between the control system and a third party FDD software. Other popular control systems do not have as well developed infrastructure for machine-to-machine (m2m) data exchange, and users can only access measurement data in bulk through comma delimited files or spread sheets.

The ubiquitous adoption amongst the building controls industry of standard communication protocols such as BACnet and LONtalk offers an additional avenue to data extraction from control systems; using a common protocol, FDD software may directly query measurement data from individual building controllers. This approach to data collection overcomes any systemic limitations that a control system may have towards database-level electronic exchange of measurement data however it also introduces other potential drawbacks. Interviews with control system experts, and research into the appropriate literature (58), suggests that open protocols such as BACnet and LONtalk are not supported nearly as well as we might desire by control companies. The use of open communication protocols for direct queries on building controllers has often resulted in situations where the building controller's response time, or scanning rate, is made sluggish. While open communication protocols appear to facilitate transparency of data exchange between different products, their level of support by building control companies does not seem sufficient to enable facile deployment of FDD across the building stock (58).

4.2.2 Cost of Data Collection

While building control systems are a ready vehicle for collecting FDD data, their existence within a building does not guarantee that sufficient measurements are being made to support any level of FDD. Interviews with experts and review of pertinent literature indicates that the rich instrumentation of buildings is not the prerogative of building control installations; instead the opposite is true where just enough measurements are made to support the lowest cost, yet still functional control architecture. It is likely that FDD systems in buildings will require additional installation of sensors beyond what is normally found in a low-cost controls installation, depending on the desired depth and breadth of FDD.

The cost of installing instrumentation in buildings is a significant hurdle to the scalable deployment of FDD throughout the building stock, as well as the retrofit of building control systems that could significantly impact building energy efficiency. Interviews with building control experts, as well as a review of pertinent literature (57) indicates that the total cost of a building control system, furnished with front-end graphics and deployed over the entire HVAC system, is amortized at roughly $1,000 per measurement or control point. While the price range for typical HVAC sensors like duct-immersion temperature or humidity probes is on the order of $10 to $100[2], the inclusive cost of installation, start-up, engineering and commissioning escalates the per-point cost of building measurement to roughly $1,000. Of that $1,000 point cost, roughly 25% is the cost of the sensor and related control hardware, 10% may be attributed to profit, and the remaining 65% is for the labor associated with installation and setup[3].

Stand alone sensor overlays for the purpose of FDD can be cheaper than $1,000 per point, but the lower point cost typically comes with a sacrifice of control capability; stand-alone sensor overlays do not necessarily contribute to the control of a building. Likewise, the labor cost associated with installing certain useful sensors will always be much greater than the cost of the sensor itself. Take for example the measurement of outdoor air flow in an air-handler, or the measurement of chilled water temperature for a chilled water supply main that is equal to or

[2] Price checks at www.kele.com, a popular supplier of HVAC sensor products
[3] Personal communications with control contractors from Schneider Electric, Carrier Corporation and Evco Mechanical, Inc.

greater than 3" in diameter; in the latter case a temperature measurement well must be installed on the chilled water supply main, and in the former case an air flow station must be installed in the air-handler's ductwork. Both of these activities could require shutting down of systems, draining of pipes, welding, cutting, piping, and other expensive labor activities.

The costs and complexities of collecting building data extend far beyond the cost of acquiring sensors and data acquisition hardware; instrumentation in buildings is a systems engineering problem whose planning and execution underpin the success of FDD.

4.2.3 What to Measure

From one perspective we may consider deploying dense sensor networks in order to create maps of energy transmission and consumption in a building; the usefulness of that approach has not yet been born out in practice or research. A more pragmatic approach to building measurement is to start with the un-answered questions and interests of the building managers, and work backwards from that desirable information to determine what additional measurements are needed in the building. Since our research has focused on FDD in AHUs, we asked building maintenance personnel at MIT what they would like to know about their AHUs. The top most interest of MIT personnel in the operations of their AHUs is the quantity of electrical and thermal energy consumed by that equipment. Complimentary to their inquiries about AHU energy consumption, MIT personnel also wanted to know if their AHUs could potentially consume less energy without sacrificing occupant comfort or safety.

In order to satisfy the interests of MIT personnel, we created the following list of desirable, standard AHU instrumentation:

- Electrical sub-metering of supply and return fans
    - Preferably full electric power measurement; this is often integrated in contemporary variable-frequency drive products that are used for VAV AHUs
- Thermal sub-metering of heating and cooling coils

    - Air-side measurements are typically cheaper than other options, and require air flow measurements across the heat exchanger, as well as air-inlet and exhaust temperature measurements.
    - Humidity measurements across the heat exchanger are also important for cooling applications.
    - Measurements of valve stem positions on the heat exchangers are also useful for detecting leaking valves
- Mixing box characterization
    - Air-flow measurement of the outdoor air intake, as well as the supply air discharge
    - Temperature and humidity measurements on the outdoor and return air flow streams
    - Temperature and humidity of the mixed air conditions

Electrical and thermal sub-metering of AHUs has straightforward uses in educating building management about the energy consumption of their air-handling equipment. Additional measurements of the valve stem position on AHU heat exchangers also helps to identify when valves are leaking. Finally, the measurement of outdoor air fraction, especially for AHUs that can operate in an economizing mode, is very useful for identifying prevalent and costly energy in-efficiencies like a broken economizer cycle.

4.2.4 Instrumentation Options

Temperature and humidity measurements in occupied building spaces are commonplace and generally easy to implement. In an AHU, however instrumentation design and implementation is made more complex by the need to make accurate measurements over large airflows and energy intensive processes. Mixed air temperature measurement in a mixing box, for example, is classically a difficult measurement to make accurately because mixing boxes are prone to poor mixing of air that is derived from multiple sources and at various temperatures. Likewise, air flow measurements, especially at the outdoor air intake manifold are difficult to implement because of the low face velocity of the outdoor air intake flow rate and typical lack of space for ductwork that supports an accurate measurement of air flow. The accurate measurement of bulk air flow typically requires a straight run of ductwork that is multiples in

length of the duct's hydraulic diameter, and sufficiently small cross sectional area to yield a face velocity of over 100 FPM (59).

In the case of outdoor air flow, ASHRAE research project 980 focused on alternative means by which the outdoor air fraction could be measured for a VAV AHU, other than directly measuring the outdoor and supply air flow rates. ASHRAE RP 980 explored the use of CO2 species and enthalpy balances to infer outdoor air fraction. These alternative methods utilized extra CO2 sensors and temperature sensors, respectively in order to obviate the need for an outdoor air flow station. The use of enthalpy balance to measure outdoor air fraction required higher accuracy measurements of mixed air temperature, which ultimately rendered the approach less useful because of the difficulties associated with making such measurements in practice. In contrast, the CO2 species balance was found to be more practical, lower cost and accurate than the enthalpy balance method, however only during periods where the building was sufficiently occupied that the return air CO2 concentration differed from the outdoor CO2 concentration by about 100 PPM.

4.2.5 Sensor Measurement Uncertainty

The uncertainties of HVAC sensor measurements are usually included as the rated accuracy for a manufacturer's sensor product. Sensor accuracies are typically expressed as a relative percent error of the sensor's measurement, or measurement range. For example, a typical HVAC temperature probe may have an operating range of -30 °F to 230 °F, and a rated accuracy of 0.2% over that range; the manufacturer may also quote this accuracy as +/- 0.5 °F. Alternatively, a manufacturer may rate the accuracy of their sensors simply as a given percentage, for example 0.2%, without specifying that this relative percent error is effective over a range of values or for each individual measurement. In the latter case, the sensor accuracy is improved at the higher extreme of the measurement range, and it asymptotically approaches the accuracy that is associated with a relative percent error over the entire measurement range.

Temperature	+/- 0.3 °F for probes, +/- 3 °F for averaging sensors
Humidity	+/- 2 to 5% of measurement value
Air-flow	+/- 2 to 5% of measurement value
CO_2 Concentration	+/- 2 to 5% of measurement value
Pressure	+/- 1 to 3% of measurement value

Table 2 Measurement accuracies for typical HVAC instrumentation[4]

We can use the findings in table 1 to estimate the impact of measurement uncertainty on the calculation of energy consumption by an AHU. For example, consider an AHU with 20,000 CFM of air flow, and a steam coil that incurs a 30 °F rise for the air flowing over the coil. The enthalpy rate of the steam coil is calculated in mBTU/hr (10^6 BTU per hour) according to (60)

$$H_s = \frac{1.1 \times CFM \times \Delta T}{10^6}$$

Equation 4-1

and is found to be 0.66 mBTU/hr. Using the accuracies in table 1, we can also compute an upper and lower limit estimate of the steam coil's enthalpy rate as 0.76 mBTU/hr and 0.56 mBTU/hr, respectively. Combing all of these results, we can express the energy calculation with its associated relative percent uncertainty as 0.66 mBTU/hr +/- 15%. If the AHU's steam coil were to operate in this capacity for an entire work week (5 days, 24 hours each), and we assumed that the cost of heating was equivalent to the cost of natural gas, roughly $1.50 per therm (one therm is equal to 100,000 BTU), then the thermal energy cost of the AHU for that time period would be $1,188 +/- $178.

If we interpret HVAC sensor measurement uncertainty as the measure of a statistical distribution, then the above results can form the basis for a probabilistic, rather than a deterministic inference about the cost of AHU operation. The probabilistic interpretation is preferable because it accommodates for the inherent uncertainty of HVAC measurements, and facilitates statistical testing of FDD hypotheses. A deterministic approach to evaluating HVAC operation wholly ignores the fundamental fact that measurements are never perfect, and in doing

[4] Values were taken from www.Kele.com for popular HVAC instrumentation products

so presents building managers with an incomplete picture of how their systems could be operating.

4.3 Propagation of Uncertainty in HVAC analysis

Uncertainty in HVAC analysis can arise from a variety of sources, including the natural precision error of measurement, bias of measurements, and uncertainty of models. In the following sections we will explore a method for propagating measurement uncertainty through HVAC FDD analysis, and expands its capability to include other forms of uncertainty or confidence that are important to the final FDD inference.

4.4 Modes of uncertainty

4.4.1 Random Measurement Error

Random measurement error is a familiar form of uncertainty amongst scientists and engineers; it underpins the accuracy and precision ratings of sensors. HVAC sensor accuracies are tabulated by their manufacturers in order to help HVAC engineers select a sensor that is appropriate for its intended application. High quality HVAC sensors can be purchased with accuracy certifications that are traceable to NIST (The National Institute of Standards and Technology), but the vast majority of products are subject to less stringent certification standards, if they are certified according to any standard at all. In many cases manufacturers will publish accuracies for their sensors without including any description of how that accuracy was ascertained or made credible.

4.4.2 Systematic Measurement Error

Measurement bias, or systematic error, is a form of uncertainty that does not contribute to the scatter of measurement data. Bias error is typically fixed and attributable to finite limitations of measurement (such as the number of bits in an analog-to-digital converter), but in other cases it can change over long periods of time resulting in measurement drift. Commercial HVAC sensors typically come with an estimate of drift magnitude versus time in order to inform HVAC engineers of the intended lifetime of the sensor product. Humidity sensors are notorious within the HVAC industry for drifting measurements; the growing bias error is often attributed to fouling of the capacitative elements within the sensor that are used to measure the water content of air. Unlike random error that can be easily quantified by taking numerous samples in a static system, bias error is almost always unknown unless measurements can be compared between a calibrated sensor and one under scrutiny. Sensor calibration in the building commissioning process typically consists of that comparative process where measurements from installed HVAC sensors are compared against similar measurements drawn from temporarily installed, laboratory-calibrated instruments.

4.4.2.1 Model uncertainty

In most engineering analysis of HVAC systems, the fundamental principles of physics are used to describe, size, and analyze the operation of an HVAC system. Inherently, engineers do not perceive "uncertainty" in their application of fundamental thermal-fluid analysis to HVAC systems. On the other hand, the classic technique of "over-sizing" equipment during the specification of HVAC systems is a broad method by which HVAC engineers embrace the uncertain accuracy of their calculations. Since HVAC systems are designed based on a model of the building, whether that be a detailed energy model, load calculations or design-day estimates, HVAC engineers respond to the unknown factors of the model by increasing the load capacity of their designs.

In FDD, data-driven and physical modeling techniques must both consistently accommodate for measurement uncertainty because such models are used to continuously predict the operation of a system. FDD measurement uncertainty is a mainstream component of black and grey box modeling techniques because statistics provides a straightforward toolset for

managing empirical uncertainty. In that same vein, measurement uncertainty is also a common consideration for physical models, again because statistics has a broad tool set for managing empirical errors.

Model uncertainty of the physical, grey or black box model itself, as in a mathematical description of the model's shortcomings, is a less common consideration in FDD models of systems. The inaccuracies of physical models in particular are not often directly considered because the underlying assumption of a physical model is that its description of a system is sufficient for the desired FDD inference. In practice, however, we found in the literature review of chapter 2 that physical models of HVAC equipment very rarely capture the entire spectrum of possible physics that can influence FDD inference; experience has shown that inaccurate models result in extensive filtering and handling of extraneous fault signals. The typical fashion for handling such model uncertainty is simply to state the limits of the models; for example, a model may only apply for certain temperature ranges or hours of the day.

A more implicit method of handling model uncertainty may facilitate the comparison of identical engineering results that stem from multiple physical models. For example, the energy consumption of a boiler that is fired by natural gas may be measured or predicted in a variety of ways:

- Sub-metering of natural gas flow
- Sub-metering of thermal BTU output
- Count of degree days or hours (reference measurements of indoor and outdoor temperatures)
- Status measurement on the natural gas valve (61)
- Acoustic signature of combustion (62)
- CO_2 and CO measurements on the boiler exhaust

All of these physical measurements can be converted by some model into an estimate of the energy that is consumed by the boiler over a period of time. Intuitively, as engineers we may identify that sub-metering of natural gas flow provides the most direct estimate of energy consumption (or alternatively, utility cost), but in many cases the cost of implementing that measurement is prohibitively high. And even if we did measure the flow of natural gas into the

boiler, there are several possible models to convert that gas flow into energy consumption; the simplest model is to assume that all of the natural gas is converted in thermal output. Clearly that simple model is flawed, however, because the heat exchange process is known to be imperfect within the boiler and there could be small leaks of fuel out of the boiler.

Alternatively, lower cost measurements such as acoustics on the combustion chamber or status on the valve are preferable from a financial perspective to the installation of a gas flow meter. But the models for converting those lower cost measurements into energy estimates may be more even inaccurate than the models that are used with gas flow measurements. The increased model inaccuracy is partially due to the less direct method of measuring the flow of energy; acoustics and valve status measure properties that are related to energy consumption, where as gas flow measures the flow of chemical energy itself.

This thought experiment exposes the need for an analytical FDD framework that includes model uncertainty and permits the comparison of engineering results between different models that use different measurements. For example, we might be comfortable using a low cost measurement such as valve status to yield a rough estimate of the boiler's energy consumption, instead of an expensive measurement with a more accurate estimate, if the more rough result came with an estimate of its model uncertainty.

A fundamental hypothesis of this thesis is that characterization of both model and measurement uncertainty in FDD is what can enable probability to become the common language of comparison between different building systems. Ultimately, such a probabilistic framework would support our desired probabilistic inference, and yield a FDD system that may be agnostic to what sensors or models are actually used to describe the building. Bayesian networks are a common approach to handling model and parameter uncertainties within a unified framework (63), but in the ensuing discussion we will discuss a simpler approach that follows from the engineering application of experimental uncertainty.

4.4.3 Methods for accommodating uncertainty in analysis

The scalable deployment of FDD requires consideration of the costs and practicalities of deploying sensors in buildings and drawing inference about complex and uncertain systems. In the real world, the best models and measurements will not always be available for use in every situation, but characterizing model and measurement uncertainty may provide a basis for using whatever model or measurement is most practical for a given situation.

4.4.3.1 Mathematical treatment of experimental uncertainties

Our treatment of measurement uncertainty follows directly from the classic exposition of Coleman and Steele on the subject of engineering application of experimental uncertainty (64). In their publication, Coleman and Steele review the underlying assumptions, derivations and extensions of the International Standard Organization's (ISO) "Guide to the Expression of Uncertainty in Measurement", otherwise known as ISO/IEC Guide 98:1995. ISO Guide 98, at the time of its publication, served as the basis for standards bodies such as ANSI/ASME and AIAA to review and update their own procedures for characterizing and communicating experimental uncertainties. Today, ISO Guide 98 is still a significant standard that is used across numerous industries to communicate and characterize experimental uncertainty.

Measurement uncertainties in ISO Guide 98 and in Coleman and Steele are broken down into two basic modes: bias and precision error, or, alternatively systematic error and data scatter. We shall use the nomenclature found in Coleman and Steele and refer to bias error as β_k and precision error as $\epsilon_k$where the subscript k enumerates the sequential index of a measurement on the variable X. In nearly all experiments or measurement processes, the measured values of different variables are combined using a data reduction equation (DRE) to yield some useful information; in the case of FDD, the DRE may be an energy function that outputs energy consumption and the measured variables (for an AHU) may be air flow and temperature change.

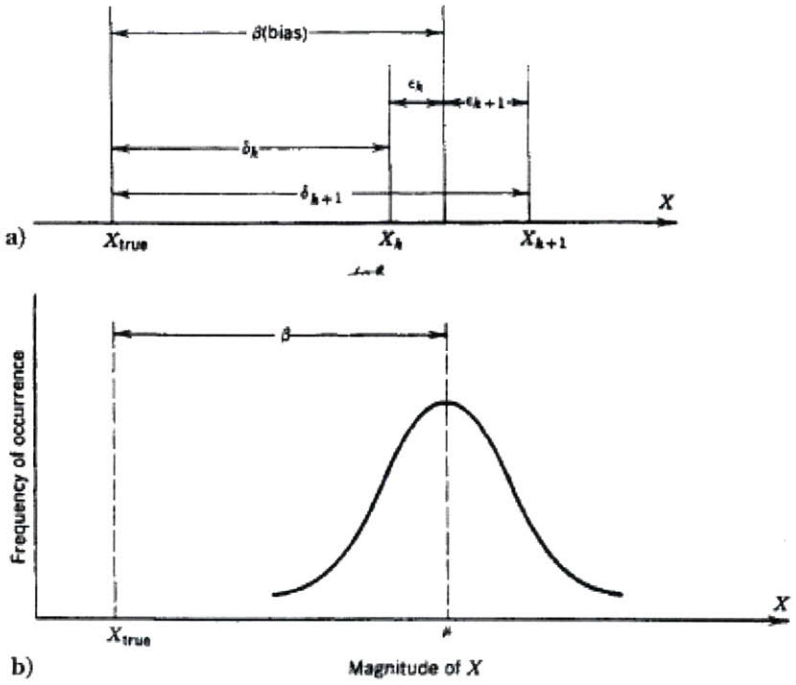

Figure 30 Modes of uncertainty for measuring a variable, X: a) two discrete readings b) infinite number of readings (adapted from Coleman and Steele, 1995). In both cases, the bias error, β_k represents a systematic difference between the averaged measured value of X and its true value. The precisions error, ϵ_k is a measure of the data scatter between the mean measurement of X, and its kth sample.

The interaction between the DRE and its parent measurements is presented mathematically as

$$r = f(x, y)$$

Equation 4-2

for two measurement variables, x and y, and the DRE output, r. The kth result of the DRE may be expanded in a Taylor series around its true value, $r_{true}$to show the mathematical interplay between uncertain parent measurements, x_k and y_k, their true values, x_{true} and y_{true} and the output value uncertainty, r_k

$$r_k = r_{true} + \frac{\partial r}{\partial x}(x_k - x_{true}) + \frac{\partial r}{\partial y}(y_k - y_{true}) + R_2$$

Equation 4-3

where R_2 is the expansion remainder that includes higher partial derivatives of the parent measurement variables.

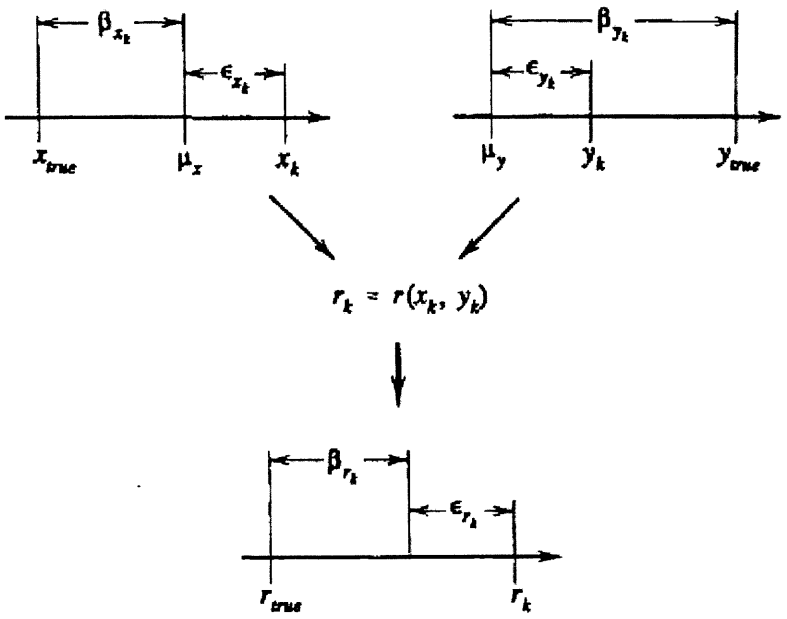

Figure 31 Propagation of measurement bias and precision errors into experimental results

The higher order terms of the Taylor expansion may be neglected if we assume that the measurement errors themselves are small or if the DRE itself is linear. In the former case, we expect the higher order terms to approach zero faster than the first and zero order terms, and hence the residual term, R_2 is negligible in comparison to the other terms. Following Coleman and Steele, we will assume that the measurement errors are sufficiently small to neglect higher order terms in the Taylor expansion.

As more measurements are made on the system, the output from the DRE will naturally embody the bias and precision errors that exists in its parent measurements; the effects of propagation are shown schematically in figure 2. Rewriting the Taylor expansion using the total error δ for the k^{th} determination of the result r, and using the notation $\theta_x = {}^{\partial r}/_{\partial x}$, the DRE error can be written as

$$\delta_{rk} = \theta_x(\beta_{xk} + \epsilon_{xk}) + \theta_y(\beta_{yk} + \epsilon_{yk})$$

Equation 4-4

The variance of the total error δ for some large number N of results r is defined by

$$\sigma_{\delta_r}^2 = \lim_{N\to\infty} \frac{1}{N} \sum_{k=1}^{N} \delta_{rk}^2$$

Equation 4-5

Substituting the definition for total error variance into the DRE Taylor expansion, as well as similar definitions of variance for the bias and precision errors of the individual measurements and taking the limit as N approaches infinite yield

$$\sigma_{\delta_r}^2 = \theta_x^2\sigma_{\beta_x}^2 + \theta_y^2\sigma_{\beta_y}^2 + 2\theta_x\theta_y\sigma_{\beta_y}\sigma_{\beta_x} + \theta_x^2\sigma_{\epsilon_x}^2 + \theta_{\epsilon_y}^2\sigma_y^2 + 2\theta_x\theta_y\sigma_{\epsilon_y}\sigma_{\epsilon_x}$$

Equation 4-6

In reality the population variances are never known and they must be estimated. Defining u_c^2 as an estimate of the distribution of total errors in the result, b^2as the estimate of the variance of a bias error distribution and S^2as an estimate of the variance of a precision error distribution, we can write

$$u_c^2 = \theta_x^2 b_x^2 + \theta_y^2 b_y^2 + 2\theta_x\theta_y b_{xy} + \theta_x^2 S_x^2 + \theta_y^2 S_y^2 + 2\theta_x\theta_y S_{xy}$$

Equation 4-7

where the variance terms with mixed subscripts indicate covariant terms between measurement variables.

The derivation so far is distribution free and seeks to reveal how parent measurement uncertainties are propagated into analytical results. The choice of distribution is forced by the desire to obtain an uncertainty U_r from the result variance u_c with some confidence level. The transition from variance to result uncertainty is made in the ISO guide as well as in Coleman and Steele by the selection of a coverage factor K, which supports

$$U_r = Ku_c$$

Equation 4-8

For normally distributed errors, which is typically assumed via arguments based on the Central Limit Theorem, the coverage factor K is equivalent to the value of the t statistic for a given number of degrees of freedom and level of confidence. The appropriate degrees of freedom for a system of measurements can be computed either from the Welch-Satterthwaite formula (65) or according to large sample and engineering approximations that are enumerated in Coleman and Steele and in their cited literature. In either case, the value of the coverage factor for normally distributed variances approaches 2.0 for increasing degrees of freedom and a confidence interval of 95%. Coleman and Steele show that degrees of freedom equal to or greater than 10 are sufficient to allow the use of a coverage factor of 2.0. Furthermore, Coleman and Steele also show that for DREs composed of multiple variables, the product of the degrees of freedom for each individual variable can be substituted as the total degrees of freedom for the estimate of the coverage factor. Because of these engineering and large sample arguments, the ISO standard as well as Coleman and Steele and other cited authors suggest that a coverage factor of 2.0 can typically be used to compute a 95% confidence interval for analytical results that are based on uncertain measurements.

The results from Coleman and Steel can be reduced to a generalized equation for handling bias and precision errors in the computation of a 95% confidence interval, U_r for engineering analyses

$$U_r^2 = B_r^2 + P_r^2$$

Equation 4-9

where B_r and P_r are the total bias and precision uncertainties for an arbitrary analysis. The bias and precision uncertainties for a total analysis are subsequently defined as

$$B_r^2 = \sum_{i=1}^{j} \theta_i^2 B_i^2 + 2 \sum_{i=1}^{j-1} \sum_{k=i+1}^{j} \theta_i \theta_k \rho_{bik} B_i B_k$$

Equation 4-10

and

$$P_r^2 = \sum_{i=1}^{j} \theta_i^2 P_i^2 + 2 \sum_{i=1}^{j-1} \sum_{k=i+1}^{j} \theta_i \theta_k \rho_{Sik} P_i P_k$$

Equation 4-11

where i and k are indices over the number of variables involved in the analysis, and ρ_{bik} and ρ_{Sik} are the correlation coefficients between pairs of bias and precision errors, respectively. Finally, the precision error of an individual variable, P_i is related to its sample variance, S_i according to

$$P_i = t_i S_i$$

Equation 4-12

where t_i is the value of the t statistic for the desired level of confidence and degrees of freedom of the variable. Based on the prior arguments, the t statistic is often assumed to be equal to 2.0 for a 95% confidence interval.

4.4.3.2 Application of measurement uncertainty to single sample measurements

HVAC systems are dynamic and responsive to their environments, which makes it challenging to collect multiple measurement samples that are representative of a single state of the system. From a physical perspective, the time interval over which buildings can transition between states may be limited by the time constants of energy exchange in the building. Thermal interactions in buildings generally occur over a period of minutes, which suggests that sub-minute measurement sampling can yield statistically significant populations of data to represent building states. From a practical perspective, however, control loops for building control systems operate at the scan rate of their processors, and representative forcing functions such as PID

loops can modulate at sub-minute periods. In general this means that building control systems can react very quickly to small thermal interactions, despite the slower response of bulk building physical properties such as air or wall temperatures. Since many of our FDD measurements can be outputs from the building control system, or physical measurements about the building space, it becomes difficult in practice to define a sampling period for building systems other than an instantaneous single sample.

Coleman and Steele as well as ISO Guide 68 admit that most practical experimental systems only permit single-sample instantaneous measurements of the state of the system, much like we may be limited to in buildings. If it were possible to create statistically significant sample sizes for each state of a system (perhaps with 10 measurements or more, per state), then a sample precision and possibly bias error could be computed for each state sample and used in the above equations to create a confidence interval. In the presence of single sample measurements, however, Coleman and Steele and ISO Guide 68 indicate that the selection of precision and bias errors should be selected for each measurement according to the best relevant information available at the time of measurement.

Following these recommendations, we will in this thesis use the manufacturer's quoted accuracy to generate estimates of the precision error for our measurements made about building systems. Specifically, we will make a fundamental assumption that the manufacturer's quoted relative percent error for their sensor represents a 95% confidence interval about the sensor's measurement, under a normal distribution that is centered at the sensor measurement. For example, under this assumption a temperature measurement of 65 °F, from a sensor with a manufacturer's quoted accuracy of +/- 3 °F (typical for an averaging sensor in an AHU) will be interpreted as representing a normal distribution of the possible real values for temperature, with sample mean of 65 °F, and sample standard deviation of 1.5 °F. This assumption is effectively a form of bootstrapping in order to form a probabilistic inference over a time series of data. In terms of the equations above, we will assume that t_i is equal to 2.0 for a 95% confidence interval, and that P_i is provided by the manufacturer as their sensor accuracy. The validity of this assumption should be considered on a case-by-case basis for analysis on uncertain measurements

however it is a simple program that could be used to quickly build probabilistic FDD for HVAC systems.

Momentarily ignoring the bias error and assuming independent variables and measurements, we can now use the above equations to propagate measurement uncertainty through our FDD analysis. Furthermore, by assigning a statistical interpretation to a manufacturer's published sensor accuracy we can present the results of FDD as a statistical distribution instead of a deterministic number.

4.4.3.3 Model uncertainties and Bias Error

Coleman and Steele present the estimate of a measurement bias limit, B_i as a root-sum-square combination over elemental bias limits, B_n

$$B_i = \left[\sum_{n=1}^{M} (B_i)_n^2\right]^{1/2}$$

Equation 4-13

for M possible elemental biases. According to Coleman and Steele and ISO Guide 68, elemental biases can include a wide variety of experimental contributions including information theoretic limits of data acquisition systems, manufacturer's specifications on drift, experimental insights from calibration studies and much more (66). In general, the ISO as well as Coleman and Steele specify that the experience of the investigator is central to the selection of important elemental biases; many biases can exist in an experimental system, but not all of them carry a significant contribution to the outcome of the experiment.

One possible elemental bias that may be included in the estimate of the bias limit is the bias of the DRE model itself; for example it may be possible to characterize the DRE such that under certain circumstances we expect its result to not be very accurate (as in the model itself is known to be inaccurate for certain regimes of variable measurements). Inclusion of the DRE bias as en elemental bias may serve as a bridge between different mathematical models that attempt

to estimate the same result using different measurements. In terms of the prior thought experiment of estimating boiler energy consumption, inclusion of DRE biases within the estimates of confidence intervals would allow us to compare energy consumption estimates derived from measurement of natural gas flow against less accurate energy consumption estimates that are derived from measurements on the status of the natural gas valve. If both results are expressed as a statistical distribution, then a hypothesis test can be used to determine whether the results are effectively equivalent to within a specific degree of confidence.

There may be a variety of approaches to encapsulating the uncertainty of models within a bias term however we have considered the simple engineering approach of using upper and lower bound extremes of plausible experimental results. This approach to estimating model biases is drawn from the classical engineering technique of bounded-value estimation. For example, thermal-fluid analysis of engineering problems are often complicated with multiple modes of heat exchange and transient phenomena, but we often simplify such problems by making bulk estimates of heat transfer coefficients, physical properties, and flow geometries. These types of simplifications are intended to peal-back real-world complexities and permit engineers to get a "ball-park" estimate of how a system is working, without having to invest significant time in a more accurate analysis that would otherwise require numerical simulation or approximation. When engineers create such "ball park" analyses they must also compute upper and lower bound "brackets" on their gross estimate in order to couch their results within the uncertainty of their analysis. Our approach, which is one of many possibilities, is simply to assign a confidence interval across the upper and lower bound brackets of the engineering analysis. Admittedly this is an "expert" approach to expressing model uncertainty since it relies on the confidence that an expert perceives in their bounded model. Despite the drawbacks of expert systems (67), this approach does allow us to use the simple mathematics of experimental uncertainty in order to simultaneously manage measurement and model uncertainties.

4.5 Probabilistic Inference

Since the results of our engineering analysis are now statistical distributions rather than deterministic numbers, we can use the tools found in statistics and probability to perform rigorous FDD inference. Expert rules and other methods mentioned in chapter 2 can now be re-cast in terms of probability so that the results of discrimination and classification can be communicated to within certain levels of confidence. Confidence levels in FDD analysis can then be used to measure the strength of a fault signal and filter out FDD results that are less certain.

4.5.1 Example of cost prioritization

Earlier in this chapter we computed the thermal-cost of operation for a simple AHU; when measurement uncertainties were included in that computation we found that the cost of operation was $1,188 +/- $178, but we did not prescribe any interpretation for the +/- $178. Applying the ideas from experimental uncertainty to those same results allows us to now interpret the included uncertainty of the calculation as a 95% confidence interval for a normal distribution centered at $1,188. If we desired, we could now compute the probability, P, of the cost of operation exceeding a certain threshold, UD, according to

$$P = \int_{UD}^{\infty} \frac{2}{U_{tot}\sqrt{2\pi}} e^{-2(x-M)^2 / U_{tot}^2} dx$$

Equation 4-14

where M is the center of normal distribution, which is $1,188 in this example, and U_{tot} is the result's total uncertainty, which is $178 in this example. Using this approach, we can now prioritize the investigation into different pieces of equipment by the probability that their energy cost exceeds a certain threshold.

4.6 Taking action under uncertainty

A large majority of FDD analysis and results, as they were presented in the literature review of chapter 2, are often based on absolute terms; expert rules are either violated or not-violated, decision trees consist of binary yes or no questions, and dollar losses are expressed as a single number. But as we have seen in this chapter the world of measurement and data in buildings is fraught with uncertainties, whether they stem from random errors or systematic biases. Furthermore, building managers and stakeholders are used to that uncertainty and building management processes have clearly grown to accommodate it.

FDD in buildings may benefit from embracing building model and data uncertainty in order to be more flexible with practical deployment and also to communicate results in terms that building managers and stakeholders can better appreciate. The engineering application of uncertainty provides a convenient and simple mathematical framework that FDD can use to manage building model and measurement uncertainties. The added benefit of this analytical approach is that numerical results are delivered as statistical distributions rather than deterministic values, which allows FDD systems to naturally leverage the large tool set of probability and statistics in order to perform rigorous analysis.

5 Experimental work on MIT equipment

This chapter of the thesis will describe our experimental methods for testing the FDD system, as well as the methods by which the system was deployed on all of the AHUs involved in this research. The chapter is broken into two primary sections; the first section will focus on AHUs 9 and 10 in MIT building 46, and the latter section will focus on the other three AHUs.

5.1 Introduction

An FDD system was designed, developed and deployed on five AHUs across three separate buildings on the MIT campus. For one of those AHUs, AHU 9 in MIT building M46, mechanical and software-control faults were purposefully applied to it and the FDD system was

tested to see if it could detect those faults and evaluate their impact on AHU energy consumption. Another three AHUs, AHUs 2 and 3 in MIT building 56, and AHU 2 in MIT building 16, were also used to test the FDD system because of their known pre-existing conditions of simultaneous heating and cooling. These last three AHUs served as control variables in our experiment to test if the FDD system could detect and evaluate pre-existing conditions, as oppose to the conditions that we purposefully applied on the other tested AHUs. The three control AHUs were also chosen because MIT had already deployed on them a fault detection service that is offered by a local company, Cimetrics Inc. MIT has had a long-standing relationship with Cimetrics, and we were interested in comparing the results of our FDD system against what was already identified by Cimetrics and known to the MIT facilities department. The FDD system was also deployed on AHU 10 in MIT building 46, however no purposeful faults were applied to that AHU; AHU 10 is a twin unit to AHU 9, and it served as the control for our experiments where we purposefully deployed mechanical and software-control faults.

5.2 AHUs 9 and 10 in MIT Building 46

MIT Building 46, otherwise known as the McGovern Institute for Brain Research and the Picower Brain and Cognitive Research Center, is a mixed-use laboratory and educational space that was completed in 2006. The building includes over 250,000 square feet of research, conference and teaching space including a primate research laboratory, animal vivarium, nearly 1,500 VAV air distribution boxes, hundreds of laboratory fume hoods and a seven story open-air atrium with a glass ceiling. The building's HVAC system includes over two dozen AHUs, most of which are dedicated outdoor air units (DOAs), of varying sizes and complexity. The building's ventilation system was designed so that laboratory animals could be transported anywhere in the building at any time, resulting in essentially the largest laboratory-grade ventilation system on the MIT campus. In addition to the high air exchange rates used throughout the building, some of the AHUs are also designed with demand-based ventilation systems that utilize local and return air measurements of CO_2 concentration to modulate outdoor air flow.

5.2.1 Physical Description of AHUs 9 and 10

AHUs 9 and 10, which were included in our research, are two of the AHUs in MIT building 46 that utilize demand-based ventilation; they are also two out of a small minority of AHUs in the building that re-circulate air and possess an economizing cycle (the nature and implications of AHU economizing cycles were discussed previously in chapter three). AHUs 9 and 10 are of a twin design, both with a name plate supply air flow rating of 50,000 CFM, static duct pressure rating of 1.4 inches of water column, steam heating coils rated at 1 million BTUs per hour, and chilled water coils rated at 167 Tons. AHU 9 serves four conference rooms, one of which is rated for 70 occupants and the others at 20 occupants each, as well as part of the main atrium of the building. AHU 10 serves the remaining portion of the atrium, as well as the main reception and administrative office of the building, which together hold roughly 100 occupants.

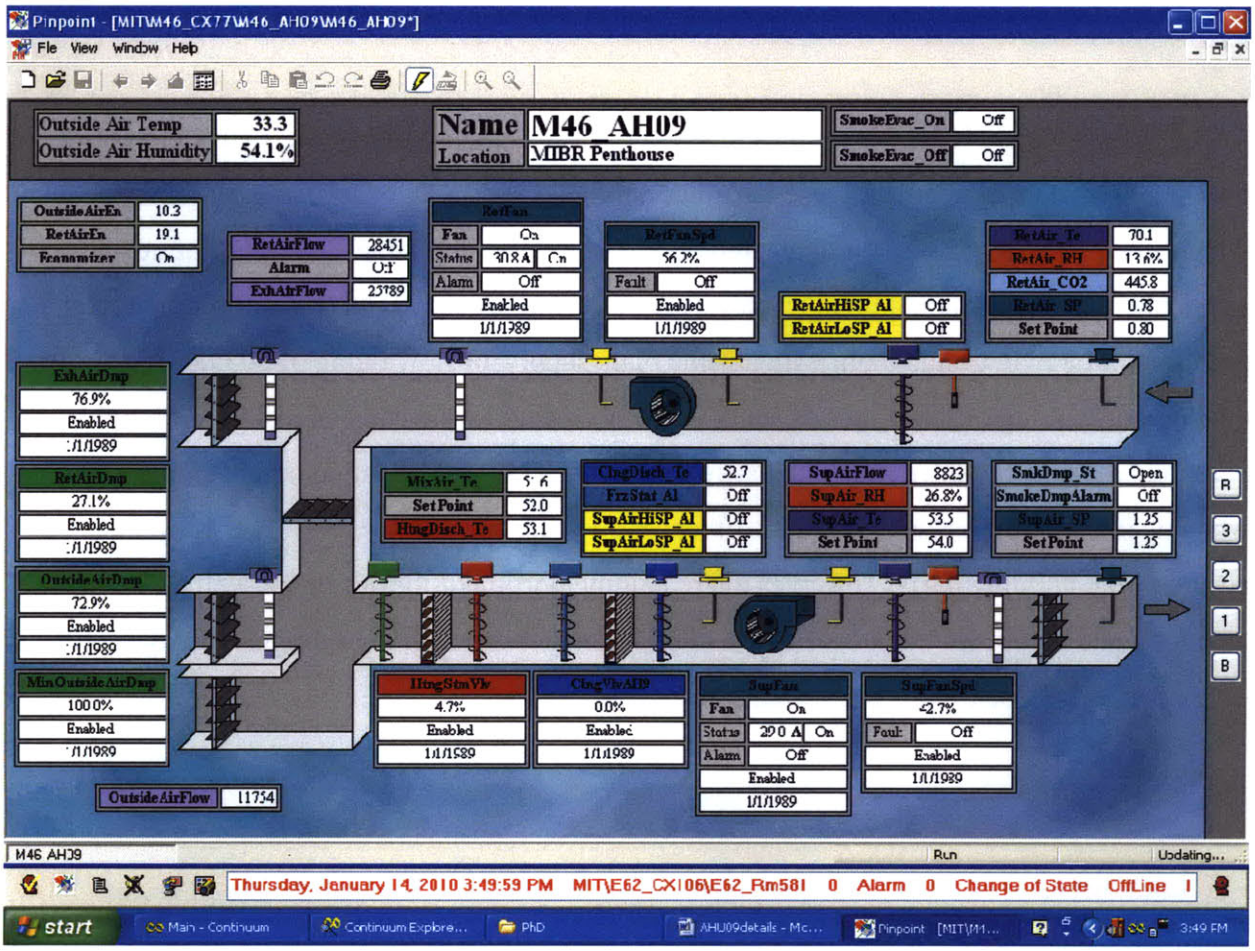

Figure 32 Screenshot from the building control system for MIT building 46, detailing a schematic of AHU 09; the schematic for AHU 10 is identical. The design of these AHUs is reminiscent of the schematic VAV AHU design that was included as figure 1 in chapter 3 of this thesis

The schematic diagrams of AHUs 9 and 10 are identical, and are represented in Figure 32 by a screenshot of the building control interface for AHU 9. The outdoor air intake of the AHUs is split into two sections; a maximum outdoor intake with continuously adjustable dampers, and a minimum outdoor air intake with binary open or closed dampers. The minimum outdoor air dampers occupy roughly 30% of the cross sectional area of the total outdoor air intake manifold, corresponding to a minimum outdoor air fraction of 30% across all supply air flow rates for both AHUs 9 and 10.

The schematic in Figure 32 is slightly misleading because it portrays the measurement of the outdoor air flow only at the maximum outdoor air inlet to the mixing box; in reality the

outdoor air flow measurement station is located in the throat of the outdoor air intake manifold that leads to both sets of outdoor air dampers. In general, AHUs 9 and 10 are better instrumented than most VAV AHUs; for example Figure 32 shows that all four of the AHU air flows are measured: outdoor, return, supply, and exhaust. Measurement of all four AHU air flows is not commonly found in conventional AHU installations, but was included with this installation because of the critical role that ventilation plays in the safe use of the laboratory building. The magnitudes of all four air flows are controlled by the interplay between the supply and return air fans, and the exhaust, outdoor, and recirculation air dampers. The economizer section of the AHUs is composed of the mixing box and the set of three control dampers.

The heating and cooling coils of the AHUs are also well instrumented with discharge temperature sensors following each coil, as well as a mixed air temperature sensor in the mixing box and supply air temperature sensor following the discharge of the supply air fan. Additionally, static pressure sensors are included on the supply and return air sides of the AHUs in order to control the rotational speed of their respective fans. Humidity sensors are also included on the supply, return and outdoor air flows in order to facilitate enthalpy computations for each flow and subsequent enthalpy control over the economizer cycle. Other features included in the control graphic snapshot include set-points and valve positions which are output and stored values from the building control system, as well as measured status labels for the fans.

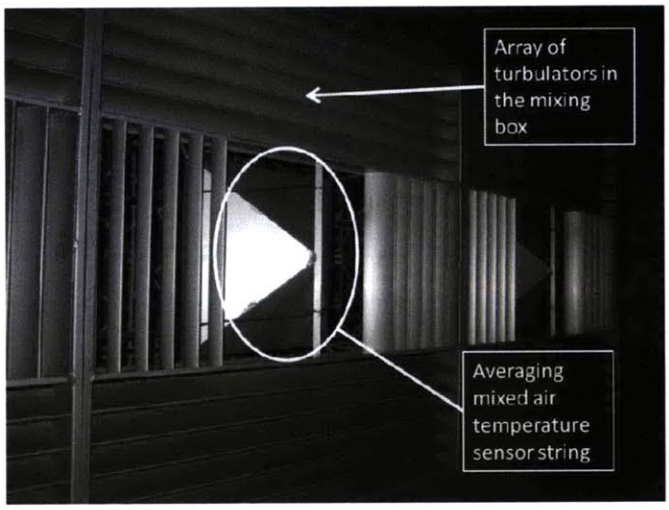

Figure 33 Picture of the mixing box interior showing the use of turbulators to improve mixing of return and outdoor air; the averaging mixed air temperature string is also shown traversing the turbulator surface. As a length scale, the length of a horizontal turbulator fin is 4 feet

The building control system for MIT building 46 was provided by Schneider Electric and installed by the MIT branch office for Schneider Electric Systems Integration. The control system for MIT building 46 includes over 5,000 measurement and control points, distributed over several hundred hardware controllers, and several dozen graphical interfaces. The graphics interface and control platform belong to the continuum product line of the building controls division of Schneider Electric; roughly two-thirds of the MIT campus operates on this control platform, with over 200,000 measurement and control points across the campus. The sophistication of the Schneider Electric control interface is representative of features and functions that are commonly found amongst most building control systems.

5.2.2 Relevant Controls Programming

The basic control elements for AHUs 9 and 10 are embodied within their respective sequence of operations document; this is a descriptive document that is kept by MIT facilities as part of the operations and maintenance (O&M) manual for the equipment. The O&M manual is a typical resource that building managers have for each piece of equipment or major system that is in their buildings; O&M manuals exist to support repair and maintenance efforts over the

lifetime of building equipment. The key control sequence elements that influence day-to-day AHU energy consumption have been summarized here:

1. Supply and return fans shall be on at all times

2. The variable speed drives on the supply and return fans shall modulate the fans to maintain 1.25 inches of water column static duct pressure in their respective ducting

3. The unit will enter economizer mode when the outside air enthalpy is less than the return air enthalpy. The economizer will be locked out when the outside air enthalpy equals or exceeds the return air enthalpy; in that case the maximum outdoor air dampers will close, the exhaust air dampers will go to their minimum position of 15%, and the return air dampers will be fully open.

4. In economizer mode, the exhaust, return and outdoor air dampers will modulate to maintain a mixed air temperature of 52 ^{o}F; the return air dampers will close as the exhaust and outdoor air dampers open

5. CO_2 control over outdoor air flow takes precedence over the mixed air temperature control and economizer lock-out. A minimum outdoor air flow of 20,000 CFM will be maintained for the minimum CO_2 threshold concentration of 800 ppm; as the measured return or space CO_2 concentration exceeds 800 ppm, the minimum outdoor air flow rate will proportionally increase in value.

6. The AHU's steam pre-heat coil and chilled water cooling coil are controlled by separate PID control loops in order to maintain a supply air temperature of 55 ^{o}F.

The control sequences for these two are-handlers are more sophisticated than similarly sized air-handlers that are found in other buildings; for example, most buildings do not include CO_2 concentration as a ventilation control signal. MIT Building 46 uses CO_2 concentration-based ventilation control as a means of carefully managing the indoor air quality (IAQ) of the building. Instead of CO_2 control signals, most buildings have a time-based schedule that reduces the supply air flow rate of the air-handler when the building is supposed to be largely un-occupied (typically in the evening till the early morning for conventional office buildings). To compensate for periods where the pre-programmed ventilation and occupancy schedule is inaccurate, occupants are often provided with manual override switches to temporarily re-activate the ventilation system; these override switches or buttons are typically integrated into

the local room thermostat. On the other hand, occupants may not realize that their IAQ has decreased until it reaches uncomfortable levels, and in research laboratory spaces with animal quarter and noxious chemicals, poor IAQ can pose serious health risks to building occupants.

In order to actively control the energy consumption associated with ventilation, as well as maintain high IAQ, the dynamic control over outdoor air ventilation rates in MIT building 46 is embedded within the programming of the system as well as in occupant override switches. The related control methods and set-points for the economizing cycle and discharge air conditions in AHUs 9 and 10 are more customary amongst other VAV AHU implementations. Likewise, the static duct pressure control limits and scheme for modulating the supply and return air flows are common for VAV AHUs that are controlled to meet static duct pressure constraints.

5.2.3 Instrumentation and Data Acquisition

5.2.3.1 Measurements

Over two dozen physical and virtual points exist for AHUs 9 and 10 on their building controller. Some of these physical and virtual points will be important to our FDD system, while others do not yet appear useful. Without yet developing a FDD model of the system, our experiences from chapter three suggest that physical and virtual control points that are relevant to thermal energy calculations will be useful to the FDD system. The schematic in figure 1 suggests an initial list of measurement points that can be used to compute AHU thermal energy consumption:

- Outdoor, supply, return, and mixed air temperatures
- Heating and cooling coil discharge air temperatures
- Outdoor, return and supply air relative humidities
- Outdoor, return, supply and exhaust air flow rates

Other data sources shown in Figure 32 may also prove useful to detecting faulty operation of the AHU, or computation of their electrical energy consumption, including:

- Outdoor, exhaust and recirculation damper position
- Heating and cooling coil valve position
- Supply and return static pressure
- Return and other CO_2 measurements
- Supply and return fan amperage

5.2.3.2 Data Collection

The collection of data for all of the experiments included in this thesis was performed by the Schneider Electric building control system itself. At the equipment level, measurements and control points were affected by the controller product that was installed on the equipment; in all cases, these were Infinet/TAC brand products from Schneider Electric. The controllers on AHUs 9 and 10 were i2-920 system controllers by Schneider Electric, with 12-bit digital-to-analog conversion for 16 universal sensor inputs and 8 control outputs, analog voltage resolution of 2.5 mV over the range of 0 to 10 V, analog voltage accuracy of +/- 7.5 mV, 1 MB of SRAM and 2 MB FLASH memory.

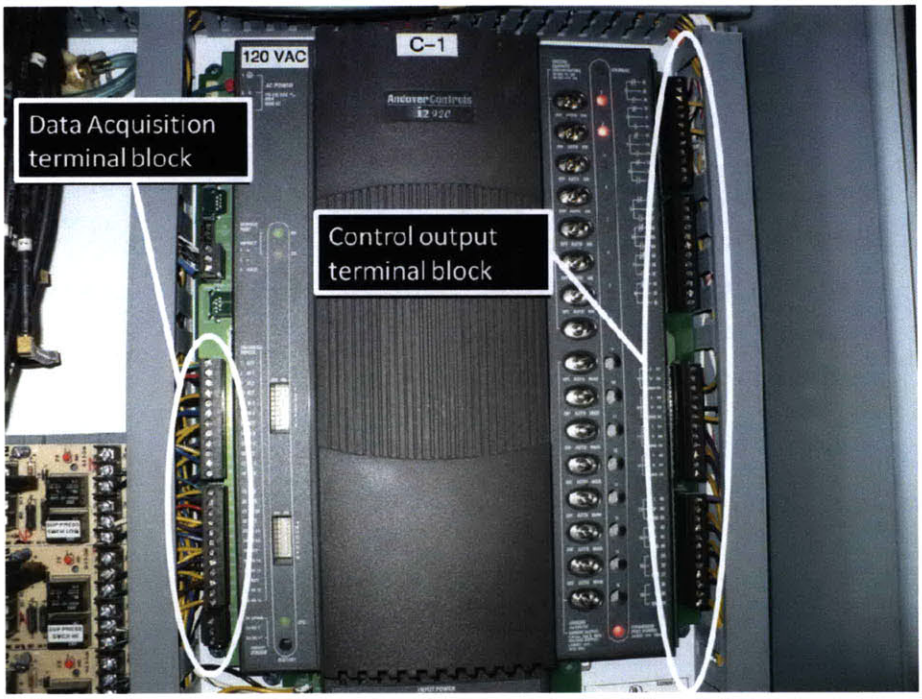

Figure 34 Picture of the i2-920 system controller for AHU 9 in MIT building 46

In addition to on-board memory and implementation of control programming, the controllers also serve as a terminal block for collecting and aggregating control point data. Point data can arrive at the controller's terminal block in a variety of forms, typically including but not limited to:

- Analog voltage (typically 0 to 10 V)
- Analog current (typically 4 to 20 mA)
- Circuit resistance (typically 3 to 20 kΩ)
- Digital pulse (typically 5 V peak-to-peak)

The collection and storage of building control data in the Schneider Electric product is also facilitated through their software. All control points, whether they are numerical computations, input values or output signals can be logged over predetermined time intervals at the controller or database level. The building controller itself has a finite memory of 100 MB that can store, and will overwrite point data; for a typical AHU with 20 control points, polled regularly at 5 minute intervals, the local controller can store roughly 15 days worth of data. The limitations on local data storage at the controller level are overcome if the building control data can be regularly sent to a much larger storage database. At MIT and in many other buildings that

have Schneider Electric control products, the local building controllers are connected over an Ethernet network to a central database that can support extended logging of building control data.

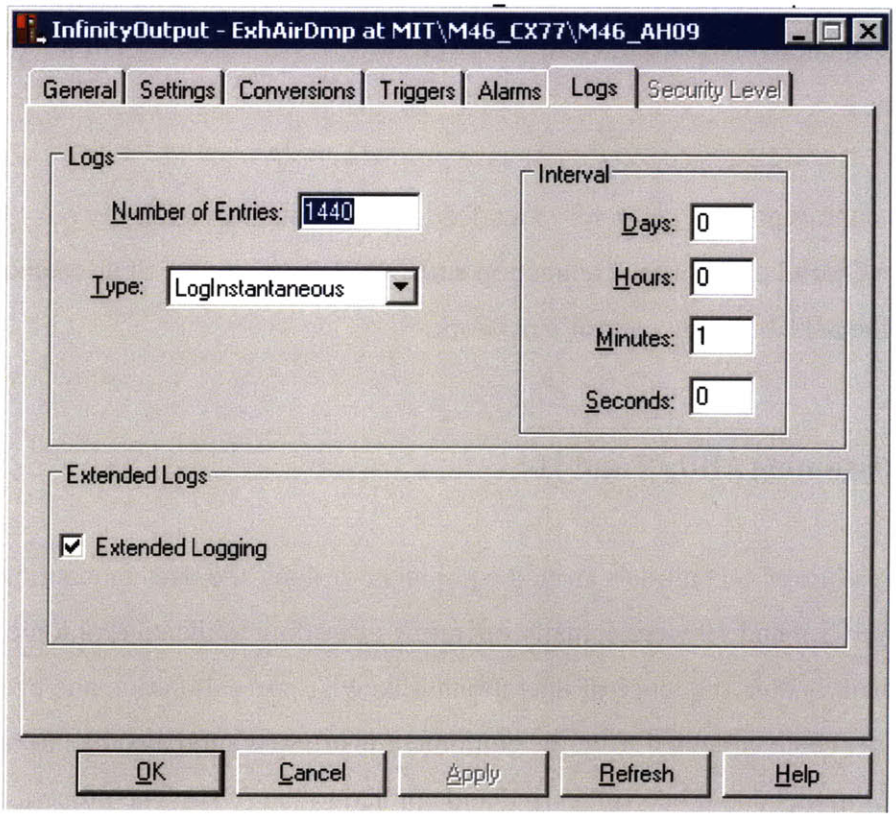

Figure 35 Screenshot of the setup screen for implementing extended logging on an Infinity point; the data logging feature is supported on MIT's building control system for all types of control points.

Figure 35 shows the setup screen for implementing extended data logging on the exhaust air damper position for AHU 9 in MIT building 46. The logging setup allows users to specify the interval of data logging, as well as its typology; data logging can be instantaneous at a certain interval, or averaged over that interval. The user can also specify how many total data points should be collected per day, and whether the data should be stored over the network at the central database (this is termed extended logging) or just locally on the building controller.

Despite the capacity for a network database server to store seemingly limitless amounts of data, of the nearly 200,000 points that are on the building control network at MIT, only about

380 of them are currently setup for extended logging; less than half of a percent of the control infrastructure at MIT is setup for any extended data logging. On the one hand the volume of possible building control data at MIT seems like a tremendous untapped resource for better management and maintenance of MIT's buildings, however the reality is that few personnel at MIT facilities have the expertise, let alone the time, to use that data to improve the campus. Consequently, data is rarely collected and reviewed by MIT personnel except for emergency or forensic purposes.

Once a control point is setup for extended logging, data harvesting by our FDD system is achieved via an ODBC data source connection and T-SQL query to the SQL database that rests at the heart of the MIT building control's network.

5.2.4 Commissioning AHUs 9 and 10

Before creating FDD models around equipment designs and data acquisition, we needed to verify that AHUs 9 and 10 were actually operating according to the control logic that was prescribed in their written sequence of operations. Likewise, the calibration and even existence of building control data collected from the equipment instrumentation also had to be verified before any FDD model could be synthesized and applied to the AHU. The process of verifying control logic, checking sensor calibration and visually inspecting the integrity of the equipment is collectively termed commissioning; the goal of commissioning is to verify that the AHU is operating in the manner in which it was originally designed and intended to operate. We commissioned AHUs 9 and 10 in our research program according to the following simple procedure:

1. Verify the enthalpy control over the AHU's economizer cycle
2. Compare outdoor and supply air flows against a calibrated flow meter
3. Examine duct work for mechanical deformation, possibly resulting from over or under pressurization
4. Check the integrity of the pneumatic connections and piping between the controller and equipment actuators
5. Cycle the AHU's dampers to ensure that they work properly

6. Verify that the heating and cooling coil control valves can be controlled to open and fully close
7. Verify the balance of air flows between the four air flows that are measured about the AHU

Building commissioning is a profession unto itself, and our abbreviated commissioning list does not reflect the much deeper and thorough investigation that is associated with a dedicated commissioning and retro commissioning process (68). Our primary interest in AHU commissioning was to understand how the AHU was programmed to operate; without that knowledge we could not create an accurate FDD model of the AHU. In addition to learning how the AHU was truly programmed to operate, we also wanted to gauge the level of AHU sensor calibration and equipment integrity. Over the course of this research project, like in most commercial settings, there was never an opportunity to remove sensors from the equipment in order to test their calibration under laboratory conditions. Our efforts to test instrument calibration were on par with conventional commissioning practice; measurement data from the AHU was compared against a hand-held or temporarily installed sensor with a known calibration, and also reviewed by hand for any practical indications of gross inaccuracy. The process of commissioning AHUs 9 and 10, especially with the respect to verifying actuator controllability, eventually mixed with and transitioned into the development and deployment of an FDD system on the AHUs. For that reason, aspects of commissioning such as verifying damper actuation, air flow balancing and control valve controllability were included as FDD modules themselves; the results of that particular work will be discussed in the next chapter of this thesis.

The first step in commissioning AHUs 9 and 10 was to examine their control programming. We quickly determined from an inspection of the control logic that the AHUs were not programmed exactly according to the control architecture that was specified in their written sequence of operations. In particular, the AHUs were programmed to use outdoor and return air temperature instead of enthalpies (as was mentioned in the sequence of operations) to control their economizing cycle. This was not a surprising finding; most economizing cycles in VAV AHUs are based on temperature comparisons rather than enthalpy comparisons, usually because the system installation does not include humidity sensors on the outdoor and return air

flows. In the present case, AHUs 9 and 10 are equipped with humidity sensors and the sequence of operations does specify enthalpy control over the economizer, however the enthalpy-derived control programming was not implemented in practice. It is unknown to us why the controls programming did not include enthalpy control despite its inclusion in the written sequence of operations; it may be possible that the sequence of operations that was kept on record by the MIT facilities department did not reflect any last minute changes that were made in the controls specifications. The discovery of that discrepancy between programming practice and specification underscores the need to commission equipment before deploying an FDD system. We learned throughout this research project that the contents of building documentation such as engineering plans, operating manuals and sequences of operation are not guaranteed to accurately reflect physical installations.

We found other control programming errors in the control logic for AHUs 9 and 10, most important of which was as an erroneous conversion factor between the measured supply air flow velocity and computed volumetric flow rate. While this programming error has no consequence on the control of the AHU, it would dramatically alter the outcome of data computations within an FDD model of the AHU. We found in both AHUs 9 and 10 that the conversion factor between the supply flow velocity measurements and volumetric flow computation was off by a factor of 2; the building control system was reporting volumetric supply air flow rates that were only half as much as they actually were.

Our physical inspection of the AHUs found no observable mechanical, pneumatic or electrical faults. Since the building itself had only been occupied for less than 3 years, we did not expect the equipment to show any visible signs of degradation. We worked with a technician from the building management office to examine electrical connections, pneumatic tubing, mechanical linkages, air filters and fan belts, all of which appeared to be in acceptable working order. While there were no obvious physical failures about the equipment, we relied on our FDD system models to identify any software or hardware faults that were not identified during the commissioning process.

5.2.5 AHU Models

5.2.5.1 Introduction

The fundamental thermal energy model for a generic VAV AHU was derived in chapter 3 and is also applicable here:

$$TE_{AHU} = \dot{m}_{OA}(h_{OA} - h_{RC}) + \dot{m}_{SA}(h_{RC} - h_{SA})$$

Equation 5-1

where positive and negative values of TE_{AHU} correspond to heating and cooling energy consumption by the AHU, respectively, measured in BTUs per hour. The mass flow rates of outdoor and supply air flow rates are represented by $\dot{m}_{OA}$ and $\dot{m}_{SA}$ respectively, measured in pounds of dry air per hour. The enthalpies of the outdoor, recirculated and supply air flows are represented by h_{OA}, h_{RC}, and h_{SA} respectively, measured in BTUs per pound of air. Based on our list of measurements for AHUs 9 and 10 in the previous section, the building control system yields either direct measurements of each quantity in equation 1 (such as air flows), or measurements of quantities that can help to evaluate the variables in equation 1 (such as temperatures and humidities to yield enthalpies).

While Equation 5-1 can tell us how much thermal energy is consumed by AHUs 9 and 10, it does not immediately provide insight into how much thermal energy *should* be consumed by AHUs 9 and 10 under optimal operating conditions. Since the AHU consumes thermal energy in order to condition mixtures of recirculated and outdoor air flow to meet supply air conditions, we can predict how much energy should be consumed by these AHUs by predicting how much outdoor and recirculated air they should receive.

Chapter 3 explored this same modeling exercise for the AHUs at the Iowa Energy Center and in that case we derived relationships for the outdoor air fraction of the supply air flow based on the AHUs' mixed air temperature control logic. AHUs 9 and 10 in MIT building 46 have similar mixed air temperature controls to the AHUs that are at the IEC, however the analysis for

the AHUs at MIT are made more complex by the presence of additional ventilation control logic that is based on measured CO_2 concentrations.

5.2.5.2 Mixed Air Temperature Control

The thermal energy consumption of the air-handler, as computed in equation 1, is dependent on several key variables, including the mass flow rates of supply and outdoor air, and the enthalpy of the return, outdoor, and supply air flows. According to the control sequence of operations, the outdoor air flow into the MIT AHUs is a function of the mixed air temperature conditions in the mixing box and the CO_2 concentration in the return and space air. Regardless of the control logic, however, the analysis in chapter 3 revealed that the flows of recirculated and outdoor air through VAV AHUs are controlled by the positions of the AHU's mixing box dampers. The model for the outdoor and recirculated air flows in AHUs 9 and 10 can therefore be derived from the control logic that links their mixed air temperature and CO_2 sensor inputs to the positions of their exhaust, outdoor and recirculated air dampers.

Upon close examination of the control logic in AHUs 9 and 10 we found the programming that computes the control signal for the mixing box dampers according to CO_2 concentrations and mixed air temperature conditions:

```
1     Max_CO2 = maximum(RetAir_CO2, Rm3310_CO2, Rm4199_CO2, Rm4300A_CO2, Rm5199_CO2,
2     Rm6199_CO2)
3     CO2Mode = Range.Fn((Max_CO2 > CO2_Se), 0, 1, CO2Mode, 0, 0.2, Timer[2])
4
5     If CO2Mode = On then
6       OutsideAirFlo_Se = Ratio.Fn(CO2Signal, 0, 1, OutsideAirFloMin, SupAirFlow)
7       MixAirDmpSignal = maximum(MixAirTeSignal, OutsideAirFloSig)
8     Else
9       OutsideAirFlo_Se = OutsideAirFloMin
10      MixAirDmpSignal = MixAirTeSignal
11    Endif
```

The first two lines of code show that the CO_2-based control signal, "Max_CO2" is actually dependent on several measurements of CO_2 concentration; five from rooms that are served by the air-handler, and one from the return air flow. The maximum measured CO_2 concentration from of all of those sources is used create the "Max_CO2" signal which controls whether the air-handler ventilates the building according to CO_2 concentration logic or mixed air temperature logic

The third line of code executes a comparison between the maximum measured CO_2 concentration and a set-point concentration of CO_2, the result of which is used to activate the "CO2MODE" for the air-handler; the set-point CO_2 concentration, CO2_Se, is equal to 800 ppm[6]. The comparison in line 3 is implemented with a function called Range.Fn, which also accepts several other arguments; further investigation into Range.Fn revealed that its extra arguments define a timer that prevents the CO2MODE from activating unless the Max_CO2 signal exceeds the CO2_Se for a certain period of time.

Once the CO2MODE has been defined as either "on" or "off", its value is used in a conditional statement, code lines 4 through 10, to switch between ventilation that is controlled by CO_2 concentration and mixed air temperature conditions. When the CO2MODE is off (code lines 7 through 10), the control signal for the dampers, "MixAirDmpSinal" is equal to the mixed air temperature control signal, "MixAirTeSignal". During periods when the CO2MODE is on (code lines 5 through 7), the damper control signal will equal the maximum of two different signals, "MixAirTeSignal" and "OutsideAirFloSig". In cases where the "CO2MODE" is on, line 5 of the control program shows that the program also computes a new set-point value for the outdoor air flow rate, "OutsideAirFlo_Se"; this new set-point value is used to generate an updated "OutdoorAirFloSig" for the active "CO2MODE" state.

[5] Code snippet taken from a program entitled "AHMisc.Pr" on the controller for AHU 9 in MIT building 46. The snippet is intentionally included in order to show the reader what control programming looks like in the Schneider product line.

[6] The values for various AHU set points was also stored in the control program "AHMisc" on AHU 9 in MIT building 46

The mixed air temperature control signal, "MixAirTeSignal", and outside air flow signal, "OutsideAirFloSig" are both computed from their own respective proportional-integral-derivative (PID) control programs[7]. In both cases, the PID control programs generate a single corrective control signal based on a measurement value, a set-point for that measurement value and three gain constants (one each for the proportional, integral and derivate gains). In the case of "OutsideAirFloSig", the set-point value for the PID calculation is also computed in the main control program (see line 5 of the code snipped included above). The PID programs for MixAirTeSignal and OutsideAirFlowSig output a signal equal to zero for measurement values that are equal to or less than their respective set-points. For measurement values that exceed their respective set-points, the MixAirTeSignal and OutsideAirFloSig are positive and range from 0 to 1.

For example, if the "CO2MODE" is off and the measured mixed air temperature is 65 °F while its set-point is 52 °F, then the "MixAirTeSignal" PID program would generate a positive control signal between 0 and 1. In that case and according to line 10 of the included code snippet, the MixAirDmpSignal which drives the damper actuators would be equal to the MixAirTeSignal. Examination of the actuation programs for the exhaust, outdoor and recirculation dampers revealed that the MixAirDmpSignal is directly used to drive the exhaust and outdoor air dampers, but its unit compliment, or one minus the MixAirDmpSignal, is used to drive the recirculation dampers. This follows our expectations from chapter 3 where we showed that normal mixing box operation requires that the recirculation dampers continue to close as the outdoor air dampers continue to open.

In the case of AHUs 9 and 10, the exhaust and outdoor air dampers are both directly modulated by the MixAirDmpSignal, however the exhaust damper actuator also includes a constant positive bias on the control signal in order to prevent those dampers from fully closing when the control signal value is equal to 0. The outdoor air dampers do not include a bias on their damper control signal because the minimum outdoor air dampers will continue to remain fully open even when the adjustable outdoor air dampers are fully closed.

[7] The controller for AHU 9 includes several individual PID loop control programs, including ones for the mixed air temperature signal and outside air flow signal.

The PID programs for both "MixAirTeSignal" and "OutsideAirFlo" are evaluated once each time the air-handler controller cycles through its complete set of programs. The cycle rate of the air-handler is termed its scan rate, and is typically on the order of 10 to 60 seconds[8]; the scan rate is sufficiently slow to allow the physical actuators on the air-handler to respond to their control signals. A detailed model for outdoor air flow into the AHU may include the actual PID loop control logic over "MixAirTeSignal" and "OutsideAirFlo", however in that case our model of the AHU would also have to include sub-component models for actuator response, hysteresis, and other complicated dynamics[9]. It is not clear whether such a detailed dynamic model of an AHU yields superior insight into the impact or presence of gross energy in-efficiencies like simultaneous heating and cooling or imbalanced air-flows.

Since our goal is to simplify the FDD modeling process such that FDD systems can be rapidly deployed, we will make a fundamental assumption that the dynamics of satisfactory AHU control response are much faster acting than the equipment pathologies that incur the greatest energy in-efficiencies. Another way of stating this assumption is that the time period over which important energy in-efficiencies exist is much longer than the time needed for AHU actuators to properly respond to their control signals. This assumption does not discount the impact of poor PID loop tuning or the slow response of actuators that suffer from mechanical degradation; on the contrary we are simply putting a minimum time threshold for AHU dynamics that contribute to energy in-efficiency. This assumption may be stated mathematically as

$$\tau_{in-efficiency} \gg \tau_{dynamics}$$

Equation 5-2

where we expect the proper response time of equipment actuators $\tau_{dynamics}$ to be on the order of 10's of seconds, so that $\tau_{in-efficiencies}$ is on the order of minutes, if not longer. Practically speaking, our modeling assumption forces an instantaneous reaction time between AHU systems

[8] The limiting factor for controller scan rate is the response time of the equipment actuators; typical fast acting pneumatic actuators have a response time of 10 to 15 seconds.

[9] ASHRAE 1020 RP made extensive use of modular models of actuators and systems, but this approach has been shown to require extensive training data in order to fit parametric component models

and their disturbances, and creates a lower bound of roughly 1 minute on the building data sampling rate needed for FDD.

If we follow the assumption that the PID control response occurs instantaneously, then the outdoor air flow into the AHU under mixed air temperature control conditions follows directly from the control sequence of operations. The results can be summarized in the following set of relations, which were derived previously in chapter 3 of this thesis:

$$\dot{\forall}_{OA} = \left(\frac{T_{MA,sp} - T_{RA}}{T_{OA} - T_{RA}}\right)\dot{\forall}_{SA} \; for \; T_{OA} < T_{RA}$$

Equation 5-3

$$\dot{\forall}_{OA} = 0.3 \times \dot{\forall}_{SA} \; for \; T_{OA} \geq T_{RA}$$

Equation 5-4

$$\dot{\forall}_{SA} \times 0.3 \leq \dot{\forall}_{OA} \leq \dot{\forall}_{SA}$$

Equation 5-5

where $\dot{\forall}_{OA}$ and $\dot{\forall}_{SA}$ are the AHU's outdoor and supply air flow, respectively, measured in CFM. $T_{MA,sp}$, T_{OA}, and T_{RA} are the mixed air temperature set-point, outdoor air temperature, and return air temperature, respectively, measured in degrees Fahrenheit.

Equation 5-2 and Equation 5-3 define the outdoor air flow into the AHU for a variety of weather conditions. Equation 5-3 and Equation 5-4 follow from the physical constraint that the minimum outdoor air flow cannot decrease below 30% of the supply air flow, because of the outdoor air flow attributed to the minimum outdoor air dampers. Equation 5-4 also shows that the maximum possible outdoor air flow is when the supply air delivers 100% outdoor air.

We can use the prediction of outdoor air flow from Equation 5-2, Equation 5-3, and Equation 5-4 in Equation 5-1, along with actual measurements of the supply air flow and other temperatures to create a prediction of the thermal energy consumption of the AHU. The energy in-efficiency of the AHU can then be inferred from the comparison of its predicted and measured energy consumption. In the absence of CO_2 ventilation controls, Equation 5-2 and Equation 5-3 would suffice as a model to predict the outdoor air flow into the AHU as a function of weather; AHUs 9 and 10, however require additional modeling constraints to accommodate their CO_2 based ventilation logic.

5.2.5.3 CO_2-based Ventilation Controls

The code snippet included in the previous section shows that the CO_2-based ventilation control logic in AHUs 9 and 10 is a function of several possible CO_2 measurements. When any one of those measurements exceeds a critical concentration threshold of 800 PPM, the AHU controller will increase the outdoor flow in order to reduce all CO_2 concentration measurements to below that threshold. Variations in outdoor air flow due to this control action are difficult to model from first principles because the dynamic CO_2 concentrations in the building are a function of the space geometries, ventilation rates, and real-time occupancy. We attempted to model the impact of the CO_2-based ventilation logic on the AHU's outdoor air flow through a steady state species balance across the spaces served by the AHUs.

The largest of the spaces served by AHUs 9 and 10 is the central atrium; it is seven stories tall and has a floor seating area for roughly 50 occupants. Air-handler 9 also serves five conference rooms, the largest of which seats roughly 70 occupants in 11,000 cubic feet while the other four each seat another 25 occupants in 3,800 cubic feet. The atrium itself occupies roughly 450,000 cubic feet, half of which is ventilated by air-handler 9, and the other half by air-handler 10. The total space and maximum occupancy served by air-handler 9 is roughly 250,000 cubic feet, and 220 people, respectively; our estimates of volume and occupancy were created by inspecting the mechanical drawing set for the building, as well as onsite measurement of room geometries and occupancy.

The conference and class rooms served by air-handler 9 have a much larger ratio of occupancy to air flow rate, and we expect that the CO_2 measurements in those rooms, if anywhere, are the most likely to exceed 800 PPM. This hypothesis is supported by the CO_2 sensor placement for AHU 9; of its six CO_2 measurements, one CO_2 sensor exists in each of the five rooms served by the AHU, while the remaining sensor exists in the return air ductwork leading to the AHU's mixing box. Due to its size and low-occupancy, the atrium return air serves to blend with and dilute the CO_2 concentration of the return air derived from the conference and class rooms. As a result of that dilution we do not expect the total return air CO_2 measurement at the air-handler to often yield a significant CO_2 ventilation control signal. Likewise, since the atrium return air dilutes the total CO_2 concentration of the AHU's return air, it also helps to reduce the CO_2 concentration in the recirculation air that is delivered back to the space. The total CO_2 concentration of the air-handler supply air is a function of the ratio between recirculation and outdoor air flow rates, and their respective CO_2 concentrations.

The ventilation design engineers for MIT building 46 adhered to ASHRAE standard 62.2-2007 in order to specify the minimum fresh air ventilation rates that prevent the room air CO_2 concentrations from exceeding 1,000 PPM. According to the ASHRAE standard, about 15 CFM of fresh air per occupant should be supplied to conference rooms and classrooms in order to maintain comfortable indoor air quality; this rate is a blend of the per-person and per-square foot fresh air ventilation ratings for an adult educational facility. At that rate, the minimum fresh air ventilation for the large and small conference rooms that are served by air-handler 9 should be roughly 1,050 CFM and 375 CFM, respectively. Upon examining the mechanical design drawing set for the building as well as the individual conference rooms, we found that the design flow rates for the large and small conference rooms was actually 4,200 CFM and 1,000 CFM, respectively. The ratio of the minimum outdoor air ventilation specified by ASHRAE to the actual design flow rates of the building suggests that the minimum outdoor air fraction for air handler 9 is roughly .31; this matches our previous observation that the minimum outdoor air dampers for AHU 9 occupy about 30% of the cross-sectional area of the AHU's outdoor air intake duct, and hence they supply a minimum outdoor air fraction of roughly 0.3.

The worst-case scenario for attaining high CO_2 concentrations in the conference rooms is when all of them are filled to their occupant capacity, assuming ideal mixing with no supply air short circuiting to the room exhaust. This scenario results in the highest possible return air CO_2 concentrations, and when counting recirculation air flow, the scenario also yields the highest CO_2 concentrations for the air that is supplied back to the rooms, per quantity of outdoor air flow. We can create a steady-state model of room-level CO_2 concentrations for this worst-case scenario by performing a species balance for CO_2 around the air-handler, atrium and conference rooms, and return air plenum:

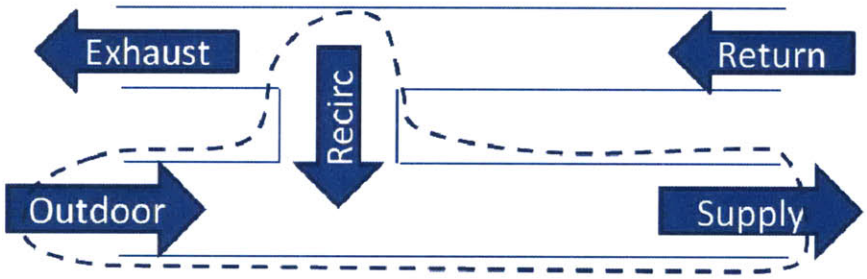

Figure 36 Control volume analysis of CO2 species balance around air-handler

Equation 5-6

The species balance around the air-handler relates the supply air CO_2 concentration to the outdoor air and recirculation air CO_2 concentrations and flow rates, respectively.

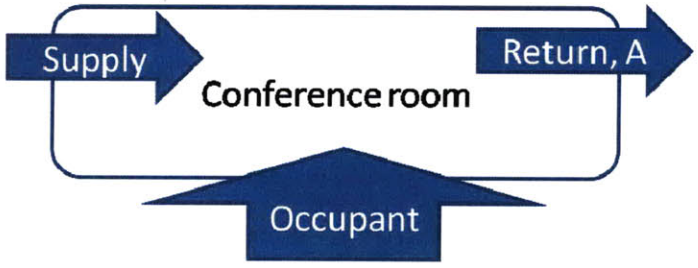

Figure 37 control volume analysis of CO2 species balance around set of conference rooms

Equation 5-7

We can lump all of the conference and class rooms together into a single element because the worst case scenario requires that the all of those rooms be filled to their occupant capacity, and all of the rooms are known to have the same design ratio of air flow to maximum occupancy. Lumping all of the conference rooms together yields a single model component that relates the supply air CO_2 concentration to the CO_2 generated by the conference room occupants and the CO_2 concentration of the air leaving those conference rooms.

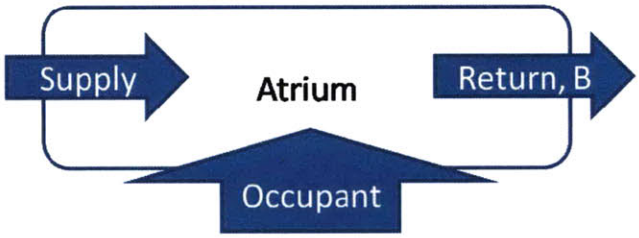

Figure 38 control volume analysis of CO2 species balance around atrium

Equation 5-8

In a similar fashion, the atrium model also considers the CO_2 balance between the supply air, return air and human occupants. In order to simplify the model and remove the influence of AHU 10 on the CO_2 concentration of the atrium we will also assume here that AHU 9 provides all of the fresh air ventilation to meet the needs of atrium occupants. This approximation can yield an upper bracket on the impact of occupant derived CO_2 in the atrium if we expect the return air CO_2 concentration from the office space served by AHU 10 to have a negligible CO_2 contribution relative to that of the conference rooms. This is a reasonable assumption because the conference rooms served by AHU 9 have a smaller ratio of outdoor air flow to occupants than will the office space served by AHU 10, when the conference rooms are at their maximum occupancy.

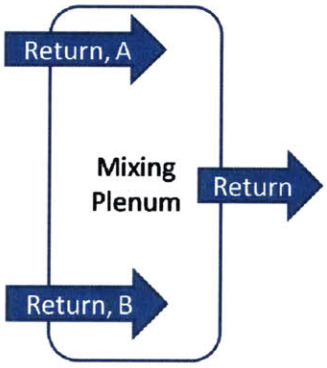

Figure 39 control volume around mixing box plenum for atrium and conference room returns

Equation 5-9

The final component of the model is the plenum return where the return air from the conference rooms is blended with the return air from the atrium to yield a total CO_2 concentration for the air that returns to the AHU.

The goal of this analysis is to identify combinations of supply and outdoor air flow rates that yield CO_2 concentrations in the conference rooms that exceed 800 PPM under the worst case conditions of saturated room occupancy and room-level design flow rates (8200 CFM total). The set of four equations from the model sub-components is cast with four unknown variables, Csa, Cra, Cra1, Cra2, or the supply, return, conference room return and atrium return air CO_2 concentrations, respectively.

The other parameters in the model include constants or derived values; we will assume that the CO_2 output of a sedentary human, , is 0.00706 CFM, or roughly 200 mili-liters per minute, and that the background outdoor air CO_2 concentration is 380 ppm. Additional assumptions include that the supply and return air flow to the conference rooms are equal to the sum of the design air flows for all of the conference rooms, 8,200 CFM, and that the flow of air to the atrium is equal to the total supply air flow rate minus the supply air flow rate to the conference rooms. Combining these assumptions, parameters and equations yields the following

linear system of equations that we can solve to find conference room CO_2 concentration as a function of total supply and outdoor air flow rate:

$$\dot{\forall}_{oa}c_{oa} + (\dot{\forall}_{sa} - \dot{\forall}_{oa})c_{ra} = \dot{\forall}_{sa}c_{sa}$$

Equation 5-10

$$\dot{\forall}_{sa1}c_{ra1} = \dot{\forall}_{sa1}c_{sa} + n_p c_{pp}$$

Equation 5-11

$$(\dot{\forall}_{sa} - \dot{\forall}_{sa1})c_{sa} + n_{peaople,a}c_{pp} = (\dot{\forall}_{sa} - \dot{\forall}_{sa1})c_{ra2}$$

Equation 5-12

$$\dot{\forall}_{sa1}c_{ra1} + (\dot{\forall}_{sa} - \dot{\forall}_{sa1})c_{ra2} = \dot{\forall}_{sa}c_{ra}$$

Equation 5-13

Re-arranging these linear equations into matrix form allows us to solve the set of equations for any given pair of outdoor and supply air flows. Solving these equations for a range of supply air flow rates, 15,000 to 50,000 and outdoor air fractions of 0.15 to 1.0 yields a set of surface contours that define the worst-case CO_2 concentration in the conference rooms for that ordered pair:

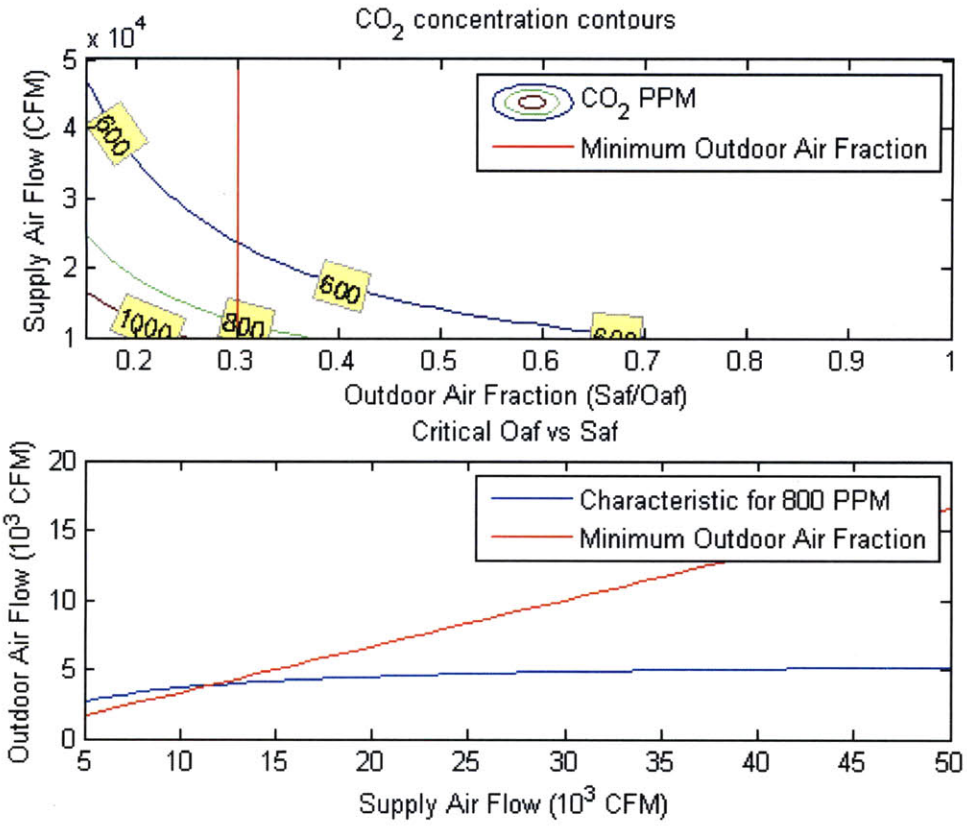

Figure 40 (upper graph) CO_2 concentration contours per supply air flow and outdoor air fraction, and (lower graph) critical curve between supply and outdoor air flow rate for a CO_2 concentration of 800 PPM. In both graphs a red-line has been included to show the minimum outdoor air flow that is possible for the AHU

The upper graph in Figure 40 shows several contours for conference room CO_2 concentrations under the worst case conditions prescribed in the model, and for various combinations of AHU outdoor air fraction and supply air flow. The outdoor air fraction was used as the ordinate in this graph in order to facilitate easier comparison against the minimum outdoor air fraction for the AHU. The red line represents the minimum outdoor air fraction for the AHU, and bisects the space into operating points that are possible and impossible; operating points to the left of the line are impossible because the minimum outdoor air dampers physically prevent the AHU's outdoor air fraction from decreasing below 0.3. The graph shows a small space of operating points for outdoor air fractions between 0.3 and 0.37 and supply air flows between 10,000 and 12,000 CFM where under the worst case conditions of this model, the conference room air CO_2 concentration may equal or exceed 800 PPM. The space of possible points that satisfies this constraint grows larger as the supply air flow decreases, however it is not likely that

the supply air flow would decrease far below 10,000 CFM given that the supply air fan has name-plate capacity of 50,000 CFM and the conference rooms receive 8200 CFM. Even in cases where the modeling assumptions made around the atrium and elsewhere were less accurate, it appears that only a small space of typical operating conditions leads to ventilation control based on CO_2 concentrations. This may imply that the CO_2 based ventilation controls were intended more as a safety feature in case laboratory animals were transported through the atrium or other parts of the building.

The lower graph in Figure 40 shows the same information as the upper graph, but in terms of outdoor and supply air flow rates instead of outdoor air fraction and supply air flow rate; the blue curve represents the critical relationship between outdoor and supply air flow rates that bisects the space of operating points according to whether they yield conference room air CO_2 concentrations that exceed 800 PPM, or not. AHU operating points that exist on the blue line yield conference room air CO_2 concentrations that equal 800 PPM, while operating points above the blue line yield concentrations that are less than 800 PPM. The red line in that graph represents the minimum outdoor air fraction for the AHU; operating points below the red line are physically impossible due to the construction of the AHU. Once again we can see that only a small region of operating points satisfies the constraint on minimum outdoor air fraction and exist below the critical threshold to yield undesirable conference room air CO_2 concentrations.

The blue line in the lower graph of figure 40 can be analytically determined by setting the conference room CO_2 concentration, Cra1, to 800 ppm and solving the set of equations for the outside air flow as a function of the supply air flow:

$$\dot{\forall}_{OA} = \frac{5775 * \dot{\forall}_{SA}}{5775 + \dot{\forall}_{SA}}$$

Equation 5-14

For any given supply air flow rate, Equation 5-14 yields the critical outdoor air flow rate at which the conference room CO_2 concentration reaches 800 ppm; below this critical outdoor air flow rate the conference room CO_2 concentration will increase steadily beyond 800 PPM or even

grow without bound. For example, if all of the conference rooms are occupied at full capacity and the supply air flow is slightly less than half of its maximum rating at 20,000 CFM, then the outdoor flow rate must be at least 4,481 CFM, or 22% outdoor air fraction in order to maintain conference room CO_2 concentrations below 800 PPM. If the conference rooms are not fully occupied, or the supply air flow rate is less than its full capacity, then the outdoor air flow rate can be less than 4,481 CFM and still maintain conference room CO_2 levels below 800 PPM.

The documented sequence of operations for AHUs 9 and 10 specifies that the minimum outdoor air flow rate for periods where the CO_2 concentration signal achieves 800 PPM shall be 20,000 CFM, and that the outdoor air flow rate is supposed to increase from that value proportionally as the CO_2 concentration signal exceeds 800 PPM. Our investigation into the AHU controller programming, however, shows that the CO_2 control logic cannot directly influence the outdoor air flow rate but instead can only drive the outside air dampers to a more fully opened position. This means that at best the CO_2 control logic could drive the outdoor air fraction to unity and cause the AHU to supply the building with 100% outdoor air. While the implemented CO_2 control logic does not necessarily satisfy the claims written in the sequence of operations, our worst-case scenario model suggests that those written requirements are in fact excessive and unnecessary. Our model shows that even with low supply air flow rates there exists only a small space of supply and outdoor air operating points that can trigger the CO_2 control logic. Even then, that space of points only exists under worst case conditions and they can often be avoided without using more than 10,000 CFM of outdoor air flow. In the event that those conditions are satisfied, however, we have also generated a mathematical model of room air CO_2 concentration to help reform our predictions of outdoor air fraction under those circumstances.

5.2.5.4 Summary of Model for AHU Energy Measurement and Prediction

The thermal energy consumption of our experimental VAV AHUs is given by Equation 5-1, which is reproduced here for convenience of presentation:

$$TE_{AHU} = \dot{m}_{OA}(h_{OA} - h_{RC}) + \dot{m}_{SA}(h_{RC} - h_{SA})$$

We can estimate the actual thermal energy consumption of a VAV AHU by evaluating Equation 5-1 with measurements of outdoor and supply air mass flow rates, and their corresponding outdoor, return and supply air enthalpies. We can also use the control logic, and weather and building conditions to predict the ideal AHU outdoor air flow rate. The mixed air temperature controls and CO_2-based ventilation controls yield the following set of relationships to predict the ideal outdoor air flow rate for AHU 9 in MIT building 46:

$$\dot{\forall}_{OA} = \left(\frac{T_{MA,sp} - T_{RA}}{T_{OA} - T_{RA}}\right)\dot{\forall}_{SA} \; for \; T_{OA} < T_{RA}$$

Equation 5-15

$$\dot{\forall}_{OA} = 0.3 \times \dot{\forall}_{SA} \; for \; T_{OA} \geq T_{RA}$$

Equation 5-16

$$\dot{\forall}_{SA} \times 0.3 \leq \dot{\forall}_{OA} \leq \dot{\forall}_{SA}$$

Equation 5-17

$$\dot{\forall}_{OA,crit} = \frac{5775 * \dot{\forall}_{SA}}{5775 + \dot{\forall}_{SA}}$$

Equation 5-18

$$if \; \dot{\forall}_{OA} < \dot{\forall}_{OA,crit}, then \; \dot{\forall}_{OA} = \dot{\forall}_{OA,crit}$$

Equation 5-19

Based on the model derived in the ventilation control section, we expect these relationships to hold over the majority of the AHU's operation. Equation 5-18 defines a critical relationship between outdoor and supply air flow for AHU 9 that demarcates operating conditions which under worst case conditions could lead to air CO_2 concentrations that exceed 800 PPM. Equation 5-19 integrates that critical relationship into the overall model of outdoor air flow for the AHU simply by requiring that the AHU's outdoor air flow never decrease below the critical outdoor air flow. This approach to combining the ventilation models rests on the assumption that the instantaneous PID loop response to the CO_2 ventilation signal would prevent

the outdoor air flow from ever decreasing below the critical outdoor air flow. We also assume that the control response to the CO_2 ventilation signal will not force the outdoor air to flow be very different from what would otherwise be predicted by the mixed air temperature control. These assumptions are based on the conclusions drawn from the worst case scenario model which suggested that the CO_2-based ventilation control would only have a significant presence at very low outdoor air fractions and supply air flow rates. Even in those situations, the model suggests that that increased demand for outdoor air flow in order to satisfy CO_2 conditions would not far exceed the demands for outdoor air flow placed by the mixed air temperature control.

Having combined the ventilation and mixed air temperature controls to yield a uniform prediction of AHU outdoor air flow, we can now use equation 1 to yield an estimate of the ideal thermal energy consumption of AHU 9.

5.3 FDD System

The detection of faults in AHUs 9 and 10 follow from the model describing outdoor air flow of the AHUs combined with the uncertainty analysis presented in the previous chapter. The Taylor expansion for propagating experimental uncertainty was used with the AHU models to estimate parameters for the distribution of AHU thermal energy consumption.

Energy consumption parameter distributions estimates were made on both the predicted and measured operation of the AHUs. In addition to statistical tests on the homogeneity of energy measurement and prediction populations, we also created a series of expert rules to detect and diagnose the presence of equipment faults within those populations.

5.3.1 FDD Expert Rules

We used FDD expert rules to target the detection and diagnosis of two principle AHU pathologies in our experiments: simultaneous heating and cooling, and imbalanced air flows. We

chose those faults because of their known prevalence within the commercial building stock, and their typically large impact on energy consumption and indoor air quality.

We evaluated the presence of simultaneous heating and cooling by computing the average probability that the heating and cooling coils may both be active at the same time:

$$P_{SimHC} = \langle P\left(\frac{\dot{m}_{SA}\Delta h_{hc}}{\dot{Q}_{HC}} \geq 0.03\right), P\left(\frac{\dot{m}_{SA}\Delta h_{cc}}{\dot{Q}_{CC}} \geq 0.03\right)\rangle$$

Equation 5-20

where $\dot{Q}_{HC}$ and $\dot{Q}_{CC}$ refer to the capacity of the heating and cooling coils, respectively. The probability of either the heating or cooling coil being active is given by the probability that the heating or cooling enthalpy exchange for the AHU, when scaled by its respective coil capacity, exceeds 0.03. The threshold of 0.03 was chosen based on our experience that 3% of a heating or cooling coil's capacity defines the start of meaningful heat transfer between the working fluids of a heat exchanger. The probability computations use the mean and standard deviation of the population that are defined in chapter 4 according to the Taylor expansion for propagating experimental uncertainty, as well as the underlying assumption that all experimental populations follow a normal distribution. Assuming that the probabilities of coil operation are independent, then the mean product of those probabilities is found for an arbitrary evaluation time series, and reported as a gross indicator of whether simultaneous heating and cooling occurred in that time interval. The advantage of our approach to FDD expert rules relative to the prior art is that it embraces and propagates the uncertainty of measurements and models, and automatically scales to the size of the AHU equipment.

In order to detect air flow directions that do not match the expected orientations of intake and exhaust on the equipment, we estimated the AHU supply air flow rate from the other three air flow measurements made about the AHU, and then compared that estimate of the supply air flow rate against its actual measurement value; the probability of flow reversals was calculated in a similar fashion to how we calculated the probability of heating and cooling coil activity.

$$\dot{\forall}_{SA}^{*} = \dot{\forall}_{RA} - \dot{\forall}_{EA} + \dot{\forall}_{OA}$$

Equation 5-21

$$P_{Imbalance} = \langle P\left(\left|\frac{\dot{\forall}_{SA} - \dot{\forall}_{SA}^{*}}{\dot{\forall}_{SA,max}}\right| \geq 0.05\right)\rangle$$

Equation 5-22

The supply air flow rate estimated from the other flow measurements made about the AHU is defined as $\dot{\forall}_{SA}^{*}$, and is measured in CFM; the nameplate rating of the supply air flow is designated as $\dot{\forall}_{SA,max}$ and is also measured in CFM. The threshold of 0.05 was chosen based on our expert experience that deviations between measured and expected supply air flow rates are meaningful when they exceed 5% of the rated capacity of the supply fan. Once again, this probabilistic form of expert rules is advantageous because it automatically scales to the size of the AHU.

In both cases of expert rules, the result is always a number between 0 and 1, which provides a convenient standard for communicating the urgency of a fault. For example we could raise fault alarms for all faults whose probability of existence exceeds 75%, or some other user defined criteria. Within this framework, text-based inference or diagnosis of fault signals can also be based on the *strength* or magnitude of the probability estimate relative to unity. Beyond text-based inference, the interpretation of these expert rules is also uniform; the numerical result is always the average probability that a condition is satisfied or exceeded over the prescribed time interval.

5.3.2 FDD Financial Estimates

Our approach to assigning financial value to fault signals is based on user-defined parameters for financial investments. Using appropriate conversion factors between energy values and energy unit-costs, we can compute the value of discrepancies between the predictions and measurements of AHU energy consumption, and also the probability that the value of those discrepancies exceeds user defined investment goals:

Equation 5-23

where is the user-defined financial criteria, is the distribution of measured energy consumption, is the distribution of predicted energy consumption, and is the critical energy quantity that corresponds to the user defined financial criteria. Using this probabilistic format we can communicate the value of divergent AHU behavior in terms of the odds that correcting that behavior can yield a certain financial return.

5.3.3 Fault Experiment Program

Several different AHU faults were planned and executed on AHU 9 in MIT building 46 during our research. Hardware faults were selected for our experiments based on their ease of implementation, ability to not cause irreparable damage to the AHU, low risk of generating occupant hazards, and their prevalence in buildings.

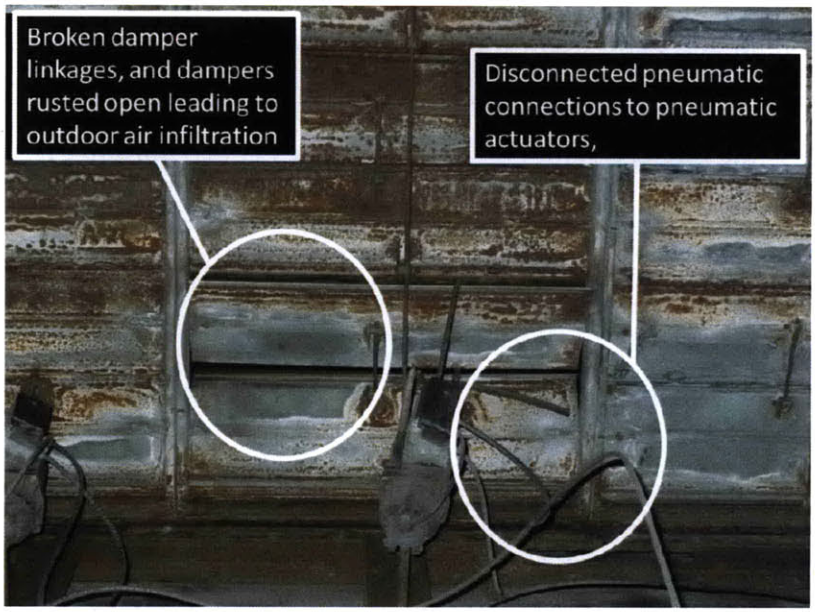

Figure 41 Picture of an AHU mixing box that is suffering from broken damper linkages, rusted open dampers, and disconnected pneumatic actuators. The picture was taken during a tour of 14 AHUs across the MIT campus; we observed similar pathologies in a number of AHUs across the campus.

The practical constraints of imposing mechanical faults on an otherwise properly functioning AHU led us to focus on mixing box mechanical faults. In particular, the MIT facilities department was comfortable with our tampering with the damper linkages and actuators for the outdoor and recirculation dampers in the mixing box of AHU 9 in MIT building 46. We were also interested in tampering with the AHU's mixed air temperature sensor and heating coil valve in order to simulate the effects of other common AHU pathologies, but MIT was not as comfortable with our imposing those types of mechanical faults on their equipment.

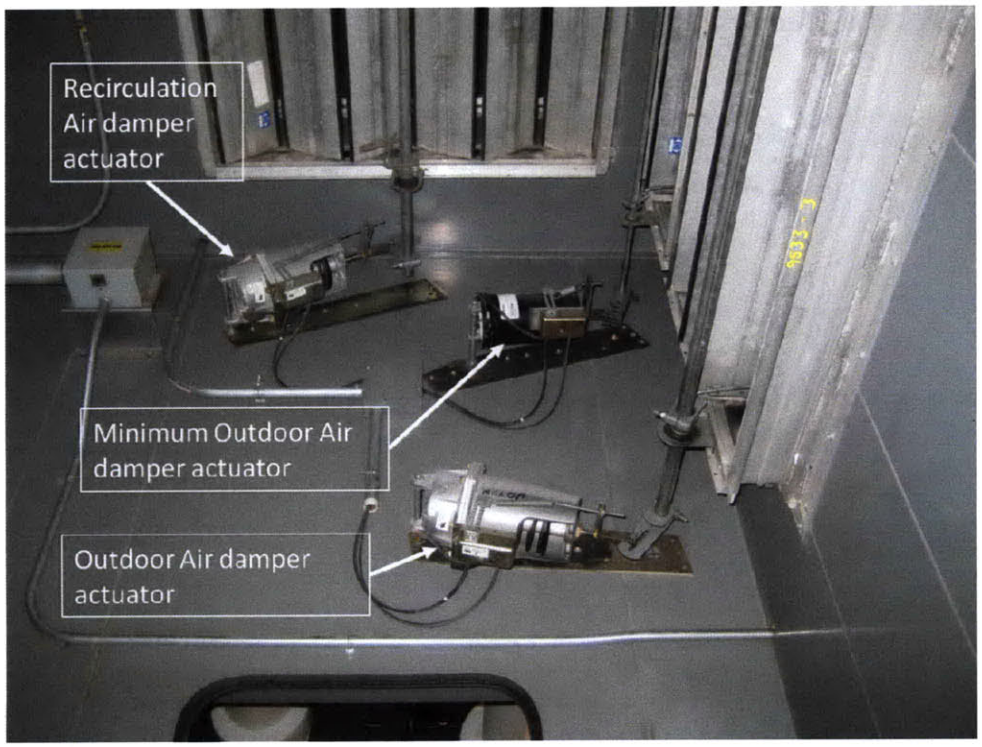

Figure 42 Interior picture of the mixing box for AHU 9, showing the outdoor and recirculation air dampers, and their corresponding actuators, and actuator linkages.

A wide variety of AHU mixing box pathologies were represented in our research by disengaging the damper linkages between the outdoor and recirculation air dampers and actuators, and the pneumatic supply lines to those same damper actuators. Broken damper linkages and malfunctioning pneumatic actuators can result from a number of AHU issues, ranging from ageing and natural degradation of equipment, to spontaneous mechanical failures or accidents during routine maintenance.

By removing the pneumatic supply pressure to the damper actuators, the recirculation and outdoor air dampers defaulted to their fully open and closed positions, respectively. This type of reaction is representative of mechanical faults that yield dampers stuck in certain positions; the potential impact of such faults was explored earlier in chapter 3. The outdoor and recirculation dampers also possessed multiple sets of linkages to their respective actuators; by incrementally removing those linkages we were able to test different levels of fault severity for leaking and broken dampers. Once the damper linkages were removed, the dampers could also be manually positioned to any particular starting orientation; over time, however, the orientation of free floating dampers was a function of the air flow over them.

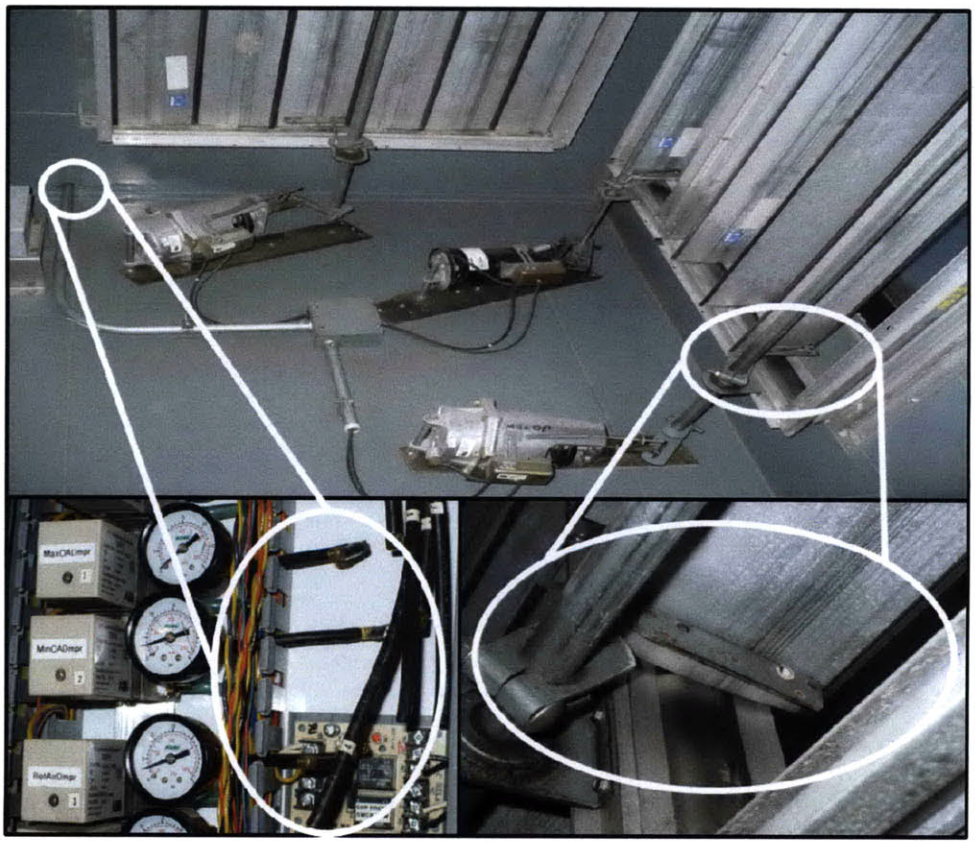

Figure 43 Cascade of mixing box images and close-up pictures of fault implementations; the lower left picture shows the disconection and capping of the pneumatic supply to the outdoor and recirculation dampers at the pneumatic supply manifold for the AHU. The lower right picture shows the removal of a linkage between the outdoor air damper and its actuator.

To compensate for our limited ability to apply diverse hardware faults on the AHU, we also altered the control system software programming to simulate pathologies that could not be

rendered physically. In particular, we altered the damper control programming to eliminate the mixed air temperature control and CO_2-based ventilation control; the changes in software led to periods of heating and cooling that would have otherwise not occurred.

```
1       If CO2Mode = On then
2         OutsideAirFlo_Se = Ratio.Fn(CO2Signal, 0, 1, OutsideAirFloMin, SupAirFlow)
3         If TrainingSwitch = On then
4           If (Timeofday >= 7:00am & Timeofday < 1:00pm) then
5             MixAirDmpSignal = (1 / 35) * ((Hour - 7) * (60 / Mint) + floor(Minute / Mint))
6           Endif
7           If (Timeofday >= 1:00pm & Timeofday < 7:00pm) then
8             MixAirDmpSignal = 1 - (1 / 35) * ((Hour - 13) * (60 / Mint) + floor(Minute / Mint))
9           Endif
10          If (Timeofday < 7:00am ! Timeofday >= 7:00pm) then
11            MixAirDmpSignal = maximum(MixAirTeSignal, OutsideAirFloSig)
12          Endif
13        Else
14          MixAirDmpSignal = maximum(MixAirTeSignal, OutsideAirFloSig)
15        Endif
16      Else
```

[10]

The software code changes were made to the same snippet of code that was shown earlier in this chapter. Instead of allowing the dampers to be controlled by the MixAirTeSignal or OutsideAirFloSig as indicated in the original code, we modified the program to cycle the damper positions as a function of time for a certain period of the day. In particular, our software modification forced the MixAirDmpSignal to follow a linear ramp from 0 to 1, in 15 minute increments, starting at 0 at 7:00 am and ending at 1 at 1:00 pm. Following that positive ramp, we also forced the MixAirDmpSignal to ramp back down to 0 in the same 15 increments, starting at 1 at 1:00 pm, and ending at 0 at 7:00 pm. Outside of the hours of 7:00 am to 7:00 pm, the MixAirDmpSignal was allowed to follow its nominal control programming. In addition to causing unnecessary periods of heating and cooling (which was desired as a fault characteristic), the new code also facilitated our commissioning of damper actuation and air flow measurements.

[10] Code snippet taken from a program entitled "AHMisc.Pr" on the controller for AHU 9 in MIT building 46

These hardware and software faults were applied intermittently over a period of three months, from January 2010 through March 2010. The FDD system was operational over that time period and was used to detect and evaluate the impact of those faults. The results of the FDD system were presented to several different members of MIT's repair and maintenance and project engineering staff in order to ascertain how MIT could actually use our FDD system to help improve the energy efficiency of buildings across the campus.

5.4 AHUs 2 and 3 in MIT Building 56, and AHU 2 in MIT Building 16

The other three AHUs involved in our experiment belonged to MIT buildings 16 and 56; these two buildings are adjoining and to an onlooker unfamiliar with MIT's buildings, they appear as one large building. Together, the buildings include seven above-ground floors with approximately 100,000 square feet of laboratory, teaching and office spaces. The bottom two floors of the building include classrooms and computer rooms, while the remaining above-ground floors are replete with biological laboratories and office space. The basement and sub-basement levels of the building include a chemical hardware stock room, offices, and mechanical rooms.

5.4.1 Physical Characteristics

AHU 2 in MIT building 16 is one of the largest AHUs on campus, with a rated supply air capacity of 100,000 CFM. AHUs 2 and 3 in MIT building 56 are much smaller in scale, each with a rated supply air capacity of 10,000 CFM. All three of these AHUs are dedicated outdoor air systems, meaning that they only condition and supply 100% outdoor air to the building.

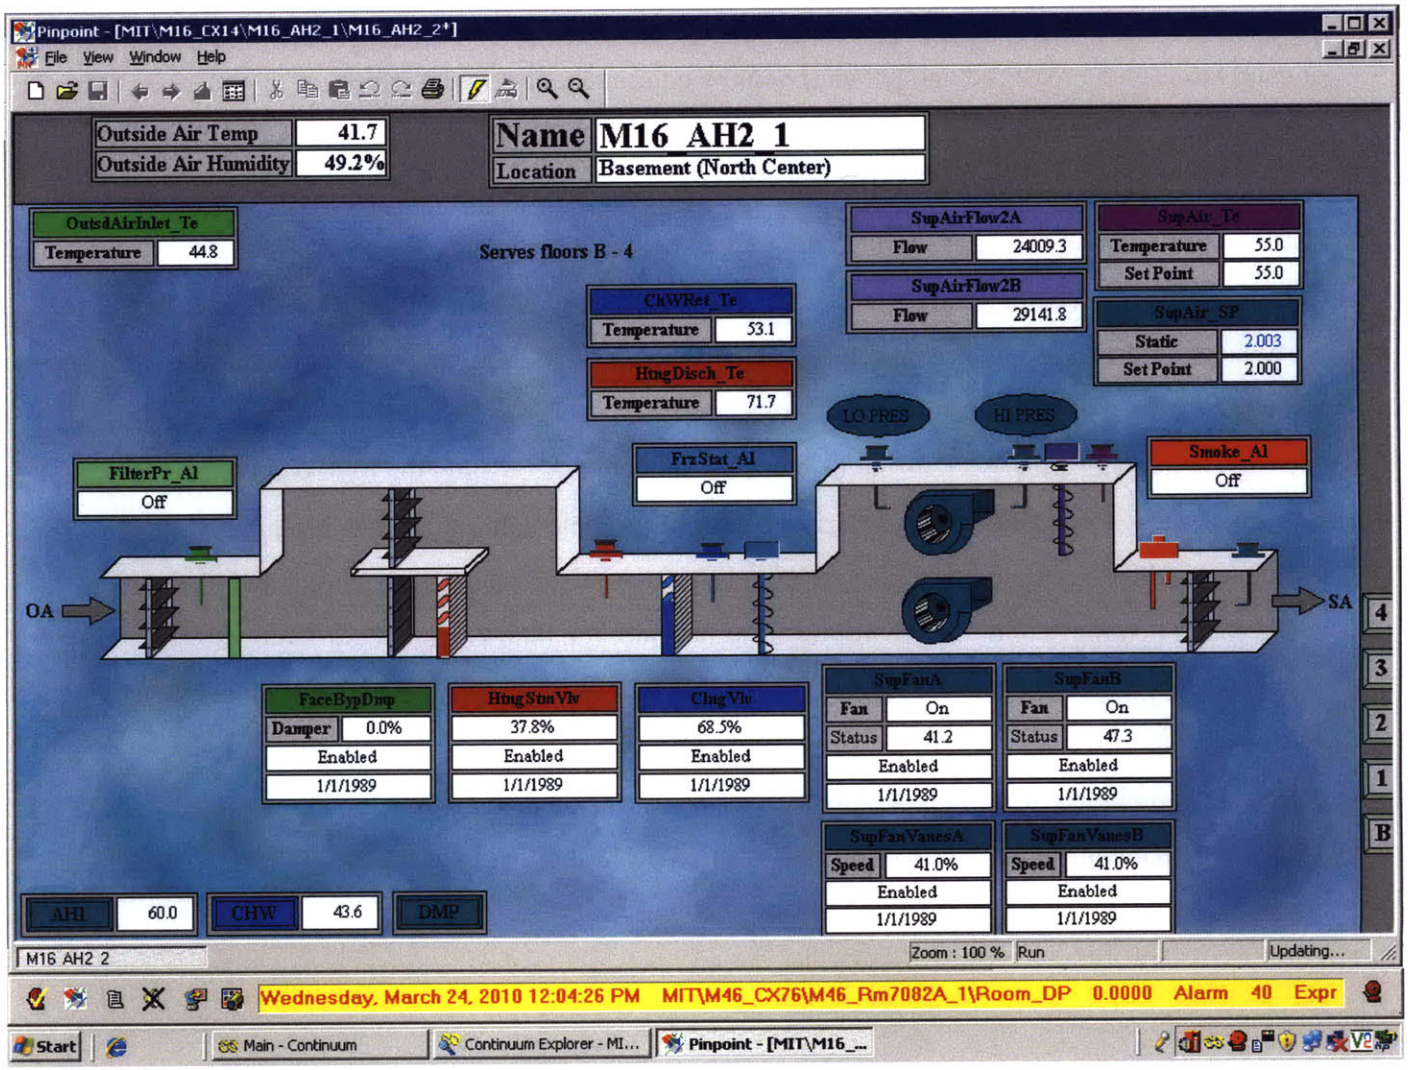

Figure 44 Schematic diagram of AHU 2 in MIT building 16; the image is a screenshot from the building control system interface to the AHU

A schematic diagram of AHU 2 in MIT building 16 is shown in Figure 44; the image in the figure shows that the AHU has two supply air fans, and a unique heating coil arrangement that involves coil-face bypass dampers. The supply air fans in this particular AHU have a constant rotational speed however the discharge air flow rate from the AHU can be modulated by the orientation of directional vanes; the vanes are not shown on the schematic in Figure 44, however they exist at the entrance to the fan-portion of the AHU, and they control the pressure drop of air entering the fans. This particular method of modulating air-flow is less energy efficient than using variable frequency drives, which are used in the AHUs in MIT building 46. The percent turn down of the control vanes for both supply air fans can be read directly from the

image of the equipment control interface that is in Figure 44. Likewise, the supply air flow rate of each fan is also measured and presented in the equipment control interface.

The heating coil for this AHU design integrates a face-bypass damper system with steam heating coils in order to better regulate how much heat is injected into the outdoor air intake stream. While the schematic in Figure 44 shows the bypass dampers and steam heating coils as portions of separated ductwork, in reality the two pieces of equipment are integrated within the same flow path. The bypass dampers articulate between vertical, cylindrical-finned steam heating elements in a clam-shell-like fashion; by opening and closing, the bypass dampers alter how much air is forced across a cylindrical steam heating element rather than around it. Fortunately, the building control system includes a temperature sensor at the inlet and discharge of the steam heating section which allows us to measure the total change in air temperature across the heating-portion of the AHU. The cooling coil of the AHU is a conventional chilled water coil like the ones encountered in MIT building 46. The discharge temperature of air from both coils is included in the equipment interface, as well as the commanded percent opening of their respective control valves, and the commanded percent opening of the steam heat bypass dampers.

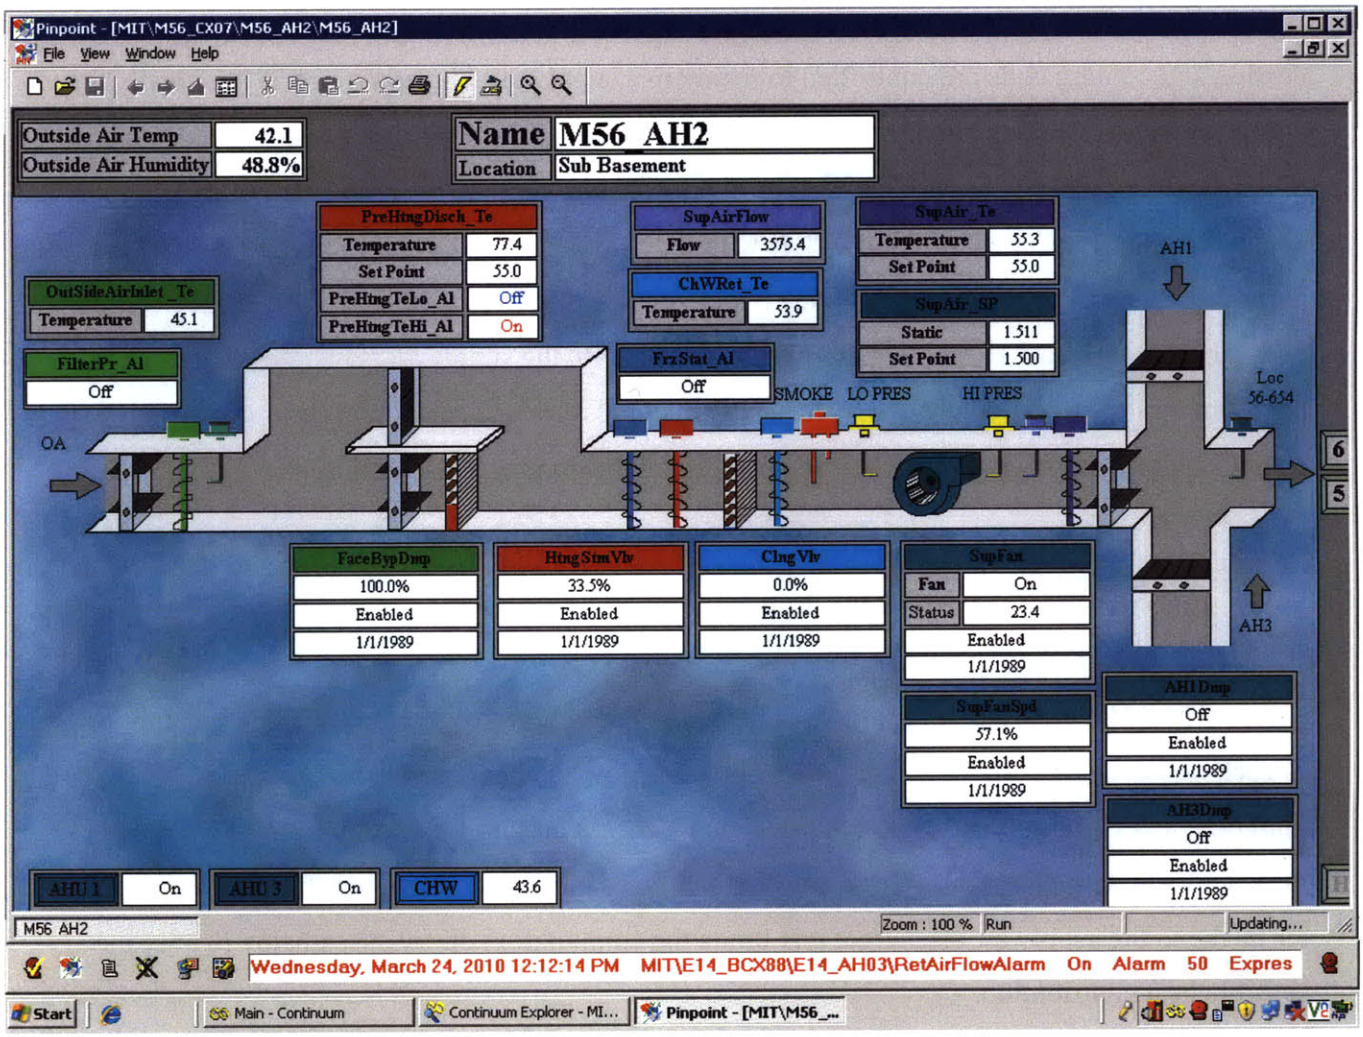

Figure 45 Schematic diagram of AHU 2 in MIT building 56; the image is a screenshot from the building control system interface to the AHU

AHUs 2 and 3 in MIT building 56 are of identical designs, and a schematic of their equipment is shown in the control interface screenshot in Figure 45. Similar to AHU 2 in MIT building 16, AHUs 2 and 3 in MIT building 56 have a set of coil-face bypass dampers that are integrated with their steam heating coil. Unlike the former AHU however, AHUs 2 and 3 in MIT building 56 only have one supply fan each, whose rotational rate is controlled by a variable frequency drive.

5.4.2 Relevant Controls Programming

The relevant controls programming for these three test AHUs is straightforward; the heating and cooling coils are controlled to maintain a supply air temperature of 55 °F. Ideally this would translate to either minimum heating or cooling energy expenditure in order to modulate the raw outdoor air temperature to meet the supply air discharge temperature set-point. The heating enthalpy exchange of these AHUs with their outdoor air flows is controlled both by the percent opening of their heating coil face-bypass dampers, as well as the position of the heating coil valve. Similar to AHUs 9 and 10 in MIT building 46, the supply air flow rates of these additional test AHUs is controlled by the static duct pressure and static duct pressure set-point for the equipment; this is the conventional control algorithm for VAV AHUs.

5.4.3 Instrumentation and Data Acquisition

The AHU measurements that are relevant to computing the thermal energy consumption of each of these AHUs include their total supply air flow rate, outdoor air temperature and humidity, and coil discharge temperatures. The discharge air humidity for these AHUs is not measured which implies that we cannot compute the change in water content from the outdoor air intake to the supply air discharge; this limits the computation of dry air enthalpy to only the sensible heat contribution.

All of the data used for performing FDD on these three AHUs was ultimately derived from the Schneider Electric building control system, however, a historical data set from 2009 was included in our analysis as well as current data from our experimental research period, January 2010 through March 2010. The historical data set was provided by Cimetrics Inc., who collected the data in 2009 from the Schneider Electric building control system in order to provide MIT with FDD services on those same AHUs. Our data collection during the research period was performed in the same exact fashion as was executed for AHUs 9 and 10 in MIT building 46.

5.4.4 Commissioning

Commissioning of these AHUs largely followed from examination of the FDD reports that Cimetrics Inc. had provided to MIT in 2009 during their analysis and inspection of the AHUs. Our discussion with MIT facilities revealed that no new work had been performed on those AHUs following the FDD work done by Cimetrics, and we therefore assumed that the instrumentation and equipment was already qualified for use in our analysis. The historical data collected from Cimetrics about those AHUs was identically used in Cimetrics' analysis and hence we assumed that it was suitable for our use as well. To be certain of their current state, however, temperature and air flow data collected from those AHUs during our experimental period was manually inspected for any gross anomalies or inaccuracies.

5.4.5 AHU Models

The thermal energy models for these three AHUs are nearly identical and straightforward to derive from a control volume analysis since the AHUs are dedicated outdoor air systems (DOAs):

$$TE_{DOAs\ AHU} = \dot{m}_{SA}(h_{SA} - h_{OA})$$

Equation 5-24

where $\dot{m}_{SA}$ is the total supply air flow rate, measured in CFM, and $h_{SA}$and h_{OA} are the enthalpy of the supply and outdoor air conditions, respectively, measured in BTUs per pound of dry air. The only difference amongst these three AHUs is that AHU 2 in MIT building 16 has two supply air fans whose independent supply air flow rates must be summed to yield the total AHU supply air flow; the other two AHUs each have a single supply air fan whose flow is identically the AHU's supply air flow rate.

5.4.6 FDD System

The specific FDD goal for these three AHUs was to identify and evaluate periods in which they exhibited simultaneous heating and cooling. This target was chosen because Cimetrics had already identified and evaluated the extent of simultaneous heating and cooling

that was occurring in these AHUs in 2009, and we wanted to verify that our FDD system could at least reproduce the results of Cimetrics' analysis using their historical data. Since MIT facilities had not yet corrected the simultaneous heating and cooling that was known to exist in those AHUs, we also wanted to make sure that our FDD system could still identify and evaluate those same problems with data drawn during our experimental period. Finally, in order to test the ability of our expert rules to generalize across equipment types, we applied the exact same expert rule used to identify simultaneous heating and cooling in AHU 9 in MIT building 46 to the test AHUs in MIT buildings 16 and 56. The financial implications of simultaneous heating and cooling in these three test AHUs was also evaluated in the same fashion as it was for the other test AHU.

5.5 Summary of Experimental Setup

We tested our FDD system on several different faults in five separate AHUs that represent three different AHU equipment designs. Our specific FDD expert rules and financial analysis are cast in a probabilistic and dimensionless framework, which, despite their reliance on an expert's choice of thresholds, may still generalize across equipment typologies and size. Prominent equipment faults such as simultaneous heating and cooling were tested in both current and historical data sets, portions of which were already labeled with existing fault diagnoses. The results of these experiments are included in the next chapter and address the efficacy of our fault detection system both from an absolute perspective for faults that we purposefully implemented on equipment, and a relative perspective for faults that had already been identified by another FDD service.

6 Results

6.1 Introduction:

Experimental results are presented in this chapter as screenshots of the output from the FD&E system that was developed to identify faults on the MIT campus. The results are presented in this fashion in order to underscore the research effort put into designing how technical information is communicated through the software's user-interface to less technical users of that software. Presenting the results as part of the FD&E system output also facilitates a discussion on how technical features of the analysis correspond to specific elements in the user-interface. Feedback from MIT facilities personnel on the software interface and design features is also included at the end of this chapter.

6.2 FD&E Software

6.2.1 Software architecture and setup at MIT

The architecture of the FD&E software developed in this thesis reflects the interactive components that are needed in order to collect, store, analyze and report on building data; figure 1 illustrates those components and their integration into the software architecture.

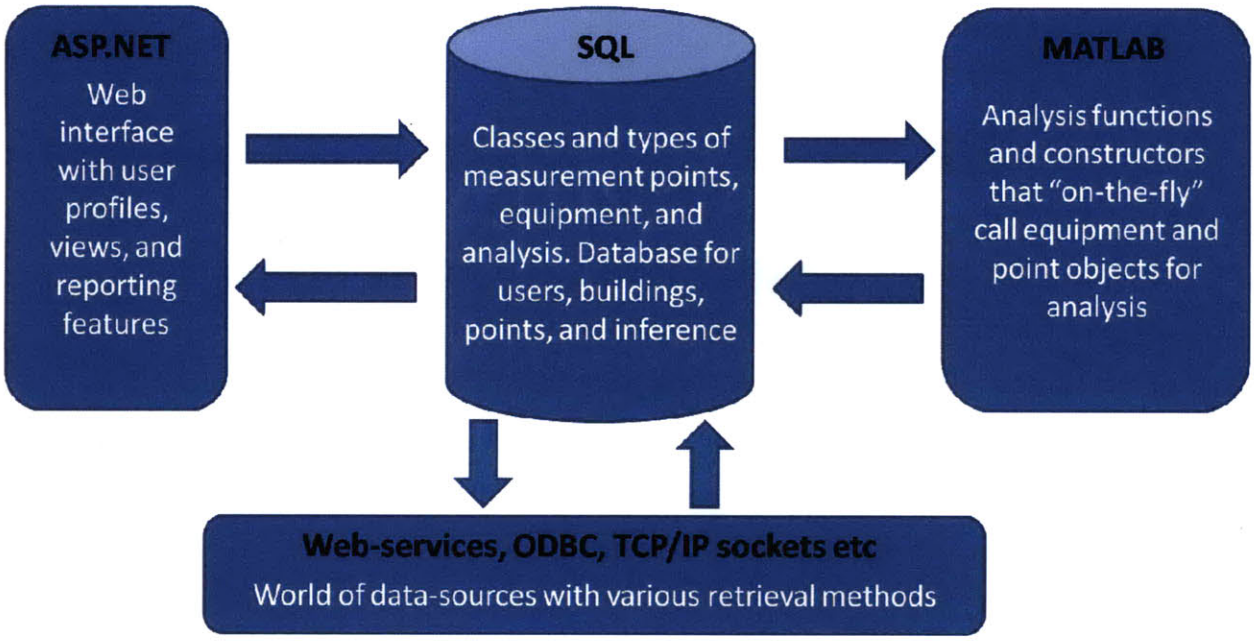

Figure 46 Schematic of the software architecture that underpins the FD&E system developed in this thesis

A class library was created to manage data acquisition from diverse data sources at multiple levels of the OSI network model of communications (69), including ODBC data source objects, web-services, TCP/IP sockets, and web-page scraping. A large database written in SQL was used to store data collected from the network and coordinate interaction between the other elements in the architecture. Additional class libraries were written in MATLAB to define equipment and point-measurement objects, as well as constructors for assembling equipment analysis. Finally, the user-interface was written in ASP.net in order to provide a fully web-based application.

Setting up this system to run at MIT required several additional pieces of information:

- Unit costs for heating and cooling energy
- Targeted dollar value of annual pay-back on correcting faults detected
- Roster of equipment and points to be monitored

The previous chapter explained in greater detail the efforts taken to model and commission equipment and measurement points once they were chosen for monitoring. MIT facilities provided blended heating and cooling costs of $1.50 per therm of heating and $0.13 per TON-hour of cooling; these values reflect the extensive co-generation and multi-plant utility

system that exists on the campus. Conversations with MIT facilities personnel also suggested that the software target the detection of faults that if corrected, could yield an annual savings of $10,000 or more.

6.2.2 Software interface

The FD&E software interface employs a variety of features that communicate the results of data analysis according to user defined fault detection settings. One approach is a zero-to-ten ranking system, or energy prioritization, of the results of energy inefficiency analyses. The zero-to-ten ranking is calculated for energy inefficiency analyses by multiplying the number 10 and the probability that the dollar value of any detected inefficiencies exceeds the annual payback target. For analyses that do not forecast losses over the entire year but instead compute the dollar value of detected inefficiencies over a specific time interval, the energy priority of those results is found by multiplying ten and the probability that the dollar value of inefficiencies in that time interval equals or exceeds the annual target, proportionally reduced to that time interval. For example, an analysis may find that there is a 20% probability that an AHU wastes $5,000 or more over a 4 week period due to simultaneous heating and cooling. The targeted annual payback of correcting faults detected is $10,000, but over a 4-week period that target is proportionally reduced to roughly $770. In this case, the probability that the dollar value of the identified AHU inefficiency over a 4 week period exceeds $770 is very close to unity; multiplying the probability by ten yields an energy priority of roughly 10; the highest possible priority. A possible future improvement to the prioritization of energy loss estimates over arbitrary time intervals is to extrapolate the results of that time interval over an entire year via regression on weather parameters such as degree days or outdoor air temperature. Not all energy analysis may lend themselves to such extrapolation and so the simpler approach described above was taken in the current software development.

But there are many more things that can go wrong with a building than just energy inefficiency, and so the software includes other prioritization systems for maintenance issues and occupant health and safety issues. In this case, the maintenance prioritization system currently reports the product of the number ten and the probability that certain maintenance conditions are

met by the equipment. For example if the expert rule for simultaneous heating and cooling identifies the probability of simultaneous heating and cooling over an interval as 0.88, then the maintenance priority for that analysis over that time interval is 8.80. This design for a maintenance priority system is very rudimentary and does not consider the costs of maintenance or other factors that maintenance personnel would use for prioritization. The presentation of maintenance priorities in the current version of the software is meant primarily to illustrate the concept of forming maintenance priorities as a means of sorting monitoring data. Likewise, the occupant health and safety prioritization system was not developed as part of this thesis, but was included as a place holder to illustrate that such a system should also be used to sort monitoring data.

Figure 47 includes a screenshot of the software interface that reveals some of the features discussed above.

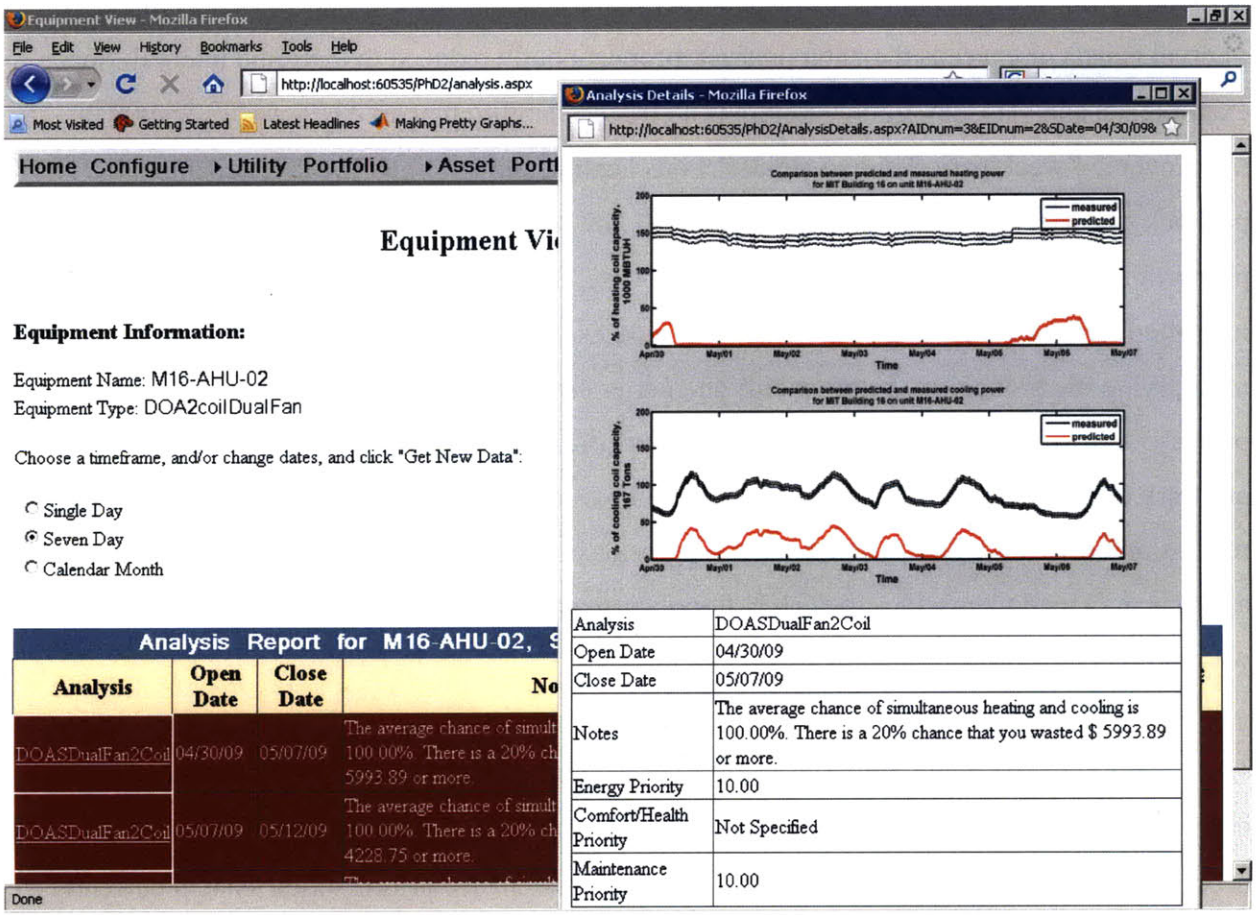

Figure 47 Screenshot of the software interface, showing features that link the underlying analysis to user defined settings

The screenshot in Figure 47 shows two elements of the interface; in the background is the main page of the interface that includes a spreadsheet-like report for a roster of analyses performed over a specific date range, and in the foreground is a pop-up window showing the detailed results for a specific analysis performed over a particular date range. The main page allows users to navigate lists of faults detected over monthly, weekly, or daily date ranges, sorted by energy, maintenance or health and safety priorities. The pop up window shows the same information that is included in each row of the spreadsheet, but only for one date range of analysis (one row of the spreadsheet). The name of the analysis performed is included in both interfaces, as is the starting and closing date for the analysis; Figure 47 includes a pop up window for the analysis "DOASDualFan2Coil", which, as the name implies, is for dedicated outdoor air system that have two fans and two coils.

The "notes" section of the interface includes text that is automatically generated by the software to help users interpret the results of analysis. In the case shown in Figure 47, the software has identified that the probability of simultaneous heating and cooling over that date range is 100%. The analysis also reports that there is a 20% chance that the equipment has wasted almost $6,000 over that date range; this is the upper tail of the distribution of dollar losses attributed to the inefficiency of the equipment. Below the notes section are the priority rankings of the results; the energy priority is a ten because the dollar value of the inefficiencies detected has a high likelihood of exceeding the targeted annual payback. Below that, the maintenance priority is also a ten because the probability of simultaneous heating and cooling is 100%. While the software does not diagnose the cause of the simultaneous heating and cooling, it does notify system users that the condition exists and that it has a high priority ranking; it is up to the system user to decide whether or not there is sufficient motivation to further investigate the possible fault. The graphs included in the foreground of the screenshot in Figure 47 are included as a visual aid to help users better interpret the analysis results and perhaps diagnose the specific origin of the possible fault. The features of those graphs will be discussed in the following sections as part of the broader presentation of experimental results.

6.3 Experiments in Building 46: AHUs 9 and 10

The previous chapter discussed an experimentation program in AHU 9 of MIT building 46 that included the purposeful application of software and hardware faults on the AHU. The hardware mechanical faults applied to AHU 9 included removal of the pneumatic supply pressure to the AHU's recirculation and outdoor air dampers, as well as removal of the damper linkages between those dampers and their actuators. The software fault on the AHU included a change in control programming to drive the mixing box damper signal as a function of time instead of its nominal dependence on CO_2 concentration in the building or mixed air temperature in the mixing box. These purposeful faults as well as other inherent faults of the equipment were all identified and evaluated by the software. None of the faults on AHU 9 were found to have significant financial value relative to MIT's targeted annual payback of correcting faults detected, however they do illustrate the ability of the software to detect and prioritize various faults.

6.3.1 Results from MIT Building 46

Figure 48 shows a close up view of the software pop-up screen that facilitates inspection of the energy efficiency analysis results for AHU 9. The upper graph in the figure shows the measured and predicted heating activity of the AHU's heating coil while the lower graph shows the measured and predicted activity of the AHU's cooling coil; measured activity is shown in black, while predicted activity is shown in red. Predictions on AHU performance are not forecasts of operation, but instead are the model-based idealized performance of the AHU for the given weather and operating conditions of the building over the date range of analysis.

The heating and cooling activity of the AHU, whether predicted or measured, are each shown as three lines in Figure 48; the dashed outer lines represent a normally distributed 95% confidence interval about the solid center line. For the graphs in Figure 48, the red, prediction curves more clearly show this banded structure than do the black lines for measured operation. In

this case the distribution on the predicted heating and cooling activity of the AHU is wider than the distribution of measured activity.

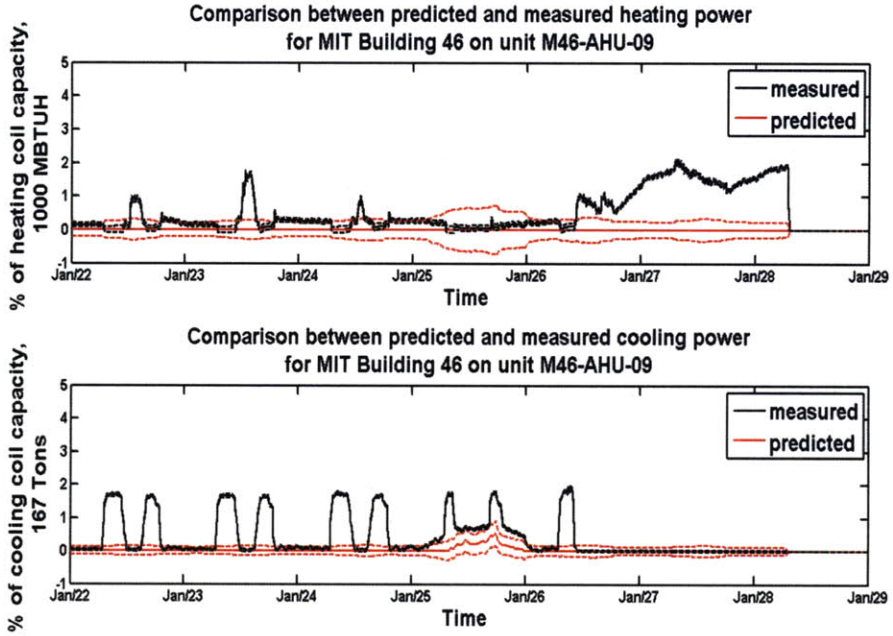

Analysis	AHUactualVidealEnergy1
Open Date	01/22/10
Close Date	01/29/10
Notes	No significant signs of simultaneous heating and cooling. There is a 20% chance that you wasted $ 26.06 or more.
Energy Priority	0.00
Comfort/Health Priority	Not Specified
Maintenance Priority	Not Specified

Figure 48 Close up view of the software pop-up screen that facilitates closer inspection of analyzed equipment performance; this particular view is for AHU 9 in MIT building 46, for a period when the damper pneumatic pressure supply was removed and the damper control programming was altered.

The vertical axis of both the heating and cooling coil performance graphs are dimensionless percentages of the corresponding heating or cooling coil capacity rates. This follows from the dimensionless expert rules defined in chapter 5; the results of analysis are presented in this dimensionless form so that they can be interpreted relative to the scale of the equipment under scrutiny. For example, the upper graph in figure 48 shows the peak heating coil activity as roughly 3% of the heating coil's rated capacity, $1x10^6$ BTUs per hour; numerically this is equivalent to 30,000 BTUs per hour, which seems like a large amount of heat exchange,

but clearly within the context of this equipment it is rather small. Likewise, the peak cooling rate shown in the lower graph of figure 3 is about 2% of the cooling coil's rated capacity; again numerically this is equivalent to about 3 tons of cooling, or roughly the peak cooling power needed by a typical New England family home in the summer time, but within the context of this equipment it is rather small.

From January 22nd, 2010, through January 26th, 2010, the control programming for the damper control signal in AHU 9 was altered to vary as a function of time rather than its nominal mixed air temperature or CO_2 concentration based control. The effects of that software change are evident in both the heating and cooling graphs shown in figure 48; the rise and fall of the cooling and heating power during that period follow the opening and closing of the mixing box dampers. Mid-day on January 26th, the damper control signal was restored to its nominal operation and the pneumatic supply pressure was removed from the recirculation and outdoor air dampers.

In the absence of supply pressure, the recirculation and outdoor air dampers go to their nominally fully closed and open positions; these are default values driven by fire safety standards. The inoperable dampers force the AHU to imbibe more cold air than it would do so otherwise on a January day, causing the measured heating coil operation to exceed its ideal predicted behavior; this response is shown in the top graph of Figure 48 starting at mid-day on January 26th. Despite their manifestation within the analysis results, the software calculates that very little financial loss is incurred by these faults; there is only a 20% chance that the faults incurred financial damages of $26.06 or more, and consequently the energy priority is 0.00.

For heating and cooling enthalpy exchange between 0 and 2% shown in Figure 48, the confidence interval of the predicted heating and cooling enthalpy exchange of the AHU often exceeds the confidence interval of the measured heating and cooling enthalpy exchange. This exemplifies a limitation of the modeling and inference approach taken in this thesis; at lower levels of enthalpy exchange, the predictions about AHU energy consumption become more uncertain. For example, on January 25th in the upper graph of Figure 48, the dotted red-line, which signifies the confidence interval of prediction, clearly exceeds the measured quantity of

heating enthalpy exchange. This effect of increased predictive uncertainty at low levels of enthalpy exchange was found in many other results from the FD&E systems, and is most likely due to the uncertainty associated with predictions on the outside air flow required by the AHU. The uncertainty of outdoor air flow predictions, especially for AHU 9 in building 46 may be large under certain flow conditions due to the possible influence of CO_2 based ventilation (see the related discussion in the preceding chapter). Despite the increased uncertainty of prediction, and hence reduced ability to detect faults at low enthalpy exchange rates, the usefulness of the FD&E system appears intact as the measured enthalpy exchange increases to more meaningful levels. Results for Buildings 16 and 56 presented later on in this chapter explore this effect at the opposite extreme with growing uncertainties at very large measured values of enthalpy exchange.

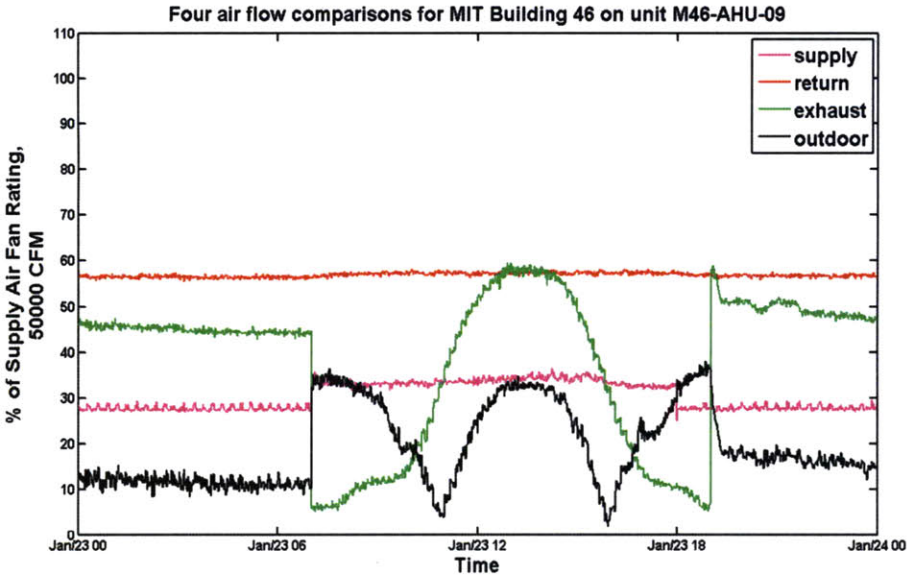

Analysis	FourAHUAirFlows
Open Date	01/23/10
Close Date	01/24/10
Notes	On average, the measured supply air flow is 30% of the supply fan rating; The mean difference between the measured and computed supply air flow is 13% of the supply fan rating; the air flow measurements or directions may be erroneous The retun air flow on average exceeds the supply air flow by 26% of the supply fan rating; the supply and return flow measurements may be erroneous or the building may have additional air sources
Energy Priority	Not Specified
Comfort/Health Priority	Not Specified

Figure 49 Results screenshot for analysis on air-flows across AHU 9 in MIT Building 46

Like the energy analysis, the results of air-flow analysis are also presented in terms of a dimensionless percentage of the supply air flow rating; that new graph is shown in Figure 49. The automated text output for the airflow analysis includes commentary on the relative magnitudes of some of the flows and inference on any anomalous flow behavior. The text highlights that the supply air flow rate is on average only 30% of the rated capacity of supply air fan; this is included to help diagnostic engineers understand the turn down ratio of the variable frequency drive on the fan. The graph and text also point out that the return air flow is much larger than the supply air flow; this is immediately surprising since AHUs are typically designed to intake less return air than they supply to a space. Finally, the supply air flow rate that is computed from the air flow balance equation in chapter 5 was compared against its measured

value in order to highlight potential discrepancies in the measurement of air flows relative to the model describing them (most notably discrepancies in the assignment of flow direction).

Discussions with the building technician for MIT building 46 revealed that the excessive return air flow to the AHU could be attributed to leaking fire dampers in the basement of the building. Other phenomena visible in the graph were harder to explain and required more investigation. For example, from 7:00 a.m. to 9:00 a.m. and again from 4:00 pm to 6:00 pm, the outdoor air flow (shown in black) appeared to exceed the exhaust air flow (shown in green), and as one of those flows increased, the other would decrease. This is an intriguing result because conservation of mass across the AHU prohibits the exhaust air flow from decreasing while the outdoor air flow increases, and the return and supply air flows remain constant. On the contrary, conservation of mass requires that the flows follow the dynamic behavior exhibited by the AHU flows from about 11:00 am to 3:00 pm on January 23rd; the exhaust and outdoor air flows should increase and decrease together while the return and supply air flows remain constant.

The entire timeframe over which the exhaust and outdoor air flows exhibited their oscillatory behavior coincides with the time frame when the altered control logic (software fault) took control of the dampers. In fact, the period of time when the outdoor air flow exceeded the exhaust air flow corresponds to a period when the damper control signals were roughly equal to or less than 50%. For that range of damper control signal, the outdoor and exhaust air dampers should have been at 50% or less open, and the recirculation damper should have been at 50% or more open. For constant values of return and supply air flow, as shown in the graph, the only plausible explanation for why the outdoor air flow could exceed the exhaust air flow and inversely follow its variation, is if the exhaust air flow were actually re-directed to flow through the outdoor air intake damper. In this case, the erroneous outdoor air flow measurement could be an artifact of the type of air flow station that was used to measure the outdoor air flow; the thermal-anemometer outdoor air flow station does not provide a direction of flow, only a magnitude, and hence flow reversal leading to discharge of air through the outdoor air intake duct is otherwise indistinguishable from air flow in the proper direction.

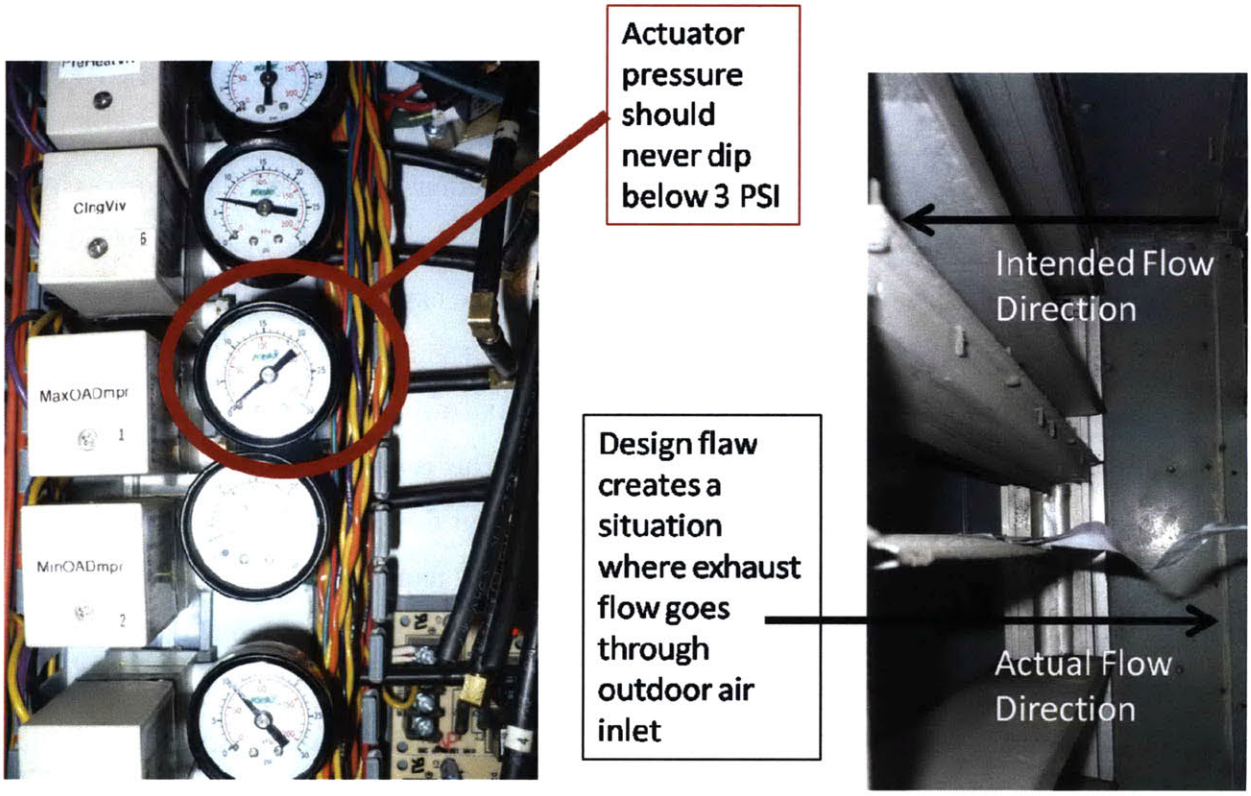

Figure 50 Pictures of faults diagnosed due to inspection of results from software analysis; the right hand image verifies suspicions of flow reversal in the mixing box, while the left hand image shows evidence of a broken pressure transducer for the outdoor air dampers.

The graphical and text results of the software analysis prompted an inspection of the equipment with the technical manager for the building. During that inspection, the flow reversal hypothesis was validated by taping a streamer to the leading edge of the outdoor air intake damper, and photographing the flow across it. The right hand photograph in Figure 50 validates the flow-reversal hypothesis; the trailing edge of the flow streamer points towards the intake of the outside air duct as oppose to the intake of the mixing box, thereby confirming suspicions that the flow was going in the wrong direction and discharging to the outside from the outdoor air intake duct.

A closer inspection of the pneumatic damper transducers for the AHU also revealed that the transducer for the maximum outdoor air damper had failed; it's pneumatic supply pressure when connected to the pneumatic supply line registered as 0 PSI, despite the fact that the minimum operating pressure for the transducer is 3 PSI. The broken transducer is shown in the left hand photograph of figure 5, with its measurement needle resting at 0 PSI; the device was

confirmed as broken after a manual perturbation of the damper actuator yielded no corrective response from the transducer. After confirming that the device was broken, the building technician issued a work order to fix the equipment; the broken transducer was replaced by MIT several days after it was identified. The flow reversal effect was confirmed both before and after the outdoor air damper pressure transducer was replaced.

The results presented so far by the software ultimately lead to the diagnosis of two unexpected faults in the AHU; a design flaw that causes flow reversal, and an equipment failure that reduced the controllability of the AHU. The serendipitous capture and resolution of these faults followed from the design intent of the FD&E software; the software tool successfully identified, ranked and described anomalous equipment behavior, and provided sufficient information for a trained building technician to diagnose, and in some cases fix, HVAC faults.

Despite these successes in discovering and fixing faults, it is not clear whether fixing the broken transducer or flow reversal would actually lead to any significant gains in energy efficiency or annual energy savings. The energy efficiency analysis of the AHU over seven days in January, which is shown in Figure 48, suggests that there is only a 20% probability of saving $26 or more in energy costs over that time interval. The result indicates that even if all faults were corrected on the AHU, at best its energy consumption over that time interval could decrease by only about $25. The data in Figure 48 and Figure 49 further suggest that even with the broken transducer for the outdoor air dampers, the AHU could still control the flow of outdoor air simply by adjusting the position of the recirculation and exhaust air dampers.

The mixing box simulations included in chapter three for northeastern weather also indicate that the AHU's outdoor and exhaust air dampers would probably only be driven to 50% or more closed for a small part of the year, primarily in the summertime. The reversal of flow in the mixing box, which occurs when the exhaust and outdoor air dampers are drawn to 50% or more closed, is actually a boon to energy conservation during the hot and humid Boston summer. By excluding outdoor air from the AHU, the flow reversal eliminates energy that otherwise would have been spent on dehumidifying or cooling air that could be more humid and hot than the return air alone. While this "fault" could favorably reduce energy consumption during part of

the year, it could potentially lead to an uncomfortable or even dangerous environment for the occupants. In the absence of fresh outdoor air, unpleasant orders, noxious vapors and CO_2 concentration will increase in the parts of the building served by AHU 9 and potentially harm the productivity and health of occupants in those spaces.

From February 1st to February 5th, the damper linkages for the outdoor and recirculation air damper were also incrementally removed and replaced. Figure 51 includes two graphs for the AHU heating and cooling activity over that time period; the latter half of the visualized cooling activity, from February 1st to February 5th, shows a small increase in cooling energy usage beyond the predicted consumption level. This seems indicative of faulty recirculation and outdoor air dampers because the fault in this case would theoretically permit more warm return air to recirculate back into the AHU than might otherwise be allowed; the increased recirculation flow would require an increase in cooling energy consumption in order to achieve discharge air conditions. Under normal circumstances, the AHU could use the cool outdoor conditions to temper the recirculation air stream without any mechanical cooling, but since the recirculation and outdoor air damper linkages had been removed, the economizing feature was disabled.

With both sets of dampers disabled, the opposite condition of excess heating (instead of excess cooling) may also have been likely; the AHU could have drawn more cold outside air than could be properly tempered with warm recirculation air alone, leading to use of the steam heating coil to properly heat the air. The final outcome of either excess heating or cooling due to faulty damper linkages is predicated on the pressure distribution within the mixing box, and the resulting air flow. The presence of excess cooling in this case is likely due to the large return airflow that is unique to AHU 9; the back pressure from the exhaust air dampers may be sufficiently large to cause the recirculation air flow to overwhelm the outdoor air flow when their respective dampers are free to move on their own.

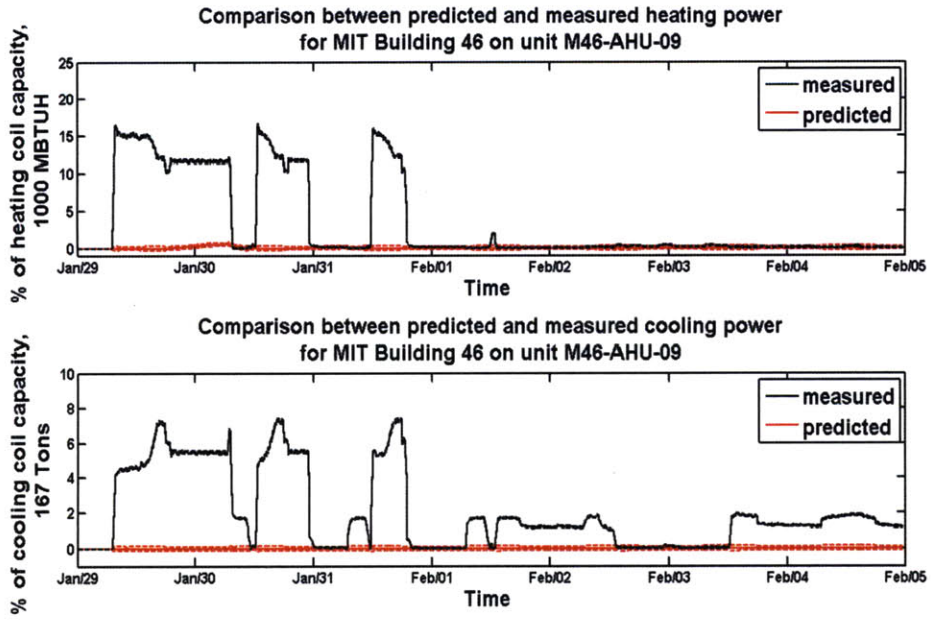

Analysis	AHUactualVidealEnergy1
Open Date	01/29/10
Close Date	02/05/10
Notes	The average chance of simultaneous heating and cooling is 25.86%. There is a 20% chance that you wasted $ 156.11 or more.
Energy Priority	7.00
Comfort/Health Priority	Not Specified
Maintenance Priority	2.59

Figure 51 Results of experimentation on AHU 9 for January 29th, 2010 through February 5th, 2010

The effects of removing the damper linkages were tested twice over the time interval shown in Figure 51; once from midday on February 1st to midday on February 2nd, and again from midday February 3rd to midday February 5th. In between, from midday February 2nd to midday February 3rd, the fault was resolved (the linkages were re-installed). The data in the lower graph of Figure 51 reflect this experimentation schedule; the excess cooling effect disappears when the damper linkages are replaced and then re-appears when they are once again removed.

In addition to visualizing the effects of broken dampers, the data in Figure 51 from January 29th though January 31st clearly shows a brief period of simultaneous heating and

cooling activity. The effect seems to disappear in the latter half of the data shown in Figure 51, and the equipment does not demonstrate the same phenomenon in any other part of the experimental data set. Furthermore, the simultaneous heating and cooling activity appears to be the result of purposeful control by the AHU controller; portions of the phenomenon appear to be in steady state, and the effects are visible over several hours at a time. The ultimate cause of this phenomenon could not be explained by inspecting the equipment or the software programming of the controller, however the analytical results suggest that it may not be worthwhile to investigate or resolve the problem; the energy priority for the week was a seven and valued at roughly $160.

Despite the demonstration of faulty dampers and simultaneous heating and cooling on the AHU, the analysis of the faults yields a modest valuation of the equipment in-efficiencies over that time interval. If the analysis had been carried out over a longer period of time that did not include more simultaneous heating and cooling, then the valuation would have been even less. In this case the software and analytical approach has once again raised some awareness about potentially fault-laden conditions, and then qualified the possible financial value of fixing those faults against user-defined investment criteria.

The analysis of AHU 10 in building 46, where no faults were ever purposefully applied during this experiment, exhibited a very similar simultaneous heating and cooling pathology to what was observed in AHU 9 above. The results of condition monitoring on AHU 10 from January 22nd, 2010 through January 29th, 2010, are shown in figure 7, and the latter half of the data show two distinct periods of controlled simultaneous heating and cooling. Once again the origins of the fault are unknown; the mechanical equipment does not exhibit any clear deficiencies, and neither does the software programming. Still, despite acknowledging the existence of a problem, the analysis also suggests that the priority for fixing that problem is not very high; the energy priority for that analysis time period is a 5.00, and the fault is valued at roughly $105.

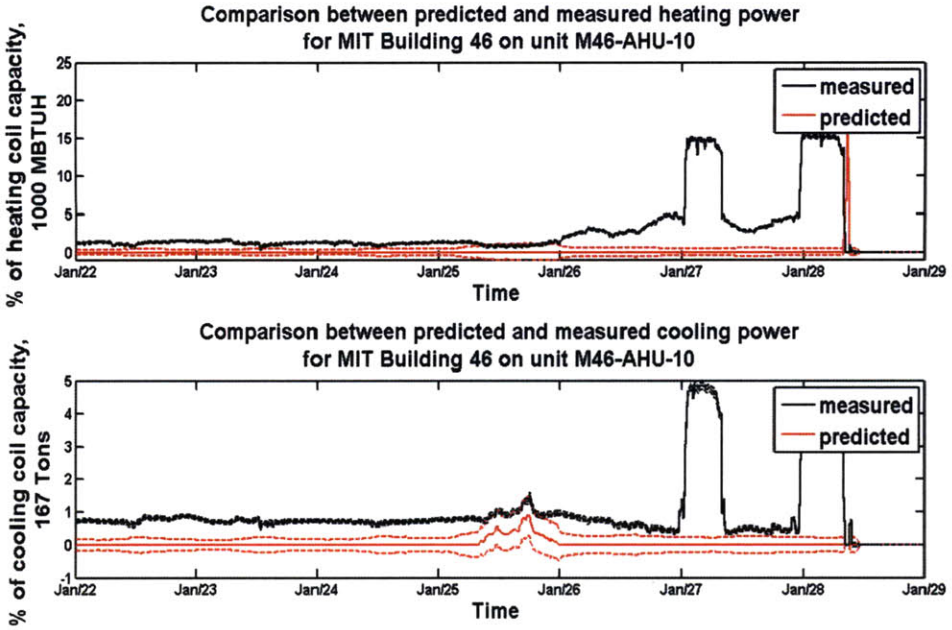

Analysis	AHUactualVidealEnergy1
Open Date	01/22/10
Close Date	01/29/10
Notes	No significant signs of simultaneous heating and cooling. There is a 20% chance that you wasted $ 105.26 or more.
Energy Priority	5.00
Comfort/Health Priority	Not Specified
Maintenance Priority	Not Specified

Figure 52 Condition monitoring on AHU 10 in building 46, from January 22nd, 2010 through January 29th, 2010

6.3.2 Summary of Results for MIT Building 46

Several faults on two AHUs in MIT building 46 were made visible through the building condition monitoring software. The analysis of that related data was also used to draw inference about the existence of simultaneous heating and cooling in the equipment. Both experimentally applied and prior existing faults were detected and evaluated by the software; in some cases the faults were immediately addressed and resolved, and in other cases the analysis showed that the faults may not be a high priority to address.

6.4 Analysis of Buildings 16 and 56

In 2009 Cimetrics Inc. was hired to provide a monthly continuous commissioning service on MIT buildings 16 and 56; they successfully identified over $250,000 of possible energy efficiency savings, the majority of which was attributed to simultaneous heating and cooling in a handful of large AHUs. The same data that Cimetrics collected and used to draw their inference in 2009 was used in this current research to compare the results of the FD&E software and analysis against Cimetrics' 3rd party results. Furthermore, current data for AHU 2 in building 16 was also collected and analyzed to determine if any changes had been on the AHU since Cimetrics first diagnosed the equipment with simultaneous heating and cooling in the spring of 2009.

6.4.1 Findings on 2009 and 2010 performance data

Analysis of Cimetrics' 2009 equipment data by the FD&E software provided similar results to what was reported by Cimetrics a year ago; the analysis yields a very high probability for simultaneous heating and cooling, and an overwhelming financial loss due to that fault.

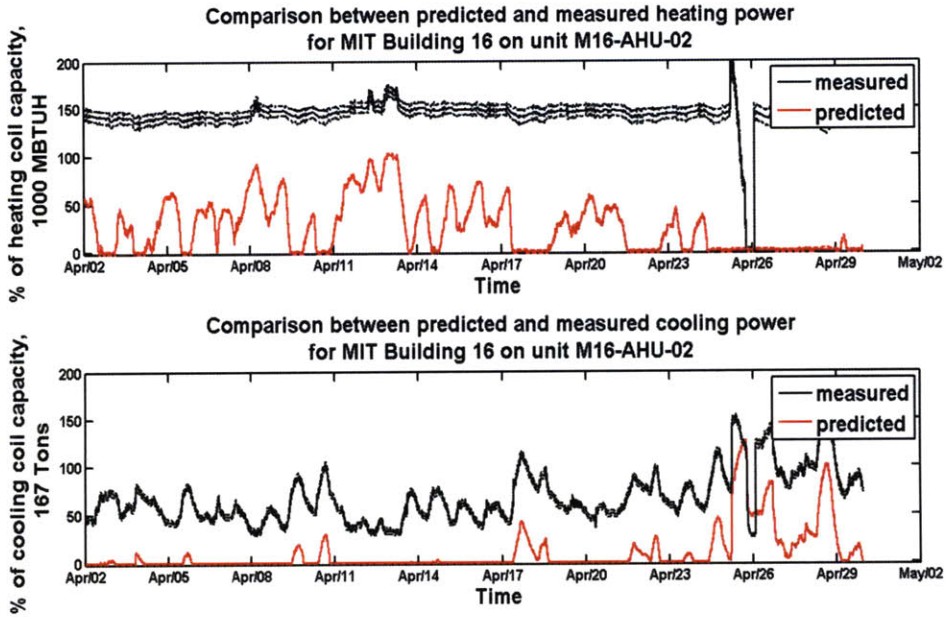

Analysis	DOASDualFan2Coil
Open Date	04/02/09
Close Date	04/30/09
Notes	The average chance of simultaneous heating and cooling is 99.13%. There is a 20% chance that you wasted $ 20147.15 or more.
Energy Priority	10.00
Comfort/Health Priority	Not Specified
Maintenance Priority	9.91

Figure 53 Analysis results on AHU 2 in MIT building 16

Figure 53 presents the results of analysis; the graphs of heating and cooling activity show a consistently large difference between the measured and predicted energy consumption of the AHU. The upper graph in Figure 53 shows a visibly consistent level of measured heating coil activity over almost the entire month of April, 2009. The large and nearly constant value of the measured heating coil enthalpy exchange over the entire month suggests that the steam valve for the coil may be stuck at 100% open. The measured heating coil activity in Figure 53 fluctuates about 150% of the coil's nominal heating capacity because the maximum enthalpy exchange rate of the coil under real operating conditions exceeds the manufacturer's original rating of the coil. Especially for steam heating systems, the real operating enthalpy exchange of the coil is a

function of many variables and may not always perfectly mach a manufacturer's specification for the coil.

The predicted heating coil activity, shown as a red line in the upper graph of Figure 53, suggests that the outdoor weather during the month was sufficiently cold to require heating, however not to the extent that it was delivered by the unit. The lower graph in Figure 53 also shows a large discrepancy between the measured and predicted cooling coil activity; the excess cooling seems to be a controlled reaction of the AHU to compensate for the excessive heating coil activity.

In addition to detecting simultaneous heating and cooling, the FD&E analysis also calculated that the excess heating and cooling activity of the AHU resulted in approximately $20,000 worth of wasted energy, just for the month of April, 2009. Since one month of equipment operation already exceeded the $10,000 target for energy efficiency investment at MIT, this fault was assigned a priority of 10. Using a regression on outdoor air temperature, Cimetrics extrapolated their data from March, 2009, to the entire year in order to forecast the annual financial loss due to this fault; they reported an annual loss of $256,686. Because the heating coil activity is so large and largely overwhelms the effects of changing weather, the annual impact of this fault could also be estimated simply by multiplying the results from one month by 12; this alternative extrapolation predicts a potential annual loss of roughly $240,000. The loss estimates from Cimetrics and the FD&E analysis differ by roughly 5%, but both approaches find the fault to exceed MIT's $10,000 efficiency payback target by more than a factor of 20.

In comparison to the results in Figure 48 for AHU 9 that show increased uncertainty of prediction at low levels of enthalpy exchange, Figure 53 shows the opposite effect at very large values of measured heating enthalpy exchange; in such instances the uncertainty of measured performance is much larger than that of predicted performance. The increased uncertainty of the measured performance at high levels of heating enthalpy exchange is directly due to the large influence of uncertainty associated with temperature measurements across the heating coils. Since sensor measurement accuracy is taken as a relative percent uncertainty by the FD&E

system, temperature measurements made at high temperatures are less certain than those made at low temperatures; this is why the confidence interval for the cooling coil enthalpy exchange in the lower graph of figure 53 is much tighter than the confidence interval for the heating coil enthalpy exchange shown in the upper graph of the same figure. Especially since the heating enthalpy exchange is over 150% of the coil's rating, we expect the air temperature change across the coil to be significantly large and in a high temperature range, leading to greater uncertainty associated with the measurement.

Despite this increased uncertainty, however, the cost differential between the predicted and measured performance is so large that the increased uncertainty has very little impact on the final inference; clearly there is a large amount of money and energy being wasted. On the other hand, however, if the measured and predicted values for enthalpy exchange were both within a similar range, perhaps both close to 100%, then the amplification of measurement uncertainty in that range would decrease the ability of the FD&E system to detect and evaluate a fault. By incorporating sensor and model uncertainties within the fault inference, we may create situations where faults are indistinguishable due to overlapping confidence intervals; so far we have not yet encountered a situation where this effect masks or falsely implies a significant level of energy waste.

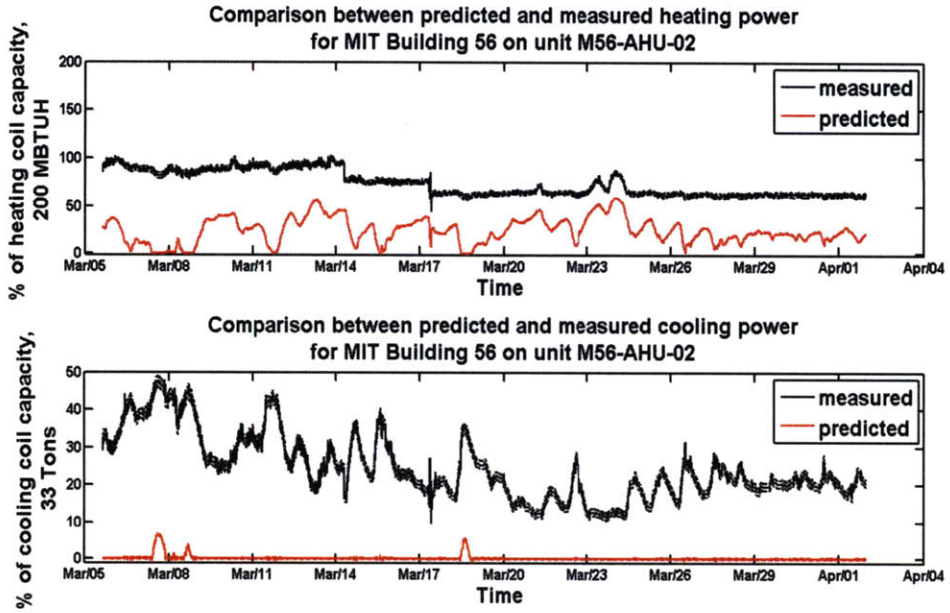

Analysis	DOASSinglelFan2Coil
Open Date	03/05/09
Close Date	04/02/09
Notes	The average chance of simultaneous heating and cooling is 100.00%. There is a 20% chance that you wasted $ 1663.53 or more.
Energy Priority	10.00
Comfort/Health Priority	Not Specified
Maintenance Priority	10.00

Figure 54 Results of FD&E analysis on AHU 2 in Building 56

Cimetrics' data for AHU 2 in MIT building 56 was also consumed by the FD&E software and once again the analytical results were similar to what Cimetrics had reported; both approaches found several thousand dollars worth of energy wasted due to simultaneous heating and cooling. Figure 54 includes the results of FD&E analysis on Cimetrics' 2009 data for that AHU, and shows heating and cooling coil activity that are indicative of simultaneous heating and cooling. AHU 2 in building 56 is roughly an order of magnitude smaller than AHU 2 in building 16, and consequently the financial impact of simultaneous heating and cooling in the former equipment is about ten times smaller than in the latter. Despite their difference in size, both AHUs have the potential to pay back $10,000 or more in energy savings if their faults were corrected; both units exhibit a high probability of meeting the investment pay back target for

MIT and therefore they both received an energy priority of 10. The results from analyzing AHU 3 in building 56 were very similar to the results for AHU 2 in the same building; both AHUs exhibited the same fault characteristics and pay back potential.

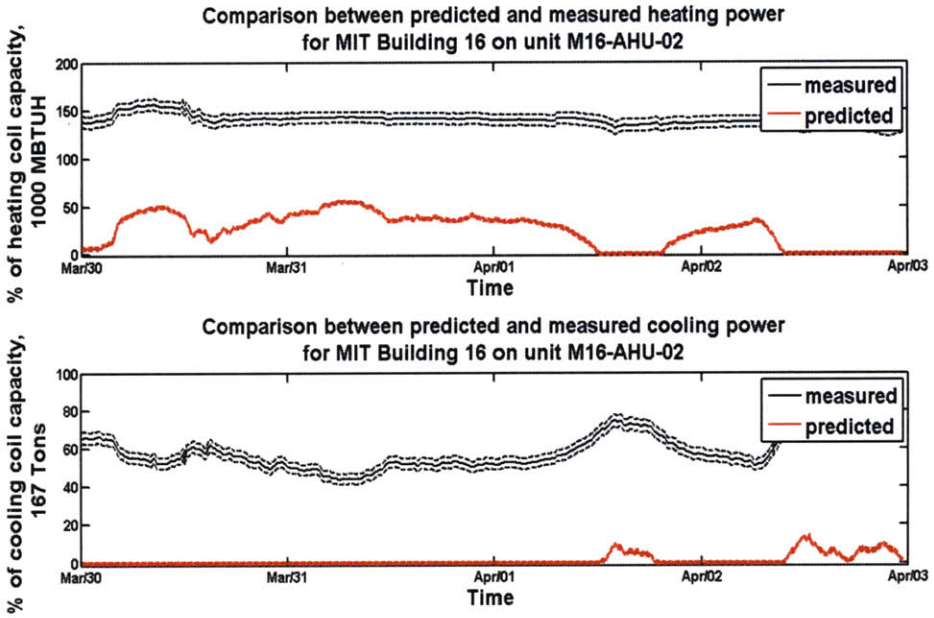

Analysis	DOASDualFan2Coil
Open Date	03/30/10
Close Date	04/03/10
Notes	The average chance of simultaneous heating and cooling is 100.00%. There is a 20% chance that you wasted $ 2855.18 or more.
Energy Priority	10.00
Comfort/Health Priority	Not Specified
Maintenance Priority	10.00

Figure 55 Results of FD&E analysis for AHU 2 in building 16, with current 2010 data

In addition to using 2009 data from Cimetrics to analyze AHU 2 in building 16, more recent data for March and April, 2010, was also collected from the equipments' building control system and analyzed by the FD&E software. The results of the analysis are included in figure 10 above for a snapshot of equipment operation during early April. Similar fault characteristics exist between the recent building control data and the 2009 Cimetrics data; the measured heating coil activity appears to remain consistently at full bore, and the financial loss extrapolated for the

whole year remains over $200,000. Evidently no changes or corrections have been made to the AHU, even after 12 months since Cimetrics first identified that the AHU was hemorrhaging hundreds of thousands of dollars per year in excess heating and cooling.

6.4.2 Summary of results for MIT buildings 16 and 56

Three AHUs with known pre-existing fault conditions were evaluated by the FD&E software tool. The tool was used to detect and evaluate the financial loss of simultaneous heating and cooling within the equipment. Results of the analysis by the FD&E tool matched previous fault detection and evaluation results provided by Cimetrics Inc.; between both methods, the predictions on financial loss due to those faults were within 5-10% of each other.

6.5 Results of Monte Carlo Simulations

6.5.1 Introduction

A core assumption in the approach taken to propagate uncertainty through the FD&E analysis is that all distributions within each step of the analysis are Gaussian. The validity of this assumption is particularly relevant to the accurate calculation of equipment energy consumption because of the large number of computations involved in its evaluation; as data is manipulated through computation, it is not guaranteed that the starting assumptions of normality are preserved to the final result of analysis.

6.5.2 MC Simulations

In order to validate that assumption, Monte Carlo (MC) simulations were conducted with bootstrapping of raw data vectors to yield a sample representation of equipment energy consumption that is free from any assumptions about propagating distributions. The intent of MC

simulation is to compare the distribution of energy consumption produced by the MC simulation against a single sample distribution of energy consumption produced by the nominal analysis with propagating distributions.

The only assumption used in the MC simulation was that the raw data vectors could be re-sampled, or bootstrapped, to create a sample population of measurements for each real measurement taken. Sample populations of each measurement were made by random sampling of a Gaussian distribution of values that was created for each real measurement; the real measurement value was taken as the mean of that distribution, and the manufacturers' rated accuracy of the sensor was taken as equivalent to two standard deviations of that distribution. Gaussian distributions were used to bootstrap the time series because of the starting assumptions about normality that are implicit to the approximation methods that underlie the propagation of uncertainty through the single-sample analysis. If the Taylor expansion in Coleman and Steele (64) were made under the influence of a different distribution, then the time-series bootstrap would need to reflect that different distribution in order for the MC simulation to compare against the single sample approximation.

The real time series of data for each measurement was bootstrapped at each timestamp to create 100 additional time series; the bootstrap created a large population from which total energy consumption could be directly evaluated for each re-sampled time series and then used to construct a distribution of possible energy consumption values over the interval of that time series. The results from the MC simulations were used in two ways to validate the accuracy of the assumptions that underlie the singe-sample propagation of uncertainty:

1. The distribution of the MC simulation results were tested for normality
2. The single-sample and MC results were tested for homogeneity

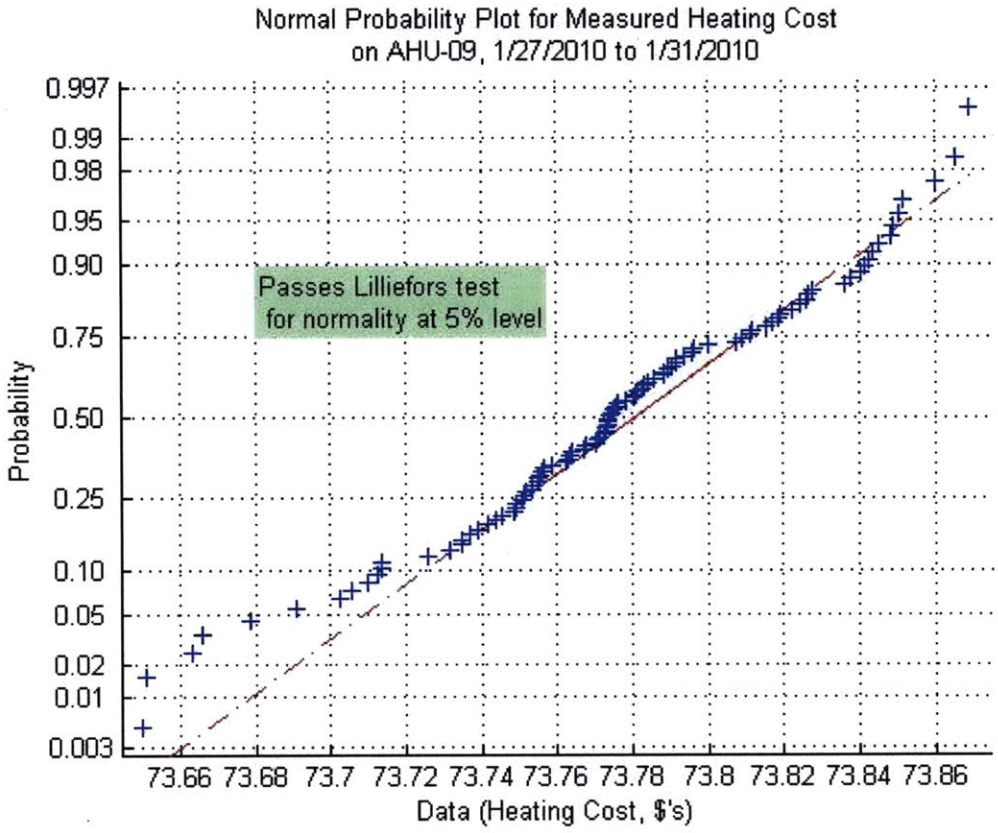

Figure 56 Q-Q plot of empirical CDF on MC simulation results; a Lilliefors test for normality was applied to the distribution of energy consumption and found the null hypothesis valid at the 5% level

The Lilliefors test for normality, a two-sided goodness-of-fit test suitable for small sample sizes with unknown parameters (70), was used to test the hypothesis that the MC results followed a normal distribution. Applied to several different sets of simulation results, the Lilliefors test consistently found that the null hypothesis was true at the 5% level. Visual inspection of the quantile plot for the empirical distribution in Figure 56 shows that the tails of the distribution for N=100 simulations begin to defy the strictly normal distribution. Repeating the MC simulations with a larger population, N=500, resulted in greater linearity at the tails, suggesting that the divergent behavior at smaller sample sizes is an artifact of simulation sample size.

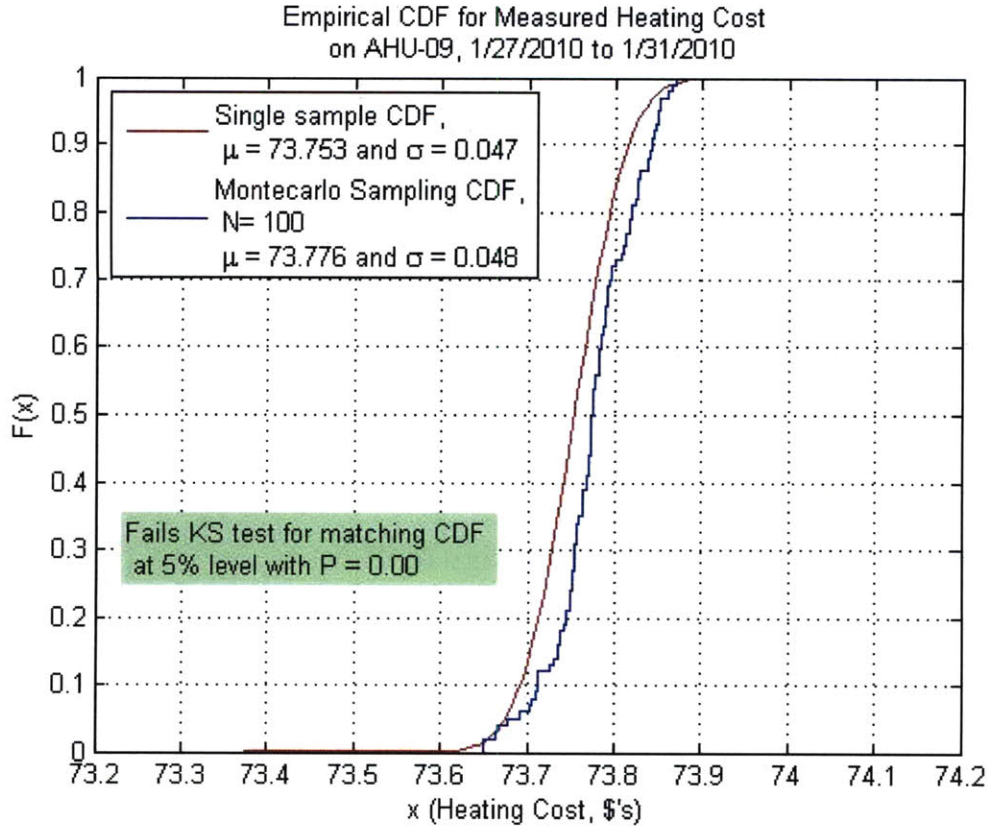

Figure 57 Two-sample Kolmogorov-Smirnov test of homogeneity between MC and single sample distributions; the KS test rejeceted the null hypothesis at the 5% level for almost all test samples

The homogeneity of the MC and single-sample distributions was tested with a two-sample Kolmogorov-Smirnov (KS) test (70). Across a large range of dollar value distributions for measured and predicted equipment energy consumption, the KS test at the 5% level consistently failed to find that the single sample and MC sample distributions had identical distributions. Testing for homogeneity in the ~$100 range is shown in Figure 57; notice that while the values for mean and standard deviation are in fact different between the distributions, their practical difference for estimating the dollar value of large equipment energy inefficiencies that are financially meaningful to system users is apparently insignificant.

6.5.3 Summary of results from MC simulations

MC simulations were carried out with population sizes extending from 100 to 500, for distributions of measured and predicted equipment energy costs of operation over the range of $1

to $1000. The distributions that resulted from MC simulations were all tested for and found to follow Gaussian distributions at the 5% level by the two-sided Lillefors test. The two-sample KS test was also used to test for the homogeneity between the MC and single-sample result distributions; for most sample sizes and range of operating costs, the KS test rejected the null hypothesis at the 5% level. Despite rejecting the hypothesized homogeneity of the distributions, visual inspection of the MC and single sample distributions suggested that discrepancies between them were practically insignificant in the search for energy inefficiencies that are likely to be valuable to system users.

6.6 Feedback from users on system design features

While most of the tangible results in this research have focused on the objective analytical results of detecting and evaluating building equipment faults, several important subjective observations were also made during user testing of the software deployed on MIT's campus. In particular, several personnel from the MIT facilities department were interviewed to gain their feedback on the usefulness of the software in detecting, evaluating and ultimately resolving gross building pathologies. Several important themes were observed in the feedback from MIT facilities personnel:

- Expressing FD&E results in the form of a probability appears to be more palatable to consumers of that information than deterministic results because the probabilistic form implies that uncertainty is implicitly embraced by the analysis. Uncertainty seems to be a significant part of the maintenance culture; facilities personnel appear to acknowledge that not all measurements can be trusted and that equipment installations are often subject to installation or programming errors. By formulating the results of analysis within a context that matches the culture of maintenance, it appears that the results of analysis are more believable by consumers of that information. Even if the results do not include all of the uncertainty that a user may expect in a building, the incorporation of just measurement and model uncertainty at least shows the user that the analysis doesn't expect to be perfect. While facilities personnel may not have the background to rigorously interpret statistical results, they are well equipped to consider the chances of risk and reward in making decisions.

- Sorting and prioritization of results according to user-defined parameters appears to be a key feature for system users; while there may be several dozen pathologies in a building, the system user often only has the time and inclination to address the top 1 or 2 most pressing issues. Furthermore, user-defined sorting of pathologies according to a dollar value of annual pay-back on fixing a problem appears to also mesh with the maintenance culture. Interactions with system users in the MIT facilities department also demonstrated that the dollar value of fixing a fault is not always the driving factor in deciding to address a fault; occupant health, safety and comfort is usually more immediately important than dollars lost to energy inefficiency.

- Automated text generation for interpreting inference is helpful for suggesting to users where they might find physical manifestations of pathologies that are detected and evaluated by the software. Furthermore, exposing the actual data graphic to system users allows them to drill-down from a spread-sheet summary of information into the process of diagnosing issues.

7 Conclusions

The overarching purpose of this research was to investigate how buildings across the United States and beyond may be rapidly and consistently assessed for large energy inefficiencies attributable to the pathological malfunction of HVAC equipment. A key assumption underlying that goal is the expectation that if building managers and contractors knew about valuable equipment malfunctions, then they could be more aggressive towards fixing gross building energy inefficiencies. The process of fixing broken buildings, however, is more complicated than simply identifying broken building energy systems, otherwise the past twenty years of successful research in the subject surely would have yielded a fitting technical solution. A detailed review of the prior art suggests that the lack of widespread use of fault detection technologies in buildings today may be partially attributed to under-emphasis of the following factors in fault detection research and development:

1. The social culture and business processes that underlie the physical act of fixing a broken building (71)

2. Measurement and model uncertainty that is inherent to inference performed on complicated systems like buildings

3. Use of simple, yet broadly applicable models to gain bulk information on gross energy inefficiencies

While the research activities in this thesis focused predominantly on the underlying physics and mathematics of designing a fault detection system, I also learned through this experience that the social culture and business process of fixing building faults are equally important design elements to consider in that system. On the surface it appears that considerations of the social culture and business process of fixing buildings have little place in the research and development of mathematical algorithms for detecting building energy inefficiency. However when the goal of those algorithms is to catalyze the fixing of a building, it may be more logical to start with the social and business cultures that dominate the fixing of buildings, and then work backwards to define the mathematics and physics that are also important in that process.

7.1 Limitations of current approach

While this thesis showed some success at detecting and evaluating faults within a building, it does not represent a complete solution to total fault detection, diagnostics and evaluation of building operation. The FD&E approach taken in this thesis is foremost limited by the need for models that describe equipment behavior. The simple thermal-fluid models used here to analyze the operation of AHUs do not provide a detailed insight into the origin of faults, but only their existence. Moreover, the modeling approaches used here are useful for detecting certain types of gross energy inefficiencies like simultaneous heating and cooling, but are not guaranteed to identify all possible faults within equipment. For example, the thermal-fluid models used here to describe AHU operation do not consider the fine-grained control action of PID loops or other

transient phenomenon which means that faults pertaining to transience or poor control loop tuning may be overlooked.

In addition to limitations placed by modeling and measurement, the probabilistic foundation of the FD&E system designed in this thesis makes the system more difficult to judge for its precise efficacy of detecting and evaluating faults. Since the system reports its findings in terms of a likelihood of fault existence or value, the actual false positive or negative rate of inference is more difficult to judge than if the system generated deterministic results. Consequently, the system may erroneously report a small likelihood of fault existence or valuation, but its true accuracy is subject to interpretation by the user. This artifact of how the system operates makes it more difficult to compare directly against other FDD systems that provide deterministic results on fault existence, diagnosis and value. Over time, with continued use of the system it will be possible to include user feedback on system accuracy in order to quantify and refine the statistical power of the system's inference.

7.2 Further Development

The key finding of this research is that the uncertainty of inference plays the central role in both the judicious algorithmic detection of equipment faults and the effective communication of the value of those faults to individuals who fix equipment malfunctions. The importance of uncertainty in communicating the results of fault detection and evaluation stems from the cultural expectation of building service agents; they know that buildings are complicated and highly uncertain, and so deterministic results appear far less credible and hence less useful than those with upfront acknowledgements of their uncertainty.

Furthermore, while a deeper investment in fault detection technology may yield much improved tools for finding and evaluating building energy inefficiencies, they will always remain just tools amongst many other tools; if our interest is in fixing buildings and improving energy efficiency, then our focus must be on how to design tools such that they are put into action and made to yield results. Through my research activities I have observed several themes of future

technical research that could potentially aid in creating and putting tools to work for us in fixing buildings:

1. Maturation of data-driven decision making within the building repair and maintenance culture

2. Aggregation and dissemination of building performance data

3. Smarter building systems that reduce the labor cost of implementing related control and fault detection software technologies

Those research themes are further developed below in a series of short, medium and long term research and development goals:

Short term:

- **Standardization of control nomenclature**: The lack of standard building system designs and nomenclature of building equipment and instrumentation is a significant hurdle to the widespread deployment of software-based building energy models and FDD services. While upfront costs and design constraints prevent standard instrumentation and equipment packages from being enforceable across the building stock, a standard protocol for naming of control points and equipment is more readily applicable simply as a control programming guide. If consistent nomenclature was applied to control points that performed the same functions across the entire building stock, then an FDD system could automatically draw and interpret data from any building control system without user programming of data composition. In that same vein, if a standard nomenclature was applied to pieces of equipment as well as their associated control points, then building energy models and FDD systems could perhaps automatically configure themselves to the customized design of the building. While we cannot and should not force buildings or their systems to adhere to a standard set of designs, it would be very useful to building energy model and FDD system software automation efforts if the nomenclature of describing buildings and their systems could be standardized. Standards bodies such as AHSRAE could develop such a nomenclature, as well as software tools to help designers and installers adhere to that nomenclature.

- **Manufacturer characterization and distribution of equipment models**: Cut-sheets on equipment and products from manufacturers that service the building industry rarely

include the detailed dynamic response of the equipment to varying input control signals. Cut sheets typically include the minimum yet most pertinent design data that engineers need in order to specify equipment for particular applications that are known at the design stage of a project. This approach to the design and engineering of building systems follows from the mentality of "up-front" cost and expediency of design; just enough equipment information is included to quickly finish a design to meet the lowest possible construction costs for the system design known at that time. Software-based FDD systems and building energy modeling tools represent a change in the mentality of building system design; these tools exist in order to manage the life-cycle costs of building, owning and managing a building. The power of these tools could be enhanced, and their deployment costs reduced, if building equipment and system manufacturers provided detailed models to describe the performance of their equipment under different dynamic and static conditions. 3rd party software-based FDD systems could in particular leverage detailed equipment models to enable more specific diagnoses and detection of equipment faults. The equipment response under various dynamic and static conditions would also be helpful to evaluate how a system should be re-commissioned when the originally intended building design is no longer applicable. To require manufacturers to provide mathematical models of every piece of their equipment is impractical and unreasonable, however, a useful starting point may be with unitary valves, pumps and fans that are used to build custom hydronic and air systems. Packaged units such as roof top units, air conditions, and chillers often come pre-packaged with on-board fault detection systems, but custom designed and installed systems have no such luxury. If fault detection engineers were equipped with manufacturer's models of valves, pumps and fans (system actuators), then they could build up "systemic" fault detection models for custom systems that are analogous to those found on "packaged" products.

- **Building Data Warehousing:** Much of this thesis has emphasized the use of data-driven decision making, and in doing so has exposed that much data in buildings is often created but never stored, sorted and converted into useful information. Since many building control platforms have extended logging capabilities, a short term innovation in the building stock would be to activate those logging features in as many buildings as possible, and tie that numerical data to categorical information like work-order and service tickets that describe equipment malfunctions and their resolution. An extensive database of equipment and system performance could be created simply by taking a campus of buildings, such as at MIT, and storing all of its historical building control data. While on its own the data is just bits on hard disk, numerous building control companies and independent software developers may use such a database to create more sophisticated FDD tools, or to verify equipment models used in building energy models. A clear lesson that I have learned from this research is that very little is precisely known and documented about how buildings operate once they are built and occupied; this is

especially pertinent to the future of commercial buildings where public and private entities have made pledges to create or approximate zero-net energy and carbon buildings. Without measuring, aggregating and reviewing real building system performance, we cannot confirm whether our buildings work the way in which we intended them to work, or learn how to improve their performance. Furthermore, by starting to warehouse and review data, we may also identify the critical set of building data that is truly needed to improve building performance. At this point, the volume and typology of data needed in order to improve building performance is relatively unknown; only by starting down this process of data collection and review can we identify the important subset of data and eventually reduce the cost of data collection by eliminating extraneous information.

Medium Term Innovations:

- **Smart sensors**: The reduced cost of ASICs could eventually allow a practical marriage of sensors with memory and processing that enables network communication of sensor data rather than analog data acquisition. One advantage of network based data acquisition is greater flexibility in sharing of sensor data amongst multiple data consumers (electronic or otherwise), but a more subtle advantage is the incorporation of sensor meta-data with measurement results. For example, sensor data today is typically comprised of voltage or amperage values that are converted at the data acquisition terminal into engineering values. Furthermore, that data acquisition terminal also handles thresholds for identifying significant changes in the measurement signal, and the nomenclature for identifying that sensor data to rest of the control system. This architecture of data collection requires state implementation within the hierarchy of data collection; a system user must program sensor characteristics at multiple levels within the system, taking up time and increasing the cost and complexity of data acquisition. If sensors were not just analog measurement instruments, but instead included a thin client for communicating sensor meta-data with measurement values, then perhaps we could eliminate state programming within the network and define it only at the network boundaries. From a practical perspective, we might include the following sensor meta-information within network packets of sensor measurement data:

    - Type or class of measurement (e.g. what does the sensor measure in the system?)
    - Engineering units of the measurements (e.g. Temperature, pressure, flow)
    - Date of sensor manufacturing or commissioning
    - Manufacturer model and serial number
    - Manufacturer's rated accuracy of the analog instrumentation, and standard by which that accuracy was determined

- Manufacturer's rated drift of the sensor, and standard by which that drift was determined
- Run-time of the sensor, or anther estimate of its health

Smart sensors represent the next step in an evolution of the building stock that would help facilitate the short term goals enumerated above while also laying the foundation for the longer-term, next generation of control systems. Furthermore, it may be possible to push certain elements of FDD out to the fringes of this new building network by loading smart sensors with some simple algorithms to detect their own health.

- **Smart actuators:** Complimentary to smart sensors, we could also use implement ASIC chips to create smart actuators with properties and advantages similar to smart sensors. For example, we might include the following meta-information in network communications between actuators, sensors and controllers:

  - Type or class of actuator (e.g. what does the actuator do)
  - Engineering units for the actuator (e.g. valve stem percent closing, motor rotational speed)
  - Date of manufacturing or commissioning
  - Manufacturer model and serial number
  - Manufacturer's model of actuator response (e.g hysteresis, valve stem position vs. voltage, etc.)
  - Run-time of the actuator, or perhaps an on-board estimate of its health.

  Like the related smart sensors, intelligent actuators can come with their own on-board FDD, while also providing a central controller with a model of the actuator's response characteristics.

- **Data mining on buildings:** If we are successful at creating a database of building performance data it may then be possible to use supervised and un-supervised learning techniques to identify useful and insightful building performance features. Specifically if the building database is labeled with categorical information such as periods of time when certain pieces of equipment were broken, or occupants were uncomfortable, or energy consumption was too high, then modern techniques for black-box modeling may identify features of numerical data that identify those categorical inferences. This thesis has discussed some of the shortcomings that befall first-principle modeling techniques, and we may overcome those limitations by subscribing to completely data-driven approaches to building modeling and inference.

Longer term innovations:

- **Perturbation based fault detection and control:** One logical extension to identifying a strong fault signal within building data is to create a closed loop system to perturb fault signal sources in order to forcibly test potential fault conditions. For example, the algorithms used in this thesis compute a probability of simultaneous heating and cooling; the existence of a stuck valve contributing to that fault signal could be tested simply by at perturbing the valve position to measure the response of the coil. If under that scenario perturbation of the valve positions yields no significant change in the fault signal, then there is substantial evidence that the fault signal is true and attributable to broken equipment.

- **Collaborative control and fault detection:** From a growing database of labeled data that represents the state of building equipment, properly normalized data and labels could be used to collaboratively detect undesirable equipment states across equipment in different buildings or even from different manufacturers. In many ways this longer term innovation seeks to replicate in a database the experiential knowledge that engineers and technicians use on a daily basis to detect, diagnose and fix faults in buildings. For example, if a database included the lifetime history of 10 different chillers within buildings in a similar climate zone, then perhaps that historical numerical and categorical data could be used to troubleshoot and diagnose another chiller within a different building in that climate zone.

- **Control Optimization:** A fresh perspective on building controls and fault detection is to look at them as two sides of the same coin. In today's control industry, software programming of control systems follows traditional control theory with error signals that drive PID loops or other similar systems. Fault detection, however, must incorporate user-definitions of faulty or otherwise undesirable system operation (for example simultaneous heating and cooling that wastes over $10,000 a year). In that way, fault detection poses system objective or loss functions that can measure the ability of a building system to meet supervisory-level numerical or categorical goals for operation. Combining the previous ideas on collaborative and perturbation based fault detection could yield a new approach to building control systems; fault detection could perhaps form the basis for finer tuning of conventional control systems to meet objective functions that are otherwise difficult to encode within the classical control architecture. For example a collaborative database could include the performance of PID-controlled heating and cooling systems for a wide variety PID-gain values in different buildings. A perturbation fault detection system could test the efficacy of those collaboratively known PID values against the PID gain values originally programmed into the control system

and determine which PID gain values were best for the building by measuring the building's response along energy, maintenance and comfort heuristics.

7.3 Final Remarks

There are several obstacles that stand in the way of applying the lessons learned in this thesis to buildings across the world and having an impact on their energy consumption. Most notable amongst those obstacles, I believe, is the need to educate building stakeholders on the potential value of using real time building data and analysis to help better manage building operations. As was discussed previously, building operations management is far more reactive than proactive, and justifiably so; occupants matter most in a building, and for a building with limited human resources, the only possible management practice is to react quickly to the critical aspects of an occupant's experience in the building. The best way to educate building stakeholders on the value of a data-driven approach to proactive building management is to lead by example; only by piloting, testing and widely publishing the results of research and development on these ideas will we manage to convince a meaningful number of stakeholders of the value in building data and its analysis.

Other obstacles, however, cannot be influenced by any further testing or development of these ideas. Most notable amongst that of type of obstacles, I believe, is the low cost of energy (at least here in the US). Despite concerns over rising energy costs, the actual financial burden of energy on most buildings pales in comparison to other operational expenses that they bear; for example commercial building owners are far more concerned about vacancy rates and tenant rent than they are about the energy cost of lights and HVAC being left on. Other building owners or managers care more about the experience that the occupant has in the building much more than the amount of energy consumed in order to yield that experience (take a restaurant or clothing store for example). Regardless of our motivations or aspirations for improving building energy efficiency, the simple reality is that energy is cheap in many places around the world, and we build buildings so that people will have a place to live out their life's goals and dreams as they see fit. Energy efficiency will not succeed if it gets in the way of how people want to live their lives, run their businesses, or enjoy their freedoms. If energy were more expensive, then energy

efficiency would play a larger role in the daily practice of using a building, however, for many people and buildings today this is just not the case. Given this tremendous obstacle it seems very clear that one of the best strategies for achieving energy efficiency is to make it the ancillary benefit of meeting other needs of a building that are more important.

Another large obstacle to using the ideas in this thesis across the building stock is the inherent business model for making buildings; so long as the people making a building are not affected by its lifetime costs then there is little motivation to include the infrastructure for improved energy efficiency. The most natural feedback loop for including lifecycle considerations within construction is for customers to demand that from their contractors. Once again though we face the complicating factor that customer might not care about the energy implications of construction, either because energy is cheap or they won't be paying the utility bills once the space is rented or purchased.

The path going forward to transition the technology in this thesis into practice includes a number of goals designed specifically to overcome these critical obstacles:

- Pilot testing and deployment of the system: the software technology must be demonstrated on a large scale to clearly show how data driven decision making in buildings is significantly more profitable and useful to building stakeholders than alternative management practices

- Multiple value propositions: the software system was designed from inception to include nominal maintenance and occupant health issues that were completely isolated from any energy implications. Ironically, we may achieve more widespread energy efficiency by not focusing on energy as the core objective of any fault detect technology, but as the ancillary benefit to some other application of FDD.

- Interaction at the ground level: people, not software, fix buildings; this means that in order to be effective, a software technology destined to fix buildings must be embraced by the fundamental group of people who fix buildings. By working closely with those people to demonstrate the value of data in buildings, we may form a strong grass roots effort to use tools like the one developed through this thesis to better manage and fix buildings around the world. In that same vein, close interaction at the ground level of fixing buildings will help focus software development on the most user-friendly and

intuitive interface for communicating and using the results of FDD; the ease of use of such software solutions are as much, if not more so important to the success of FDD than even the content of the analysis.

8 Bibliography

1. **US DoE EERE.** Buildings Energy Data Book. [Online] http://buildingsdatabook.eere.energy.gov (2009)..

2. **TIAX LLC.** *Energy Impact of Commercial Building Controls and Performance Diagnostics.* Cambridge : Tiax, LLC, 2005.

3. **Brambley, Michael.** *Advanced sensors and controls for building applications: Market assessment and potential research and developmen pathways.* s.l. : Pacific Northwest National Laboratory, 2005.

4. **Gershenfeld, Neil, Samouhos, Stephen and Nordman, Bruce.** Intelligent Infrastructure for Energy Efficiency. *Science.* 327, 2010, Vol. 5969.

5. **Frankel, M; Turner, C.** *Energy Performance for LEED New Construction.* s.l. : New Buildings Institute, 2008.

6. **Katipamula, Srinivas and Brambley, Michael.** Methods for Fault Detection, Diagnostics and Prognostics for Building Systems; A review, part I. *HVAC&R Research.* 2005, Vol. 11, 1.

7. —. Methods for Fault Detection, Diagnostics and Prognostics in Building Systems; A review, Part II. *HVAC&R Research.* 2005, Vol. 11, 2.

8. **Krarti, Moncef.** An Overview of Artificial Intelligence-based methods for Building Energy Systems. *J. Sol. Ener. Eng.* 2003, Vol. 125.

9. **Norford, L, et al.** Demonstration of Fault Detection and Diagnosis Methods for Air Handling Units (ASHRAE 1020-RP). *HVAC&R Research.* 2002, Vol. 8, 1.

10. **Kreider, J and Haberl, J.** Predicting Hourly Building Energy Use: The Great Predictor Shootout - Overview and Discussion of Results. *ASHRAE Transactions.* 1994, Vol. 100, 2.

11. List of EERE tools for Building Energy Simulation. [Online] http://apps1.eere.energy.gov/buildings/toolsdirectory/subjects.cfm/pagename=subjects/pagename menu=wholebuildinganalysis/pagenamesubmenu=energysimulation.

12. **al, Crawley et.** EnergyPlus: creating a new-generation building energy simulation program. *Eergy and Buildings.* 2001, Vol. 33.

13. **Crawley, Drury, Hand, Jon and Kummert, Michael.** Contrasting the capabilities of building energy performance simulation programs. *Proceedings of Building Simulation 2005.*

14. **Haberl, J and Themilseran, S.** Predicting Hourly Building Energy Use: The Great Energy Predictor Shootout II: Measuring Retrofit Savings. *ASHRAE Trans.* 1999, Vol. 102.

15. **Dodier, R and Henze, G.** Statistical analysis of neural networks as applied to building energy prediction. *Trans. ASME.* 2004, Vol. 124.

16. **Mackay, D.** Bayesian Non-linear Modeling for the Prediction Competition. *ASHRAE Trans.* 1994, Vol. 100, 2.

17. **Ohlsson, M, et al.** Predicting system loads with artificial neural networks. *ASHRAE Trans.* 1994, Vol. 100, 2.

18. **Feuston, B and Thurtell, J.** Generalized nonlinear regression with ensemble of neural nets. *ASHRAE Trans.* 1994, Vol. 100, 2.

19. **Karatsou, S, Santamouris, M and Geros, V.** Modeling and predicting building energy use with artificial neural networks. *Energy and Buildings.* 2006, Vol. 38.

20. **Anders, U and Korn, O.** Model selection in neural networks. *Neural Networks.* 1999, Vol. 12.

21. **Rivals, I and Personnaz, L.** Neural network construction and selection in nonlinear modeling. *IEEE Trans. Neural Networks.* 2003, Vol. 14, 4.

22. **Seem, J.** Using intelligent data analysis to detect abnormal energy consumption in buildings. *Energy and Buildings.* 2007, Vol. 39.

23. —. Pattern recognition algorithm for determining days of the week with similar energy consumption profiles. *Energy and Buildings.* 2005, Vol. 37.

24. **Lai, F, Magoules, F and Lherminier, F.** Vapnik's learning theory applied to energy consumption forecasts in residential buildings. *Int. J. Computer Mathematics.* 2008, Vol. 85, 10.

25. **Dong, B, Cao, C and Lee, S.** Applying support vector machines to predict building energy consumption in tropical region. *Energy and Buildings.* 2005, Vol. 37.

26. **Liang, J and Du, R.** Model based fault detection and diagnosis of HVAC systems using support vector machines. *Int. J. Refrig.* 2007, Vol. 30, 6.

27. **Rifkin, Ryan.** Everything old is new again: a fresh look at historical perspectives in machine learning. *PhD Thesis, Sloan School of Management, MIT.* 2002.

28. **Samanta, B.** Gear fault detection using artificial neural networks and support vector machines with genetic algorithms. *Mech. Syst. Sig. Proc.* 2004, Vol. 18.

29. **Jack, L and Nandi, A.** Fault detection using support vector machines and neaural networks augmented by genetic algorithms. *Mech. Syst. Sig. Proc.* 2002, Vol. 16, 2-3.

30. **Huang, T, Kecman, V and Kopriva, I.** *Kernel based methods for mining huge data sets.* New York : Springer, 2006.

31. **Katipamula, S, Pratt, R and Braun, J.** Building System Diagnostics and Preventative Maintenance. [book auth.] J. F. Krieder ed. *Handbook of HVAC.* Boca Raton : CRC Press, 2001.

32. **Dodier, R and Krieder, J.** Whole Building Energy Diagnostics. *ASHRAE Trans.* 1999, Vol. 105, 1.

33. **Krarti, M, et al.** Estimation of energy savings for building retrofits using neural networks. *ASME J. Solar Energy Engineering.* 1998, Vol. 120, 3.

34. **Breekweg, M, Gruber, P and Ahmed, O.** Development of a generalized neural network model to predict faults in building energy performance. *ASHRAE Trans.* 2000, Vol. 106, 1.

35. **Norford, L and Little, R.** Fault detecion and load monitoring in ventilation systems. *Int. J. Mol. Sci.* 1993, Vol. 99, 1.

36. **Ahn, C, Mitchell, J and McIntosh, I.** Model based fault detection and diagnosis for cooling towers. *ASHRAE Trans.* 2001, Vol. 107, 2.

37. **Rossi, T and Braun, J.** A statistical rule based fault detection and diagnostic method for vapor compression air conditioners. *Int. J. HVAC&R.* 1997, Vol. 3, 2.

38. **Stylianou, M.** Application of classification functions to chiller fault detection and diagnosis. *ASHRAE Trans.* 1997, Vol. 103, 1.

39. **Lee, W, House, J and Shin, D.** Fault detection of an air-handling unit using residual and recursive parameter identification method. *ASHRAE Trans.* 1997, Vol. 103, 1.

40. **Ngo, D and Dexter, A.** A robust model based approach to diagnosing faults in air-handling units. *ASHRAE Trans.* 1999, Vol. 105, 1.

41. **Braun, J.** Automated fault detection and diagnostics for vapor compression cooling equipment. *ASME Trans.* 2003, Vol. 125.

42. **Peng, X, Haves, P and Curtil, D.** A library of HVAC component models for use in automated diagnostics. *Proceedings of SimBuild 2006.* 2006.

43. **Wagner, J and Shoureshi, R.** Failure detection and diagnostics for thermofluid systems. *Jounal of Dynamic Systems, Measurement and Control.* 1992, Vol. 114, 4.

44. **Haves, P, Salsbury, T and Wright, J.** Condition monitoring in HVAC subsystems using first principles. *ASHRAE Trans.* 1996, Vol. 102, 1.

45. **Salsbury, T and Diamond, R.** Fault detection in HVAC systems using model based feedforward control. *Energy and Buildings.* 2001, Vol. 33.

46. **Kaler, G.** Embedded expert system development for monitoring packaged HVAC equipment. *ASHRAE Transactions.* 1990, Vol. 96 , 2.

47. **Kaldorf, S and Gruber, P.** Practical experiences from developing and impementing an expert system diagnostic tool. *ASHRAE Transactions.* 2002, Vol. 108, 1.

48. **Brambley, M, et al.** Automated diagnostics for outdoor air ventilation and economizers. *ASHRAE Journal.* 1998, Vol. 40, 10.

49. **Katipamula, S, et al.** Automated fault detection and diagnostics for outdoor air ventilation systems and economizers: Methodology and rsults from field testing. *ASHRAE Transactions.* 1999, Vol. 105, 1.

50. **House, J, Vaezi-Nejad, H and Whitcomb, J.** An expert rule st for fault detection in air hadling units. *ASHRAE Transactions.* 2001, Vol. 107, 1.

51. **Schein, J, et al.** A rule-based fault detection method for air handling units. *Energy and Buildings.* 2006, Vol. 38.

52. **Katipamula, S, Brambley, M and Luskay, L.** Automated proactive techniques for continuous commissioning o air handling units. *Journal of Solar Energy Engineering.* 2003, Vol. 15.

53. **Venkatasubramanian, S, et al.** A review of process fault detection and diagnosis part III: process history based methods. *Computers and chemical engineering.* 2003, Vol. 27.

54. **Najafi, M, et al.** Application of machine learning in fault diagnostics of mechanical systems. *Proceedings of the world congress on engineering and computer science.* 2008.

55. **Choi, K, et al.** Fault diagnosis in HVAC chillers. *IEEE Instrumentation and Measurement Magazine.* 2005, Vol. August.

56. **Bendapudi, S and Braun, J.** *Development and validation of a mechanistic, dynamic model for a vapor compression centrifugal liquid chiller.* West Lafayette, IN : Ray W. Herrick Labs, Purdue University, 2002.

57. **RS Means Inc.** *Light Commercial Construction Data.* s.l. : Reed Construction Data, 2007.

58. **Cimetrics Inc.** *Standards for fault detection, diagnostics, and optimization in building systems.* s.l. : Southern California Gas Company, 2007.

59. **al, Najafi et.** *ASHRAE RP 980.* s.l. : ASHRAE, 1995.

60. **Bell, Arthur.** *HVAC Equation,s Data, and Rules of Thumb.* New York : McGraw Hill, 2008.

61. *Combustion monitoring and control system. 4059385* United States.

62. *Combustion efficiency analyzer, acoustic. 4959638* United States.

63. **Kass, R and Raftery, A.** Bayes Factor. *Journal of the American Statistical Association.* 90, 1995, 430.

64. **Coleman, H and Steele, W.** Engineering application of experimental uncertainty analysis. *AIAA* . 33, 1995, 10.

65. **ISO.** *Guide to the expression of uncertainty in measurement.* Geneva, Switzerland : ISO, 1993.

66. **Coleman, H and Steele, W.** *Experimental and uncertainty analysis for engineers.* New York : Wiley, 1989.

67. **Meyer, A and Booker, J.** *Eliciting and analyzing expert judgement: a practical guide.* Philadelphia : ASA-SIAM, 2001.

68. **University, Texas A&M.** *Federal Facilities Guide to Commissioning and Retro-commissioning.*

69. **Wetteroth, Debbra.** *OSI Reference model for telecommunications.* New York City : McGraw-Hill, 2002.

70. **Hahn, Otto.** *Statistical distributions in engineering.* New York City : Wiley, 1997.

71. **Heinemeier, K, Richardson, R and Kulathumani, K.** *User and market factors that influence diagnostic tool development.* s.l. : Honeywell Inc. on behalf of LBNL, 2003.

72. **Katipamula, S and Brambley, M.** Methods for fault detection, diagnostics and prognostics for building systems - a review part I. *HVAC&R Research.* 2005, Vol. 11, 1.